CHAPTER SEVEN

Variable Frequency Operation of Induction Motors

1. INTRODUCTION

We saw in Chapter 6 the many attributes of the induction motor which have made it the preferred workhorse of industry. These include simple low-cost construction which lends itself to totally enclosed designs suitable for dirty or even hazardous environments; limited routine maintenance with no brushes; only three power connections; and good full-load efficiency. We have also seen that when operated from the utility supply there are a number of undesirable characteristics, the most notable being that there is only one speed of operation (or more precisely a narrow load-dependent speed range). In addition, starting equipment is often required to avoid excessive currents of up to six times rated current when starting direct-on-line; reversal requires two of the power cables to be interchanged; and the instantaneous torque cannot be controlled, so the transient performance is poor.

We will see in this chapter that all the good features of the mains operated induction motor are retained and all the bad characteristics detailed above can be avoided when the induction motor is supplied from a variable-frequency source; i.e. its supply comes from an inverter.

The chapter divides broadly into two parts, both dealing with the capabilities of the induction motor when supplied from an inverter. The first part (sections 2 and 3) deals with the steady-state behavior when the operating frequency is solely determined by the inverter, and is independent of what is happening at the motor. We will refer to this set-up as 'inverter-fed', and in the early days of converter-driven induction motors this was the norm, the frequency being set by an oscillator that controlled the sequential periodic switching of the devices in the inverter. We will see that by appropriate control of the frequency and voltage we are able to operate over a very wide range of the torque–speed plane, but we will also identify the factors that place limits on what can be achieved. On a steady-state basis, this arrangement proved able to compete with the d.c. drive, but even when incorporated into a closed-loop control scheme the transient performance remained inferior. It is well worth absorbing the main messages from this study because although most contemporary drives now operate on a different basis, the steady-state running conditions at the motor remain the same. Readers who were able to follow the material in Chapter 6 should find this part straightforward.

Electric Motors and Drives
http://dx.doi.org/10.1016/B978-0-08-098332-5.00007-3

In the second major part of this chapter (from section 4), we explore both 'field-oriented' and 'direct torque' methods for control of the inverter/induction motor combination. We prefer not to use the term 'inverter-fed' in these circumstances because although the motor derives its supply from an inverter, the switching of the inverter devices is determined by the state of the flux and currents in the motor, rather than being imposed from a separate oscillator. Both methods allow us to achieve hitherto unheard of levels of transient performance, but it is important to note that they only became possible because of the development of fast, cheap, digital processors that can implement the high-speed calculations necessary to model and control the motor in real time.

Understanding field-oriented control is usually regarded as challenging, even for experienced drives personnel, not least because the subject tends to be highly mathematical. However, although we will adopt a largely graphical approach to get to the heart of the matter, it is likely that most readers will find it advisable to absorb the tutorial material in sections 4 and 5 first.

For most of this chapter we will assume that the motor is supplied from an ideal balanced sinusoidal voltage source. Our justification for doing this is that although the pulse-width-modulated voltage waveform supplied by the inverter will not be sinusoidal (see Figure 7.1), the motor performance depends principally on the fundamental (sinusoidal) component of the applied voltage. This is a somewhat surprising but extremely welcome simplification, because it allows us to make use of our knowledge of how the induction motor behaves with a sinusoidal supply to anticipate how it performs when fed from an inverter.

There will be more discussion of the various control arrangements and practicalities of inverter-fed operation in Chapter 8.

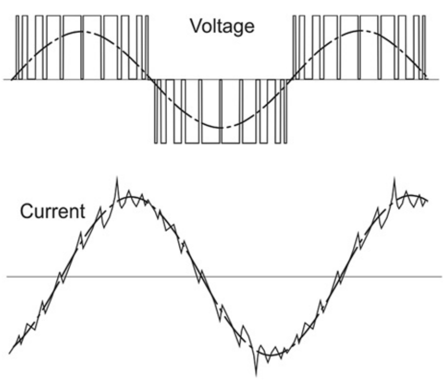

Figure 7.1 Typical voltage and current waveforms for PWM inverter-fed induction motor. (The fundamental-frequency component is shown by the dotted line.)

2. INVERTER-FED INDUCTION MOTOR DRIVES

It was explained in Chapter 6 that the induction motor can only run efficiently at low slips, i.e. close to the synchronous speed of the rotating field. The best method of speed control must therefore provide for continuous smooth variation of the synchronous speed, which in turn calls for variation of the supply frequency. This is readily achieved using a power electronic inverter (as discussed in Chapter 2) to supply the motor. A complete speed control scheme, which is illustrated with speed feedback, is shown in simplified block diagram form in Figure 7.2.

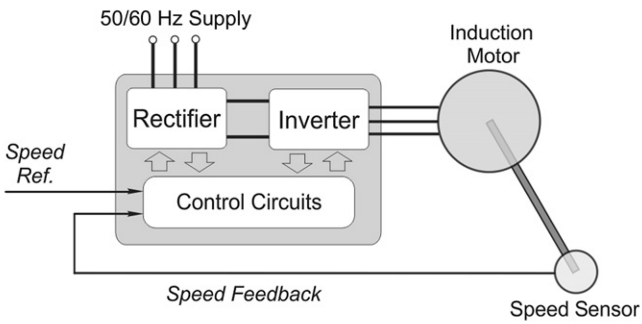

Figure 7.2 General arrangement of inverter-fed variable-frequency induction motor controlled-speed drive.

The arrangement shown in Figure 7.2 shows the motor with a speed sensor attached to the motor shaft. For all but the most demanding dynamic applications, or where full torque at standstill is a requirement, a speed sensor would not normally be required. This is good news as fitting a speed sensor to a standard induction motor involves significant additional cost and additional cabling.

We should recall that the function of the converter (i.e. rectifier and variable-frequency inverter) is to draw power from the fixed-frequency constant-voltage mains, rectify it and then convert it to variable frequency, variable voltage for driving the induction motor. Both the rectifier and the inverter employ switching strategies (see Chapter 2), so the power conversions are accomplished efficiently and the converter can be compact.

2.1 Steady-state operation – importance of achieving full flux

Three simple relationships need to be borne in mind in order to simplify understanding of how the inverter-fed induction motor behaves. First, we established in Chapter 5 that, for a given induction motor, the torque developed depends on the magnitude of the rotating flux density wave, and on the slip speed of the rotor, i.e. on the relative velocity of the rotor with respect to the flux wave. Secondly, the strength or amplitude of the flux wave depends directly on the supply voltage to the stator windings, and inversely on the supply frequency. Thirdly, the absolute speed of the flux wave depends directly on the supply frequency.

Recalling that the motor can only operate efficiently when the slip is small, we see that the basic method of speed control rests on the control of the speed of rotation of the flux wave (i.e. the synchronous speed), by control of the supply frequency. If the motor is a 4-pole one, for example, the synchronous speed will be 1500 rev/min when supplied at 50 Hz, 1800 rev/min at 60 Hz, 750 rev/min at 25 Hz, and so on. The no-load speed will therefore be almost exactly proportional to the supply frequency, because the torque at no load is small and the corresponding slip is also very small.

Turning now to what happens on load, we know that when a load is applied the rotor slows down, the slip increases, more current is induced in the rotor, and more torque is produced. When the speed has reduced to the point where the motor torque equals the load torque, the speed becomes steady. We normally want the drop in speed with load to be as small as possible, not only to minimize the drop in speed with load, but also to maximize efficiency: in short, we want to minimize the slip for a given load.

We saw in Chapter 5 that the slip for a given torque depends on the amplitude of the rotating flux wave: the higher the flux, the smaller the slip needed for a given torque. It follows that having set the desired speed of rotation of the flux wave by controlling the output frequency of the inverter we must also ensure that the magnitude of the flux is adjusted so that it is at its full (rated) value,[1] regardless of the speed of rotation. This is achieved, in principle, by making the output voltage from the inverter vary in the appropriate way in relation to the frequency.

Given that the amplitude of the flux wave is proportional to the supply voltage and inversely proportional to the frequency, it follows that if we arrange that the voltage supplied by the inverter varies in direct proportion to the frequency, the flux wave will have a constant amplitude. This simple mode of operation – where the V/f ratio is constant – was for many years the basis of the control strategy applied to most inverter-fed induction motors, and it can still be found in some commercial products.

Many inverters are designed for direct connection to the utility supply, without a transformer, and as a result the maximum inverter output voltage is limited to a value similar to that of the supply system. Since the inverter will normally be used to supply a standard induction motor designed, for example, for 400 V, 50 Hz operation, it is obvious that when the inverter is set to deliver 50 Hz, the voltage

[1] In general, operating at rated flux gives the best performance and on most loads the highest efficiency. Some commercial drives offer a mode of control in which the flux is reduced typically with the square of the speed: this can provide some benefit at low speeds for fan and pump-type loads where the magnetizing current accounts for a significant proportion of the motor losses.

should be 400 V, which is within the inverter's voltage range. But when the frequency is raised to say 100 Hz, the voltage should – ideally – be increased to 800 V in order to obtain full flux. The inverter cannot supply voltages above 400 V, and it follows that in this case full flux can only be maintained up to the base speed. Established practice is for the inverter to be capable of maintaining the 'V/f ratio', or rather the flux, constant up to the base speed (frequently 50 or 60 Hz), but to accept that at higher frequencies the voltage will be constant at its maximum value. This means that the flux is maintained constant at speeds up to base speed, but beyond that the flux reduces inversely with frequency. Needless to say the performance above base speed is adversely affected, as we will see.

Users are sometimes alarmed to discover that both voltage and frequency change when a new speed is demanded. Particular concern is expressed when the voltage is seen to reduce when a lower speed is called for. Surely, it is argued, it can't be right to operate say a 400 V induction motor at anything less than 400 V. The fallacy in this view should now be apparent: the figure of 400 V is simply the correct voltage for the motor when run directly from the utility supply, at say 50 Hz. If this full voltage were to be applied when the frequency was reduced to say 25 Hz, the implication would be that the flux would rise to twice its rated value. This would greatly overload the magnetic circuit of the machine, giving rise to excessive saturation of the iron, an enormous magnetizing current, and wholly unacceptable iron and copper losses. To prevent this from happening, and keep the flux at its rated value, it is essential to reduce the voltage in proportion to frequency. In the case above, for example, the correct voltage at 25 Hz would be 200 V.

It is worth stressing here that when considering a motor to be fed from an inverter there is no longer any special significance about the utility network frequency, and the motor can be wound for almost any base frequency. For example, a motor wound for 400 V, 100 Hz could in the above example operate with constant flux right up to 100 Hz.

3. TORQUE–SPEED CHARACTERISTICS

When the voltage at each frequency is adjusted so that the ratio of voltage to frequency (V/f) is kept constant up to base speed, a family of torque–speed curves as shown in Figure 7.3 is obtained. These curves are typical for a standard induction motor of several kW output.

As expected, the no-load speeds are directly proportional to the frequency, and if the frequency is held constant, e.g. at 25 Hz in Figure 7.3, the speed drops only modestly from no-load (point a) to full-load (point b). These are therefore good, useful open-loop characteristics, because the speed is held fairly well from no-load to full-load. If the application calls for the speed to be held precisely, this can clearly be achieved by raising the frequency so that the full-load operating point moves to point (c).

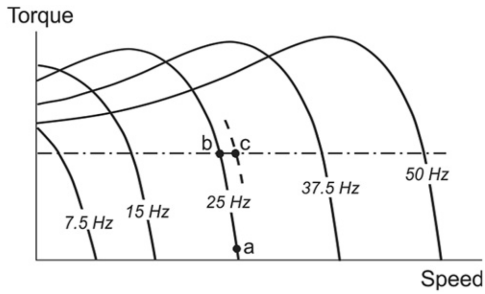

Figure 7.3 Torque–speed curves for inverter-fed induction motor with constant *V/f* ratio.

We note also that the pull-out torque and the torque stiffness (i.e. the slope of the torque–speed curve in the normal operating region) is more or less the same at all points below base speed, except at low frequencies where the voltage drop over the stator resistance becomes very significant as the applied voltage is reduced. A simple *V/f* control system would therefore suffer from significantly reduced flux and hence less torque at low speeds, as indicated in Figure 7.3.

The low-frequency performance can be improved by increasing the *V/f* ratio at low frequencies in order to restore full flux, a technique which is referred to as 'voltage boost'. In modern drive control schemes which calculate flux from a motor model (see section 8), the voltage is automatically boosted from the linear *V/f* characteristic that the approximate theory leads us to expect. A typical set of torque–speed curves for a drive with the improved low-speed torque characteristics obtained with voltage boost is shown in Figure 7.4.

The characteristics in Figure 7.4 have an obvious appeal because they indicate that the motor is capable of producing practically the same maximum torque at all speeds from zero up to the base (50 Hz) speed. This region of the characteristics is

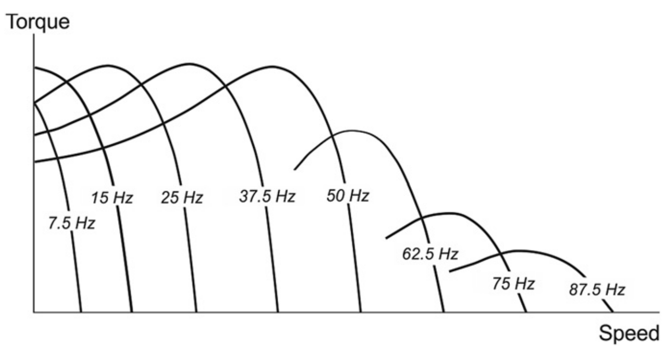

Figure 7.4 Typical torque–speed curves for inverter-fed induction motor with constant flux up to base speed (50 Hz) and constant voltage at higher frequencies.

known as the 'constant torque' region, which means that for frequencies up to base speed, the maximum possible torque which the motor can deliver is independent of the set speed. Continuous operation at peak torque will not be allowable because the motor will overheat, so an upper limit will be imposed by the controller, as discussed shortly. With this imposed limit, operation below base speed corresponds to the armature-voltage control region of a d.c. drive, as exemplified in Figure 3.9.

We should note that the availability of high torque at low speeds (especially at zero speed) means that we can avoid all the 'starting' problems associated with fixed-frequency operation (see Chapter 6). By starting off with a low frequency which is then gradually raised the slip speed of the rotor is always small, i.e. the rotor operates in the optimum condition for torque production all the time, thereby avoiding all the disadvantages of high slip (low torque and high current) that are associated with utility-frequency/direct-on-line (DOL) starting. This means that not only can the inverter-fed motor provide rated torque at low speeds, but – perhaps more importantly – it does so without drawing any more current from the utility supply than under full-load conditions, which means that we can safely operate from a weak supply without causing excessive voltage dips. For some essentially fixed-speed applications, the superior starting ability of the inverter-fed system alone may justify its cost.

