CHAPTER NINE

Synchronous and Brushless Permanent Magnet Machines and Drives

1. INTRODUCTION

Chapters 6 to 8 have described the virtues of the induction motor and how, when combined with power electronic control, it is capable of meeting the performance and efficiency requirements of many of the most demanding applications. In this chapter another group of a.c. motors is described. In all of them the electrical power that is converted to mechanical power is fed into the stator, so, as with the induction motor, there are no sliding contacts in the main power circuits. All also have stators that are identical (or very similar) to the induction motor, although some new constructional and winding techniques involving segmented construction are being applied at the lower-power end to improve the power density of the machine and/ or reduce cost.

It has to be admitted that the industrial and academic communities have served to make life confusing in this area by giving an array of different names to essentially the same machine, so we begin by looking at the terminology. The names we will encounter include:

1. Conventional *synchronous machine* with its rotor field winding (excited-rotor). This is the only machine that may have brushes, but even then they will only carry the rotor excitation current, not the main a.c. power input.
2. *Permanent magnet synchronous machine* with permanent magnets replacing the rotor field winding.
3. *Brushless permanent magnet synchronous motor* (same as (2)). The prefix 'brushless' is superfluous.
4. *Brushless a.c. motor* (same as (2)).
5. *Brushless d.c. motor* (same as (2) except for detailed differences in the field patterns). This name was coined in the 1970s to describe 'inside-out' motors that were intended as direct replacements for conventional d.c. motors, and in this sense it has some justification.
6. *Permanent magnet servo motor* (same as (2)).

We begin by looking at motors that are intended, or at least have the capability, to be operated directly off the utility supply, usually at either 50 or 60 Hz. These are known as 'synchronous' motors, and they operate at a specific and constant speed

Electric Motors and Drives
http://dx.doi.org/10.1016/B978-0-08-098332-5.00009-7

for a wide range of loads, and therefore can be used in preference to induction motors when precise (within the tolerance of the utility frequency), constant speed operation is essential: there is no load-dependent slip as is unavoidable with the induction motor. These machines are available over a very wide range from tiny single-phase versions in domestic clocks to multi-megawatt machines in large industrial applications such as gas compressors. The clock application means that utility companies have a responsibility to ensure that the average frequency over a 24-hour period always has to be precisely the rated frequency of the supply in order to keep us all on time. Ironically, in order to do this, they control the speed of very large, turbine-driven, synchronous machines that generate the vast majority of the electrical power throughout the world.

To overcome the fixed-speed limitation that results from the constant frequency of the utility supply, inverter-fed synchronous motor drives have been developed. We will see that all forms of this generic technology use a variable-frequency inverter to provide for variation of the synchronous speed, but that in almost all cases, the switching pattern of the inverter (and hence the frequency) is determined by the rotor position and not by an external oscillator. In such so-called self-synchronous drives, the rotor is incapable of losing synchronism and stalling (which is one of the main drawbacks of the utility-fed machine).

We will also see how field-oriented control can be applied to synchronous machines to achieve the highest levels of performance and efficiency with machines which have higher inherent power densities than the equivalent induction motors.

Finally, reluctance motors will be considered briefly. Their simple rotor construction appears to offer great opportunity, but we will see that their low operating power-factor restricts their star rating. It does, however, remain a commercially available product and warrants review. For completeness hysteresis motors also get a mention.

2. SYNCHRONOUS MOTORS

In Chapter 5 we saw that the 3-phase stator windings of an induction motor produce a sinusoidal rotating magnetic field in the air-gap. The speed of rotation of the field (the synchronous speed) was shown to be directly proportional to the supply frequency, and inversely proportional to the pole-number of the winding. We also saw that in the induction motor the rotor is dragged along by the field, but that the higher the load on the shaft, the more the rotor has to slip with respect to the field in order to induce the rotor currents required to produce the torque. Thus although at no-load the speed of the rotor can be close to the synchronous speed, it must always be less; and as the load increases, the speed has to fall.

In the synchronous motor, the stator windings are essentially the same as in the induction motor, so when connected to the 3-phase supply, a rotating magnetic field is produced. However, instead of having a cylindrical rotor with a cage winding, the synchronous motor has a rotor with either a d.c. excited winding (supplied via sliprings, or on larger machines an auxiliary exciter[1]), or permanent magnets, designed to cause the rotor to 'lock on' or 'synchronize with' the rotating magnetic field produced by the stator. Once the rotor is synchronized, it will run at exactly the same speed as the rotating field despite load variation, so under constant-frequency operation the speed will remain constant as long as the supply frequency is stable.

As previously shown, the synchronous speed (in rev/min) is given by the expression

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

where f is the supply frequency and p is the pole-number of the winding. Hence for 2-, 4- and 6-pole industrial motors the running speeds on a 50 Hz supply are 3000, 1500 and 1000 rev/min, while on a 60 Hz supply they become 3600, 1800 and 1200 rev/min, respectively. At the other extreme, the little motor in a timeswitch with its cup-shaped rotor with 20 axially projecting fingers and a circular coil in the middle is a 20-pole reluctance (synchronous) motor that will run at 300 rev/min when fed from a 50 Hz supply. Users who want speeds different from these discrete values will be disappointed, unless they are prepared to invest in a variable-frequency inverter.

With the synchronous machine, we find that there is a limit to the maximum (pull-out) torque which can be developed before the rotor is forced out of synchronism with the rotating field. This 'pull-out' torque will typically be 1.5 times the continuous rated torque but can be designed to be as high as 4 or even 6 times higher in the case of high-performance permanent magnet motors where, for example, high accelerating torques are needed for relatively short periods. For all torques below pull-out the steady running speed will be absolutely constant. The torque–speed curve is therefore simply a vertical line at the synchronous speed, as shown in Figure 9.1. We can see from Figure 9.1 that the vertical line extends into quadrant 2, which indicates that if we try to force the speed above the synchronous speed the machine will act as a generator.

Traditionally, utility-fed synchronous motors were used where a constant speed is required, high efficiency desirable, and power-factor controllable. They were also used in some applications where a number of motors were required to run at

1 An auxiliary exciter is simply a second, smaller machine with a 3-phase stator and rotor winding, mounted on the same shaft. A 3-phase supply on the stator is 'transformer coupled' to the rotor winding. The induced e.m.f. in the rotor is rectified and fed to the main motor field winding. The phase rotation of the supply to the auxiliary stator is opposite to that of the main motor so that when the motor comes to high speed the induced rotor e.m.f. remains high.

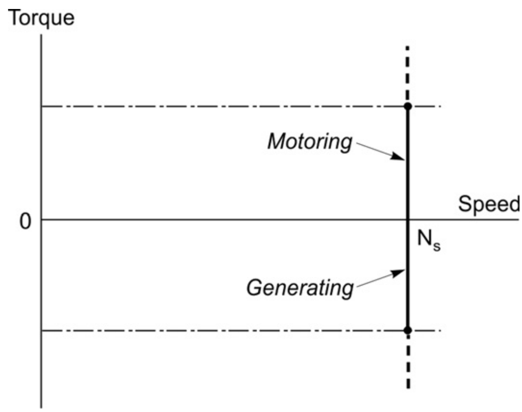

Figure 9.1 Steady-state torque–speed curve for a synchronous motor supplied at constant frequency.

precisely the same speed. However, a group of utility-fed synchronous motors could not always replace mechanical shafting[2] because while their rotational speed would always be matched, the precise relative rotor angle of each motor would vary depending on the load on the individual motor shafts.

