



Induction Motors – Operation from 50/60 Hz Supply

1. INTRODUCTION

This chapter is concerned with how the induction motor behaves when connected to a supply of constant voltage and frequency. Despite the onward march of the inverter-fed motor, this remains the most widely used and important mode of operation, the motor running directly connected to a utility supply.

The key operating characteristics are considered, and we look at how these can be modified to meet the needs of some applications through detailed design. The limits of operation are investigated for the induction machine operating as both a motor and a generator. Methods of speed control which are not dependent on changing the frequency of the stator supply are also explored. Finally, while the majority of industrial applications utilize the 3-phase induction motor, the role played by single-phase motors is acknowledged with a review of the types and characteristics of this variant.

2. METHODS OF STARTING CAGE MOTORS

2.1 Direct starting – problems

Our everyday domestic experience is likely to lead us to believe that there is nothing more to starting a motor than closing a switch, and indeed for most low-power machines (up to a few kW) – of whatever type – that is indeed the case. By simply connecting the motor to the supply we set in train a sequence of events which sees the motor draw power from the supply while it accelerates to its target speed. When it has absorbed and converted sufficient energy from electrical to kinetic form, the speed stabilizes and the power drawn falls to a low level until the motor is required to do useful mechanical work. In these low-power applications acceleration to full speed may take less than a second, and we are seldom aware of the fact that the current drawn during the acceleration phase is often higher than the continuous rated current.

For motors over a few kW, however, it is necessary to assess the effect on the supply system before deciding whether or not the motor can be started simply by switching directly on to the supply. If supply systems were ideal (i.e. the supply voltage remained unaffected regardless of how much current was drawn) there

would be no problem starting any induction motor, no matter how large. The problem is that the heavy current drawn while the motor is running up to speed may cause a large drop in the system supply voltage, annoying other customers on the same supply and perhaps taking it outside statutory limits.

It is worthwhile reminding ourselves about the influence of supply impedance at this point, as this is at the root of the matter, so we begin by noting that any supply system, no matter how complicated, can be modeled by means of the delightfully simple Thévenin equivalent circuit shown in Figure 6.1. (We are assuming balanced 3-phase operation, so a 1-phase equivalent circuit will suffice.)




Figure 6.1 Equivalent circuit of supply system.

The supply is represented by an ideal voltage source (V_s) in series with the supply impedance Z_s . When no load is connected to the supply, and the current is zero, the terminal voltage is V_s ; but as soon as a load is connected the load current (I) flowing through the source impedance results in a volt-drop, and the output voltage falls from V_s to V , where

$$V = V_s - IZ_s \quad (6.1)$$

For most industrial supplies the source impedance is predominantly inductive, so that Z_s is simply an inductive reactance, X_s . Typical phasor diagrams relating to a supply with a purely inductive reactance are shown in Figure 6.2: in (a) the load is also taken to be purely reactive, while the load current in (b) has the same magnitude as in (a) but the load is resistive. The output (terminal) voltage in each case is represented by the phasor labeled V .

For the inductive load (a) the current lags the terminal voltage by 90° , while for the resistive load (b) the current is in phase with the terminal voltage. In both cases the volt-drop across the supply reactance (IX_s) leads the current by 90° .

The first point to note is that, for a given magnitude of load current, the volt-drop is in phase with V_s when the load is inductive, whereas with a resistive load the volt-drop is almost at 90° to V_s . This results in a much greater fall in the magnitude of the output voltage when the load is inductive than when it is resistive. The second – obvious – point is that the larger the current, the more the drop in voltage.

Unfortunately, when we try to start a large cage induction motor we face a double-whammy because not only is the starting current typically five or six times




Figure 6.2 Phasor diagrams showing the effect of supply-system impedance on the output voltage with (a) inductive load and (b) resistive load.

rated current, but it is also at a low power-factor, i.e. the motor looks predominantly inductive when the slip is high. (In contrast, when the machine is up to speed and fully loaded, its current is perhaps only one-fifth of its starting current and it presents a predominantly resistive appearance as seen by the supply. Under these conditions the supply voltage is hardly any different from no-load.)

Since the drop in voltage is attributable to the supply impedance, if we want to be able to draw a large starting current without upsetting other consumers it would clearly be best for the supply impedance to be as low as possible, and preferably zero. But from the supply authority viewpoint a very low supply impedance brings the problem of how to cope in the event of an accidental short-circuit across the terminals. The short-circuit current is inversely proportional to the supply impedance, and tends to infinity as Z_s approaches zero. The cost of providing the switchgear to clear such a large fault current would be prohibitive, so a compromise always has to be reached, with values of supply impedances being set by the supply authority to suit the anticipated demands.

Systems with low internal impedance are known as ‘stiff’ supplies, because the voltage is almost constant regardless of the current drawn. (An alternative way of specifying the nature of the supply is to consider the fault current that would flow if the terminals were short-circuited: a system with a low impedance would have a high fault current or ‘fault level’.) Starting on a stiff supply requires no special arrangements and the three motor leads are simply switched directly onto the mains. This is known as ‘direct-on-line’ (DOL) or ‘direct-to-line’ (DTL) starting. The switching will usually be done by means of a contactor, incorporating fuses and other overload/thermal protection devices, and operated manually by local or remote pushbuttons, or interfaced to permit operation from a programmable controller or computer.

In contrast, if the supply impedance is high (i.e. a low fault level) an appreciable volt-drop will occur every time the motor is started, causing lights to dim and interfering with other apparatus on the same supply. With this 'weak' supply, some form of starter is called for to limit the current at starting and during the run-up phase, thereby reducing the magnitude of the volt-drop imposed on the supply system. As the motor picks up speed, the current falls, so the starter is removed as the motor approaches full speed. Naturally enough the price to be paid for the reduction in current is a lower starting torque, and a longer run-up time.

Whether or not a starter is required depends on the size of the motor in relation to the capacity or fault level of the supply, the prevailing regulations imposed by the supply authority, and the nature of the load.

The references above to 'low' and 'high' supply impedances must therefore be interpreted in relation to the impedance of the motor when it is stationary. A large (and therefore low impedance) motor could well be started quite happily direct-on-line in a major industrial plant, where the supply is 'stiff', i.e. the supply impedance is very much less than the motor impedance. But the same motor would need a starter when used in a rural setting remote from the main power system, and fed by a relatively high impedance or 'weak' supply. Needless to say, the stricter the rules governing permissible volt-drop, the more likely it is that a starter will be needed.

Motors which start without significant load torque or inertia can accelerate very quickly, so the high starting current is only drawn for a short period. A 10 kW motor would be up to speed in a second or so, and the volt-drop may therefore be judged as acceptable. Clutches are sometimes fitted to permit 'off-load' starting, the load being applied after the motor has reached full speed. Conversely, if the load torque and/or inertia are/is high, the run-up may take many seconds, in which case a starter may prove essential. No strict rules can be laid down, but obviously the bigger the motor, the more likely it is to require a starter.

2.2 Star/delta (wye/mesh) starter

This is the simplest and most widely used method of starting. It provides for the windings of the motor to be connected in star (wye) to begin with, thereby reducing the voltage applied to each phase to 58% (i.e. $1/\sqrt{3}$) of its direct-on-line value. Then, when the motor speed approaches its running value, the windings are switched to delta (mesh) connection. The main advantage of the method is its simplicity, while its main drawbacks are that the starting torque is reduced (see below), and the sudden transition from star to delta gives rise to a second shock – albeit of lesser severity – to the supply system and to the load. For star/delta switching to be possible both ends of each phase of the motor windings must be

brought out to the terminal box. This requirement is met in the majority of motors, except small ones which are usually permanently connected in delta.

With a star/delta starter the current drawn from the supply is approximately one-third of that drawn in a direct-on-line start, which is very welcome, but at the same time the starting torque is also reduced to one-third of its direct-on-line value. Naturally we need to ensure that the reduced torque will be sufficient to accelerate the load, and bring it up to a speed at which it can be switched to delta without an excessive jump in the current.

Various methods are used to detect when to switch from star to delta. Historically, in manual starters, the changeover is determined by the operator watching the ammeter until the current has dropped to a low level, or listening to the sound of the motor until the speed becomes steady. Automatic versions are similar in that they detect either falling current or speed rising to a threshold level, but for all but large motors they operate after a preset time.