Beyond the base frequency, the flux ('V/f ratio') reduces because V remains constant. The amplitude of the flux wave therefore reduces inversely with the frequency. Under constant flux operation, the pull-out torque always occurs at the same absolute value of slip, but in the constant-voltage region the peak torque reduces inversely with the square of the frequency and the torque–speed curve becomes less steep, as shown in Figure 7.4.

Although the curves in Figure 7.4 show what torque the motor can produce for each frequency and speed, they give no indication of whether continuous operation is possible at each point, yet this matter is of course extremely important from the user's viewpoint, and is discussed next.

3.1 Limitations imposed by the inverter – constant power and constant torque regions

A primary concern in the inverter is to limit the currents to a safe value as far as the main switching devices and the motor are concerned. The current limit will be typically set to the rated current of the motor, and the inverter control circuits will be arranged so that no matter what the user does the output current cannot exceed this safe (thermal) value, other than for clearly defined overload (e.g. 120% for 60 seconds) for which the motor and inverter will have been specified and rated. (For some applications involving a large number of starts and stops, the motor and drive may be specially designed for the specific duty.)

In modern control schemes (sections 8 and 9) it is possible to have independent control of the flux- and torque-producing components of the current, and in this way the current limit imposes an upper limit on the permissible torque. In the region below base speed, this will normally correspond to the rated torque of the motor, which is typically about half the pull-out torque, as indicated by the shaded region in Figure 7.5. Note that this is usually a thermal limit imposed by the motor design.

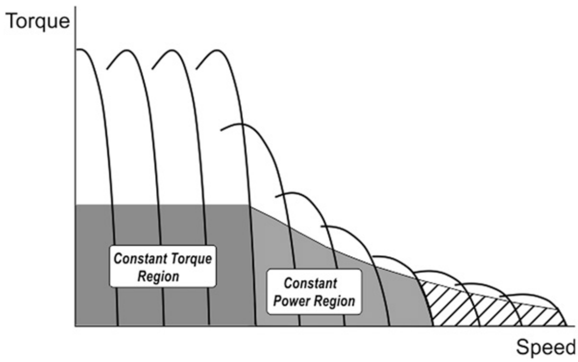

Figure 7.5 Constant torque, constant power and high-speed motoring regions.

Above base speed, it is not possible to increase the voltage and so the flux reduces inversely with the frequency. Since the stator (and therefore rotor) currents are again thermally limited (as we saw in the constant torque region), the maximum permissible torque also reduces inversely with the speed, as shown in Figure 7.5. This region is consequently known as the 'constant power' region. There is of course a close parallel with the d.c. drive here, both systems operating with reduced or weak field in the constant power region. In the constant power region, the flux is reduced and so the motor has to operate with higher slips than below base speed to develop the full (rated) rotor current and correspondingly reduced torque.

The voltage drop over the stator leakage inductance (see Appendix 2) increases with frequency. At typically twice base speed the extent of this voltage drop reduces the available voltage to such an extent that it is no longer possible for the motor to provide constant power operation, as indicated by the cross-hatched area in Figure 7.5.

3.2 Limitations imposed by the motor

The traditional practice in d.c. drives is to use a motor specifically designed for operation from a thyristor converter. The motor will have a laminated frame, will probably come complete with a tachogenerator, and – most important of all – will have been designed for through ventilation and equipped with an auxiliary air blower. Adequate ventilation is guaranteed at all speeds, and continuous operation with full torque (i.e. full current) at even the lowest speed is therefore in order.

By contrast, it is still common for inverter-fed systems to use a standard industrial induction motor. These motors are usually totally enclosed, with an external shaft-mounted fan which blows air over the finned outer case (and an internal stirring fan to circulate air inside the motor to avoid spot heating). They are designed first and foremost for continuous operation from the fixed frequency utility supply, and running at base speed.

As we have mentioned earlier, when such a motor is operated at a low frequency (e.g. 7.5 Hz), the speed is much lower than base speed and the efficiency of the cooling fan is greatly reduced. At the lower speed the motor will be able to produce as much torque as at base speed (see Figure 7.4) but in doing so the losses in both stator and rotor will also be more or less the same as at base speed, so it will overheat if operated for any length of time.

However, induction motors bearing the name 'inverter grade' or similar are readily available. As well as having reinforced insulation systems (see Chapter 8), they have been designed to offer a constant torque operating range below rated speed, typically down to 30% of base speed, without the need for an external cooling fan. In addition they may be offered with an external cooling fan to allow operation at constant (rated) torque down to standstill.

3.3 Four-quadrant capability

So far in this chapter it is natural that we have concentrated on motoring in quadrant 1 of the torque–speed plane (see Figure 3.12), because this is where the machine will spend most of its time running, but it is important to remind ourselves that the induction motor is equally at home as a generator, a role that it will frequently perform, even with an ordinary load, when a reduction in speed is called for. We should also recall that in this part of the chapter we are studying variable-frequency operation at the fundamental level, so we should bear in mind that in practice details of the control strategy will vary from drive to drive.

We can see how intermittent generation occurs with the aid of the torque–speed curves shown in Figure 7.6. These have been extended into quadrant 2, i.e. the negative-slip region, where the rotor speed is higher than synchronous, and a braking torque is exerted.

The family of curves indicates that for each set speed (i.e. each frequency) the speed remains reasonably constant because of the relatively steep torque–slip characteristic of the cage motor. If the load is increased beyond rated torque, an internal current limit comes into play to prevent the motor from reaching the unstable region beyond pull-out. Instead, the frequency and speed are reduced, and so the system behaves in a similar way to a d.c. drive.

Sudden changes in the speed reference are buffered so that the frequency is gradually increased or decreased. If the load inertia is low and/or the ramp time sufficiently long, the acceleration will be accomplished without the motor entering

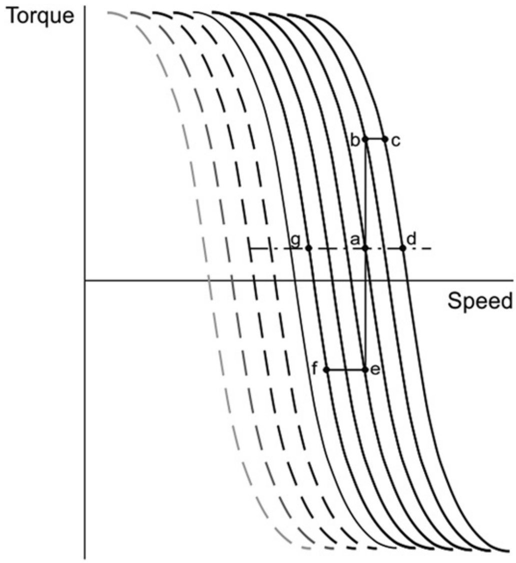

Figure 7.6 Acceleration and deceleration trajectories in the torque–speed plane.

the current-limit region. On the other hand, if the inertia is large and/or the ramp time was very short, the acceleration will take place as discussed below.

Suppose the motor is operating in the steady state with a constant load torque at point (a), when a new higher speed corresponding to point (d) is demanded. The frequency is increased, causing the motor torque to rise to point (b), where the current has reached the allowable limit. The rate of increase of frequency is then automatically reduced so that the motor accelerates under constant-current conditions to point (c), where the current falls below the limit: the frequency then remains constant and the trajectory follows the curve from (c) to settle finally at point (d).

A typical deceleration trajectory is shown by the path *aefg* in Figure 7.6. The torque is negative for much of the time, the motor operating in quadrant 2 and regenerating kinetic energy. Because we have assumed that the motor is supplied from an ideal voltage source, this excess energy will return to the supply automatically. In practice, however, we should note that many drives do not have the capability to return power to the a.c. supply, and the excess energy therefore has to be dissipated in a resistor inside the converter. (The resistor is usually connected across the d.c. link, and controlled by a chopper. When the level of the d.c. link voltage rises, because of the regenerated energy, the chopper switches the resistor on to absorb the energy. High inertia loads which are subjected to frequent deceleration can therefore pose problems of excessive power dissipation in this ‘dump’ resistor.)

To operate as a motor in quadrant 3 all that is required is for the phase sequence of the supply to be reversed, say from ABC to ACB. Unlike the utility-fed motor,

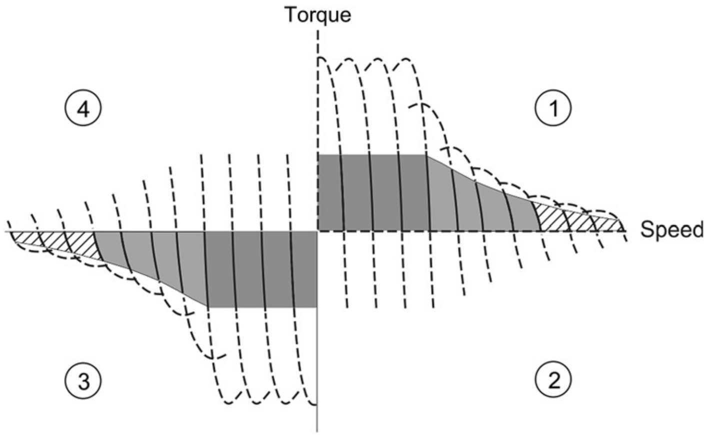

Figure 7.7 Operating regions in all four quadrants of the torque–speed plane.

there is no need to swap two of the power leads because the phase sequencing can be changed easily at the low-power logic level in the inverter. With reverse phase sequence, a mirror image set of 'motoring' characteristics is available, as shown in Figure 7.7. The shaded regions are as described for Figure 7.5, and the dashed lines indicate either short-term overload operation (quadrants 1 and 3) or regeneration during deceleration (quadrants 2 and 4).

Note that unlike the d.c. motor control strategies we examined in Chapter 4, neither the motor current, nor indeed any representation of torque, plays a role in the motor control strategy discussed so far (except when the current hits a limit, as discussed above).

4. INTRODUCTION TO FIELD-ORIENTED CONTROL

We now begin the second part of this chapter, which explores the contemporary approach to control of the inverter/induction motor combination. Field-oriented (or vector) control allows the induction motor/inverter combination to outperform conventional industrial d.c. drives, and its progressive refinement since the 1980s represents a major landmark in the history of electrical drives. It is therefore appropriate that its importance is properly reflected in this book, because one of our aims is to equip readers with sufficient understanding to allow them to converse intelligently with manufacturers and suppliers.

4.1 Outline of remainder of this chapter

Up to now, we have been able to cover topics without recourse to any very demanding mathematics, relying instead on physical explanations and diagrams, and our aim is to continue this approach. However, anyone who has consulted an article or textbook on the subject of field-oriented control will quickly become aware that most treatments involve extensive use of matrices and transform theory, and that many of the terms used will not be familiar to anyone not already schooled in the analysis of electrical machines. Fortunately, from our point of view, it is nevertheless possible to understand the underlying basis of field-oriented control via a relatively simple graphical approach, but even for this we have to make use of some ideas (such as flux linkage, and space phasors) that we have not discussed previously, so these are presented in the remainder of this section.

In section 5 we take a fresh look at the production of torque in induction motors, this time with the motor being supplied with controlled *currents* from a voltage-source inverter. We begin by taking a physical viewpoint that yields simple pictures that turn out to be accompanied by surprisingly simple formulae for torque. More importantly, the subsequent discussion in sections 6 and 7 makes clear what has to be done to achieve 'ideal' dynamic control of torque, something that was considered impossible until power electronics arrived.

The practical implementation of field-oriented control of torque is explored in section 8 through a detailed examination of the modus operandi of a typical sensorless control scheme. And finally, we look briefly at direct torque control, an alternative control strategy, in section 9.

In the remainder of this section we provide an introduction to some of the graphical and circuit-based techniques that we will need in order to understand torque control. The aim is to familiarize ourselves with the methodology and techniques that are used, after which we can sidestep the actual analysis and instead highlight the lessons that emerge from the torque-modeling exercise.

Readers who are comfortable with the distinction between transient and steady-state conditions and familiar with space phasors, transformation between reference frames, and the circuit modeling of electrical machines, may wish to skip this (inevitably rather long) treatment.

4.2 Transient and steady states in electric circuits

Field-oriented control allows us to obtain (almost) instantaneous (step) changes in torque on demand, and it does this by jumping directly from one steady-state condition to another. This simple statement is seldom given the prominence it deserves, but it is a simple truth, to be recalled whenever there is a danger of being bamboozled by a surfeit of technospeak.

Given the very poor inherent response of the induction motor to sudden changes in load or utility supply (see, for example, Figure 6.7), it will come as no surprise when we learn later that this sudden transition between steady states can only be achieved by precise control of the magnitude, frequency and instantaneous phase of the stator currents. As will emerge, the key requirement of a successful sudden transition is that it must not involve a step change in the stored energy of the system.

As an introduction to the underlying principle of changing from one steady state to another without any transient, we can look at the behavior of a series resistor and inductor circuit fed by an ideal voltage source (Figure 7.8). This is much simpler than the induction motor (it only has one energy storage element – the inductor) but it demonstrates the key requirement to be satisfied for transient-free switching.

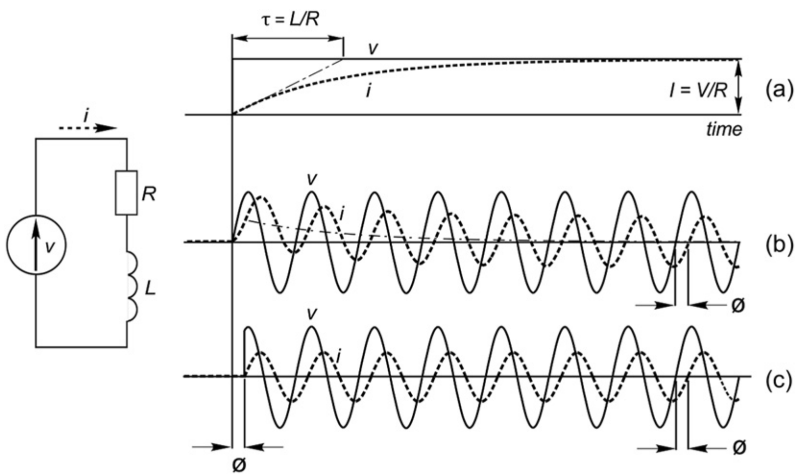

Figure 7.8 Transition between steady states in series *R–L* circuit.

First, we look at the current when the voltage is a step at $t = 0$ (Figure 7.8(a)). The steady-state current is simply V/R, but the current cannot rise instantaneously because that would require the energy stored in the inductor ($\frac{1}{2}Li^2$) to be supplied in zero time, which corresponds to an impulse of infinite power. So in addition to the steady-state term $i_{ss} = V/R$, there is a transient term given by $i_{tr} = -(V/R)e^{(-t/\tau)}$, where the time-constant, $\tau = L/R$. The total current is the sum of the steady-state and transient components, as shown by the lower dotted line in Figure 7.8(a).