We will now look briefly at the various types of synchronous motor, mentioning the advantages and disadvantages of each. The excited-rotor type is considered first because of its importance in large sizes but also because its behavior can be readily analyzed, and its mechanism of operation illuminated, by means of a relatively simple equivalent circuit. We then consider the permanent magnet type that is the most numerous in the drives arena, from which we can also benefit by dint of the use of the same equivalent circuit. Parallels are drawn with both the d.c. motor and the induction motor to emphasize that despite their obvious differences, most electrical machines also have striking similarities.

2.1 Excited-rotor motors

The rotor of a conventional synchronous machine carries a 'field' winding which is supplied with direct current either via a pair of sliprings on the shaft, or via an auxiliary brushless exciter on the same shaft. The field winding is designed to produce an air-gap field of the same pole-number and spatial distribution (usually sinusoidal) as that produced by the stator winding. The rotor may be more or less cylindrical, with the field winding distributed in slots (Figure 9.2(a)), or it may have

2 Drive shafts were used in early textile factories and the like. A single mechanical shaft would be fed through the factory or machine and various functions would be connected to the same shaft via belts. The connected equipment would then run up and down together. With toothed belts, position synchronism could be achieved.

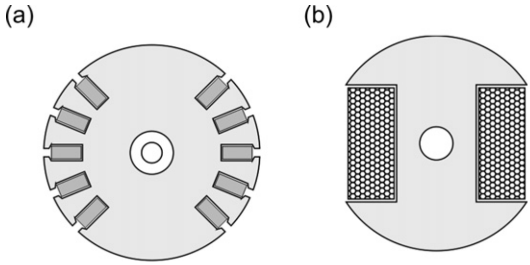

Figure 9.2 Rotors for synchronous motors. (a) 2-pole cylindrical, with field coils distributed in slots; (b) 2-pole salient pole with concentrated field winding.

projecting ('salient') poles around which the winding is concentrated (Figure 9.2(b)). A cylindrical-rotor motor has little or no reluctance (self-aligning) torque, so it can only produce torque when current is fed into the rotor. On the other hand, the salient-pole type also produces some reluctance torque even when the rotor winding has no current. In both cases, however, the rotor 'excitation' power is relatively small, since all the mechanical output power is supplied from the stator side.

Excited-rotor motors are used in sizes ranging from a few kW up to many MW. The large ones are effectively alternators (as used for power generation) but used as motors. Wound-rotor induction motors (see Chapter 6) can also be made to operate synchronously by supplying the rotor with d.c. through the sliprings.

The simplest way to visualize the mechanism of torque production is to focus on a static picture, and consider the alignment force between the stator and rotor field patterns. When the two are aligned with N facing S, the torque is zero and the system is in stable equilibrium, with any displacement to right or left causing a restoring torque to come into play. If the fields are distributed sinusoidally in space, the restoring torque will reach a maximum when the poles are misaligned by half a pole-pitch, or 90°. Beyond 90° the torque reduces with angle, giving an unstable region, zero torque being reached again when N is opposite N.

When the motor is running synchronously, we can use much the same mental picture because the field produced by the 3-phase alternating currents in the stator windings rotates at precisely the same speed as the field produced by the d.c. current in the rotor. At no-load there is little or no angular displacement between the field patterns, because the torque required to overcome friction is small. But each time the load increases, the rotor slows momentarily before settling at the original speed but with a displacement between the two field patterns that is sufficient to furnish the torque needed for steady-state running. This angle is known as the 'load angle', and we can actually see it when we illuminate the shaft of the motor with a supply-frequency stroboscope: a reference mark on the shaft is seen to drop back by a few degrees each time the load is increased. (Note that this is a characteristic of a synchronous machine when connected directly to the utility supply. We will see later in this chapter that when under inverter control with position feedback, precise angular position can be maintained regardless of load, which is important when coordinated motion is required, for example in X–Y tables where a precise path/contour must be followed.)

2.2 Permanent magnet motors

The synchronous machines considered so far require two electrical supplies, the first to feed the field/excitation and the second to supply the stator. Brushless permanent magnet machines have magnets attached to the rotor to provide the field, and so only a stator supply is required. The principle is illustrated for 2-pole and 4-pole surface-mounted versions in Figure 9.3, the direction in which the magnets have been magnetized being represented by the arrows. Motors of this sort have typical output ranging from about 100 W up to perhaps 500 kW, though substantially higher ratings have been made. (We should also mention that the magnets are sometimes buried within the rotor iron, which gives rise to an additional reluctance torque component. Although we discuss reluctance motors briefly later in this chapter, we will only deal with surface-mounted permanent magnet motors here.)

The advantages of the permanent magnet type are that no supply is needed for the rotor and the rotor construction can be robust and reliable. The disadvantage is that the excitation is inherently fixed, so the designer must either choose the shape and disposition of the magnets to match the requirements of one specific load, or seek a general-purpose compromise. Control of power-factor via excitation is of course no longer possible. Within these constraints the brushless permanent magnet synchronous motor behaves in very much the same way as its excited-rotor sister.

We will defer further discussion of some important constructional details of permanent magnet motors to a later section, and concentrate now on the tools available to help us understand the operating characteristics of synchronous motors.

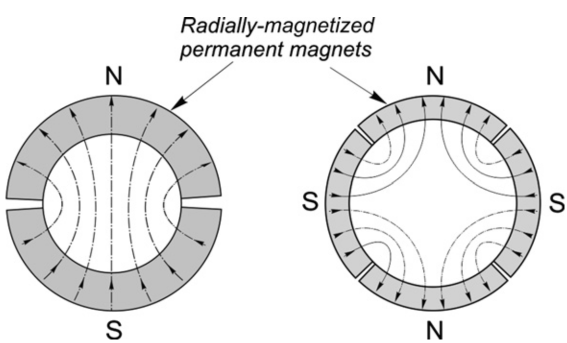

Figure 9.3 Permanent magnet synchronous motor rotors: 2-pole (left); 4-pole (right).

3. EQUIVALENT CIRCUITS OF SYNCHRONOUS MOTORS

Predicting the current and power-factor drawn from the supply by a cylindrical-rotor or permanent magnet synchronous motor is possible by means of the per-phase a.c. equivalent circuit shown in Figure 9.4. To arrive at such a simple circuit inevitably means that approximations have to be made, but we are seeking only a broad-brush picture, so the circuit is perfectly adequate.

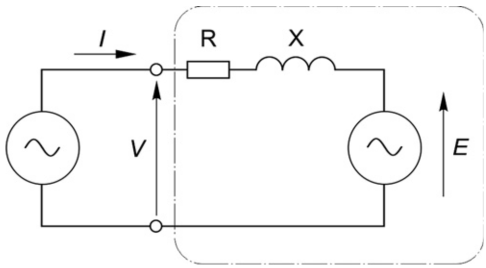

Figure 9.4 Equivalent circuit for synchronous machine.