2.3 Autotransformer starter

A 3-phase autotransformer is usually used where star/delta starting provides insufficient starting torque. Each phase of an autotransformer consists of a single winding on a laminated core. The incoming supply is connected across the ends of the coils, and one or more tapping points (or a sliding contact) provide a reduced voltage output, as shown in [Figure 6.3](#).

The motor is first connected to the reduced voltage output, and when the current has fallen to the running value, the motor leads are switched over to the full voltage.

If the reduced voltage is chosen so that a fraction α of the line voltage is used to start the motor, the starting torque is reduced to approximately α^2 times its direct-on-line value, and the current drawn from the mains is also reduced to α^2 times its direct value. As with the star/delta starter, the torque per ampere of supply current is the same as for a direct start.




Figure 6.3 Autotransformer starter for cage induction motor.

2.4 Resistance or reactance starter

By inserting three resistors or inductors of appropriate value in series with the motor, the starting current can be reduced by any desired extent, but only at the expense of a disproportionate reduction in starting torque.

For example, if the current is reduced to half its direct-on-line value, the motor voltage will be halved, so the torque (which is proportional to the square of the voltage – see later) will be reduced to only 25% of its direct-on-line value. This approach is thus less attractive in terms of torque per ampere of supply current than the star/delta method. One attractive feature, however, is that as the motor speed increases and its effective impedance rises, the volt-drop across the extra impedance reduces, so the motor voltage rises progressively with the speed, thereby giving more torque. When the motor is up to speed, the added impedance is shorted out by means of a contactor.

2.5 Solid-state soft starting

This method is now the most widely used. It provides a smooth build-up of current and torque, the maximum current and acceleration time are easily adjusted, and it is particularly valuable where the load must not be subjected to sudden jerks. The only real drawback over conventional starters is that the mains currents during run-up are not sinusoidal, which can lead to interference with other equipment on the same supply.

The most widely used arrangement comprises three pairs of back-to-back thyristors (or triacs) connected in series with the three supply lines, as shown in [Figure 6.4\(a\)](#).

Each thyristor is fired once per half-cycle, the firing being synchronized with the utility supply and the firing angle being variable so that each pair conducts for a varying proportion of a cycle. Typical current waveforms are shown in [Figure 6.4\(b\)](#): they are clearly not sinusoidal but the motor will tolerate them quite happily.

A wide variety of control philosophies can be found, with the degree of complexity and sophistication being reflected in the price. The cheapest open-loop systems simply alter the firing angle linearly with time, so that the voltage applied to the motor increases as it accelerates. The ‘ramp-time’ can be set by trial and error to give an acceptable start, i.e. one in which the maximum allowable current from the supply is not exceeded at any stage. This approach is reasonably satisfactory when the load remains the same, but requires resetting each time the load changes. Loads with high static friction are a problem because nothing happens for the first part of the ramp, during which time the motor torque is insufficient to move the load. When the load finally moves, its acceleration is often too rapid. The more advanced open-loop versions allow the level of current at the start of the ramp to be chosen, and this is helpful with ‘sticky’ loads.




Figure 6.4 (a) Thyristor soft-starter, (b) typical motor current waveforms.

More sophisticated systems – usually with on-board digital controllers – provide for tighter control over the acceleration profile by incorporating an inner current-control loop. After an initial ramping up to the start level (over the first few cycles), the current is held constant at the desired level throughout the accelerating period, the firing angle of the thyristors being continually adjusted to compensate for the changing effective impedance of the motor. By keeping the current at the maximum value which the supply can tolerate, the run-up time is minimized. As with the open-loop systems the velocity-time profile is not necessarily ideal, since with constant current the motor torque exhibits a very sharp rise as the pull-out slip is reached, resulting in a sudden surge in speed. Some systems also include a motor model which estimates speed and allows the controller to follow a ramp or other speed-time profile.

Prospective users need to be wary of some of the promotional literature: claims are sometimes made that massive reductions in starting current can be achieved

without corresponding reductions in starting torque. This is nonsense: the current can certainly be limited, but as far as torque per line amp is concerned soft-start systems are no better than series reactor systems, and not as good as the auto-transformer and star/delta methods. Caution should also be exercised in relation to systems that use only one or two triacs: these are fairly common in smaller sizes (<50 kW). Although they do limit the current in one or two phases as compared with direct-on-line starting, the unbalanced currents distort the air-gap flux and this gives rise to uneven (pulsating) torque.

2.6 Starting using a variable-frequency inverter

Operation of induction motors from variable-frequency inverters is discussed in Chapters 7 and 8, but it is appropriate to mention here that one of the advantages of inverter-fed operation is that starting is not a problem because it is usually possible to obtain rated torque from standstill up to rated speed without drawing an excessive current from the supply. None of the other starting methods we have looked at has this ability, so in some applications it may be that the comparatively high cost of the inverter is justified solely on the grounds of its starting and run-up potential.

3. RUN-UP AND STABLE OPERATING REGIONS

In addition to having sufficient torque to start the load it is obviously necessary for the motor to bring the load up to full speed. To predict how the speed will rise after switching on we need the torque–speed curves of the motor and the load, and the total inertia.

By way of example, we can look at the case of a motor with two different loads (Figure 6.5). The solid line is the torque–speed curve of the motor, while the broken lines represent two different load characteristics. Load (A) is typical of a simple hoist, which applies constant torque to the motor at all speeds, while load




Figure 6.5 Typical torque–speed curve showing two different loads which have the same steady running speed (N).

(B) might represent a fan. For the sake of simplicity, we will assume that the load inertias (as seen at the motor shaft) are the same.

The speed–time curves for run-up are shown in [Figure 6.6](#). Note that the gradient of the speed–time curve (i.e. the acceleration) is obtained by dividing the accelerating torque T_{acc} (which is the difference between the torque developed by the motor and the torque required to run the load at that speed) by the total inertia.




Figure 6.6 Speed–time curves during run-up, for motor and loads shown in [Figure 6.5](#).

In this example, both loads ultimately reach the same steady speed, N (i.e. the speed at which motor torque equals load torque), but B reaches full speed much more quickly because the accelerating torque is higher during most of the run-up. Load A picks up speed slowly at first, but then accelerates hard (often with a characteristic ‘whoosh’ produced by the ventilating fan) as it passes through the peak torque speed and approaches equilibrium conditions.

It should be clear that the higher the total inertia, the slower the acceleration, and vice versa. The total inertia means the inertia as seen at the motor shaft, so if gearboxes or belts are employed the inertia must be ‘referred’, as discussed in [Chapter 11](#).

An important qualification ought to be mentioned in the context of the motor torque–speed curves shown by the solid line in [Figure 6.5](#). This is that curves like this represent the torque developed by the motor when it has settled down at the speed in question, i.e. they are the true steady–state curves. In reality, a motor will generally only be in a steady–state condition when it settles at its normal running speed, so for most of the speed range the motor will be accelerating.

In particular, when the motor is first switched on, there will be a transient period as the three currents gradually move towards a balanced 3–phase pattern. During this period the torque can fluctuate wildly at the supply frequency, and, with brief negative excursions, typically as shown in [Figure 6.7](#) which relates to a small unloaded motor. During the transient period the average torque may be very low (as in [Figure 6.7](#)) in which case acceleration only begins in earnest after the first few cycles. In this particular example the transient persists long enough to cause an overshoot of the steady–state speed and an oscillation before settling.




Figure 6.7 Run-up of unloaded motor, showing torque transients persisting for the first few cycles of the supply.

Fortunately, the average torque during run-up can be fairly reliably obtained from the steady-state curves (usually available from the manufacturer), particularly if the inertia is high and the motor takes many cycles to reach full speed, in which case we would consider the torque–speed curve as being ‘quasi-steady state’.

3.1 Harmonic effects – skewing

A further cautionary note in connection with the torque–speed curves shown in this and most other books relates to the effects of harmonic air-gap fields. In Chapter 5 it was explained that despite the limitations imposed by slotting, the stator winding m.m.f. is remarkably close to the ideal of a pure sinusoid. Unfortunately, because it is not a perfect sinusoid, Fourier analysis reveals that in addition to the predominant fundamental component, there are always additional unwanted ‘space harmonic’ fields. These harmonic fields have synchronous speeds that are inversely proportional to their order. For example, a 4-pole, 50 Hz motor will have a main field rotating at 1500 rev/min, but in addition there may be a 5th harmonic (20-pole) field rotating in the reverse direction at 300 rev/min, a 7th harmonic (28-pole) field rotating forwards at 214 rev/min, etc. These space harmonics are minimized by stator winding design, but not all can be eliminated.