Now consider a more relevant situation, where we wish the current to jump suddenly from a steady state at one frequency (in this case zero amplitude at zero frequency (i.e. d.c.) for $t < 0$) to a sinusoidal steady state for $t > 0$.

Figure 7.8(b) shows what happens if we make the sudden transition in the applied voltage (from zero d.c.) at a point where the new voltage waveform is zero but rising, i.e. at $t = 0$. We note that the current does not immediately assume its steady state, but displays the characteristic decaying transient, lasting for several cycles before the steady state is reached, with the current finally lagging the voltage by an angle ϕ. Examination of the steady-state current waveform shows that the current is negative as the voltage rises through zero, so this particular attempt to jump straight into the steady state is clearly doomed from the outset because it would have required the circuit to anticipate the arrival of the voltage by having a negative current already in existence!

The fundamental reason for the transient adjustment in Figure 7.8(b) is that we are seeking an instantaneous increase in the energy stored in the inductor from its initial value of zero, which is clearly impossible. It turns out that if we want to avoid the transient, we must make the jump without requiring a change in the stored energy, which in this example means at the point when the current passes through zero, as shown in Figure 7.8(c). The voltage that has to be applied therefore starts abruptly at a value $V \sin \phi$, as shown, and the current immediately enters its steady state, with no transient term.

We will see later that the principle of not disturbing the stored energy is the key requirement for obtaining step changes in torque from an induction motor.

4.3 Space phasor representation of m.m.f. waves

The space phasor (or space vector) provides a shorthand graphical way of representing sinusoidally distributed spatial quantities such as the m.m.f. and flux waves that we explored in Chapter 5. It avoids us having to consider individual currents by focusing on their combined effects, and thus makes things easier to understand.

We begin by taking a fresh look at the rotating stator m.m.f., making the reasonable assumption that each of the 3-phase windings produces a sinusoidally distributed m.m.f. with respect to distance around the air-gap, which in turn implies that the winding itself is sinusoidally distributed (rather than sitting in clearly defined groups of coils as in the real machine discussed previously). For convenience, we will consider a 2-pole winding.

We can represent the relative position of the windings *in space* as shown in Figure 7.9. In Figure 7.9(a) phases B and C are on open-circuit so that we can focus on how the m.m.f. of phase A is represented. When the current in phase A is positive (i.e. current flows into the dotted end), we have chosen to represent its sinusoidal m.m.f. pattern by a vector along the axis of the winding and pointing away from it (Figure 7.9(a)), and so when the current is negative the vector points towards the coil (Figure 7.9(b)). The length of the vector is directly proportional to the instantaneous value of the current, as indicated by the relative sizes of the two vectors.

In Figure 7.9(c), phase A has its maximum positive current, while phases B and C both have negative currents of half the maximum value. Because each m.m.f. is distributed sinusoidally in space, we can find their resultant (R) using the approach that is probably more familiar in the context of a.c. circuits, i.e. by adding the three components vectorially. In this particular example, the resultant m.m.f. (Figure 7.9(d)) is co-phasal with the m.m.f. of phase A, but one and half times larger.

We can now use the approach outlined above to find the resultant m.m.f. when the windings are supplied with balanced 3-phase currents of equal amplitude but

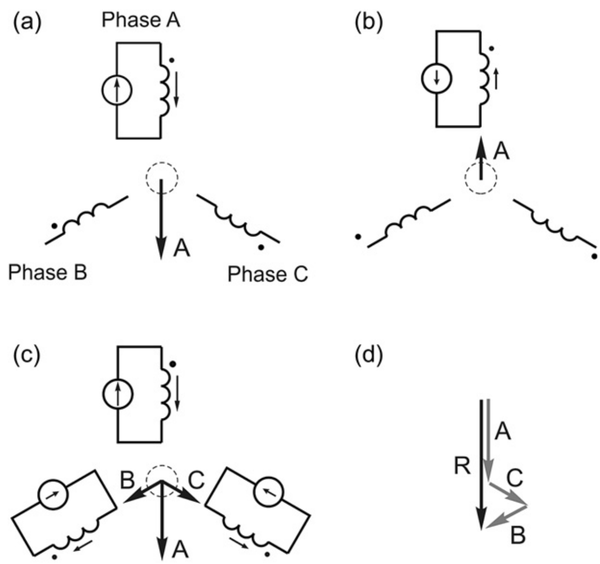

Figure 7.9 Space phasor representation of m.m.f. waves.

displaced in time by one-third of a cycle (i.e. 120°). The axes of the phases are displaced in *space* as shown in Figure 7.9, and the three currents are shown as functions of *time* in the upper part of Figure 7.10. Four consecutive times are identified, separated by one-twelfth of a complete cycle, or 30° in angle terms.

The lower part of the diagram represents the m.m.f.s in a space phasor diagram. At each instant the three individual phase m.m.f.s are shown in magnitude and position, together with the resultant m.m.f. At time t_0, for example, the situation is the same as in Figure 7.9, with phase A at maximum positive current and phases B and C having equal negative currents of half the magnitude of that in phase A; at t_1 phase B is zero while phases A and C have equal but opposite currents; and so on.

The four sketches suggest that the resultant m.m.f. describes an arc of constant radius, and it can easily be shown analytically that this is true. So although each phase produces a pulsating m.m.f. along its axis, the overall m.m.f. is constant in amplitude (with a value equal to 1.5 times the phase peak), and it rotates at a uniform rate, completing one revolution per cycle if the field is 2-pole (as here), half a revolution if 4-pole, etc. This is in line with our findings in Chapter 5.

We should note that although we have developed the idea of space phasors by focusing on steady-state sinusoidal operation, the approach is equally valid for any set of instantaneous currents, and is therefore applicable under transient conditions, for example during acceleration when the instantaneous frequency of the currents may change continuously.

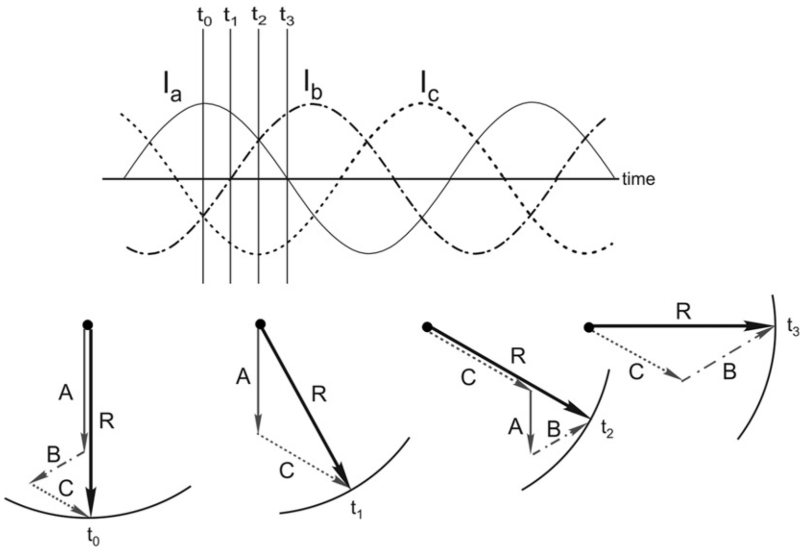

Figure 7.10 Resultant m.m.f. space phasor for balanced 3-phase operation at four discrete times, each separated by one-twelfth of a cycle (i.e. 30°).

An alternative way of representing the *resultant m.m.f. pattern* produced by a set of balanced 3-phase windings follows naturally from the discussion above. We imagine a hypothetical *single m.m.f. vector* that has a constant magnitude but *rotates relative to the stator* at the synchronous speed. This turns out to be an exceptionally useful mental picture when we come to study the behavior of the inverter-fed induction motor, because the currents are then under our control and we are able to specify precisely the magnitude, speed and angular position of the stator m.m.f. vector in order to achieve precise control of torque.

4.4 Transformation of reference frames

In the previous section we saw that the resultant m.m.f. was of constant amplitude and rotated at a constant angular velocity with respect to a reference frame fixed to the stator. As far as an observer in the stationary reference frame is concerned, the same m.m.f. could equally well be produced by a set of sinusoidally distributed windings fed with constant (d.c.) current and mounted on a structure that rotated at the same angular velocity as the actual m.m.f. wave. On the other hand, if we were attached to a reference frame rotating with the m.m.f., the space phasor would clearly appear to us to be constant.

Transformations between reference frames have long been used to simplify the analysis of electrical machines, especially under dynamic conditions, but until fast signal-processing became available that approach was seldom used for live control

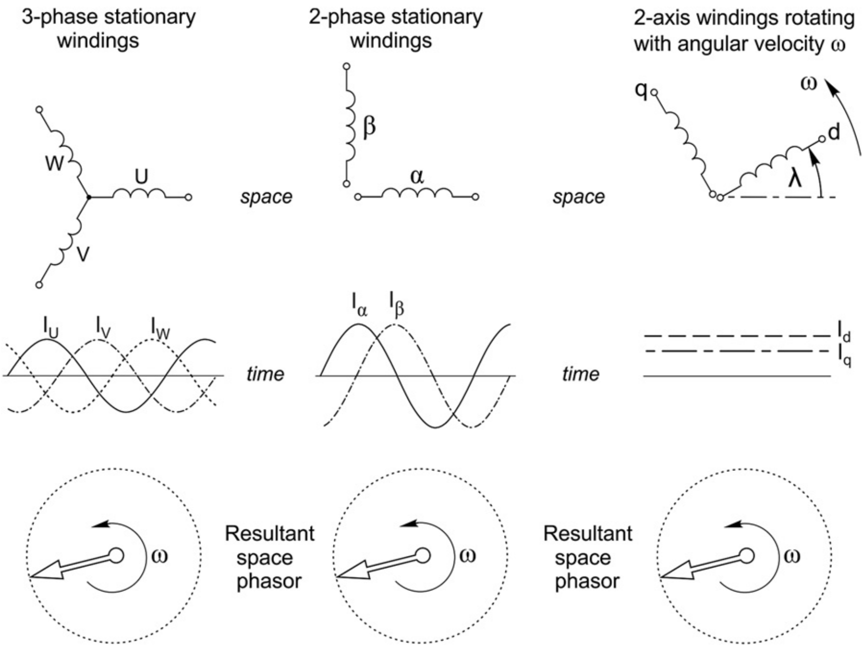

Figure 7.11 Transformation from 3-phase stationary axis reference frame to two-axis (d–q) rotating reference frame.

purposes. We will see later in this chapter that the method is used in field-oriented control schemes to transform the stator currents into a rotating reference frame locked to the rotating rotor flux space phasor, thereby making them amenable for control purposes.

Transformation is usually accomplished in two stages, as shown in Figure 7.11.

The first stage involves replacing the three windings by two in quadrature, with balanced sinusoidal currents of the same frequency but having a 90° phase shift. In this case the '$\propto\beta$' stationary reference frame has phase α aligned with phase U. To produce the same m.m.f., the two windings need either more turns or more current, or a combination of both. This is known as the Clarke transformation.

The second stage (the Park transformation) is more radical as the new variables I_d and I_q are in a rotating reference frame, and they remain constant under steady-state conditions, as shown in Figure 7.11. Again we need to specify the turns ratio and/or the current scaling. (Strictly speaking there is no need for the intermediate (2-phase) transformation, because we can transform directly from 3-phase to two-axis, but we have included it because it is often mentioned in the literature.)

It should be clear that the magnitude of the currents I_d and I_q will depend on the angle λ, which is the angle between the two reference frames at a specified instant, typically at $t = 0$. As far as we are concerned, it is sufficient to note that there are well-established formulae relating the input and output variables, both for the forward transformation (U, V, W to d, q) and for the inverse transformation, so it is straightforward to construct algorithms to perform the transformations. We should also note that while we have considered the transformation of sinusoidal currents, the technique is equally valid for instantaneous values.

4.5 Circuit modeling of the induction motor

Up to now in this book we have developed our understanding by starting with a physical picture of the interactions between the magnetic field and current-carrying conductors, but we quickly realized that in the case of the d.c. machine (and the utility-fed induction motor) there was a lot to be gained by making use of an 'equivalent circuit' model, particularly in terms of performance prediction. In so doing we were representing all the distributed interactions of the motor by way of their ultimate effect as manifested at the electrical terminals and the mechanical 'terminal', i.e. the output shaft.

As long ago as the early nineteenth century it was known that the a.c. transformer could be analyzed as a pair of magnetically linked coils, and it did not take long to show that all of the important types of a.c. electrical machine can also be analyzed by regarding them as a set of circuits, the electrical parameters (resistance, inductance) of which were either measured or calculated. The vital difference compared with the static transformer is that in the machine, the coils on the rotor move with respect to those on the stator, thereby causing a variation in the extent of the magnetic interaction between the rotor and stator. This variation turns out to be the essential requirement for the machine to produce torque and to be capable of energy conversion.

4.6 Coupled circuits, induced e.m.f. and flux linkage

By 'coupled circuits' we mean two or more circuits, often in the form of multi-turn coils sharing a magnetic circuit, where the magnetic flux produced by the current in one coil not only links with its own winding, but also with those of the other coils. The coupling medium is the magnetic field, and as we will see the electrical effect of the coupling is manifested when the flux changes.

We know from Faraday's law that when the magnetic flux (ϕ) linking a coil changes, an e.m.f. (e) is induced in the coil, given by

$$e = -N\frac{d\phi}{dt},$$

i.e. the e.m.f. is proportional to the rate of change of the flux. (The minus sign indicates that if the induced e.m.f. is allowed to drive a current, the m.m.f. produced by the current will be in opposition to that which produced the original changing flux.) This equation only applies if all the flux links all N turns of the coil, the situation most commonly approached in transformer windings that share a common magnetic circuit, and are thus very tightly coupled.

We have seen that windings for induction motors are distributed, and the flux wave produced by the current in the winding is also distributed around the air-gap. As a result not all of the flux produced by one winding links with all of its turns, and we have to perform a summation (integration) of all the 'turns times flux that links them' contributions to find the 'total effective self flux linkage', which we denote by the symbol psi (ψ). The e.m.f. induced when the self-produced flux linkage changes in, say, a stator winding (subscript S) is then given by

$$e_s = \frac{d\psi_S}{dt}$$

In an induction motor there are three distributed windings on the stator, and either a cage or three more distributed windings on the rotor, and some of the flux produced by current in any one of the windings will link all of the others. We term this 'mutual flux linkage', and often denote it by a double suffix: for example, the symbol ψ_{SR} is the mutual flux linkage between a stator winding and a rotor winding.