In this circuit X (known as the synchronous reactance, or simply the reactance) represents the effective inductive reactance of the stator phase winding; R is the stator winding resistance; V is the applied voltage; and E is the e.m.f. induced in the stator winding by the rotating field produced either by the d.c. current on the rotor or by the permanent magnet. (For the benefit of readers who are familiar with the parameters of the induction motor, it should be pointed out that X is equal to the sum of the magnetizing and leakage reactances, but because the effective air-gap in synchronous machines is usually larger than in induction motors, their per-unit synchronous reactance is usually lower than that of an induction machine with the same stator winding.)

The similarity between this circuit and that of the d.c. machine (Figure 3.6) and the induction motor (Figure 5.8) is clear, and it stems from the fact that these machines all produce torque by the interaction of a magnetic field and current-carrying conductors (the so-called '*BIl*' effect). In the case of the d.c. machine the inductance was seen not to be important under steady-state conditions because the current was steady (i.e. d.c.), and the resistance emerged as the dominant parameter. In the case of the synchronous motor the current is alternating, so not surprisingly we find that the reactance is the dominant impedance and resistance plays only a minor role, except in the case of small motors.

4. OPERATION FROM CONSTANT-VOLTAGE, CONSTANT-FREQUENCY (UTILITY) SUPPLY

At this point, readers who are not familiar with a.c. circuit theory can be reassured that they will not be seriously disadvantaged by skipping the rest of this section. But although no seminal truths are to follow, discussion of the equivalent circuit and the associated phasor diagram greatly assists the understanding of motor behavior, especially its ability to operate over a range of power-factors.

4.1 Excited-rotor motor

Our aim is to find what determines the current drawn from the supply, which from Figure 9.4 clearly depends on all the parameters therein. But for a given machine operating from a constant-voltage, constant-frequency supply, the only independent variables are the load on the shaft and the d.c. current (the excitation) fed into the rotor, so we will look at the influence of both, beginning with the effect of the load on the shaft.

The speed is constant and therefore the mechanical output power (torque times speed) is proportional to the torque being produced, which in the steady state is equal and opposite to the load torque. Hence if we neglect the losses in the motor, the electrical input power is also determined by the load on the shaft. The input power per phase is given by $VI \cos \phi$ where I is the current and the power-factor angle is ϕ. But V is fixed, so the in-phase (or real) component of input current $I \cos \phi$ is determined by the mechanical load on the shaft. We recall that, in the same way, the current in the d.c. motor (Figure 3.6) was determined by the load. This discussion reminds us that although the equivalent circuits in Figures 9.4 and 3.6 are very informative, they should perhaps carry a 'health warning' to the effect that the single most important determinant of the current (the load torque) does not actually appear explicitly on the diagrams.

Turning now to the influence of the d.c. excitation current, at a given supply frequency (i.e. speed) the utility-frequency e.m.f. (E) induced in the stator is proportional to the d.c. field current fed into the rotor. If we wanted to measure this e.m.f. we could disconnect the stator windings from the supply, drive the rotor at synchronous speed by an external means, and measure the voltage at the stator terminals, performing the so-called open-circuit test. If we were to vary the speed at which we drove the rotor, keeping the field current constant, we would of course find that E was proportional to the speed. We discovered a very similar state of affairs when we studied the d.c. machine (Chapter 3): its induced motional or 'back' e.m.f. (E) turned out to be proportional to the field current, and to the speed of rotation of the armature. The main difference between the d.c. machine and the synchronous machine is that in the d.c. machine the field is stationary and the armature rotates, whereas in the synchronous machine the field system rotates while the stator windings are at rest: in other words, one could describe the synchronous machine, loosely, as an 'inside-out' d.c. machine.

We also saw in Chapter 3 that when the unloaded d.c. machine was connected to a constant-voltage d.c. supply, it ran at a speed such that the induced e.m.f. was (almost) equal to the supply voltage, so that the no-load current was almost zero.

When a load was applied to the shaft, the speed fell, thereby reducing *E* and increasing the current drawn from the supply until the motoring torque produced was equal to the load torque. Conversely if we applied a driving torque to the shaft of the machine, the speed rose, *E* became greater than *V*, current flowed out to the supply and the machine acted as a generator. These findings are based on the assumption that the field current remains constant, so that changes in *E* are a reflection of changes in speed. Our overall conclusion was the simple statement that if *E* is less than *V*, the d.c. machine acts as a motor, while if *E* is greater than *V*, it acts as a generator.

The situation with the synchronous motor is similar, but now the speed is constant and we can control *E* independently via control of the d.c. excitation current fed to the rotor. We might again expect that if *E* was less than *V* the machine would draw in current and act as a motor, and vice versa if *E* was greater than *V*. But we are no longer dealing with simple d.c. circuits in which phrases such as ‘draw in current’ have a clear meaning in terms of what it tells us about power flow. In the synchronous machine equivalent circuit the voltages and currents are a.c., so we have to be more careful with our language and pay due respect to the phase of the current, as well as its magnitude. Things turn out to be rather different from what we found in the d.c. motor, but there are also similarities.

4.2 Phasor diagram and power-factor control

To see how the magnitude of the e.m.f. influences behavior we can examine the phasor diagrams of a synchronous machine operating as a motor, as shown in Figure 9.5.

The first point to clarify is that our sign convention is such that motoring corresponds to positive electrical input power to the machine. The power is given by $VI \cos \phi$, so when the machine is motoring (positive power) the angle ϕ lies in

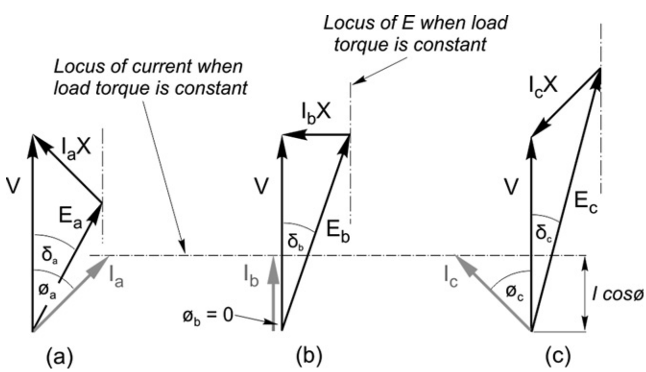

Figure 9.5 Phasor diagrams for synchronous motor operating with constant load torque, for three different values of the rotor (excitation) current.

the range ±90°. If the current lags or leads the voltage by more than 90° the machine will be generating.

Figure 9.5 shows three phasor diagrams corresponding to low, medium and high values of the induced e.m.f. (E), the shaft load (i.e. mechanical power) being constant. As discussed above, if the mechanical power is constant, so is $I \cos \phi$, and the locus of the current is therefore shown by the horizontal dashed line. The load angle (δ), discussed earlier, is the angle between V and E in the phasor diagram. In Figure 9.5, the voltage phasor diagram embodies Kirchhoff's law as applied to the equivalent circuit in Figure 9.4, i.e. $V = E + IR + jIX$, but for the sake of simplicity R is neglected so the phasor diagram simply consists of the volt-drop IX (which leads the current I by 90°) added to E to yield V.