If the rotor has a very large number of bars it will react to the harmonic field in much the same way as to the fundamental, producing additional induction-motor torques centered on the synchronous speed of the harmonic, and leading to unwanted dips in the torque speed, typically as shown in [Figure 6.8](#).

Users should not be too alarmed as in most cases the motor will ride through the harmonic during acceleration, but in extreme cases a motor might, for example,




Figure 6.8 Torque–speed curve showing the effect of space harmonics, and illustrating the possibility of a motor ‘crawling’ on the 7th harmonic.

stabilize on the 7th harmonic, and ‘crawl’ at about 214 rev/min, rather than running up to 4-pole speed (1500 rev/min at 50 Hz), as shown by the dot in [Figure 6.8](#).

To minimize the undesirable effects of space harmonics the rotor bars in the majority of induction motors are not parallel to the axis of rotation, but instead they are skewed along the rotor length, (typically by around one or two slot-pitches, see later [Figure 8.4](#)). This has very little effect as far as the fundamental field is concerned, but can greatly reduce the response of the rotor to harmonic fields.

Because the overall influence of the harmonics on the steady-state curve is barely noticeable, and their presence might worry users, they are rarely shown, the accepted custom being that ‘the’ torque–speed curve represents the behavior due to the fundamental component only.

3.2 High inertia loads – overheating

Apart from accelerating slowly, high inertia loads pose a particular problem of rotor heating which can easily be overlooked by the unwary user. Every time an induction motor is started from rest and brought up to speed, the total energy dissipated as heat in the motor windings is equal to the stored kinetic energy of the motor plus load. Hence with high inertia loads, very large amounts of energy are released as heat in the windings during run-up, even if the load torque is negligible when the motor is up to speed. With totally enclosed motors the heat ultimately has to find its way to the finned outer casing of the motor, which is cooled by air from the shaft-mounted external fan. Cooling of the rotor is therefore usually much worse than that of the stator, and the rotor is thus most likely to overheat during high inertia run-ups.

No hard and fast rules can be laid down, but manufacturers usually work to standards which specify how many starts per hour can be tolerated. Actually, this information is useless unless coupled with reference to the total inertia, since doubling the inertia makes the problem twice as bad. However, it is usually assumed that the total inertia is not likely to be more than twice the motor inertia, and this is certainly the case for most loads. If in doubt, the user should consult the

manufacturer, who may recommend a larger motor than might seem necessary if it was simply a matter of meeting the full-load power requirements.

3.3 Steady-state rotor losses and efficiency

The discussion above is a special case which highlights one of the less attractive features of induction machines. This is that it is never possible for all the power crossing the air-gap from the stator to be converted to mechanical output, because some is always lost as heat in the rotor circuit resistance. In fact, it turns out that at slip s the total power (P_r) crossing the air-gap always divides, so a fraction sP_r is lost as heat, while the remainder $(1 - s)P_r$ is converted to useful mechanical output (see also Appendix 2).

Hence when the motor is operating in the steady-state the energy-conversion efficiency of the rotor is given by

$$\eta_r = \frac{\text{Mechanical output power}}{\text{Power into rotor}} = (1 - s) \quad (6.2)$$

This result is very important, and shows us immediately why operating at small values of slip is desirable. With a slip of 5% (or 0.05), for example, 95% of the air-gap power is put to good use. But if the motor was run at half the synchronous speed ($s = 0.5$), 50% of the air-gap power would be wasted as heat in the rotor.

We can also see that the overall efficiency of the motor must always be significantly less than $(1 - s)$, because in addition to the rotor copper losses there are stator copper losses, iron losses and windage and friction losses. This fact is sometimes forgotten, leading to conflicting claims such as ‘full-load slip = 5%, overall efficiency = 96%’, which is clearly impossible.

3.4 Steady-state stability – pull-out torque and stalling

We can check stability by asking what happens if the load torque suddenly changes for some reason. The load shown by the lower broken line in Figure 6.9 is stable at speed X : for example, if the load torque increased from T_a to T_b , the load torque




Figure 6.9 Torque-speed curve illustrating stable operating region (OXYZ).

would be greater than the motor torque, so the motor would decelerate. As the speed dropped, the motor torque would rise, until a new equilibrium was reached, at the slightly lower speed (Y). The converse would happen if the load torque reduced, leading to a higher stable running speed.

But what happens if the load torque is increased more and more? We can see that as the load torque increases, beginning at point X, we eventually reach point Z, at which the motor develops its maximum torque. Quite apart from the fact that the motor is now well into its overload region, and will be in danger of overheating, it has also reached the limit of stable operation. If the load torque is further increased, the speed falls (because the load torque is more than the motor torque), and as it does so the shortfall between motor torque and load torque becomes greater and greater. The speed therefore falls faster and faster, and the motor is said to be ‘stalling’. With loads such as machine tools (a drilling machine, for example), as soon as the maximum or ‘pull-out’ torque is exceeded, the motor rapidly comes to a halt, making an angry humming sound. With a hoist, however, the excess load would cause the rotor to be accelerated in the reverse direction, unless it was prevented from doing so by a mechanical brake.

4. TORQUE–SPEED CURVES – INFLUENCE OF ROTOR PARAMETERS

We saw earlier that the rotor resistance and reactance influenced the shape of the torque–speed curve. Both of these parameters can be varied by the designer, and we will explore the pros and cons of the various alternatives. To limit the mathematics the discussion will be mainly qualitative, but it is worth mentioning that the whole matter can be dealt with rigorously using the equivalent circuit approach, as discussed in Appendix 2.

We will deal with the cage rotor first because it is the most important, but the wound rotor allows a wider variation of resistance to be obtained, so it is discussed later.

4.1 Cage rotor

For small values of slip, i.e. in the normal running region, the lower we make the rotor resistance the steeper the slope of the torque–speed curve becomes, as shown in Figure 6.10. We can see that at the rated torque (shown by the horizontal dotted line in Figure 6.10) the full-load slip of the low-resistance cage is much lower than that of the high-resistance cage. But we saw earlier that the rotor efficiency is equal to $(1 - s)$, where s is the slip. So we conclude that the low resistance rotor not only gives better speed holding, but is also much more efficient. There is of course a limit to how low we can make the resistance: copper allows us to achieve a lower resistance than aluminum, but we can’t do any better than fill the slots with solid copper bars.




Figure 6.10 Influence of rotor resistance on the torque–speed curve of a cage motor. The full-load running speeds are indicated by the vertical dotted lines.

As we might expect there are drawbacks with a low resistance rotor. The direct-on-line starting torque is reduced (see Figure 6.10), and worse still the starting current is increased. The lower starting torque may prove insufficient to accelerate the load, while increased starting current may lead to unacceptable volt-drops in the supply.

Altering the rotor resistance has little or no effect on the value of the peak (pull-out) torque, but the slip at which the peak torque occurs is directly proportional to the rotor resistance. By opting for a high enough resistance (by making the cage from bronze, brass or other relatively high-resistivity material) we can if we wish arrange for the peak torque to occur at or close to starting, as shown in Figure 6.10. The snag in doing this is that the full-load efficiency is inevitably low because the full-load slip will be high (see Figure 6.10).

There are some applications for which high-resistance motors were traditionally well suited, an example being for metal punching presses, where the motor accelerates a flywheel which is used to store energy. In order to release a significant amount of energy, the flywheel slows down appreciably during impact, and the motor then has to accelerate it back up to full speed. The motor needs a high torque over a comparatively wide speed range, and does most of its work during acceleration. Once up to speed the motor is effectively running light, so its low efficiency is of little consequence. (We should note that this type of application is now often met by drives, but because the induction motor is so robust and long-lived, significant numbers of ‘heritage’ installations will persist, particularly in the developing world.)

High-resistance motors are sometimes used for speed control of fan-type loads, and this is taken up again later when we explore speed control.

To sum up, a high rotor resistance is desirable when starting and at low speeds, while a low resistance is preferred under normal running conditions. To get the best of both worlds, we need to be able to alter the resistance from a high value at starting to a lower value at full speed. Obviously we can’t change the actual resistance of the

cage once it has been manufactured, but it is possible to achieve the desired effect with either a ‘double cage’ or a ‘deep bar’ rotor.