In the same way that an e.m.f. is induced in a winding when its self-produced flux changes, so also are e.m.f.s induced in all other windings that are mutually coupled to it. For example, if the flux produced by the stator winding changes, the e.m.f. in the rotor (subscript R) is given by

$$e_R = \frac{d\psi_{SR}}{dt}$$

4.7 Self and mutual inductance

The self and mutual flux linkages produced by a winding are proportional to the current in the winding. The ratio of flux linkage to the current that produces it is therefore a constant, and is defined as the inductance of the winding. The self inductance (L) is given by

$$L = \frac{\text{Self flux linkage}}{\text{Current}} = \frac{\psi_S}{i_S},$$

while the mutual inductance (M) is defined as

$$M_{SR} = \frac{\text{Mutual flux linkage}}{\text{Current}} = \frac{\psi_{SR}}{i_S}$$

The self and mutual inductances therefore depend on the design of the magnetic circuit and the layout of the windings.

We can now recast the expressions for e.m.f. derived above so that they involve the rates of change of the currents and the inductances, rather than the fluxes. This is a very important simplification because it allows us to represent the distributed effects of the magnetic coupling in single lumped-parameter electric circuit terms. The self-induced and mutually induced e.m.f.s are now given by

$$e_S = L\frac{di_S}{dt},$$

and

$$e_R = M_{SR}\frac{di_S}{dt}$$

4.8 Obtaining torque from a circuit model

We represent the two sets of 3-phase distributed windings of the induction motor by means of the six fictitious 'equivalent' coils shown in Figure 7.12. (We are using the well-proven fact that a cage rotor behaves in essentially the same way as one with a wound rotor, as explained in Chapter 5.) The three stator coils remain stationary, while those on the rotor obviously move when the angle θ changes.

Because the air-gap is smooth, and the rotor is assumed to be magnetically homogeneous, all the self inductances are independent of the rotor position, as are the mutual inductances between pairs of stator coils and between pairs of rotor coils. Symmetry also means that the mutuals between any two stator or rotor phases are the same.

However, it is obvious that the mutual inductance between a stator and a rotor winding will vary with the position of the rotor: when stator and rotor windings are aligned, the flux linkage will be maximum, and when they are positioned at right angles, the flux linkage will be zero. With windings that are distributed so as to produce sinusoidal flux waves, the mutual inductances vary sinusoidally with the angle θ.

To a circuit theory practitioner, it is this variation of mutual inductance with position that immediately signals that torque production is possible. In fact it is straightforward (if somewhat intellectually demanding) to show that the torque

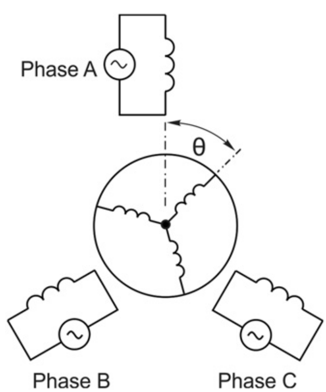

Figure 7.12 Coupled-circuit model of 3-phase induction motor.

produced when the sets of windings in Figure 7.12 carry currents is given by the rather fearsome-looking expression

$$T = \sum i_S i_R \frac{dM_{SR}}{d\theta}$$

What this means is that to find the total torque we have to find the summation of nine terms, each of which represents a contribution to the total torque from one of the nine stator–rotor pairs. So we need the instantaneous value of each of the six currents, and the rate of change of inductance with rotor position for each stator–rotor pair. For example, the term representing the contribution to torque made by stator coil A interacting with rotor coil B is given by

$$T_{SARB} = i_{SA} i_{RB} \frac{dM_{SARB}}{d\theta}$$

In practice we can use various expedients to simplify the torque expression, for example we know that mutual inductance is a reciprocal property, i.e. $M_{AB} = M_{BA}$, and we can exploit symmetry, but the important thing to note here is that it is a straightforward business to find the torque from the circuit model, provided that we know the currents and the inductances.

4.9 Finding the rotor currents

The rotor currents are induced, and to find them we have to solve the set of six equations relating them to the applied stator voltages, using Kirchhoff's voltage law.

So, for example, the voltage equation below relating to rotor phase A includes a term representing the resistive volt-drop, another representing the self-induced e.m.f. and five others representing the mutual coupling with the other windings. There are two more rotor equations and three similar ones for the stator windings.

$$v_{RA} = i_{RA} R_R + L_{RA} \frac{di_{RA}}{dt} + M_{RARB} \frac{di_{RB}}{dt} + M_{RARC} \frac{di_{RC}}{dt} + M_{RASA} \frac{di_{SA}}{dt} + M_{RASB} \frac{di_{SB}}{dt} + M_{RASC} \frac{di_{SC}}{dt}$$

In the induction motor the rotor windings are usually short-circuited, so there is no applied voltage and the left-hand side of each rotor equation is zero.

If we have to solve these six simultaneous differential equations when the stator terminal *voltages* are specified (typical of utility-fed constant-frequency conditions), we have a very challenging task that demands computer assistance, even under steady-state conditions. However, when the stator *currents* are specified (as we will see is the norm in an inverter-fed motor under vector control), the equations can be solved much more readily. Indeed under steady-state locked rotor conditions we can employ an armory of techniques such as *j* notation and phasor diagrams to solve the equations by hand.

We have now seen in principle how to predict the torque, and how to solve for the rotor currents, when the stator currents are specified. So we are now in a position to see what can be learned from a study of the known outcomes under two specific conditions.

In the next section, we look at how the torque varies when the stator windings are fed with a balanced set of 3-phase a.c. currents of constant amplitude but variable frequency, and the rotor is stationary. Although this is not of practical importance, it is very illuminating, and it points the way to the second and much more significant mode of operation, in which the net rotor flux is kept constant at all frequencies; this is dealt with in section 6.

5. STEADY-STATE TORQUE UNDER CURRENT-FED CONDITIONS

Historically there was little interest in analysis under current-fed conditions because we had no means of direct control over the stator currents, but the inverter-fed drive allows the stator currents to be forced very rapidly to whatever value we want, regardless of the induced e.m.f.s in the windings. Fortunately, we will see that knowing the currents from the outset makes quantifying the torque very much easier, and it also allows us to derive simple quantitative expressions that indicate how the machine should be controlled to achieve precise torque control, even under dynamic conditions.

To simplify our mental picture we will begin with the rotor at rest, and we will assume that we have a wound rotor with balanced 3-phase windings that for the moment are open-circuited, i.e. that no current can flow in them. With balanced 3-phase stator currents of amplitude I_s we know from the discussion above that the traveling stator m.m.f. wave can be represented by a single space phasor that rotates at the synchronous speed, and that in the absence of any currents in the rotor (and neglecting saturation of the iron) the flux wave will be in phase with the m.m.f. and proportional to it. This is shown Figure 7.13(a): in this sketch the rotor and stator are stationary, but all the patterns rotate at synchronous speed.

On the left of Figure 7.13(a) is a graphical representation of the sinusoidal distribution of resultant current around the stator at a given instant, and the corresponding flux pattern (dashed lines). Note that there is no rotor current. On the right of Figure 7.13(a) is a phasor that can represent both the stator m.m.f. and what we will call the resultant mutual flux linkage, both of which are proportional to the stator current. The expression 'mutual flux linkage' in Figure 7.13(a) thus represents the total effective flux linkages with the rotor due to the stator traveling flux wave. In circuit terms, this flux linkage is proportional to the mutual inductance between the stator and rotor windings (M), and to the stator current (I_S), i.e. MI_S.

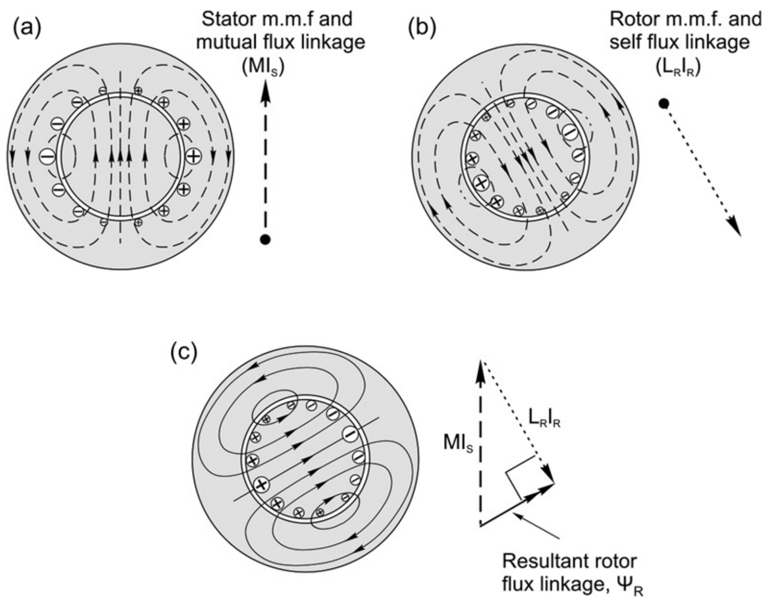

Figure 7.13 Space phasors of m.m.f. and flux linkage under locked-rotor conditions.

Now we short-circuit the rotor windings, and solve the set of equations for the rotor currents in the steady state. In view of the symmetry it comes as no surprise to find that they also form a balanced 3-phase set, at the same frequency as those of the stator, but displaced in time-phase. The resultant pattern of currents in the rotor is shown on the left of Figure 7.13(b), together with the flux pattern (dashed lines) that they would set up if they acted alone. Note that the stator currents that are responsible for inducing the rotor currents have been deliberately suppressed in Figure 7.13(b), because we want to highlight the rotor's reaction separately.

The m.m.f. due to the rotor currents is represented by the phasor shown on the right of Figure 7.13(b), and again this can also serve to represent the rotor self flux linkages ($L_R I_R$) attributable to the induced currents. It is clear that the time phase shift between stator and rotor currents causes a space phase shift between stator and rotor m.m.f.s, with the rotor m.m.f. broadly tending to oppose the stator m.m.f. If the rotor had zero resistance, the rotor m.m.f. would directly oppose that of the stator. The finite rotor resistance displaces the angle, as shown in Figure 7.13(b). We will see shortly that this phase angle varies widely and is determined by the frequency.

To find the resultant m.m.f. acting on the rotor we simply add the stator and rotor m.m.f. vectors, as shown in Figure 7.13(c). It is this m.m.f. that produces the resultant flux at the rotor, and we can therefore also use it to represent the net rotor flux linkage (ψ_R). The flux pattern at the rotor is shown by the solid lines in Figure 7.13(c, left). (But we should note that the number of flux lines shown in Figure 7.13(a–c) is not intended to reflect the relative magnitudes of flux densities, which, if saturation was not present, would be higher in the two upper sketches.) We should also note that, as expected, the behavior is independent of the rotor position angle θ, because the rotor symmetry means that when viewed from the stator, the rotor always looks the same overall.

Close examination of Figure 7.13(c) reveals an extremely important fact. The resultant rotor flux vector (ψ_R) is perpendicular to the rotor current vector. This means that the rotor current wave (shown on the left of Figure 7.13(c)) is oriented in the ideal position in space to maximize the torque production, because the largest current is coincident with the maximum flux density. If we look back to Figures 3.1 and 3.2, we will see that this is exactly how the flux and current are disposed in the d.c. machine, the N pole facing the positive currents and the S pole opposite the negative currents.

When we evaluate the torque under these conditions, a very simple analytical result emerges: the torque turns out to be given by the product of the rotor flux linkage and the rotor current, i.e.

$$T = \psi_R I_R$$

The similarity of this expression to the torque expression for a d.c. machine is self-evident, and further underlines the fundamental unity of machines exploiting the '*BIl*' mechanism discussed in Chapter 1. We note also that in Figure 7.13(c, right), one side of the right-angle triangle is ψ_R while the other is proportional to the rotor current I_R. Hence the area of the triangle is proportional to the torque, which makes it easy to visualize how torque varies with frequency, which we look at shortly.

(The keen reader may recall that the mental pictures we employed in Chapter 5 were based on the calculation of torque from the product of the air-gap flux wave and the rotor current wave, and that these were not in phase, except at very low slip frequency. The much simpler picture which has now been revealed – in which the flux and current waves are always ideally disposed as far as torque production is concerned – arises because we have chosen to focus on the resultant rotor flux linkage, not the air-gap flux: we are discussing the same mechanism as in Chapter 5, but the new viewpoint has thrown up a much simpler picture of torque production.)

We will see later that the rotor flux linkage is the central player in the field-oriented or vector control methods that now dominate in inverter-fed drives. To make full use of the flux-carrying capacity of the rotor iron, and to achieve step changes in torque, we will keep the amplitude of ψ_R constant, and we will explore this shortly. But first we will look at how the torque depends on slip when the amplitude of the stator current is kept constant.

5.1 Torque vs slip frequency – constant stator current

An alert reader might question why the title of this section includes reference to slip frequency, when we have specified so far that the rotor is stationary, in which case the effective slip is 1 and the frequency induced in the rotor will always be the same as the stator frequency. The reason for referring to slip frequency is that, as far as the reaction of the rotor is concerned, the only thing that matters is the relative speed of the traveling stator field with respect to the rotor.

So if we study the static model with an induced rotor frequency of 2 Hz, the torque that we predict can represent locked rotor conditions with 2 Hz on the stator; or the rotor running with a slip of 0.1 with 20 Hz on the stator; or a slip of 0.04 with 50 Hz on the stator; and so on. In short, under current-fed conditions, our model correctly predicts the rotor behavior, including the torque, when we supply the stator windings at the slip frequency. (Note that all other aspects of behavior on the stator side are not represented in this model.)

The variation in the flux linkage triangle with slip frequency, assuming that the amplitude of the stator current is constant, is shown in Figure 7.14. The locus of the resultant rotor flux linkage as the slip is varied is shown by the semi-circles. The right-hand side relates to low values of slip frequency, where the rotor self flux linkage is much less than the stator mutual flux linkage, so the resultant rotor flux is not much less than when the slip is zero. In other words, at low slips the presence of the rotor currents has little effect on the magnitude of the resultant flux, as we saw in Chapter 5. Low-slip operation is the norm in controlled drives. The left-hand drawing relates to high values of slip, where the large induced currents in the rotor lead to a rotor m.m.f. that is almost able to wipe out the stator m.m.f., leaving a very small resultant flux in the rotor. We will not be concerned with this end of the diagram in an inverter drive.

There is a simple formula for the angle ϕ, which is given by

$$\tan\phi = \omega_s \tau \tag{7.1}$$

where $\tau = L_R/R_R$, the rotor time-constant.