Figure 9.5(a) represents a condition where the field current has been set so that the magnitude of the induced e.m.f. (E) is less than V. This is called an 'under-excited' condition, and as can be seen the current is lagging the terminal voltage and the power-factor is cos ϕ_a, lagging. When the field current is increased (increasing the magnitude of E) the magnitude of the input current reduces and it moves more into phase with V: the special case shown in Figure 9.5(b) shows that the motor can be operated at unity power-factor if the field current is suitably chosen. Finally, in Figure 9.5(c), the field current is considerably higher (the 'overexcited' case), which causes the current to increase again but this time the current leads the voltage and the power-factor is cos ϕ_c, leading. We see that we can obtain any desired power-factor by appropriate choice of rotor excitation: this is a freedom not afforded to users of induction motors, and arises because in the synchronous machine there is an additional mechanism for providing excitation, as we will now discuss.

When we studied the induction motor we discovered that the magnitude and frequency of the supply voltage V governed the magnitude of the resultant flux density wave in the machine, and that the current drawn by the motor could be considered to consist of two components. The real (in-phase) component represented the real power being converted from electrical to mechanical form, so this component varied with the load. On the other hand, the lagging reactive (quadrature) component represented the 'magnetizing' current that was responsible for producing the flux, and it remained constant regardless of load.

The stator winding of the synchronous motor is the same as the induction motor, so it is to be expected that the resultant flux will be determined by the magnitude and frequency of the applied voltage. This flux will therefore remain constant regardless of the load, and there will be an associated requirement for magnetizing m.m.f. But now we have two possible means of providing the excitation m.m.f, namely the d.c. current fed into the rotor and the lagging component of current in the stator.

When the rotor is underexcited, i.e. the induced e.m.f. E is less than V (Figure 9.5(a)), the stator current then has a lagging component to make up for the shortfall in excitation needed to yield the resultant field that must be present as determined by the terminal voltage V. With more field current (Figure 9.5(b)), however, the rotor excitation alone is sufficient and no lagging current is drawn by the stator. And in the overexcited case (Figure 9.5(c)), there is so much rotor excitation that there is effectively some reactive power to spare and the leading power-factor represents the export of lagging reactive power that could be used to provide excitation for induction motors elsewhere on the same system.

To conclude our look at the excited-rotor motor we can now quantify the qualitative picture of torque production we talked about earlier, by noting from the phasor diagrams that if the mechanical power (i.e. load torque) is constant, the variation of the load angle (δ) with E is such that E sin δ remains constant. As the rotor excitation is reduced, and E becomes smaller, the load angle increases until it eventually reaches its maximum of 90°, at which point the rotor will lose synchronism and stall. This means that there will always be a lower limit to the excitation required for the machine to be able to transmit the specified torque. This is just what our simple mental picture of torque being developed between two magnetic fields, one of which becomes very weak, would lead us to expect.

4.3 Permanent magnet motor

Although the majority of permanent magnet motors are supplied from variable-frequency inverters, some are directly connected to the utility supply, and we can explore their behavior using the equivalent circuit shown in Figure 9.4. Because the permanent magnet acts as a source of constant excitation, we no longer have control over the magnitude of the induced e.m.f. E, which now depends on the magnet strength and the speed, the latter being fixed by the utility frequency. So now we only have the load torque as an independent variable, and, as we saw earlier, because the supply voltage is constant, the load torque determines the in-phase or work component of the stator current $I \cos \phi$, as indicated in the phasor diagrams in Figure 9.5.

In order to identify which of the three diagrams in Figure 9.5 applies to a particular motor we need to know the motional e.m.f. (E) with the rotor spinning at synchronous speed and the stator open-circuited. If E is less than the utility voltage, diagram (a) applies; the motor is said to be underexcited; and it will have a lagging power-factor that worsens with load. Conversely, if E is greater than V (the overexcited case), diagram (b) or (c) is typical, and the power-factor will be leading.

4.4 Starting

It should be clear from the discussion of how torque is produced that unless the rotor is running at the same speed as the rotating field, no steady torque can be produced. If the rotor is running at a different speed, the two fields will be sliding past each other, giving rise to a pulsating torque with an average value of zero. Hence a basic synchronous machine is not self-starting, and some alternative method of producing a run-up torque is required.

Most synchronous motors, designed for direct connection to the utility supply, are therefore equipped with some form of rotor cage, similar to that of an induction motor, in addition to the main field winding. When the motor is switched onto the supply, it operates as an induction motor during the run-up phase, until the speed is just below synchronous. The excitation is then switched on, and as long as the load is not too high, the rotor is able to make the final acceleration and 'pull in' to synchronism with the rotating field. Because the cage is only required during starting, it can be short-time rated, and therefore comparatively small. Once the rotor is synchronized, and the load is steady, no currents are induced in the cage, because the slip is zero. The cage does, however, come into play when the load changes, when it provides an effective method for damping out the oscillations of the rotor as it settles at its new steady-state load angle.

Large motors will tend to draw a very heavy current during run-up, perhaps six or more times the rated current, for many tens of seconds, or longer, so some form of reduced voltage starter is often required (see Chapter 6). Sometimes, a separate small or 'pony' motor is used simply to run up the main motor prior to synchronization, but this is only feasible where the load is not applied until after the main motor has been synchronized.

5. VARIABLE-FREQUENCY OPERATION

Just as we have seen in Chapters 7 and 8 for the induction motor, once we interpose a power electronic converter between the utility supply and the machine we introduce new levels of performance and lose most of the inherent drawbacks which we find when the motor is directly connected to the utility supply.

Most obviously, a variable-frequency converter frees the synchronous machine from the fixed-speed constraint imposed by utility-frequency operation. The obvious advantage over the inverter-fed induction motor is that the speed of the synchronous motor is exactly determined by the supply frequency whereas the induction motor always has to run with a finite slip. On the down side, we lose the ability of the excited-rotor motor to vary the power-factor as seen by the utility supply.

In principle, a precision frequency source (oscillator) controlling the inverter switching is all that is necessary to give precise speed control with a synchronous motor, while speed feedback is essential to achieve accuracy with an induction motor. In practice, however, we seldom use open-loop control, where the voltage and frequency are generated within the inverter and are independent of what the motor does. Instead, field-oriented control, almost identical to that described for the inverter-fed induction motor, predominates. The principal advantage of field-oriented control is that it allows us to control the torque and flux components of the stator current independently, and in the case of the synchronous motor it prevents the motor from losing synchronism with the traveling field by locking the supply frequency to the speed of the rotor.

However, in the steady state, an observer looking at the stator voltage and current would see steady-state sinusoidal waveforms, and would be unaware of the underlying control mechanism. We can therefore study the steady-state behavior using the equivalent circuit in much the same way as we did with the utility-fed permanent magnet motor. We will continue to ignore resistance because this makes the phasor diagrams much simpler to understand without seriously compromising our conclusions.

The approach that we take differs somewhat from our previous discussion in this chapter by laying more emphasis on the relation between the torque and the fluxes in the motor. We imagine the flux produced by the magnet and the flux produced by the stator as if they existed independently, although in reality there is only one resultant flux. Intuitively we can see that because the fluxes rotate in synchronism, the magnitude of the torque will depend on the product of the two field strengths and the angle between them: when aligned, the torque is zero, and when perpendicular, it is maximum. This is equivalent to saying that the torque is maximum when the stator current wave is aligned with the magnet flux wave, which is the traditional '*BIl*' picture.