4.2 Double cage and deep bar rotors

Double cage rotors have an outer cage made of relatively high-resistivity material such as bronze, and an inner cage of low resistivity, usually copper, as shown on the left in [Figure 6.11](#).

The inner cage of low-resistance copper is sunk deep into the rotor, so that it is almost completely surrounded by iron. This causes the inner bars to have a much higher leakage inductance than if they were near the rotor surface, so that under starting conditions (when the induced rotor frequency is high) their inductive reactance is very high and little current flows in them. In contrast, the bars of the outer cage (of higher resistance bronze) are placed so that their leakage fluxes face a much higher reluctance path, leading to a low leakage inductance. Hence under starting conditions, rotor current is concentrated in the outer cage, which, because of its high resistance, produces a high starting torque.

At the normal running speed the roles are reversed. The rotor frequency is low, so both cages have low reactance and most of the current therefore flows in the low-resistance inner cage. The torque–speed curve is therefore steep, and the efficiency is high.

Considerable variation in detailed design is possible in order to shape the torque–speed curve to particular requirements. In comparison with a single-cage rotor, the double cage gives much higher starting torque, substantially less starting current, and marginally worse running performance.

The deep bar rotor has a single cage, usually of copper, formed in slots which are deeper and narrower than in a conventional single-cage design. Construction is simpler and therefore cheaper than in a double-cage rotor, as shown on the right in [Figure 6.11](#).

The deep bar approach ingeniously exploits the fact that the effective resistance of a conductor is higher under a.c. conditions than under d.c. conditions. With a typical copper bar of the size used in an induction motor rotor, the difference in effective resistance between d.c. and say 50 or 60 Hz (the so-called ‘skin effect’) would be negligible if the conductor were entirely surrounded by air. But when it is almost completely surrounded by iron, as in the rotor slots, its effective resistance at mains frequency may be two or three times its d.c. value.




Figure 6.11 Double cage (left) and deep bar (right) rotors.

Under starting conditions, when the rotor frequency is equal to the supply frequency, the skin effect is very pronounced, and the rotor current is concentrated towards the top of the slots. The effective resistance is therefore increased, resulting in a high starting torque from a low starting current. When the speed rises and the rotor frequency falls, the effective resistance reduces towards its d.c. value, and the current distributes itself more uniformly across the cross-section of the bars. The normal running performance thus approaches that of a low-resistance single-cage rotor, giving a high efficiency and stiff torque–speed curve. The pull-out torque is, however, somewhat lower than for an equivalent single-cage motor because of the rather higher leakage reactance.

Most small and medium motors are designed to exploit the deep bar effect to some extent, reflecting the view that for most applications the slightly inferior running performance is more than outweighed by the much better starting behavior. A typical torque–speed curve for a general-purpose medium-size (55 kW) motor is shown in Figure 6.12. Such motors are unlikely to be described by the maker specifically as ‘deep bar’ but they nevertheless incorporate a measure of the skin effect and consequently achieve the ‘good’ torque–speed characteristic shown by the solid line in Figure 6.12.

The current–speed relationship is shown by the dashed line in Figure 6.12, both torque and current scales being expressed in per-unit (p.u.). This notation is widely used as a shorthand, with 1 p.u. (or 100%) representing rated value. For example, a torque of 1.5 p.u. simply means one and a half times rated value, while a current of 400% means a current of four times rated value.

4.3 Starting and run-up of slipping motors

By adding external resistance in series with the rotor windings the starting current can be kept low but at the same time the starting torque is high. This was the major advantage of the wound-rotor or slipping motor, which made it well suited for loads




Figure 6.12 Typical torque–speed and current–speed curves for a general-purpose industrial cage motor.




Figure 6.13 Torque–speed curves for a wound-rotor (slipring) motor showing how the external rotor-circuit resistance (R) can be varied in steps to provide an approximately constant torque during acceleration.

with heavy starting duties such as stone-crushers, cranes and conveyor drives, for most of which the inverter-fed cage motor is now preferred.

The influence of rotor resistance is shown by the set of torque–speed curves in Figure 6.13.

Typically the resistance at starting would be selected to give full-load torque together with rated current from the utility supply. The starting torque is then as indicated by point A in Figure 6.13.

As the speed rises, the torque would fall more or less linearly if the resistance remained constant, so in order to keep close to full torque the resistance is gradually reduced, either in steps, in which case the trajectory ABC, etc. is followed (Figure 6.13), or continuously so that maximum torque is obtained throughout. Ultimately the external resistance is made zero by shorting out the sliprings, and thereafter the motor behaves like a low-resistance cage motor, with a high running efficiency.

5. INFLUENCE OF SUPPLY VOLTAGE ON TORQUE–SPEED CURVE

We established earlier that at any given slip, the air-gap flux density is proportional to the applied voltage, and the induced current in the rotor is proportional to the flux density. The torque – which depends on the product of the flux and the rotor current – therefore depends on the square of the applied voltage. This means that a comparatively modest fall in the voltage will result in a much larger reduction in torque capability, with adverse effects which may not be apparent to the unwary until too late.

To illustrate the problem, consider the torque–speed curves for a cage motor shown in Figure 6.14. The curves (which have been expanded to focus attention on




Figure 6.14 Influence of stator supply voltage on torque-speed curves.

the low-slip region) are drawn for full voltage (100%), and for a modestly reduced voltage of 90%. With full voltage and full-load torque the motor will run at point X, with a slip of say 5%. Since this is the normal full-load condition, the rotor and stator currents will be at their rated values.

Now suppose that the voltage falls to 90%. The load torque is assumed to be constant so the new operating point will be at Y. Since the air-gap flux density is now only 0.9 of its rated value, the rotor current will have to be about 1.1 times rated value to develop the same torque, so the rotor e.m.f. is required to increase by 10%. But the flux density has fallen by 10%, so an increase in slip of 20% is called for. The new slip is therefore 6%.

The drop in speed from 95% of synchronous to 94% may well not be noticed, and the motor will apparently continue to operate quite happily. But the rotor current is now 10% above its rated value, so the rotor heating will be 21% more than is allowable for continuous running. The stator current will also be above rated value, so if the motor is allowed to run continuously, it will overheat. This is one reason why all large motors are fitted with protection which is triggered by over-temperature. Many small and medium motors do not have such protection, so it is important to guard against the possibility of undervoltage operation.

Another potential danger arises if the supply voltage is unbalanced, i.e. the three line-to-line voltages are unequal. This is most likely where the supply impedance is high and the line currents are unequal due to unbalanced loads elsewhere on the system.

Electrical engineers use the technique of 'symmetrical components' to analyze unbalanced 3-phase voltages. It is a method whereby the effect of unbalanced voltages on, say, a motor is quantified by finding how the motor behaves when subjected, independently, to three sets of balanced voltages that together simulate the actual unbalanced voltages.

The first balanced set is the positive sequence component: its phase-sequence is the normal one, e.g. UVW or ABC; the second is the negative sequence, having the phase sequence WVU; and the third is the zero sequence, in which the 3-phase

voltage components are co-phasal. There are simple analytic formulae for finding the components from the original three unbalanced voltages.

If the supply is balanced, the negative sequence and zero sequence components are both zero. Any unbalance gives rise to a negative sequence component, which will set up a rotating magnetic field traveling in the opposite direction to that of the main (positive sequence) component, and thus cause a braking torque and increase the losses, especially in the rotor. The zero sequence components produce a stationary field with three times the pole-number of the main field, but this can only happen when the motor has a star point connection.

Perhaps the most illuminating example of the application of symmetrical components as far as we are concerned is in the case of single-phase machines, where a single winding produces a pulsating field. This is an extreme case of unbalance because two of the phases are non-existent! It turns out that the positive sequence and negative sequence components are equal, which leads naturally to the idea that the pulsating field can be resolved into two counter-rotating fields. We will see later in this chapter that we can extend this picture to understand how the single-phase induction motor works.

Returning to 3-phase motors, a quite modest unbalance can be serious in terms of overheating. For example, to meet international standards motors are expected to tolerate only 1% negative sequence voltage continuously. Large motors often have negative sequence protection fitted, while small ones will rely on their thermal protection device to prevent overheating. Alternatively, if continuous operation under unbalanced conditions is required the motor must be de-rated significantly, e.g. to perhaps 80% of full-load with a voltage unbalance of 4%.