We noted earlier that the torque is proportional to the area of the triangle, so it should be clear that the peak torque is reached when the slip increases from the

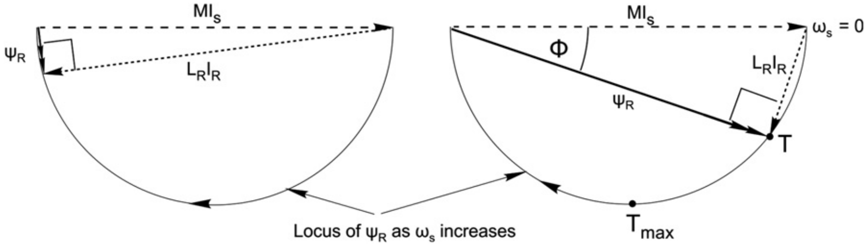

Figure 7.14 Locus of rotor flux linkage space phasor as slip frequency varies.

point T and moves to T_{max}. At this point, $\phi = \pi/4$ and the slip frequency is given by $\omega_S = 1/\tau = R_R/L_R$. Under these constant-current conditions, the slip at which maximum torque occurs is much less than under constant-voltage conditions, because the rotor self inductance is much larger than the rotor leakage inductance.

6. TORQUE VS SLIP FREQUENCY – CONSTANT ROTOR FLUX LINKAGE

As already mentioned, it is clear that to make full use of the flux-carrying capacity of the rotor iron, we will want to keep the amplitude of the rotor flux ψ_R constant. Given that the majority of the rotor flux links the stator (see Figure 7.13(c)), keeping the rotor flux constant also means that for most operating conditions, the stator flux is also more or less constant, as we assumed in Chapter 5.

In this section we explore how steady-state torque varies with slip when the rotor flux is maintained constant: this is illuminating, but much more importantly it prepares us for the final section, which deals with how we are able to achieve precise control of torque even under dynamic conditions.

We can see from Figure 7.14 that to keep the rotor flux constant we will have to increase the stator current with slip. This is illustrated graphically in Figure 7.15, in which the rotor flux linkage ψ_R is shown vertically to make it easier to see that it remains constant. In the left-hand diagram, the slip is very small, so the induced rotor current and the torque (which is proportional to the area of the triangle) are both small. The rotor flux is more or less in phase with the applied stator flux linkage because the 'opposing' influence of the rotor m.m.f. is small.

In the middle and right-hand diagrams the slip is progressively higher, so the induced rotor current is larger and the stator current has to increase in order to keep the rotor flux constant.

There is a simple analytical relationship that gives the required value of stator current as a function of slip, but of more interest is the relationship between the

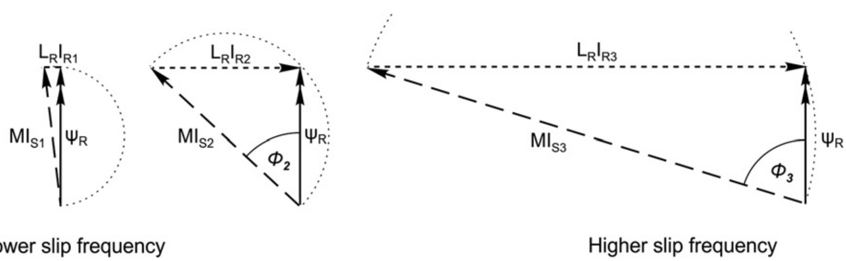

Figure 7.15 Constant rotor flux linkage space phasors at low, medium and high values of slip, showing variation of stator current required to keep rotor flux constant.

induced rotor current and the slip. From Figure 7.15, we can see that the tangent of the angle ϕ is given by

$$\tan\phi = \frac{L_R I_R}{\psi_R}$$

Combining this with equation (7.1) we find that the rotor current is given by

$$I_R = \left(\frac{\psi_R}{R_R}\right)\omega_{slip} \tag{7.2}$$

The bracketed term is constant; therefore the rotor current is directly proportional to the slip. Hence the horizontal sides of the triangles in Figure 7.15 are proportional to slip, and since the vertical side is constant, the area of each triangle (and torque) is also proportional to slip. To emphasize this simple relationship, the right-hand diagram in Figure 7.15 has been drawn to correspond to a slip three times higher than that of the middle one, so the horizontal side of the right-hand sketch is three times as long, and the area of the triangle (and torque) is trebled.

We note that when the rotor flux is maintained constant, the torque–speed curve becomes identical to that of the d.c. motor. In this respect the behavior differs markedly from that under both constant-voltage and constant-current conditions, where a peak or pull-out torque is reached at some value of slip. With constant rotor flux there is no theoretical limit to the torque, but in practice the maximum will be governed by thermal limits on the rotor and stator currents.

For those who prefer the physical viewpoint it is worth noting that the results discussed in this section could have been deduced directly from Figure 7.13(c), which indicates that the resultant rotor flux and rotor current waves are always aligned (i.e. the peak flux density coincides with the peak current density) so that if the flux is held constant, the torque is proportional to the rotor current. The rotor current is proportional to the motionally induced e.m.f., which in turn depends on the velocity of the flux wave relative to the rotor, i.e. the slip speed.

6.1 Flux and torque components of stator current

If we resolve the stator flux-linkage phasor MI_s into its components parallel and perpendicular to the rotor flux, the significance of the terms 'flux component' and 'torque component' of the stator current becomes obvious (Figure 7.16).

We can view the 'flux' component as being responsible for setting up the rotor flux, and this is the component that we must keep constant in order to maintain the working flux of the machine at a constant value for all slips. It is clearly analogous to the field current that sets up the flux in a d.c. motor.

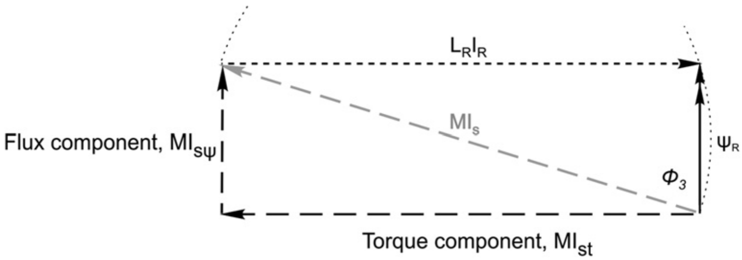

Figure 7.16 Flux and torque components of stator current.

The other ('torque') component (which is proportional to the rotor current) can be thought of as being responsible for nullifying the opposing effect of the rotor current that results when the rotor conductors are 'cut' by the traveling flux wave. This current component is therefore seen as the counterpart of the armature or work current in the d.c. motor.

Looking back to the left-hand diagram in Figure 7.15, we see that at small slips (light load) the stator current is small and practically in phase with the flux; this is what we referred to as the magnetizing current in previous chapters. At higher slips, the stator current is larger, reflecting that it now has a torque or 'work' component in addition to its magnetizing component, which again accords with our findings in previous chapters.

6.2 Establishing the flux

In the previous discussion we assumed that steady-state conditions prevailed, with the rotor flux wave remaining of constant magnitude and rotating relative to the rotor at the slip speed. We now look at how the flux wave was established, and we will see that the reaction of the rotor is very different from its subsequent steady-state behavior.

We start with the rotor at rest, no current in any of the windings, and hence no flux. With reference to Figure 7.9, we suppose that we supply a step (d.c.) current into phase A, which will split with half exiting from each of phases B and C, and producing a stationary sinusoidally distributed m.m.f. that, ultimately, will produce the flux pattern labeled 'final state' in Figure 7.17.

But of course the rotor windings are short-circuited, with no flux through them, and closed electrical circuits behave like many things in the physical world in that they react to change by opposing it. In this context, if the flux linking a winding changes, Faraday's law tells us there will be an induced e.m.f. The direction of the e.m.f. is such that if it acts in a closed circuit and produces a current, the m.m.f. produced by that current will be in opposition to the 'incoming' m.m.f./flux. (Formally this is expressed by Lenz's law, and sometimes by the use of a negative sign preceding the e.m.f. equation that quantifies Faraday's law.)

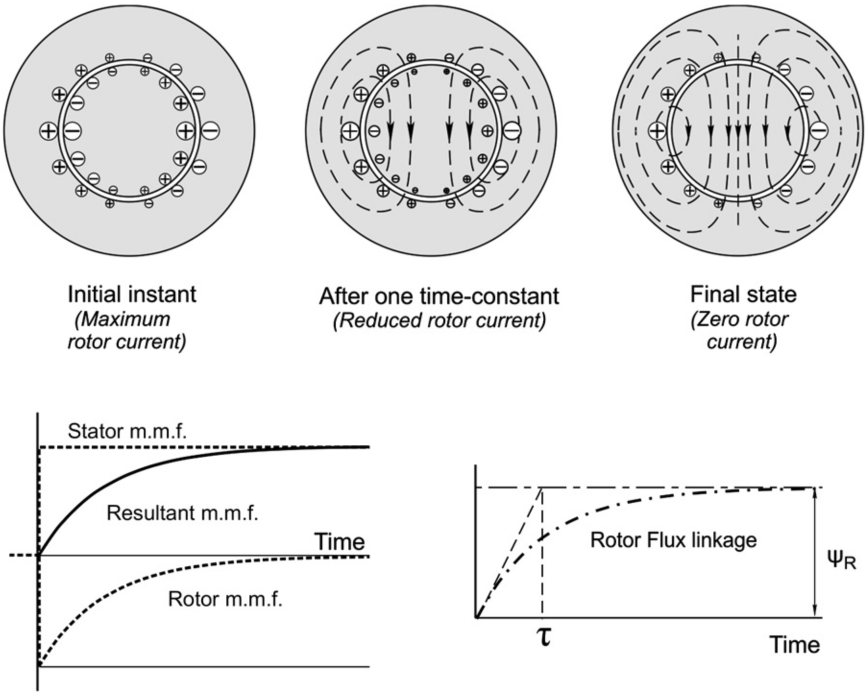

Figure 7.17 Diagrams illustrating the build-up of rotor flux when a step of stator m.m.f. occurs.

So when the stator m.m.f. phasor suddenly comes into existence, the immediate reaction of the rotor is the production of a negative stationary rotor m.m.f. pattern, i.e. in direct opposition to the stator m.m.f.: this is labeled 'initial instant' in Figure 7.17. Instantaneously, the magnitude of the rotor m.m.f. is such as to keep the rotor flux linkage at zero, as it was previously.

However, because of the rotor resistance, the rotor current needs a voltage to sustain it, and the voltage can only be induced if the flux changes. So the rotor flux begins to increase, rising rapidly at first (high e.m.f.) then with ever-decreasing gradient leading to lower current and lower rotor m.m.f. The response is a first-order one, governed by the rotor time-constant, so after one time-constant (middle sketch) the flux linking the rotor reaches about 63% of its final value, while the rotor current has fallen to 37% of its initial value. Finally, the rotor's struggle to prevent the flux changing comes to an end and the rotor flux linkage reaches a steady value determined by the stator current. If the resultant rotor flux linkage is the target value for steady-state running (ψ_R), the corresponding stator current is what we previously referred to as the 'flux component'.

The physical reason why it takes time to build the flux is that energy is stored in a magnetic field, so we cannot suddenly produce a field because that would require an impulse of infinite power. If we want to build up the flux more rapidly, we can put in a bigger step of stator current at first, so that the flux heads for a higher final value than we really need, and then reduce the current when we get close to the flux we are seeking.

We began this section with d.c. current in the stator, which in effect corresponds to zero slip frequency, all the field patterns being stationary in space. Because there is no relative motion involved, there is no motional e.m.f. and hence no torque. The 'torque component' only comes into play when there is relative motion between the rotor and the rotor flux wave, i.e. when there is slip. Obviously, to cause rotation the frequency must be increased, and as we have seen in the previous section the stator current then has to be adjusted with slip and torque to keep the rotor flux linkage constant.

Finally, it is worth revisiting Figure 7.16 briefly to reconcile what we have discussed in this section with our picture of steady-state operation, where the rotor currents are at slip frequency. On the left we have the fictitious 'flux component' of stator current, which remains constant in magnitude and aligned with the rotor flux linkage, ψ_R, along the so-called direct axis. When we first established this flux, the rotor reacted as we have discussed above, but after a few time-constants the flux settled to a constant value along the direct axis in the direction of the flux component of the stator m.m.f. This is why the arrows on ψ_R point in the same direction as the stator flux producing component.

However, we note from Figure 7.16 that the rotor flux linkage phasor ($L_R I_R$) is always equal in magnitude to the torque component of the stator mutual flux linkage phasor, but, as shown by the arrows, it is in the opposite direction. There is therefore no resultant m.m.f. or tendency for flux to develop along this, the so-called quadrature axis. This is what we would expect in the light of the previous discussion, where we saw that the reaction of mutually coupled windings to any suggestion of change is for currents to spring up so as to oppose the change. In the literature when, as here, the 'torque' current does not affect the flux, the axes are said to be 'decoupled'.

7. DYNAMIC TORQUE CONTROL

If we want to obtain a step increase in torque, we have to change the rotor flux or the rotor current instantaneously, so as to jump instantaneously from one steady-state operating condition to another. We have seen above that because a magnetic field has stored energy associated with it, it is not possible to change the rotor flux linkage instantaneously. In the case of the induction motor, any change in the rotor flux is governed by the rotor time-constant, which will be as much as 0.25 seconds for even a modest motor of a few kW rating, and much longer for large motors. This is not acceptable when we are seeking instantaneous changes in torque.

The alternative is to keep the flux constant, and make the rotor current change as quickly as possible.

In the previous section, our aim was to grow the rotor flux, which, because of its stored energy, took a while to reach the steady state. However, if we keep the rotor flux linkage constant (by ensuring that the flux component of the stator current is constant and aligned with the flux) we can cause sudden changes to the motionally induced rotor current by making sudden changes in the torque component of the stator current.

We achieve sudden step changes in the stator currents by means of a fast-acting closed-loop current controller. Fortunately, under transient conditions the effective inductance looking in at the stator is quite small (it is equal to the leakage inductance), so it is possible to obtain very rapid changes in the stator currents by applying high, short-duration impulsive voltages to the stator windings. In this respect the stator current controller closely resembles the armature current controller used in the d.c. drive.

When a step change in torque is required the magnitude, frequency, and phase of the stator currents are changed (almost) instantaneously in such a way that the rotor current jumps suddenly from one steady state to another. But in this transition it is only the torque component of stator current that is changed, leaving the flux component aligned with the rotor flux. There is therefore no change in the magnitude of the rotor flux wave and no change in the stored energy in the field, so the change can be accomplished almost instantaneously.

We can picture what happens by asking what we would see if we were able to observe the stator m.m.f. wave at the instant that a step increase in torque was demanded. For the sake of simplicity, we will assume that the rotor speed remains constant, and consider an increase in torque by a factor of three (as between the middle and right-hand sketches in Figure 7.15), in which case we would find that:

- the stator m.m.f. wave suddenly increases its amplitude;
- it suddenly accelerates to a new synchronous speed, so the slip increases by a factor of three;
- it jumps forward to retain its correct relative phase with respect to the rotor flux; i.e. the angle between the stator m.m.f. and the rotor flux increases from ϕ_2 to ϕ_3.