Because both flux distributions are sinusoidal, the torque depends on the sine of the angle between them (λ). The stator field strength depends on the current, so maximum torque will be obtained when the current is perpendicular to the magnet flux in the phasor diagram, as shown in later figures (9.7) and (9.8).

When we discussed the utility-connected excited-rotor motor in section 4, we were reminded that with the voltage and frequency fixed, the resultant (stator) flux was constant, and that if the rotor excitation was low, extra magnetizing current would be drawn from the utility supply, and the power-factor would be lagging. When the rotor excitation was high there was a surplus of excitation and a leading current would be exported to the utility. By suitable adjustment of the rotor current we could achieve a power-factor of unity for any value of the load torque. We then saw that for a utility-fed permanent magnet motor, in which the rotor excitation is constant, the stator current adjusted itself to satisfy the requirement for the resultant flux to be constant, but as a result we had no control over the power-factor.

With an inverter-fed motor we gain control of both the stator voltage and frequency, so that together with the load torque we now have three independent variables in the case of the permanent magnet motor, or four for the excited-rotor machine. The majority of inverter-fed synchronous motor drives, and almost all below 200 kW, employ permanent magnet motors, so we will concentrate on their behavior for the remainder of this section.

5.1 Phasor diagram – nomenclature and basic relationships

The general diagram (Figure 9.6) is for an underexcited case; i.e. at the speed in question, the open-circuit e.m.f. (E) is less than the terminal voltage. We will discuss what each phasor represents first, and then turn to the relationships that allow the diagram to be produced.

E is the open-circuit e.m.f. produced by the magnet flux (ϕ_{mag}): it is proportional to the magnet flux and the speed, which is proportional to the stator frequency ω. It is convenient for us to use this as our reference phasor because once the frequency is specified, the magnitude of E is known, so we can start the phasor diagram with the known E.

Recalling that phasors rotate anticlockwise, and that the projection of any phasor onto the vertical axis represents the instantaneous value, we can deduce that at the instant shown in Figure 9.6, the induced e.m.f. is at its maximum. We also know from Faraday's law that the induced e.m.f. depends on the rate of change of the flux, which is greatest when the flux is passing through zero, so in Figure 9.6 the instantaneous flux is zero, as shown. In general, therefore, the flux lags the corresponding induced e.m.f. by 90°, as shown.

ϕ_{arm} is the flux that would be set up if the armature (stator) current existed alone. We are assuming no saliency, so the magnitude and direction of this flux depend only on the magnitude and phase of the stator current.

ϕ_{res} is the resultant ('the') flux. With resistance ignored, this flux is effectively determined by the applied voltage and frequency, as we have seen several times earlier in this book, so in the phasor diagram, the magnitude of this flux is proportional to, and in quadrature with, V.

The self-induced e.m.f. from the stator flux is modeled in the equivalent circuit by means of the voltage across the stator inductive reactance X, equal to ωL, where L is the self inductance of the winding. In the phasor diagram the voltage IX leads

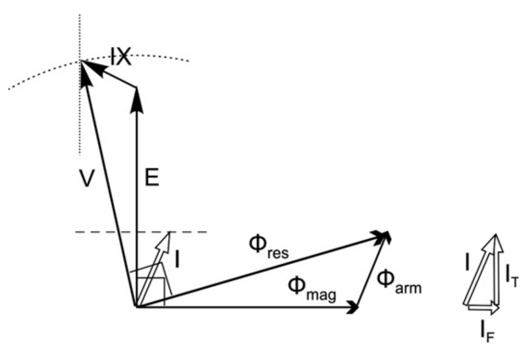

Figure 9.6 Phasor diagram for permanent magnet synchronous motor.

the current by 90°. The e.m.f. could equally well have been represented by an e.m.f. source similar to the E that represents the magnet-induced e.m.f., but historically the circuit representation is preferred because it is a compact way of representing both the magnitude and time-phase of the induced voltage with respect to the current.

V is the applied stator voltage, here regarded as an independent variable, and I is the resulting stator current. The stator phasor equation is $V = E + IX$ (or $V = E + jIX$ for those who prefer the complex notation).

The other independent variable is the load torque, which is our next consideration in developing the diagram.

In the steady state, the motor torque must equal the load torque. The motor torque is proportional to the product of the current, the magnet flux ϕ_{mag}, and the sine of the angle between them, so, knowing the magnet strength, we can find the in-phase or torque component of the stator current (I_T). The locus of the current I is then the horizontal dashed line, and the locus of V then becomes the vertical dotted line as shown in Figure 9.6. So when we finally specify the magnitude of V (shown by the dotted arc), the intersection of the arc and vertical line fixes the phase of the stator current and finalizes the diagram.

The current has two components as shown on the right in Figure 9.6. The component in phase with E is the useful or torque component I_T, while the component in phase with the magnet flux is the flux component I_F. The area of the flux triangle is proportional to the product of I_T and ϕ_{mag}, and thus the area provides an immediate visual indication of the torque.

We note that by control of V, we can alter the flux component (I_F), but the steady-state torque component is determined by the load. In particular, we can adjust V so that the stator current is in phase with E (i.e. the flux component is zero), so that we get the maximum value of torque per ampere of stator current. In this condition the armature (stator) flux is perpendicular to the magnet flux, as discussed above, and is explored in the next section.

Before we move on, we should note that because we specified an applied voltage just a bit bigger than E, the current turned out to be of modest amplitude and at a reasonable power-factor angle. However, if we had specified a much higher V, we would find that the current would have had a much larger lagging flux component (with the same torque component, determined by the load), and a much worse power-factor. And conversely, a much smaller applied voltage would result in a large leading power-factor current. Neither of these conditions is desirable because of the increased copper loss.

5.2 Field-oriented control

In a later section we will look at the torque–speed capabilities of the inverter-fed permanent magnet motor, and we will find that, as with the d.c. drive and the induction motor drive, there is a so-called constant torque region that extends up to base speed, within which full rated torque is available on a continuous basis. And in common with the other drives, at higher speeds there is a field weakening region where the maximum available torque is reduced. However, the usual constraints imposed by the maximum supply voltage and allowable continuous stator current may result in more serious restrictions on operation in this region than we have seen with other drives, including the likelihood that only intermittent operation will be possible.

Given the number of parameters involved and their variation between motor designs, it is only possible for us paint a broad picture, so we will look at one hypothetical machine and use it to provide an insight into some of the issues involved. We will focus on three conditions, two in the constant-torque region and one in field weakening. All three phasor diagrams are to the same scale to allow easy comparisons to be made.

We discussed field-oriented control in the context of the induction motor in Chapter 7. In principle for both induction and permanent magnet motors the strategy involves independent control over the torque and flux components of the stator current. But whereas the induction motor has to have both components (because it has no other source of flux), the permanent magnet motor only needs the torque component because its flux is provided by the magnets on the rotor. At speeds up to base speed, therefore, field control of the permanent magnet motor involves control of the stator current (via control of the stator voltage) to keep the flux component zero, while at the same time the torque component matches the input torque demand to the controller (see later Figure 9.10). During acceleration or deceleration, the torque demand will be high, and the controller will maintain the correct torque component while the speed is changing, until the target speed is reached when the torque component will be reduced to the appropriate steady-state value.