6. GENERATING

Having explored the torque–speed curve for the normal motoring region, where the speed lies between zero and just below synchronous, we must ask what happens if the speed is above the synchronous speed, i.e. the slip is negative. The unambiguous answer to this question is that the machine will switch from motoring to generating. Strangely, as mentioned previously, the authors have often encountered users who express deep scepticism, or even outright disbelief at the prospect of induction machines generating, so it is important that we attempt to counter what is clearly a widely held misconception.

A typical torque–speed curve for a cage motor covering the full range of speeds which are likely in practice is shown in [Figure 6.15](#).

We can see from [Figure 6.15](#) that the decisive factor as far as the direction of the torque is concerned is the slip, rather than the speed. When the slip is positive the torque is positive, and vice versa. The torque therefore always acts so as to urge the rotor to run at zero slip, i.e. at the synchronous speed. If the rotor is tempted to




Figure 6.15 Torque–speed curve over motoring region (slip between 0 and 1), braking region (slip greater than 1) and generating region (negative slip).

run faster than the field it will be slowed down, while if it is running below synchronous speed it will be urged to accelerate forwards. In particular, we note that for slips greater than 1, i.e. when the rotor is running backwards (i.e. in the opposite direction to the field), the torque will remain positive, so that if the rotor is unrestrained it will first slow down and then change direction and accelerate in the direction of the field.

6.1 Generating region

For negative slips, i.e. when the rotor is turning in the same direction, but at a higher speed than the traveling field, the ‘motor’ torque is in fact negative. In other words, the machine develops a torque that opposes the rotation, which can therefore only be maintained by applying a driving torque to the shaft. In this region the machine acts as an induction generator, converting mechanical power from the shaft into electrical power into the supply system. Cage induction machines are used in this way in wind-power generation schemes, as described in a later section.

It is worth stressing that, just as with the d.c. machine, we do not have to make any changes to an induction motor to turn it into an induction generator. In both cases, all that is needed is a source of mechanical power to turn the rotor faster than it would run if there was zero load or friction torque. For the d.c. motor, its ideal no-load speed is that at which its back e.m.f. equals the supply voltage, whereas for the induction motor, it is the synchronous speed.

On the other hand, we should be clear that, unlike the d.c. machine, the induction machine can only generate when it is connected to the supply. If we disconnect an induction motor from the utility supply and try to make it generate simply by turning the rotor we will not get any output because there is nothing to

set up the working flux: the flux (excitation) is not present until the motor is supplied with magnetizing current from the supply. It seems likely that this apparent inability to generate in isolation is what gave rise to the myth that induction machines cannot generate at all: a widely held view, but wholly incorrect!

There are comparatively few applications in which utility-supplied motors find themselves in the generating region, though as we will see later it is quite common in inverter-fed drives. We will, however, look at one example of a utility supply-fed motor in the so-called regenerative mode to underline the value of the motor's inherent ability to switch from motoring to generating automatically, without the need for any external intervention.

Consider a cage motor driving a simple hoist through a reduction gearbox, and suppose that the hook (unloaded) is to be lowered. Because of the static friction in the system, the hook will not descend on its own, even after the brake is lifted, so on pressing the 'down' button the brake is lifted and power is applied to the motor so that it rotates in the lowering direction. The motor quickly reaches full speed and the hook descends. As more and more rope winds off the drum, a point is reached where the lowering torque exerted by the hook and rope is greater than the running friction, and a restraining torque is then needed to prevent a runaway. The necessary stabilizing torque is automatically provided by the motor acting as a generator as soon as the synchronous speed is exceeded, as shown in [Figure 6.15](#). The speed will therefore be held at just above the synchronous speed, provided of course that the peak generating torque (see [Figure 6.15](#)) is not exceeded.

6.2 Self-excited induction generator

In previous sections we have stressed that the rotating magnetic field or excitation is provided by the magnetizing current drawn from the supply, so it would seem obvious that the motor could not generate unless a supply was provided to furnish the magnetizing current. However, it is possible to make the machine 'self-excite' if the conditions are right, and given the robustness of the cage motor this can make it an attractive proposition, especially for small-scale isolated installations.

We saw in Chapter 5 that when the induction motor is running at its normal speed, the rotating magnetic field that produces the currents and torque on the rotor also induced balanced 3-phase induced e.m.f.s in the stator windings, the magnitude of the e.m.f.s being not a great deal less than the voltage of the utility supply. So to act as an independent generator what we want to do is to set up the rotating magnetic field without having to connect to an active voltage source.

We discussed a similar matter in Chapter 3, in connection with self-excitation of the shunt d.c. machine. We saw that if enough residual magnetic flux remained in the field poles after the machine had been switched off, the e.m.f. produced when the shaft was rotated could begin to supply current to the field winding, thereby increasing the flux, further raising the e.m.f. and initiating a positive feedback (or

bootstrap) process which was ultimately stabilized by the saturation characteristic of the iron in the magnetic circuit.

Happily, much the same can be achieved with an isolated induction motor. We aim to capitalize on the residual magnetism in the rotor iron, and, by turning the rotor, generate an initial voltage in the stator to kick-start the process. The e.m.f. induced must then drive current to reinforce the residual field and promote the positive feedback to build up the traveling flux field. Unlike the d.c. machine, however, the induction motor has only one winding that provides both excitation and energy converting functions, so given that we want to get the terminal voltage to its rated level before we connect whatever electrical load we plan to supply, it is clearly necessary to provide a closed path for the would-be excitation current. This path should encourage the build-up of magnetizing current – and hence terminal voltage.

‘Encouraging’ the current means providing a very low impedance path, so that a small voltage drives a large current, and, since we are dealing with a.c. quantities, we naturally seek to exploit the phenomenon of resonance, by placing a set of capacitors in parallel with the (inductive) windings of the machine, as shown in Figure 6.16.

The reactance of a parallel circuit consisting of pure inductance (L) and capacitance (C) at angular frequency ω is given by $X = \omega L - 1/\omega C$, so at low and high frequencies the reactance is very large, but at the so-called resonant frequency ($\omega_0 = 1/\sqrt{LC}$), the reactance becomes zero. Here the inductance is the magnetizing inductance of each phase of the induction machine, and C is the added capacitance, the value being chosen to give resonance at the desired frequency of generation. Of course the circuit is not ideal because there is resistance in the windings, but nevertheless the inductive reactance can be ‘tuned out’ by choice of capacitance, leaving a circulating path of very low resistance. Hence by turning the rotor at the speed at which the desired frequency is produced by the residual




Figure 6.16 Self-excited induction generator. The load is connected only after the stator voltage has built up.

magnetism (e.g. 1800 rev/min for a 4-pole motor to generate 60 Hz), the initial modest e.m.f. produces a disproportionately high current and the flux builds up until limited by the non-linear saturation characteristic of the iron magnetic circuit. We then get balanced 3-phase voltages at the terminals, and the load can be applied by closing switch S (Figure 6.16).

The description above gives only a basic outline of the self-excitation mechanism. Such a scheme would only be satisfactory for a very limited range of driven speeds and loads, and in practice further control features are required to vary the effective capacitance (typically using triac control) in order to keep the voltage constant when the load and/or speed vary widely.

6.3 Doubly-fed induction machine for wind-power generation

The term doubly-fed refers to an induction machine in which both stator and rotor windings are connected to an a.c. power source: we are therefore talking about wound rotor (or slipring) motors, where the rotor windings are accessed through insulated rotating sliprings.

Traditionally, large slipring machines were used in Kramer¹ drives to recover the slip energy in the rotor and return it to the supply, so that efficient operation was possible at much higher slips than would otherwise have been possible. Some Kramer drives remain, but in the twenty-first century by far the major application for the doubly-fed induction machine is in wind-power generation where the wind turbine feeds directly into the utility grid.

Before we see why the doubly-fed motor is favored, we should first acknowledge that in principle we could take a cage motor, connect it directly to the utility grid, and drive its rotor from the wind turbine (via a gearbox with an output speed a little above synchronous speed) so that it supplied electrical power to the grid. However, the range of speeds over which stable generation is possible would be only a few per cent above synchronous, and this is a poor match as far as the turbine characteristics are concerned. Ideally, in order to abstract the maximum power from the wind, the speed of rotation of the blades must vary according to the conditions, so being forced to remain at a more-or-less constant speed by being connected to a cage motor is not good news. In addition, when there are rapid fluctuations in wind speed that produce bursts of power, the fact that the speed is constant means that there are rapid changes in torque which produce unwelcome fatigue loading in the gearbox. What is really wanted is a generator in which the generated frequency can be maintained constant over a wider speed range, and this is where the doubly-fed system scores.