Thereafter the stator m.m.f. retains its new amplitude, and rotates at its new speed.

The rotor current experiences a step change from a steady state at its initial slip frequency to a new steady state with three times the amplitude and frequency, and there is a step increase in torque by a factor of three, as shown in Figure 7.18(a). The new current is maintained by the new (higher) stator currents and slip frequency.

We should note particularly that it is the jump in the *angular position* (i.e. *space phase angle*) that accompanies the step changes in magnitude and frequency of the stator m.m.f. phasor that allows the very rapid and transient-free control of torque. Given that the definition of a vector quantity is one which has magnitude and

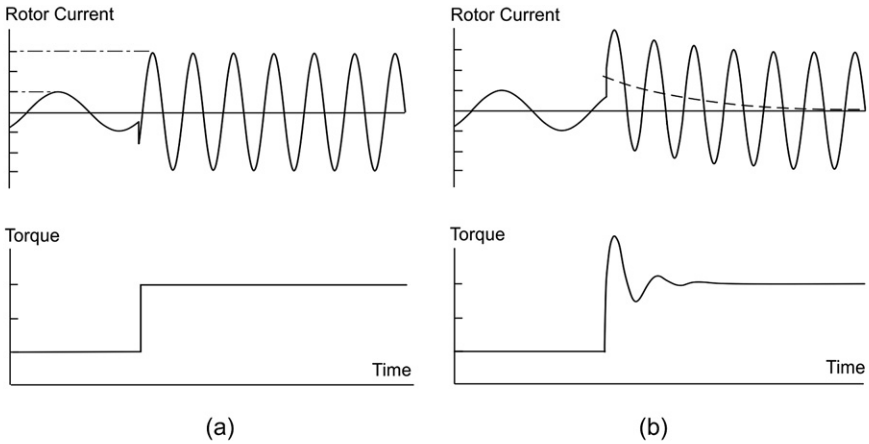

Figure 7.18 Step changes in rotor current. (a) Transient-free transition with correct changes to magnitude, frequency and instantaneous position of stator m.m.f. wave (i.e. vector control). (b) Same changes to magnitude and frequency, but not phase.

direction, and that the angular position of the phasor defines the direction in which it is pointing, it is clear why this technique gained the name 'vector control'.[2]

To underline the importance of the sudden change in *phase* of the stator current (i.e. the sudden jump in angular position of the stator m.m.f. in achieving a step of torque), Figure 7.18(b) shows what happens typically if only the magnitude and frequency, but not the position, are suddenly changed. The steady-state conditions are ultimately reached, but only after an undesirable transient governed by the (long) rotor time-constant, which may persist for several cycles at the slip frequency. The fundamental reason for the transient is that if the magnitude of the stator current is suddenly increased without a change of position, the flux and torque components both increase proportionately. The change in the flux component portends a change in the rotor flux (and associated stored energy), which in turn is resisted by induced rotor currents until they decay and the steady state is reached.

7.1 Summary

This section has described the underlying principles by which very rapid and precise torque control can be achieved from an induction motor, but we should remember that until sophisticated power-electronic control became possible the approach outlined here was only of academic interest. The fact that the modern inverter-fed drive is able to implement torque control and achieve such outstandingly impressive

[2] The term vector control has sometimes been misused to refer to drives that do not include field orientation. However, the term is so ubiquitous that we cannot avoid it, so when we refer to 'vector control' we mean a proper field-oriented system.

performance from a motor whose inherent transient behavior is poor, represents a major milestone in the already impressive history of the induction motor. The way in which such drives achieve field-oriented control is discussed next.

8. IMPLEMENTATION OF FIELD-ORIENTED CONTROL

An essential requirement if we are to unravel the workings of the overall scheme for field-oriented control is an understanding of the pulse-width modulation (PWM) vector modulator/inverter combination that is a feature of all such schemes, so this is covered first.

8.1 PWM controller/vector modulator

In the inverters we have looked at so far (see section 4 of Chapter 2) we have supposed that the periodic switching required to approximate a sinusoidal output was provided from a master oscillator. The frequency of the oscillator determined the frequency of the a.c. voltage applied to the motor, and the amplitude was controlled separately. In terms of space phasors this allows control of the amplitude and frequency, but not the instantaneous angular position of the voltage and current phasors. As we have seen, it is the additional ability to make instantaneous changes to the *angular position* of the output phasor that is the key to dynamic torque control, and this is the key feature provided by the 'vector modulator'.

We now explore what the inverter can produce in terms of its output voltage phasor. We recall that there are six devices (switches) in three legs (see Figure 2.17), and to avoid a short-circuit both switches in one leg must not be turned on at the same time. If we make the further restriction that each phase winding must at all times be connected to one or other of the d.c. link terminals, there are only eight possible combinations, as shown in Figure 7.19.

The six switching combinations labeled 1–6 each produce an output voltage phasor of equal amplitude but displaced in phase by 60° as shown in the lower part of each diagram, while the final two combinations have all three terminals joined together so the voltage is zero. The six unit vectors are shown with their correct relative phase, but rotated so as to bring U6 horizontal, in Figure 7.20.

Having only six states of the voltage phasor at our disposal is clearly not satisfactory, because we need to exert precise control over the magnitude and position of the voltage phasor at any instant, so this is where the 'time modulation' aspect comes into play. For example, if we switch rapidly between states U1 and U6, spending the same amount of time with each, we will effectively have synthesized a voltage phasor lying half way between them, and of magnitude U1 cos 30° (or 86.7% of U1), as shown by the vector U(x) in the upper-right part of Figure 7.20. If we spend a higher proportion of the time on U1 and the remainder on U2, we could produce the vector U(y). As long as we spend

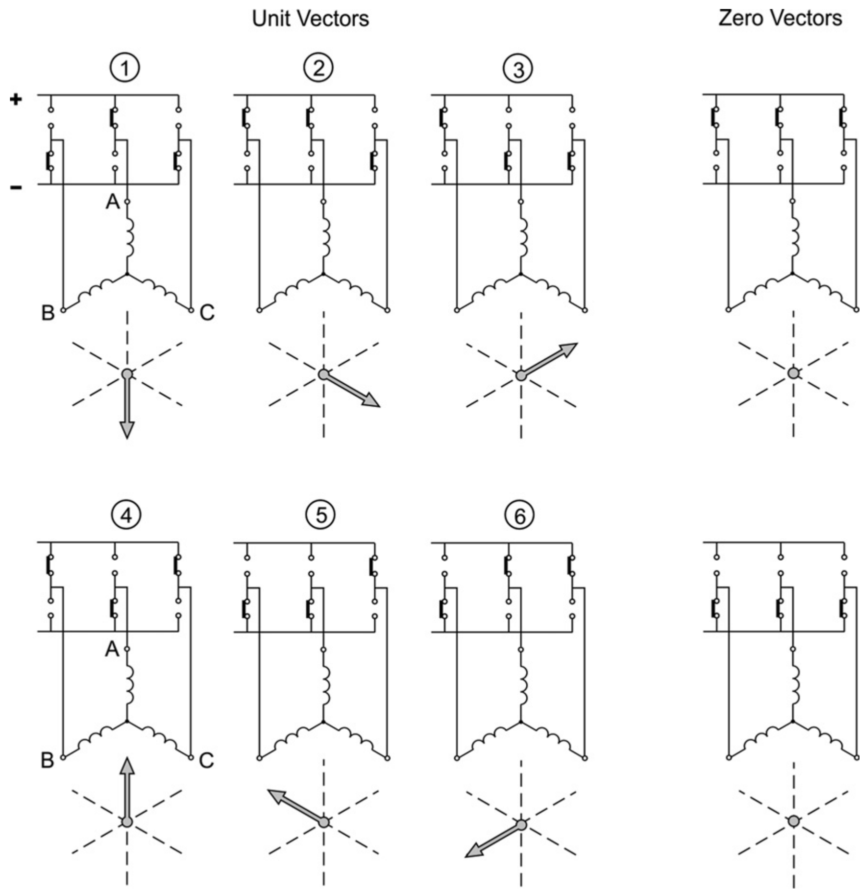

Figure 7.19 Voltage phasors for all acceptable combinations of switching for a 3-phase inverter.

the whole of the sample time on either U1 or U6, we will end up somewhere along the line joining U1 to U6.

We have used the terms 'switch rapidly' and 'the time' without specifying what they mean. In practice, we would expect the modulating frequency to be perhaps a few kHz up to the low tens of kHz, so 'the time' means one cycle at this frequency, say 100 microseconds at 10 kHz. So for as long as we wished the voltage phasor to remain at U(x), we would spend 50 μs of each sample period alternately connected to U1 and U6.

Recalling that ideally we want to be able to choose both magnitude and position it is clearly not satisfactory to be constrained to the outer edges of the hexagon. So now we bring the zero vector into play. For example, suppose we wish the voltage phasor to be U(z), as in the lower sketch. This is composed of (0.5)U5

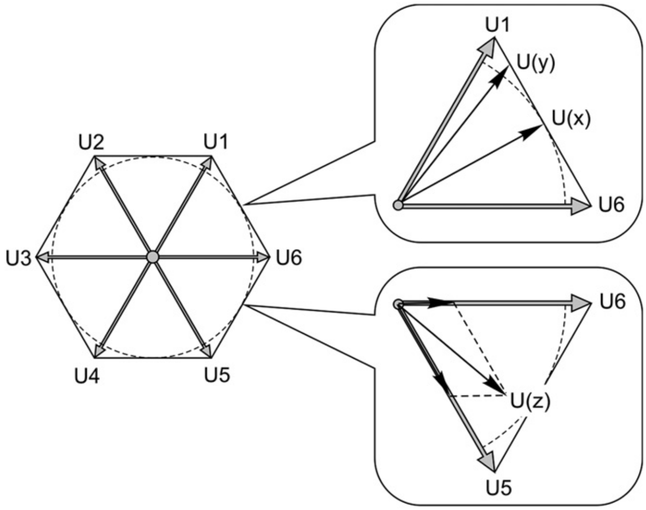

Figure 7.20 Synthesis of intermediate voltage phasors in vector modulator.

plus (0.3)U6. Hence in each modulating cycle of 100 μs, we will spend times of 50 μs on U5, 30 μs on U6 and 20 μs with one of the zero states.

The precise way in which these periods are divided within one cycle of the modulating frequency is a matter of important detail in relation to the distribution and minimization of losses between the six switching devices, but need not concern us here. Suffice it to say that it is a straightforward matter to arrange for digital software/hardware that has input signals representing the magnitude and instantaneous position of the output voltage phasor, and which selects and modulates the six switches appropriately, to create the desired output until told to move to a new location.

When we introduced the idea of space phasors earlier in this chapter, we saw that if we begin with balanced 3-phase sinusoidal voltages, the voltage phasor is of constant length and rotates at a uniform rate. Looking at it the other way round, it should be clear that if we arrange for the output of the inverter to be a voltage phasor of constant length, rotating at a constant rate, then the corresponding phase voltages must form a balanced sinusoidal set, which is what we want for steady-state running.

We conclude that in the steady state, the magnitude of the input signal to the vector modulator would have a constant amplitude and its angle would increase at a linear rate corresponding to the desired angular velocity of the output. Clearly in order to avoid having to deal with ever-increasing angles, the input signal to the modulator will reset each time a full cycle of 360° is reached, as shown in Figure 7.21.

Away from the steady-state condition, for example during acceleration, we should recall that to preserve the linear relation between torque and the stator current component (I_T), the flux component of the stator current phasor (I_F) must

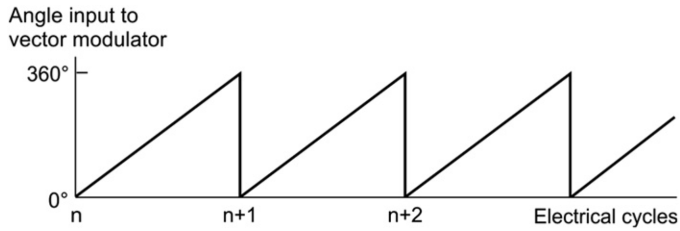

Figure 7.21 Angle reference to vector modulator corresponding to constant frequency operation of inverter. Note that the angle resets to zero at the end of each cycle.

remain aligned with the rotor flux. As we will see in the next section, this is achieved by deriving the angle input to the vector modulator directly from the absolute angular position of the rotor flux.

8.2 Torque control scheme

A simplified block diagram of a typical field-oriented torque control system is shown in Figure 7.22.

The first and most important fact to bear in mind in the discussion that follows is that Figure 7.22 represents a *torque* control scheme, and that for applications that require speed control, it will form the 'inner loop' of a closed-loop speed control scheme. The torque and flux inputs will therefore be outputs from the speed controller, as indicated in Figure 7.23.

Returning to Figure 7.22 it has to be acknowledged that it looks rather daunting, and getting to grips with it is not for the faint-hearted. However, if we examine it a bit at a time, it should be possible to grasp the essential features of its

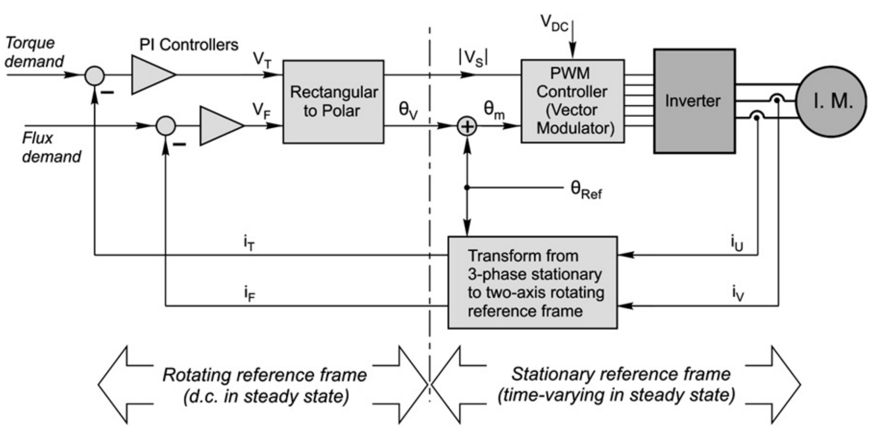

Figure 7.22 Simplified block diagram of a typical field-oriented torque control system.

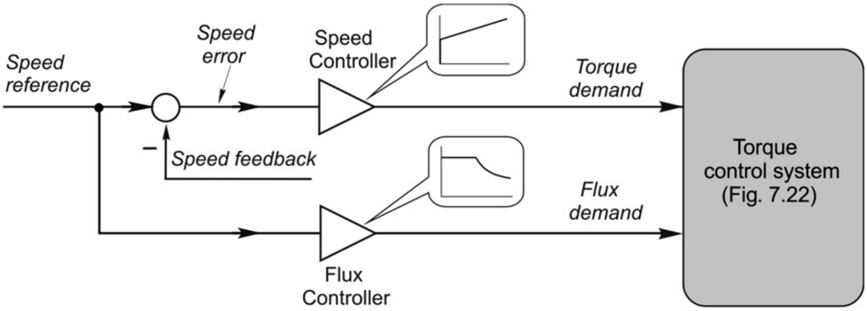

Figure 7.23 Schematic diagram of closed-loop speed control system.

operation. To simplify matters, we will focus on steady-state conditions, despite the fact that the real merit of the system lies in its ability to provide precise torque control even under transient conditions.