For convenience in calculations in relation to the phasor diagrams we will take the open-circuit e.m.f. E at the base speed (ω_B) to be 1 p.u. (per-unit). (In practice, the rated applied voltage is usually taken as 1 p.u., but there is no reason why any other value should not be chosen.) We will assume that the reactance of the winding (X) at the base speed is 0.3 p.u. (which means that at rated current (and rated frequency) the voltage across the reactance is 0.3 times the rated voltage).

5.3 Full-load

By full-load we mean that the machine is running at base speed and delivering its full rated torque. We saw above that the current depended on the applied voltage, and that, in particular, if we apply the right voltage, we can minimize the current. This is what we do in field-oriented control, so the diagram (Figure 9.7) represents this condition.

We should note that the stator current has zero flux component, and so the armature flux is perpendicular to the magnet flux, i.e. in the optimum torque-producing position. We will define the current in this situation as 1 p.u., so the

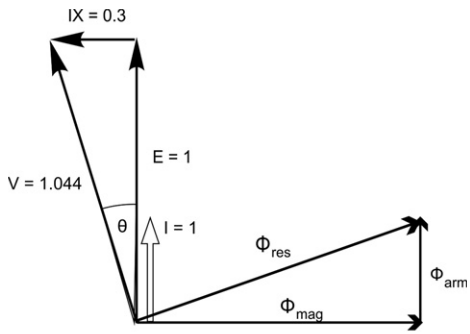

Figure 9.7 Permanent magnet motor phasor diagram – field-oriented control – full (rated) torque at base speed.

volt-drop across the reactance is $1 \times 0.3 = 0.3$ p.u. From the property of the right-angled triangle, V turns out to be 1.044 p.u., and because we have defined this as the rated power condition, we will regard 1.044 V as the maximum voltage that the inverter can produce. This condition is the most efficient state for the given current and torque, because the stator copper loss is minimized.

5.4 Full torque at half base speed (half power)

The phasor diagram is shown in Figure 9.8, again with the voltage optimized to give the most efficient stator current. The magnet flux is the same as in Figure 9.7 of course, and so for the same (rated) torque the stator current has to be 1 p.u., as before. However, the frequency is now only half ($\omega_B/2$) so the open-circuit e.m.f. is reduced to 0.5. The stator reactance is proportional to frequency, so it has also

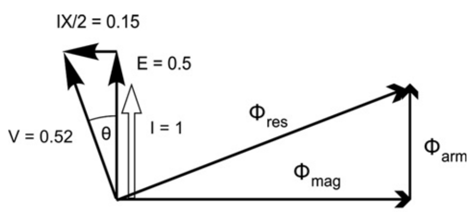

Figure 9.8 Permanent magnet motor phasor diagram – field-oriented control – full (rated) torque at half base speed.

halved, to 0.15 p.u., and the volt-drop IX becomes 0.15 p.u. also. In view of the similarity between the diagrams it is no surprise that the applied voltage turns out to be half, i.e. 0.52 p.u., and given the emphasis we have previously placed on the fact that the flux in an a.c. machine is determined by the V/f ratio, we will not be surprised to see that this is again confirmed by these results.

The angle θ between V and I (the power-factor angle) is the same at full and half speed, and since the input power is given by $VI\cos\theta$, it is clear that the input power at half speed is half of that at full speed. This is what we would expect because with resistance neglected the input power equals the mechanical power, which is half in Figure 9.8 because the torques are the same but the speed is half of that in Figure 9.7.

The two previous examples have shown how the motor can be operated to produce full rated torque up to a speed at which the full available voltage is applied; i.e. this is what we have referred to previously as the 'constant torque region'. In this region, the permanent magnet motor with field-oriented control behaves, in the steady state, in a very similar manner to a d.c. motor drive, in that the applied voltage (and frequency) is proportional to the speed and the stator current is proportional to the torque.

We saw with the d.c. drive and the induction motor drive that once we had reached full voltage and current, any further increase in speed could only be achieved at the expense of a corresponding reduction in torque, because the input power was already at its maximum or rated value. In both cases, higher speeds were obtained by entering the aptly named 'field weakening' region.

For the d.c. motor, the field flux is under our direct control, so we reduce the current in the field circuit. In the induction motor, the field is determined indirectly by the V/f ratio, so if f increases while V remains constant, the field flux reduces. The permanent magnet motor behaves in a somewhat similar fashion to the induction motor, in that if the voltage is constant at speeds above base speed, the V/f ratio reduces as the frequency (speed) is increased, so the resultant flux also reduces. However, whereas the only source of excitation in the induction motor is the stator winding, the magnet in the permanent magnet motor remains a constant (and potent!) source of excitation at all times, and so we can anticipate that in order to arrive at a reduced flux to satisfy the V/f condition, the stator current will have to nullify some of the influence of the magnet. We must therefore expect less than ideal behavior in the field weakening region, which we now examine.

5.5 Field weakening – operation at half torque, twice base speed (full power)

We will consider a situation well into the field weakening region; i.e. we will assume that in line with other drives we can expect full power, and so settle for operation at twice base speed and half rated torque. We will see that while this is achievable on an intermittent basis, it is not possible without exceeding the rated current, a conclusion that we must expect to apply to the whole of the field weakening region with a permanent magnet drive.

The condition is represented by the phasor diagram shown in Figure 9.9. The frequency is twice the base value, i.e. 2 ω_B, and so the open-circuit e.m.f. is 2 p.u. As in other drives, above base speed we apply the maximum possible stator voltage, in this case 1.04 p.u. Because the voltage cannot be greater, it is not possible for us to arrange for the current to be in phase with *E*, and we are obliged to settle for a rather poor second-best.

The load torque is half rated, which means that the torque component is 0.5 p.u. However, we note that the very large difference in voltage between *E* and *V* results in a very large stator volt-drop of $IX = 1.044$ p.u. The reactance is twice as large as at base speed because the frequency is doubled, so the current is given by $1.044/0.6 = 1.74$ p.u., which is 74% above its continuously rated value. The copper loss will therefore be increased by a factor of 1.74^2, i.e. to three times the rated value. Clearly continuous operation will not be possible because the stator will

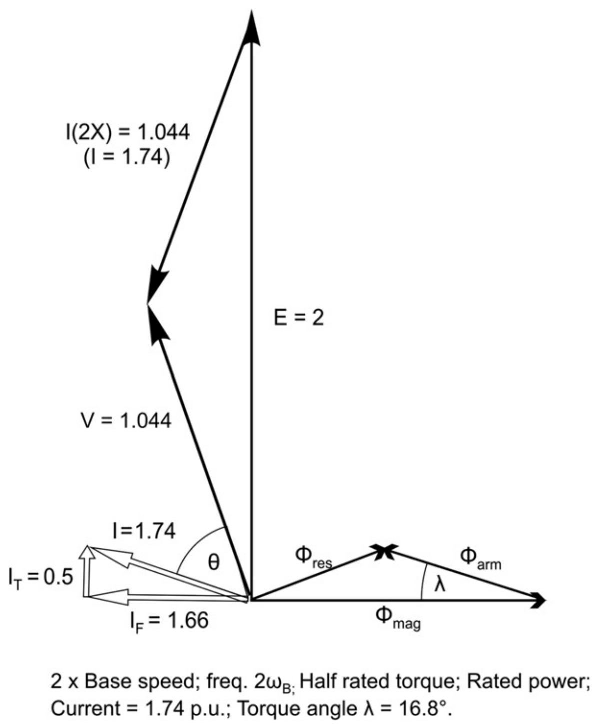

Figure 9.9 Permanent magnet motor phasor diagram – twice base speed, half rated torque, full power.

overheat, so in this field weakening condition only intermittent operation at half torque will be possible.