¹ See, for example, *The Control Techniques Drives and Controls Handbook*, 2nd edition, by W. Drury.




Figure 6.17 Doubly-fed wound-rotor induction machine for wind-power generation.

The stator is connected directly to the utility grid, while the rotor windings are also linked to the grid, but via a pair of a.c./d.c. converters, as shown in Figure 6.17. The converters – connected via a d.c. link – allow power to flow to or from the fixed-frequency grid into or out of the rotor circuit, the frequency of which varies according to the shaft speed (see below). The rating of these converters will be substantially less than the rated power output of the induction machine, depending on the speed range that is to be accepted. For example, if the speed range is to be $\pm 1.3N_s$, where N_s is the synchronous speed of the induction motor connected to the grid, and full torque is to be available over the full speed range, the rating of the converters will be only 30% of the rating of the machine. (This is a major advantage over the alternative of having a conventional synchronous generator and a frequency converter of 100% rating.)

Understanding all the details of how the doubly-fed induction generator operates is not easy, but we can get to the essence by picturing the relationship between the rotating magnetic field and the stator and rotor.

Given that the stator is permanently connected to the utility grid, we know that the magnitude and speed of the rotating magnetic field cannot vary, so the synchronous speed of a 4-pole machine connected to the 60 Hz supply will be 1800 rev/min. If we want to generate power to the grid at this speed, we feed d.c. (i.e. zero frequency) into the rotor, and we then have what amounts to a conventional synchronous machine, which can motor or generate, and its power-factor can be controlled via the rotor current (see section 2 of Chapter 9). The 4-pole field produced by the rotor is stationary with respect to the rotor, but because the rotor is turning at 1800 rev/min, it rotates at the same speed as the traveling field produced by the stator windings, the two effectively being locked together so that torque can be transmitted and power can be converted. Energy conversion is only possible at the speed of exactly 1800 rev/min.

Now suppose that the wind turbine speed falls so that the speed at the machine shaft is only 1500 rev/min. For the rotor to be able to lock on to the 1800 rev/min stator field, the field that it produces must rotate at 300 rev/min (in a positive sense) relative to the rotor, so that its speed relative to the stator is $1500 + 300 = 1800$. This is achieved by supplying the rotor with 3-phase current at a frequency of

10 Hz. Conversely, if the turbine drives the machine shaft speed at 2100 rev/min, the rotor field must rotate at 300 rev/min in a negative sense relative to the rotor; i.e. the rotor currents must again be at 10 Hz, but with reversed phase sequence.

It turns out that if the input speed is below synchronous (1500 rev/min in the example above), electrical power has to be fed into the rotor circuit. This power is taken from the utility grid, but (neglecting losses, which are small) it then emerges from the stator, together with the mechanical power supplied by the turbine. Thus we can think of the electrical input power to the rotor as merely ‘borrowed’ from the grid to allow the energy conversion to take place; of course, the net power supplied to the grid all comes from the turbine. In this example, if the turbine torque is at rated value, the overall power output will be $1500/1800P$, i.e. $5P/6$, where P is the rated power at normal (synchronous) speed. This will be made up of an output power of P from the stator, from which the rotor converters take $300/1800P$, i.e. $P/6$ into the rotor.

When the driven speed is above synchronous (e.g. 2100 rev/min), both stator and rotor circuits export power to the grid. In this case, with rated turbine torque, the power to the grid will be $2100/1800P$, i.e. $7P/6$, comprising P from the stator and $P/6$ from the rotor converter.

We should be clear that there is no magic in being able to exceed the original rated power of our machine. At full torque, the electric and magnetic loadings will be at their rated values, and the increase in power is therefore entirely due to the higher speed. We discussed this towards the end of Chapter 1.

Just as with a conventional (single-speed) synchronous machine, we can control the extent to which the stator and rotor sides contribute to the setting up of the resultant flux, and this means that we can control the grid power-factor via the rotor circuit; this can be very advantageous where there is a system requirement to export or absorb reactive volt-amperes.

7. BRAKING

7.1 Plug reversal and plug braking

Because the rotor always tries to catch up with the rotating field, it can be reversed rapidly simply by interchanging any two of the supply leads. The changeover is usually obtained by having two separate 3-pole contactors, one for forward and one for reverse. This procedure is known as plug reversal or plugging, and is illustrated in Figure 6.18.

The motor is initially assumed to be running light (and therefore with a very small positive slip) as indicated by point A on the dashed torque–speed curve in Figure 6.18(a). Two of the supply leads are then reversed, thereby reversing the direction of the field, and bringing the mirror-image torque–speed curve shown by the solid line into play. The slip of the motor immediately after reversal is




Figure 6.18 Torque-speed and speed-time curves for plug reversal of cage motor.

approximately 2, as shown by point B on the solid curve. The torque is thus negative, and the motor decelerates, the speed passing through zero at point C and then rising in the reverse direction before settling at point D, just below the synchronous speed.

The speed-time curve is shown in Figure 6.18(b). We can see that the deceleration (i.e. the gradient of the speed-time graph) reaches a maximum as the motor passes through the peak torque (pull-out) point, but thereafter the final speed is approached gradually, as the torque tapers down to point D.

Very rapid reversal is possible using plugging; for example, a 1 kW motor will typically reverse from full speed in under one second. But large cage motors can only be plugged if the supply can withstand the very high currents involved, which are even larger than when starting from rest. Frequent plugging will also cause serious overheating, because each reversal involves the ‘dumping’ of four times the stored kinetic energy as heat in the windings.

Plugging can also be used to stop the rotor quickly, but obviously it is then necessary to disconnect the supply when the rotor comes to rest, otherwise it will run up to speed in reverse. A shaft-mounted reverse-rotation detector is therefore used to trip out the reverse contactor when the speed reaches zero.

We should note that whereas in the regenerative mode (discussed in the previous section) the slip was negative, allowing mechanical energy from the load to be converted to electrical energy and fed back to the mains, plugging is a wholly dissipative process in which all the kinetic energy ends up as heat in the motor.

7.2 Injection braking

This is the most widely used method of electrical braking. When the ‘stop’ signal occurs the 3-phase supply is interrupted, and a d.c. current is fed into the stator via two of its terminals. The d.c. supply is usually obtained from a rectifier fed via a low-voltage high-current transformer.

We saw earlier that the speed of rotation of the air-gap field is directly proportional to the supply frequency, so it should be clear that since d.c. is




Figure 6.19 Torque–speed curve for d.c. injection braking of cage motor.

effectively zero frequency, the air-gap field will be stationary. We also saw that the rotor always tries to run at the same speed as the field. So if the field is stationary, and the rotor is not, a braking torque will be exerted. A typical torque–speed curve for braking a cage motor is shown in Figure 6.19, from which we see that the braking (negative) torque falls to zero as the rotor comes to rest.

This is in line with what we would expect, since there will only be induced currents in the rotor (and hence torque) when the rotor is ‘cutting’ the flux. As with plugging, injection (or dynamic) braking is a dissipative process, all the kinetic energy being turned into heat inside the motor.

8. SPEED CONTROL

We have seen that to operate efficiently an induction motor must run with a small slip. It follows that any efficient method of speed control must be based on varying the synchronous speed of the field, rather than the slip. The two factors which determine the speed of the field are the supply frequency and the pole-number (see equation (5.1)).

The pole-number has to be an even integer, so where continuously adjustable speed control over a wide range is called for, the best approach by far is to provide a variable-frequency supply. This method is very important, and is dealt with separately in Chapters 7 and 8. In this chapter we are concerned with constant-frequency (utility-connected) operation, so we are limited to either pole-changing, which can provide discrete speeds only, or slip-control which can provide continuous speed control, but is inherently inefficient.

8.1 Pole-changing motors

For some applications continuous speed control may be an unnecessary luxury, and it may be sufficient to be able to run at two discrete speeds. Among many instances

where this can be acceptable and economic are pumps, lifts and hoists, fans and some machine tool drives.