Taking the broad overview first, we can see that there are similarities with the d.c. drive with its inner current (torque) control loop (see Chapter 4), notably the stator current feedback and the use of proportional and integral (PI) controllers to control the torque and flux components of the stator current. It would be good if we could measure the flux and torque components directly, but of course the current components do not have separate existences: they are merely components of the stator current, which is what we can measure. The motor has three phases, but because we are assuming that there is no neutral connection, it suffices to measure only two of the line currents (because the sum of the three is zero). The information from these two currents allows us to keep track of the angular position of the stator current phasor with respect to the stationary reference frame (θ_S) as shown in Figure 7.24.

However, the stator current feedback signals are alternating at the frequency supplied by the inverter, and the corresponding stator space phasor is rotating at the

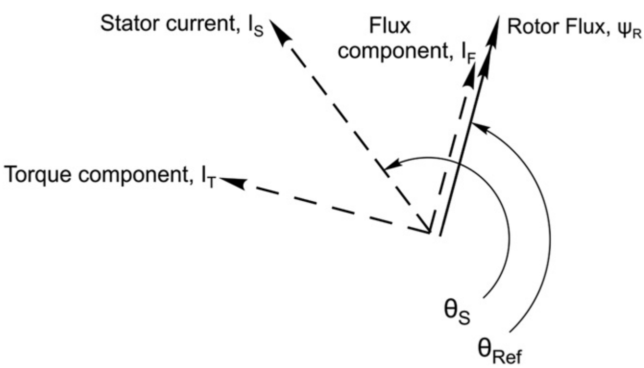

Figure 7.24 Stator current and rotor flux reference angles.

supply frequency with respect to a stationary reference frame. Before the flux and torque components of these signals (I_F and I_T) can be identified (and subsequently fed back to the PI controllers) they must first be transformed (see section 4) into a reference frame that rotates with the rotor flux. As explained previously, the rotor flux angle θ_{Ref} is therefore an essential input to the transformation algorithm, as shown in Figure 7.22.

The broken line in the middle of Figure 7.22 separates quantities defined in the stationary reference frame (on the right) from those in the rotating reference frame (on the left). In the steady state, all those on the left are d.c., while all those on the right are time-varying.

The reader might wonder why, when we follow the signal path of the current control loops, beginning on the right with the phase current transducers, there is no matching 'inverse transform' to get us back from the rotating reference frame on the left to the stationary reference frame on the right. The answer lies in the nature of the input signal to the PWM/vector modulator and inverter, which we discussed above. Let us suppose that the motor is running in the steady state, so that the output voltage phasor rotates at a constant rate with angular frequency ω. Under these conditions the rotor flux phasor also rotates with constant angular velocity ω, so the angle of the flux vector with respect to the stationary reference frame (θ_{Ref}) increases linearly with time. Also, in the steady state, the output from the PI controllers is constant, so the angle θ_V (Figure 7.22) is constant. Hence the input angle to the modulator (θ_m in Figure 7.22), which is the sum of θ_{Ref} and θ_V, is also a ramp in time, and this is what provides the rotation of the output voltage phasor. In effect, the system is self-sustaining: the primary time-varying input angle to the modulator comes from the flux position signal (which is already in the stationary reference frame), and the PI controller provides the required magnitude signal ($|V_S|$) and the additional angle (θ_V).

Turning now to the action of the PI controllers, we see from Figure 7.22 that the outputs are voltage commands in response to the differences between the feedback (actual) values of the transformed currents and their demanded values. The flux demand will usually be constant up to base speed, while the torque demand will usually be the output from the speed or position controller, as shown in Figure 7.23. The proportional term gives an immediate response to an error, while the integral term ensures that the steady-state error is zero. The outputs from the two PI controllers (which are in the form of quadrature voltage demands, V_F and V_T) are then converted from rectangular to polar form, to produce amplitude and phase signals, $|V_S|$ and θ_V, where

$$|V_S| = \sqrt{V_F^2 + V_T^2}, \text{ and } \theta_V = \tan^{-1}\frac{V_T}{V_F},$$

as shown in Figure 7.25.

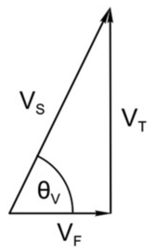

Figure 7.25 Derivation of voltage phasor from flux and torque components.

The amplitude term specifies the magnitude of the output voltage phasor (and thus the three phase voltages applied to the motor), with any variation of the d.c. link voltage (V_{dc}) being compensated in the PWM controller. The phase angle (θ_V) represents the desired angle between the stator voltage phasor and the rotor flux phasor, both of which are measured in the stationary reference frame. The angle of the rotor flux phasor is θ_{Ref}, so θ_V is added at the input to the vector modulator to yield the stator voltage phasor angle θ_m, as shown in Figure 7.22.

We can usefully conclude our look at the steady state by adding the stator voltage phasor to Figure 7.24 to produce Figure 7.26, to provide reassurance that, in the steady state, the rather different approach we have taken in this section is consistent with the classical approach taken earlier.

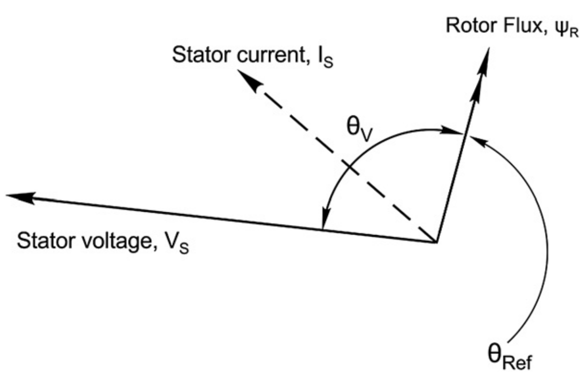

Figure 7.26 Time phasor diagram showing stator voltage and current under steady-state conditions.

8.3 Transient operation

We concluded earlier that for the motor torque to be directly proportional to the torque component of stator current, it is necessary to keep the magnitude of the rotor flux constant and to ensure that the flux component of stator current is aligned with the rotor flux. This is achieved automatically because the principal angle input to the vector modulator comes directly from the rotor flux angle (θ_{Ref}), as shown in Figure 7.22. So during acceleration, for example, the instantaneous angular velocity of the rotor flux wave will remain in step with that of the stator current phasor, so that there is no possibility of the two waves falling out of synchronism with one another.

In section 7 we discussed a specific example of how to obtain a step change in torque by making near-instantaneous changes to the magnitude, speed and position of the stator m.m.f. wave, and we are now in a position to see how this particular strategy is effected using the control scheme shown in Figure 7.22.

A step demand for torque causes a step increase in $|V_S|$ and θ_V at the output of the rectangular to polar converter in order to effect a very rapid increase in the magnitude and instantaneous position of the stator current phasor. At the same time, the algorithm that calculates the slip velocity of the flux wave (see later, equation (7.3)) yields a step increase because of the sudden increase in the torque component of stator current. The principal angular input to the vector modulator (the flux angle (θ_{Ref})) therefore changes gradient abruptly, as shown in Figure 7.27.

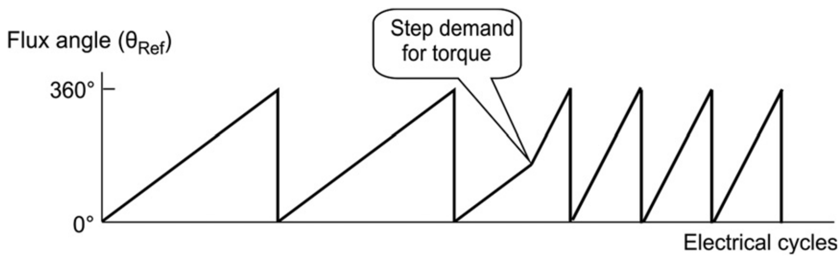

Figure 7.27 Flux angle reference showing response to a sudden step demand for increased torque.

Recalling that the steady-state stator frequency is governed by the angular velocity of the flux (i.e. $d\theta_{Ref}/dt$), this lines up with our expectation that (assuming the rotor velocity is constant) the stator frequency will increase in order to increase the slip and provide the new higher torque.

8.4 Acceleration from rest

Measured results showing how the real and transformed currents behave during acceleration from rest to base speed, are shown in Figure 7.28. These relate to a motor whose rotor time-constant is approximately 0.1 s, and cover a time of 0.5 s.

The upper diagram shows the demanded values for the transformed flux and torque components of stator current; the middle diagram shows the measured (actual) flux and torque components; and the lower shows the three phase currents.

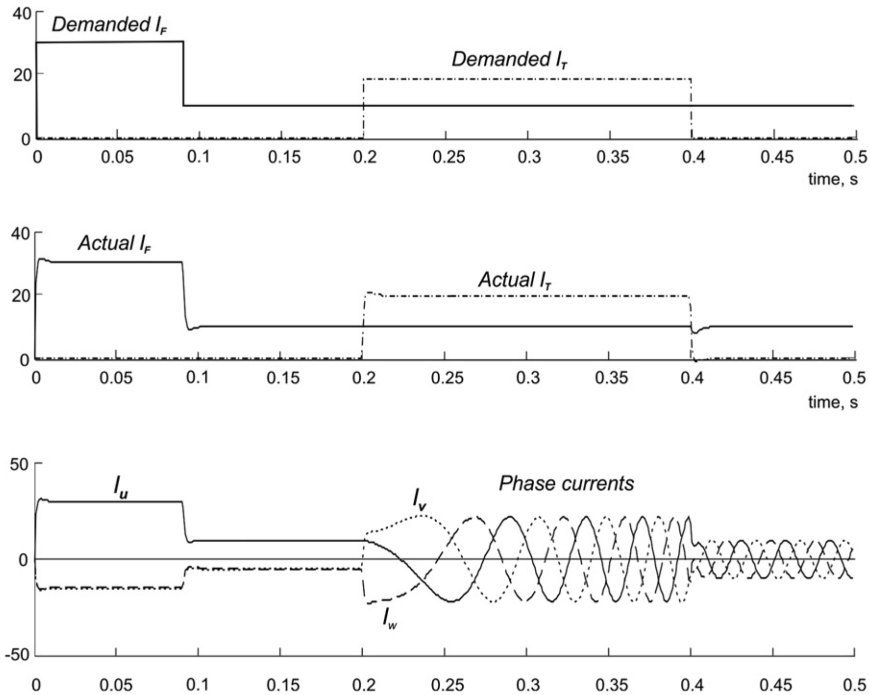

Figure 7.28 Experimental results showing build-up of flux followed by sudden demand for step increase in torque until motor reaches its target speed. *(Courtesy of Emerson – Control Techniques)*

At $t = 0$, the flux demand signal is set to its highest possible level of 30 A in order to raise the flux from zero as quickly as possible to its target value (10 A). Note how the actual transformed flux component of stator current (I_T) follows the demand signal very rapidly, with only slight overshoot. The signal I_T is the transformed version of the phase winding currents, so the fact that I_T is on target is of course indicative that the phase currents are established rapidly and held while the flux builds up, as we can see in the lower figure. During this period phase U carries a positive current of 30 A while phases V and W each carry a negative current of 15 A.

After about one time-constant the demand is reduced to 10 A, and thereafter held constant. This 'rapid forcing' ensures that by 0.2 s, the rotor flux has been fully established.

At 0.2 s, a torque producing demand is applied to accelerate the motor, and maintained until 0.4 s, when the torque producing reference is reduced to zero, and the motor stops accelerating.

We note the almost immediate and transient-free transition of the 3-phase currents from their initial steady (d.c.) values immediately prior to 0.2 s, into constant amplitude, 'smoothly increasing frequency' a.c. currents over the next 0.2 s. And then there is a similarly near-perfect transition to reduced amplitude steady-state conditions (at about 40 Hz) after 0.4 s. In the steady state, the torque component is negligible because the motor is unloaded, and the stator current consists only of the flux component, which traditionally would be referred to as the magnetizing current.

Younger readers will doubtless not require convincing of the validity of these remarkable results, but they might find it salutary to know that until the 1970s it was widely believed that such performance would never be possible.

To conclude this section we can draw a further parallel between the field-oriented induction motor and the d.c. motor. We see from Figure 7.28 that as the motor accelerates, the frequency of the stator currents increases with the speed. If we stationed ourselves on the rotor of a d.c. motor as it accelerated, the rate at which the current in each rotor coil reversed as it was commutated would also increase in proportion to the speed, though of course we are not aware of it when we are in the stationary reference frame.

8.5 Deriving the rotor flux angle

By now, the key role played by the rotor flux angle should have become clear, so finally we look at how it is obtained. It is not practical or economic to fit a flux sensor to the motor, so industrial control schemes invariably estimate the position of the flux.

We will first establish an expression for absolute rotor flux angle (θ_{Ref}) in the stationary reference frame in terms of quantities that can either be measured or estimated. Readers who find the derivation indigestible need not worry as it is the conclusions that are important, not the analytical detail.

If we let the angle of the rotor body with respect to the stationary reference frame be θ, then the instantaneous angular velocities of the rotor flux wave and the rotor itself are given by

$$\omega_{\text{flux}} = \frac{\mathrm{d}\theta_{\mathrm{Ref}}}{\mathrm{d}t}$$

$$\omega_{\text{rotor}} = \frac{d\theta}{\mathrm{d}t}$$

The rotor motional e.m.f. is directly proportional to the rotor flux linkage and the slip velocity, i.e.

$$V_{\mathrm{R}} = \psi_{\mathrm{R}}\left(\omega_{\text{flux}} - \omega_{\text{rotor}}\right),$$

and the rotor current is therefore given by

$$I_{\mathrm{R}} = \frac{\psi_{\mathrm{R}}\left(\omega_{\text{flux}} - \omega_{\text{rotor}}\right)}{R_{\mathrm{R}}}$$

The corresponding component of stator current is given (see Figure 7.16) by

$$I_{\mathrm{ST}} = \frac{L_{\mathrm{R}}}{M} I_{\mathrm{R}}$$

Combining these equations and rearranging gives

$$\frac{\mathrm{d}\theta_{\mathrm{Ref}}}{\mathrm{d}t} = \frac{MR_{\mathrm{R}}}{\psi_{\mathrm{R}} L_{\mathrm{R}}} I_{\mathrm{ST}} + \omega_{\mathrm{rotor}} = \left(\frac{M}{\tau \psi_{\mathrm{R}}}\right) I_{\mathrm{ST}} + \omega_{\mathrm{rotor}} \tag{7.3}$$

where τ is the rotor time-constant. Hence to find the rotor flux angle at time t we must integrate the expression above.