By comparing the flux triangles in Figures 9.7 and 9.9 we can see why the latter is referred to as field weakening. In Figure 9.7 our freedom to adjust V allowed us to ensure that the stator or armature flux is in quadrature with the magnet flux (i.e. the torque angle λ is 90°), leading to a slightly higher resultant flux, and maximizing efficiency by minimizing the current for the given torque.

In contrast the constraint on V means that the armature flux has a large component which is in opposition to the magnet flux, leading to a resultant flux that is much less than the magnet flux, and a very low torque angle (λ) of only 16.8°. In a sense, therefore, most of the stator current is 'wasted' in being used to oppose the magnet flux. Clearly this is not a desirable operating condition, but it is the best that we can get with the limited voltage at our disposal.

As a final check we can calculate the input power. The angle θ is given by $90 - 2\lambda = 56.4°$, so the input power ($VI \cos \theta$) is $1.044 \times 1.74 \times 0.553 = 1$ p.u., as expected since we assumed that the mechanical power was at rated value (twice base speed, half rated torque) and we ignored resistance in our calculations.

At this point we should recall that the aim throughout section 5 has been to discover how the motor parameters determine the steady-state behavior when the voltage, frequency and load vary over a range that is representative of a typical inverter-fed drive. Few readers will find it necessary to retain all of the detail (although a general awareness of the broad picture is always helpful) so those who have found it hard going can take comfort from the fact that in practice the drive will take care of everything for them. This is discussed in the next section.

6. SYNCHRONOUS MOTOR DRIVES

Inverter-fed operation of synchronous machines plays a very important and growing role in the overall drives market, as customers seek higher efficiency and higher power density than can be achieved with induction motors. We will also see later when we consider some of the available permanent magnet motor designs and control strategies that very high dynamic performance can be achieved, making this the motor of choice in many of the most demanding applications.

Inverter-fed synchronous motor drives are well established in two distinct guises. At one extreme, large (multi-MW) excited-rotor synchronous motors are used, particularly where high speeds are required or when the motor must operate in a hazardous atmosphere (e.g. in a large gas compressor). At the other (100 W to 500 kW), permanent magnet synchronous motors are used in a wide range of applications. Each of these groupings has distinct converter and control strategies and so we will look at each, beginning with the latter.

6.1 Permanent magnet motor drives

Permanent magnet motor drives are a very important and rapidly growing sector of the drives market, which is why we have devoted a large section of this chapter to a theoretical study aimed at understanding what determines their steady-state behavior. In terms of their application in drives we have little new to introduce because the converter circuits used are exactly as we have already discussed in Chapters 7 and 8 for the induction motor, and the control is also very similar.

Our discussion of field-oriented control in Chapter 7 focused heavily on the ability to control the torque-producing and field-producing components of the stator current (i_T and i_F) independently and rapidly, in order to provide virtually instantaneous control of torque. We do exactly the same for the permanent magnet synchronous machine except we normally set the demand for i_F to zero, as there is no need to provide the working flux from the stator side because it is provided by the rotor magnets.

The arrangement of a typical field-oriented control system is shown in Figure 9.10.

As with the induction motor drive the motor is supplied from a voltage source inverter, and again we are therefore controlling the stator currents via the voltage. The control strategy and the inverter are essentially the same as for an induction motor – indeed some manufacturers make life easier for users by offering a single inverter product suitable for both types of motor.

As with the induction motor control strategy, the critical issue is determining the flux position (θ_{Ref}), but in this case the task is very much easier as the flux is aligned to the rotor position. In high-performance systems a high-resolution absolute encoder would be fitted to the motor shaft to provide an accurate rotor position signal so that the flux position is known precisely at all times. However, if a position sensor is not used then the reference frame angle, θ_{Ref}, can be derived from computed motor voltages and currents in the same way as for the induction motor, but in this case we no longer have the complication of the temperature-dependent rotor time-constant.

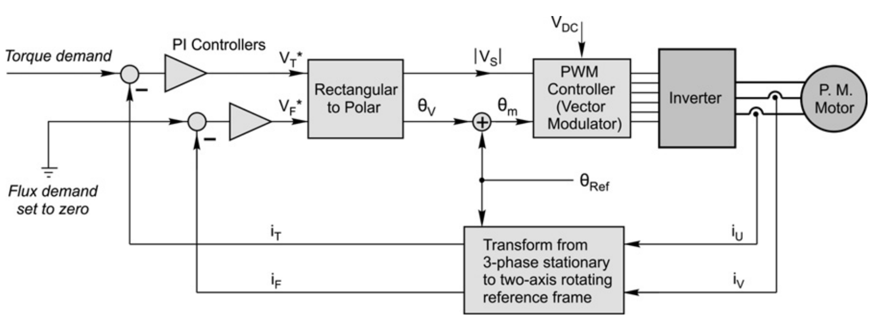

Figure 9.10 Field-oriented control scheme for permanent motor synchronous motor.

As we have discussed before, model-based schemes break down at very low frequencies because the voltage components become very small. For low-speed operation it is therefore necessary to supplement the approach with an alternative such as a position-sensing scheme using injected high-frequency currents. Or, in systems where the load is reasonably predictable, such as in a domestic washing machine, the motor can be brought up to a particular speed using open-loop switching in much the same fashion as a stepping motor (see Chapter 10) until the motional e.m.f. is large enough to be used as the control signal.

We have said with some satisfaction that because we have magnets on the rotor, we do not need to provide current to develop the field, and consequently the field current demand in Figure 9.10 is set to zero. However, if the field current demand were not set to zero, the control system would provide a flux component of stator current that could either increase or decrease the magnet's flux, depending on the polarity of the reference signal. Clearly there will be an upper limit on the flux because of saturation of the magnetic circuit, and in practice, reducing the flux is a more attractive proposition because it allows us to operate in a field weakening mode and so extend the speed range into a constant power region, as discussed at the end of section 5.

6.2 Converter-fed synchronous machine (multi-MW, excited-rotor) drive

The basic components of the drive system are shown in Figure 9.11. Current in the top of the d.c. link always flows from left to right (as shown by the direction of the thyristors), but the link voltage can be positive or negative so that energy can flow in either direction. The labels 'rectifier' and 'inverter' in Figure 9.11 indicate how each converter operates when the machine is operating as a motor, their roles being reversed when the machine is generating. An additional attraction of this set-up is that the direction of rotation of the machine is determined electronically by the switching sequence within the machine converter, and hence full four-quadrant operation is available without extra hardware.