We established in Chapter 5 that the pole-number of the field was determined by the layout and interconnection of the stator coils, and that once the winding has been designed, and the frequency specified, the synchronous speed of the field is fixed. If we wanted to make a motor which could run at either of two different speeds, we could construct it with two separate stator windings (say 4-pole and 6-pole), and energize the appropriate one. There is no need to change the cage rotor since the pattern of induced currents can readily adapt to suit the stator pole-number. Early two-speed motors did have two distinct stator windings, but were bulky and inefficient.

It was soon realized that if half of the phase-belts within each phase-winding could be reversed in polarity, the effective pole-number could be halved. For example, a 4-pole m.m.f. pattern (N-S-N-S) would become (N-N-S-S), i.e. effectively a 2-pole pattern with one large N and one large S pole. By bringing out six leads instead of three, and providing switching contactors to effect the reversal, two discrete speeds in the ratio 2:1 are therefore possible from a single winding. The performance at the high (e.g. 2-pole) speed is relatively poor, which is not surprising in view of the fact that the winding was originally optimized for 4-pole operation.

It was not until the advent of the more sophisticated pole amplitude modulation (PAM) method in the 1960s that two-speed single-winding high-performance motors with more or less any ratio of speeds became available from manufacturers. This subtle technique allows close ratios such as 4/6, 6/8, 8/10 or wide ratios such as 2/24 to be achieved. The beauty of the PAM method is that it is not expensive. The stator winding has more leads brought out, and the coils are connected to form non-uniform phase-belts, but otherwise construction is the same as for a single-speed motor. Typically six leads will be needed, three of which are supplied for one speed, and three for the other, the switching being done by contactors. The method of connection (star or delta) and the number of parallel paths within the winding are arranged so that the air-gap flux at each speed matches the load requirement. For example, if constant torque is needed at both speeds, the flux needs to be made the same, whereas if reduced torque is acceptable at the higher speed the flux can obviously be lower.

8.2 Voltage control of high-resistance cage motors

Where efficiency is not of paramount importance, the torque (and hence the running speed) of a cage motor can be controlled simply by altering the supply voltage. The torque at any slip is approximately proportional to the square of the voltage, so we can reduce the speed of the load by reducing the voltage. The method is not suitable for standard low-resistance cage motors, because their stable




Figure 6.20 Speed control of cage motor by stator voltage variation; (a) low-resistance rotor, (b) high-resistance rotor.

operating speed range is very restricted, as shown in Figure 6.20(a). But if special high-rotor resistance motors are used, the gradient of the torque–speed curve in the stable region is much less, and a somewhat wider range of steady-state operating speeds is available, as shown in Figure 6.20(b).

The most unattractive feature of this method is the low efficiency which is inherent in any form of slip-control. We recall that the rotor efficiency at slip s is $(1 - s)$, so if we run at say 70% of synchronous speed (i.e. $s = 0.3$), 30% of the power crossing the air-gap is wasted as heat in the rotor conductors. The approach is therefore only practicable where the load torque is low at low speeds, so a fan-type characteristic is suitable, as shown in Figure 6.20(b). Voltage control became feasible only when relatively cheap thyristor a.c. voltage regulators arrived on the scene during the 1970s, and although it enjoyed some success it is now seldom seen. The hardware required is essentially the same as discussed earlier for soft starting, and a single piece of kit can therefore serve for both starting and speed control.

8.3 Speed control of wound-rotor motors

The fact that the rotor resistance can be varied easily allows us to control the slip from the rotor side, with the stator supply voltage and frequency constant. Although the method is inherently inefficient it is still sometimes used because of its simplicity and comparatively low cost.

A set of torque–speed characteristics is shown in Figure 6.21, from which it should be clear that by appropriate selection of the rotor circuit resistance, any torque up to typically 1.5 times full-load torque can be achieved at any speed.

8.4 Slip energy recovery

Instead of wasting rotor-circuit power in an external resistance, it can be converted and returned to the mains supply. Frequency conversion is necessary because the rotor circuit operates at slip frequency, so it cannot be connected directly to the utility supply. These systems (which are known as the static Kramer drive) have




Figure 6.21 Influence of external rotor resistance (R) on torque–speed curve of wound-rotor motor.

largely been superseded by inverter-fed cage motors, but some survive so a brief mention is in order.

In a slip energy recovery system, the slip-frequency a.c. from the rotor is first rectified in a 3-phase diode bridge and smoothed before being returned to the utility supply via a 3-phase thyristor bridge converter operating in the inverting mode (see Chapter 4). A transformer is usually required to match the output from the controlled bridge to the mains voltage. Since the cost of both converters depends on the slip power they have to handle, this system is most often used where only a modest range of speeds (say from 70% of synchronous and above) is required, such as in large pump and compressor drives.

9. POWER-FACTOR CONTROL AND ENERGY OPTIMIZATION

In addition to their use for soft-start and speed control, thyristor voltage regulators provide a means for limited control of power-factor for cage motors, and this can allow a measure of energy economy. However, the fact is that there are comparatively few situations where considerations of power-factor and/or energy economy alone are sufficient to justify the expense of a voltage controller. Only when the motor operates for very long periods running light or at low load can sufficient savings be made to cover the outlay. There is certainly no point in providing energy economy when the motor spends most of its time working at or near full-load.

Both power-factor control and energy optimization rely on the fact that the air-gap flux is proportional to the supply voltage, so that by varying the voltage, the flux can be set at the best level to cope with the prevailing load. We can see straightaway that nothing can be achieved at full-load, since the motor needs full flux (and hence full voltage) to operate as intended. Some modest savings in losses can be achieved at reduced load, as we will see.




Figure 6.22 Phasor diagram showing improvement of power-factor by reduction of stator voltage.

If we imagine the motor to be running with a low load torque and full voltage, the flux will be at its full value, and the magnetizing component of the stator current will be larger than the work component, so the input power-factor ($\cos \phi_a$) will be very low, as shown in [Figure 6.22\(a\)](#).

Now suppose that the voltage is reduced to say half (by phasing back the thyristors), thereby halving the air-gap flux and reducing the magnetizing current by at least a factor of two. With only half the flux, the rotor current must double to produce the same torque, so the work current reflected in the stator will also double. The input power-factor ($\cos \phi_b$) will therefore improve considerably (see [Figure 6.22\(b\)](#)). Of course the slip with ‘half-flux’ operation will be higher (by a factor of four), but with a low-resistance cage it will still be small, and the drop in speed will therefore be slight.

The success (or otherwise) of the energy economy obtained depends on the balance between the iron losses and the copper losses in the motor. Reducing the voltage reduces the flux, and hence reduces the eddy current and hysteresis losses in the iron core. But as we have seen above, the rotor current has to increase to produce the same torque, so the rotor copper loss rises. The stator copper loss will reduce if (as in [Figure 6.22](#)) the magnitude of the stator current falls. In practice, with average general-purpose motors, a net saving in losses only occurs for light loads, say at or below 25% of full-load, though the power-factor will always increase.

10. SINGLE-PHASE INDUCTION MOTORS

Single-phase induction motors are simple, robust and reliable, and continue to be used in large numbers especially in domestic and commercial applications

where 3-phase supplies are not available. Although outputs of up to a few kW are possible, the majority are below 0.5 kW, and are used in straightforward applications such as refrigeration compressors, and dryers, pumps and fans, small machine tools, etc. However, these traditional applications can now benefit from the superior control and low cost provided by the simple inverter and 3-phase induction or permanent magnet motor, so the single-phase motor looks set for a niche role in future.

10.1 Principle of operation

If one of the leads of a 3-phase motor is disconnected while it is running light, it will continue to run with a barely perceptible drop in speed, and a somewhat louder hum. With only two leads remaining there can only be one current, so the motor must be operating as a single-phase machine. If load is applied the slip increases more quickly than under 3-phase operation, and the stall torque is much less, perhaps one-third. When the motor stalls and comes to rest it will not restart if the load is removed, but remains at rest drawing a heavy current and emitting an angry hum. It will burn out if not disconnected rapidly.

It is not surprising that a truly single-phase cage induction motor will not start from rest, because as we saw in Chapter 5 the single winding, fed with a.c., simply produces a pulsating flux in the air-gap, without any suggestion of rotation. It is, however, surprising to find that if the motor is given a push in either direction it will pick up speed, slowly at first but then with more vigor, until it settles with a small slip, ready to take up load. Once turning, a rotating field is evidently brought into play to continue propelling the rotor.