The mutual inductance M is a constant, and although the time-constant will vary because the rotor resistance varies with temperature, it will change relatively slowly, so we can treat it as constant, in which case the rotor flux angle is given by

$$\theta_{\mathrm{Ref}} = \int_0^t \omega_{\mathrm{rotor}} \mathrm{d}t + \frac{M}{\tau} \int_0^t \frac{I_{\mathrm{ST}}}{\psi_{\mathrm{R}}} \mathrm{d}t = \theta + \frac{M}{\tau} \int_0^t \frac{I_{\mathrm{ST}}}{\psi_{\mathrm{R}}} \mathrm{d}t$$

Note that because of the symmetry of the rotor, we only need the time-varying element of the rotor body angle (θ), not the absolute position, so the constant of integration is not required. (In contrast, for vector control of permanent magnet motors, the absolute position is important, because the rotor has saliency.)

The various methods that are used to keep track of the flux angle are what differentiate the various practical and commercial implementations of field-oriented control, as we will now see.

If we have a shaft encoder we can measure the rotor position (θ), or if we have a measured speed signal, we can derive θ by direct integration. This approach involves the fewest estimations, and therefore will normally offer superior performance, especially at low speeds, but is more costly because it requires extra transducers. We will refer to systems that use shaft feedback as 'closed-loop', but in the literature they may be also referred to as 'direct vector control'. In common with all schemes, the second term has to be estimated.

Many different methods of estimating the instantaneous parameter values are employed, but all employ a digital simulation or mathematical model of the motor/inverter system. The model runs in real time and is subjected to the same inputs as the actual motor, the model then being continuously fine-tuned so that the predicted and actual outputs match. Modern drives measure the circuit parameters automatically at the commissioning stage, and can even refine them on a near-continual basis to capture parameter variations.

The majority of vector control schemes eliminate the need for measurement of rotor position, and instead the rotor position term in equation (7.3) is also estimated from a motor model, based on the known motor voltage and currents. Rather confusingly, in order to differentiate them from schemes that do have shaft transducers, these systems are known as ‘open-loop’ or ‘indirect’ vector control. The term ‘open loop’ is a misleading one because at its heart is the closed-loop torque control shown in Figure 7.22, but it is widely used: what it really means is ‘no shaft feedback’.

The main problems of the open-loop approach occur at low speeds where motor voltages become very low and measurement noise can render the algorithms unreliable. Techniques such as high-frequency injection of diagnostic signals exist, but are yet to find acceptance in the market on standard motors. Open-loop inverter-fed induction motors are usually unsuitable for continuous operation at frequencies below 0.75 Hz, and struggle to produce full torque in this region.

An additional difficulty is that the significant variation of rotor resistance with temperature is reflected in the value of the all-important rotor time-constant τ. Any difference between the real rotor time constant and the value used by the model causes an error in the calculation of the flux position and so the reference frame becomes misaligned. If this happens, the flux and torque control are no longer completely decoupled, which results in suboptimum performance and possible instability. To avoid this, routines are included in the drive to provide ongoing estimates of the rotor time constant.

9. DIRECT TORQUE CONTROL

Direct torque control is an alternative high-performance strategy for vector/field orientation, and warrants a brief discussion to conclude our look at contemporary schemes. Developed from work first published in 1985 it theoretically provides the fastest possible torque response by employing a ‘bang-bang’ approach to maintain flux and torque within defined hysteresis bands. Like field-oriented control, it only became practicable with the emergence of relatively cheap and powerful digital signal processing.

Direct torque control avoids coordinate transformations because all the control actions take place in the stator reference frame. In addition there are no PI controllers, and a switching table determines the switching of devices in the inverter in place of the PWM approach favored for field-oriented control. These apparent advantages are offset by the need for a higher sampling rate (up to 40 kHz as compared with 6–15 kHz) leading to higher switching loss in the inverter; a more complex motor model; and inferior torque ripple. Because a hysteresis method is used the inverter has a continuously variable switching frequency which may be seen as an advantage in spreading the spectrum of acoustic noise from the motor.

We saw in the previous sections that in field-oriented control, the torque was obtained from the product of the rotor flux and the torque component of stator current. But there are many other ways in which the torque can be derived, for example in terms of the product of the rotor and stator fluxes and the sine of the angle between, or the stator flux and current and the sine of the angle between them. The latter is the approach discussed in the next section, but first a word about hysteresis control.

A good example of hysteresis control is discussed later in this book, in relation to ‘chopper drives’ for stepping motors in Chapter 9. Another more familiar example is the control of temperature in a domestic oven. Both are characterized by a simple approach in which full corrective action is applied whenever the quantity to be controlled falls below a set threshold, and when the target is reached, the power is switched off until the controlled quantity again drops below the threshold. The frequency of the switching depends on the time-constant of the process and the width of the hysteresis band: the narrower the band and the shorter the time-constant, the higher the switching frequency.

In the domestic oven, for example, the ‘on’ and ‘off’ temperatures can be a few degrees apart because the cooking process is not that critical and the time-constant is many minutes. As a result the switching on and off is not so frequent as to be irritating and wear out the relay contacts. If the hysteresis band were to be narrowed to a fraction of a degree to get tighter control of the cooking temperature, the price to be paid would be incessant clicking on and off, and shortened life of the relay.

9.1 Outline of operation

The block diagram of a typical direct torque control scheme is shown in Figure 7.29. There are several similarities with the scheme shown in Figure 7.22, notably the inverter, the phase current feedback, and the separate flux and torque demands, which may be generated by the speed controller, as in Figure 7.23.

However, there are substantial differences. Earlier we discovered that the inverter output voltage space phasor has only six active positions, and two zero states

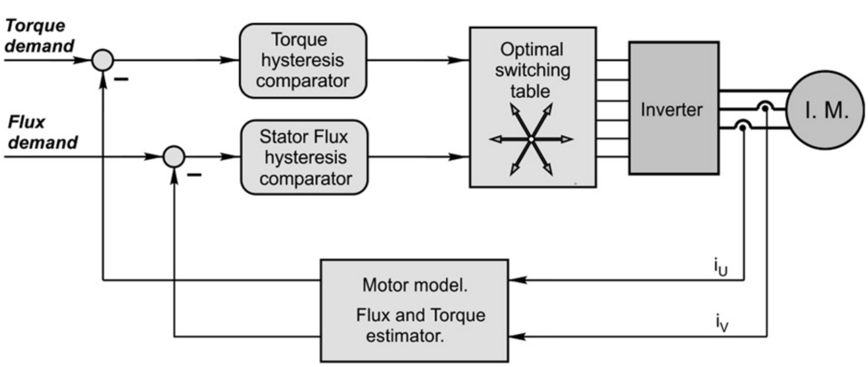

Figure 7.29 Block diagram of typical direct torque control scheme.

(see Figure 7.19), corresponding to the eight possible combinations of the six switching devices. This means that at every instant there are only eight options in regard to the voltage that we can apply to the motor terminals. In the field-oriented approach, PWM techniques are employed to alternate between adjacent unit vectors to produce an effective voltage phasor of any desired magnitude and instantaneous position. However, with direct torque control, only one of the eight intrinsic vectors is used for the duration of each sample, during which the estimated stator flux and torque are monitored.

The motor model is exposed to the same inputs as the real motor, and from it the software continuously provides updated estimates of the stator flux and torque. These are compared with the demanded values and as soon as either strays outside its target hysteresis band, a logical decision is taken as to which of the six voltage phasors is best placed to drive the flux and/or torque back onto target. At that instant the switching is changed to bring the desired voltage phasor into play. The duration of each sample therefore varies according to the rate of change of the two parameters being monitored: if they vary slowly it will take a long time before they hit the upper or lower hysteresis limit and the sample will be relatively long, whereas if they change rapidly, the sample time will be shortened and the sample frequency will increase. Occasionally, the best bet will be to apply zero voltage, so one of the two zero states then takes over.

9.2 Control of stator flux and torque

We will restrict ourselves to operation below base speed, so we should always bear in mind that although we will talk about controlling the stator flux, what we really mean is keeping its magnitude close to its normal (rated) value at which the magnetic circuit is fully utilized. We should also recall that when the stator flux is at its rated value and in the steady state, so is the rotor flux.

It is probably easiest to grasp the essence of the direct torque method by focusing on the stator flux linkage, and in particular on how (a) the magnitude of the stator flux is kept within its target limits and (b) how its phase angle with respect to the current is used to control the torque.

The reason for using stator flux linkage as a reference quantity is primarily the ease with which it can be controlled. When we discussed the basic operation of the induction motor in Chapter 5, we concluded that the stator voltage and frequency determined the flux, and we can remind ourselves why this is by writing the voltage equation for the stator as

$$V_S = I_S R_S + \frac{d\psi_S}{dt}$$

(We are being rather loose here, by treating space phasor quantities as real variables, but there is nothing to be gained by being pedantic when the message we take away will be valid.) In the interests of clarity we will make a further simplification by ignoring the resistance voltage term, which will usually be small compared with V_S. This yields

$$V_S = \frac{d\psi_S}{dt} \text{ or, in integral form, } \psi_s = \int V_S\, dt$$

The differential form shows us that the rate of change of stator flux is determined by the stator voltage, while the integral form reminds us that to build the flux (e.g. from zero) we have to apply a fixed volt-second product, with either a high voltage for a short time, or a low voltage for a long time. We will limit ourselves to the fine-tuning of the flux after it has been established, so we will only be talking about very short sample intervals of time (Δt) during which the change in flux linkage that results ($\Delta\psi_S$) is given by

$$\Delta\psi_S = V_S \Delta t$$

As far as we are concerned, ψ_S represents the stator flux linkage space phasor, which has magnitude and direction relative to the stator reference frame, and V_S represents one of the six possible stator voltage space phasors that the inverter can deliver. So if we consider an initial flux linkage vector ψ_S as shown in Figure 7.30(a), and assume that we apply, over time, Δt, each of the six possible options, we will produce six new flux-linkage vectors. The tips of the new vectors are labeled ψ_1 to ψ_6 in the figure, but only one (ψ_4) is fully drawn (dashed) to avoid congestion. There is also the option of applying zero voltage, which would of course leave the initial flux-linkage unchanged.

In (a) option 4 results in a reduction in amplitude and (assuming anticlockwise rotation) a retardation in phase of the original flux, but if the original flux linkage

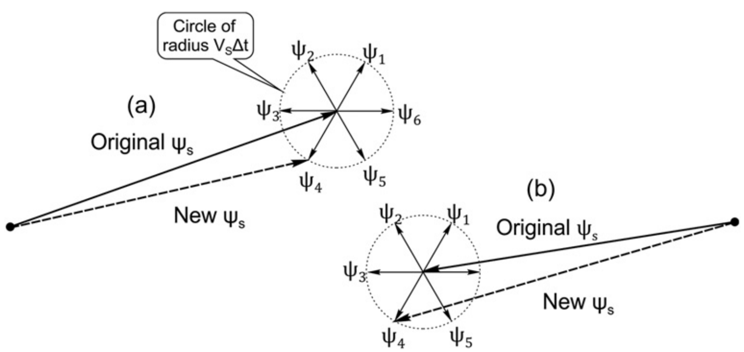

Figure 7.30 Space phasor diagram of stator flux linkage showing how the outcome of applying a given volt-second product depends on the original phase angle. In (a) the magnitude is reduced and the phase is retarded, while in (b) the magnitude is increased and the phase is advanced.

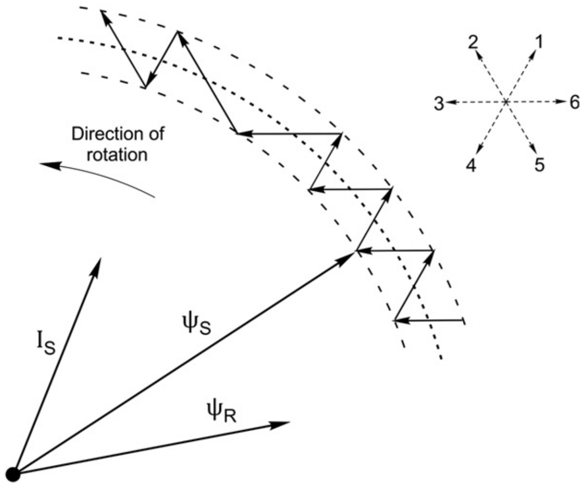

Figure 7.31 Trajectory of stator flux linkage under steady-state conditions.

had a different phase, as shown in diagram (b), option 4 results in an increase in magnitude and an advance in phase. It should be clear that outcomes vary according to initial conditions, and therefore an extensive look-up table will be needed to store all this information.

Having seen how we can alter the magnitude and phase of the stator flux, we now consider the flux linkage phasor during steady-state operation with constant speed and torque, in which case we know that ideally all the space phasors will be rotating at a constant angular velocity.

The locus of the stator flux linkage space phasor (ψ_S) is shown in Figure 7.31. In this diagram the spacing of hysteresis bands indicated by the innermost and outermost dashed lines have been greatly exaggerated in order to show the trajectory of the flux linkage phasor more clearly. Ideally, the phasor should rotate smoothly along the center dotted line.

In this example, the initial position shown has the flux linkage at the lower bound, so the first switching brings voltage vector 1 into play to drive the amplitude up and the phase forward. When the upper bound is reached, vector 3 is used, followed by vector 1 again and then vector 3. Recalling that the change in the flux linkage depends on the time for which the voltage is applied, we can see from the diagram that the second application of vector 3 lasts longer than the first. (We should also reiterate that in this example only a few switchings take place while the flux rotates through 60°: in practice the hysteresis band is very much narrower, and there may be many hundreds of transitions.)

We are considering steady-state operation, and so we would wish to keep the torque constant. Given that the flux is practically constant, this means we need to keep the angle between the flux and the stator current constant. This is where the torque hysteresis controller shown in Figure 7.29 comes in. It has to decide what switching will best keep the phase on target, so it runs in parallel with the magnitude controller we have looked at here. Each controller will output a signal for either an increase or decrease in its respective variable (i.e. magnitude or phase) and these are then passed to the optimal switching table to determine the best switching strategy in the prevailing circumstances (see Figure 7.29).

As we saw when discussing field-oriented control, it is not possible to make very rapid changes to the rotor flux because of the associated stored energy. Because the rotor and stator are tightly coupled it follows that the magnitude of the stator flux linkage cannot change very rapidly either. However, just as with field-oriented control, sudden changes in torque can be achieved by making sudden changes to the phase of the flux linkage, i.e. to the tangential component of the phasor shown in Figure 7.31.