We can understand how this comes about by first picturing the pulsating m.m.f. set up by the current in the stator winding as being the resultant of two identical traveling waves of m.m.f., one in the forward direction and the other in reverse. (This equivalence is not self-evident, but can be demonstrated by applying the method of symmetrical components, discussed earlier in this chapter.)

When the rotor is stationary, it reacts equally to both traveling waves, and no torque is developed. When the rotor is turning, however, the induced rotor currents are such that their m.m.f. opposes the reverse stator m.m.f. to a greater extent than they oppose the forward stator m.m.f. The result is that the forward flux wave (which is what develops the forward torque) is bigger than the reverse flux wave (which exerts a drag). The difference widens as the speed increases, the forward flux wave becoming progressively bigger as the speed rises while the reverse flux wave simultaneously reduces. This 'positive feedback' effect explains why the speed builds slowly at first, but later zooms up to just below synchronous speed. At the normal running speed (i.e. small slip), the forward flux is many times larger than the backward flux, and the drag torque is only a small percentage of the forward torque.

As far as normal running is concerned, a single winding is therefore sufficient. But all motors must be able to self-start, so some mechanism has to be provided to produce a rotating field even when the rotor is at rest. Several methods are employed, all of them using an additional winding.

The second winding usually has less copper than the main winding, and is located in the slots which are not occupied by the main winding, so that its m.m.f. is displaced in space relative to that of the main winding. The current in the second winding is supplied from the same single-phase source as the main winding current, but is caused to have a phase-lag, by a variety of means which are discussed later. The combination of a space displacement between the two windings together with a time displacement between the currents produces a 2-phase machine. If the two windings were identical, displaced by 90° , and fed with currents with 90° phase-shift, an ideal rotating field would be produced. In practice we can never achieve a 90° phase-shift between the currents, and it turns out to be more economic not to make the windings identical. Nevertheless, a decent rotating field is set up, and entirely satisfactory starting torque can be obtained. Reversal is simply a matter of reversing the polarity of one of the windings, and performance is identical in both directions.

The most widely used methods are described below. At one time it was common practice for the second or auxiliary winding to be energized only during start and run-up, and for it to be disconnected by means of a centrifugal switch mounted on the rotor, or sometimes by a time-switch. This practice gave rise to the term ‘starting winding’. Nowadays it is more common to find both windings in use all the time.

10.2 Capacitor run motors

A capacitor is used in series with the auxiliary winding (Figure 6.23) to provide a phase-shift between the main and auxiliary winding currents. The capacitor (usually of a few μF , and with a voltage rating which may well be higher than the mains voltage) may be mounted piggyback fashion on the motor, or located elsewhere. Its value represents a compromise between the conflicting requirements of high starting torque and good running performance.

A typical torque–speed curve is also shown in Figure 6.21; the modest starting torque indicates that capacitor run motors are generally best suited to fan-type loads. Where higher starting torque is needed, two capacitors can be used, one being switched out when the motor is up to speed.




Figure 6.23 Single-phase capacitor-run induction motor.




Figure 6.24 Single-phase split-phase induction motor.

As mentioned above, the practice of switching out the starting winding altogether is no longer favored, but many old ones remain, and where a capacitor is used they are known as ‘capacitor start’ motors.

10.3 Split-phase motors

The main winding is of thick wire, with a low resistance and high reactance, while the auxiliary winding is made of fewer turns of thinner wire with a higher resistance and lower reactance (Figure 6.24). The inherent difference in impedance is sufficient to give the required phase-shift between the two currents without needing any external elements in series. Starting torque is good at typically 1.5 times full-load torque, as also shown in Figure 6.24. As with the capacitor type, reversal is accomplished by changing the connections to one of the windings.

10.4 Shaded pole motors

There are several variants of this extremely simple, robust and reliable cage motor, which is used for low-power applications such as hair-dryers, oven fans, office equipment, and display drives. A 2-pole version from the cheap end of the market is shown in Figure 6.25.




Figure 6.25 Shaded pole induction motor.

The rotor, typically between 1 and 4 cm diameter, has a die-cast aluminum cage, while the stator winding is a simple concentrated coil wound round the laminated core. The stator pole is slotted to receive the ‘shading ring’, which is a single short-circuited turn of thick copper or aluminum.

Most of the pulsating flux produced by the stator winding by-passes the shading ring and crosses the air-gap to the rotor. But some of the flux passes through the shading ring, and because it is alternating it induces an e.m.f. and current in the ring. The opposing m.m.f. of the ring current diminishes and retards the phase of the flux through the ring, so that the flux through the ring reaches a peak after the main flux, thereby giving what amounts to a rotation of the flux across the face of the pole. This far from perfect traveling wave of flux produces the motor torque by interaction with the rotor cage. Efficiencies are low because of the rather poor magnetic circuit and the losses caused by the induced currents in the shading ring, but this is generally acceptable when the aim is to minimize first cost. Series resistance can be used to obtain a crude speed control, but this is only suitable for fan-type loads. The direction of rotation depends on whether the shading ring is located on the right or left side of the pole, so shaded pole motors are only suitable for uni-directional loads.

11. POWER RANGE

Having praised the simplicity and elegance of the induction motor, it is perhaps not surprising that they are successful over an extraordinarily wide power range from multi-megawatts down to a few tens of watts. (Indeed we cannot think of any other non-electric energy converter that can exceed a span of six orders of magnitude!)

At the upper end the limit is largely a question of low demand, there being few applications that need a shaft that delivers many tens of megawatts. But at the lower end we may wonder why there are no very small ones. Industrial (3-phase) induction motors are rarely found below about 200 W, and single-phase versions rarely extend below about 50 W, yet applications in this range are legion.

We will see that when we scale down a successful design, the excitation or flux-producing function of the windings becomes more and more demanding until eventually the heat produced in the windings by the excitation current causes the permissible temperature to be reached. There is then no spare capacity for the vital function of supplying the mechanical output power, so the machine is of no use.

11.1 Scaling down – the excitation problem

We can get to the essence of the matter by imagining that we take a successful design and scale all the linear dimensions by half. We know that in order to fully utilize the iron of the magnetic circuit we would want the air-gap flux density to be

the same as in the original design, so because the air-gap length has been halved the stator m.m.f. needs to be half of what it was. The number of coils and the turns in each coil remain as before, so if the original magnetizing current was I_m , the magnetizing current of the half-scale motor will be $I_m/2$.

Turning now to what happens to the resistance of the winding, we will assume that the resistance of the original winding was R . In the half-scale motor, the total length of wire is half of what it was, but the cross-sectional area of the wire is only a quarter of the original. As a result the new resistance is twice as great, i.e. $2R$.

The power dissipated in providing the air-gap flux in the original motor is given by $I_m^2 R$, while the corresponding excitation power in the half-scale motor is given by

$$\left(\frac{I_m}{2}\right)^2 \times 2R = \frac{1}{2} I_m^2 R$$

When we consider what determines the steady temperature rise of a body in which heat is dissipated, we find that the equilibrium condition is reached when the rate of loss of heat to the surroundings is equal to the rate of production of heat inside the body. And, not surprisingly, the rate of loss of heat to the surroundings depends on the temperature difference between the body and its surroundings, and on the surface area through which the heat escapes. In the case of the copper windings in a motor, the permissible temperature rise depends on the quality of insulation, so we will make the reasonable assumption that the same insulation is used for the scaled motor as for the original.

We have worked out that the power dissipation in the new motor is half of that in the original. However, the surface area of the new winding is only one-quarter, so clearly the temperature rise will be higher, and if all other things were equal, it will double. We might aim to ease matters by providing bigger slots so that the current density in the copper could be reduced, but as explained in Chapter 1 this means that there is less iron in the teeth to carry the working flux. A further problem arises because it is simply not practicable to go on making the air-gap smaller because the need to maintain clearances between the moving parts would require unacceptably tight manufacturing tolerances.

Obviously, there are other factors that need to be considered, not least that a motor is designed to reach its working temperature when the full current (not just the magnetizing current) is flowing. But the fact is that the magnetization problem we have highlighted is the main obstacle in small sizes, not only in induction motors but also in any motor that derives its excitation from the stator windings. Permanent magnets therefore become attractive for small motors, because they provide the working flux without producing unwelcome heat. Permanent magnet motors are discussed in Chapter 9.