In view of what we have read so far, we might have expected the switching devices in the motor inverter to be IGBTs, and to see PWM control of the output

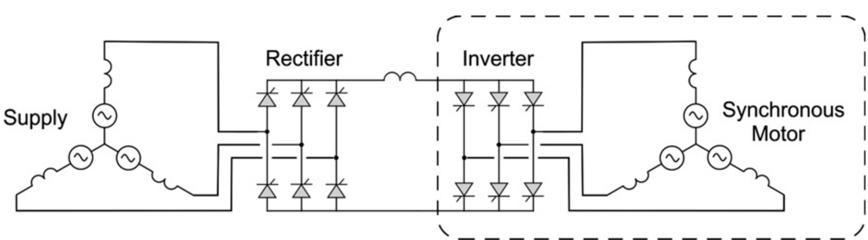

Figure 9.11 Converter-fed large synchronous machine drive.

voltage waveform, but in these large sizes, IGBT devices are prohibitively expensive, and this explains why the motor converter is exactly like that used on the supply side of a d.c. motor drive. The reason why we can use this cheaper inverter lies in the fact that, once rotating, a synchronous machine generates a.c. voltages, which can be used for the natural commutation of a converter connected to its terminals. This is why it is frequently referred to as a load commutated inverter (LCI).

In effect the motor converter behaves in the same way as it does when connected to the utility supply, and as such it is inherently capable of power flow in both directions. On the down side, we have no means of profiling the output voltage waveform, and as a result we have to accept waveforms for the motor current that are less than the perfect sinusoids that we would like.

We have already likened the combination of an inverter and synchronous motor (enclosed by the dashed line in Figure 9.11) to an 'inside-out' d.c. machine, and this view is instructive to help us to understand how this particular drive operates. We can begin by focusing on the motional-induced e.m.f.

In the synchronous motor the moving rotor carries the d.c. field winding. The rotating flux induces a sinusoidal e.m.f. in the three stator windings that is proportional to speed and field current. The firing of the devices in the motor converter is locked to the angular position of the rotor, so that each phase is connected in sequence to the d.c. link. The d.c. link therefore 'sees' a rectified 3-phase voltage that, although unidirectional, is not smooth d.c.

Conversely, the d.c. motor field is stationary, and its flux induces an alternating e.m.f. in the (many) armature coils on the rotor. The mechanical commutator and its sliding brushes rectify the e.m.f. so that at the armature terminals we get a smooth d.c.-induced voltage.

So when viewed from the d.c. link, the two are essentially the same. It is therefore not surprising that, as with a d.c. motor, the no-load speed of the synchronous motor depends on the d.c. link voltage provided by the supply-side converter, while when load is applied and speed tends to fall, the d.c. link current automatically increases until the steady state is reached. The effects of varying the d.c. excitation of the synchronous machine also mirror those in the d.c. motor, so that field weakening will lead to higher speed but reduced torque.

Of course, the ripple frequency of the rectified d.c. on the motor side of the d.c. link depends on the speed of the motor, while that on the supply side of the d.c. link will be 300 or 360 Hz depending on whether the supply is 50 or 60 Hz. So in order to prevent unwanted harmonic currents and to smooth the d.c. link current, a series inductor is included, as shown in Figure 9.11. The reactance of the inductor (ωL) is high at the ripple frequencies, but zero at d.c.

The switching strategy of the motor converter is synchronized to the back e.m.f. of the machine, each switch conducting for 120° (elect), so that the current waveform in each phase of the motor is roughly constant for 120°, and zero for the next 60°, and this pattern is repeated in the negative half-cycle. Ideally, the waveform should be sinusoidal, but the presence notably of the fifth and seventh harmonics (which leads to some unwelcome torque pulsations at six times the fundamental frequency) is the price to be paid for having a relatively unsophisticated inverter. Happily, resonances excited by the torque ripple are rare and can usually be overcome by preventing steady/continuous operation at particular speeds associated with the resonance.

It is necessary to maintain an approximately constant angular relationship between the rotor and stator m.m.f.s and hence automatically maintain the correct inverter frequency. This is an important point: the inverter does not impose a frequency upon the machine, rather the machine itself determines the frequency, and so the motor cannot lose synchronism (pole-slip). As already mentioned, the drive is accelerated by increasing the current fed to the motor, increasing the motor torque and acceleration, which thereby increases the frequency.

Once rotating, the motor terminal voltage (or better the calculated back e.m.f.) can be used to determine the rotor position, but at start-up and at very low speeds the magnitude of the voltage signals is too small to be used for control purposes or to commutate the current. A shaft-mounted absolute position sensor is therefore used to provide position information to determine when the current should be commutated from one switch to the next. The commutation itself is usually achieved by momentarily switching off the d.c. link current (by phase control of the utility-side converter) every sixth of a cycle. This allows the thyristors in the inverter to turn off so that the next pair can be fired. Above approximately 5% of rated speed the machine generates sufficient voltage for natural commutation, and subsequent control is undertaken in a similar manner to that of a d.c. drive.

As we saw earlier with the permanent magnet drive, where the load is predictable some systems actually impose a predetermined sequence of current pulses (applied sequentially to the motor phases) to 'crank' the motor up to a speed at which the back e.m.f. becomes of sufficient magnitude to be used for position sensing and commutation.

As in the d.c. drive, the a.c. supply power-factor is poor at low speeds, but on the plus side, full four-quadrant operation is possible without additional mechanical switching.

Applications for this type of drive fall into two main categories: first, as a starting mechanism for very large synchronous machines, the converter then being rated only for a fraction of the machine rating – the main motor is started off load, synchronized to the utility supply and then the load is applied; and secondly, as large high-power (and sometimes high-speed) variable-speed drives for a variety of applications. Power ratings, typically from 2 to 100 MW at speeds up to 8000 r.p.m. are available, though at powers up to 5 MW voltage source inverters are now proving to be more popular. It is also important to note that with the advent of very high-voltage thyristors, high-voltage drives can be readily designed. Supply voltages up to 12 kV are typical, but systems over 25 kV are in service where the high voltage converter technology is similar to that used for HVDC power-system converters. As powers increase, in order to keep both motor windings and interconnecting power cabling manageable it is important to increase operating voltages and thereby moderate the current levels.

Some manufacturers of synchronous machines of much more modest ratings (e.g. a few tens of kW) also offer this type of converter technology, but they tend to be niche products.

The detailed design of this type of high-power drive is highly specialized and well outside the scope of this book. The impact on the supply of such high-power converters is also an area where a great deal of detailed consideration is necessary.

7. PERFORMANCE OF BRUSHLESS MOTORS

Throughout this section we will use the terms 'brushless' and 'permanent magnet' to be synonymous when they are applied to an electric motor. We will also avoid the confusing term 'brushless d.c. motor' because, as already mentioned, such motors are always supplied with alternating currents.

However, we should mention that the inherent electromagnetic properties of a permanent magnet motor can be quantified by its 'motor constant' in much the same way as a d.c. motor (Chapter 3). If we spin the rotor of a permanent magnet machine at angular velocity ω, the r.m.s. value of the sinusoidal e.m.f. induced in each phase is given by $E = k\omega$, and if we supply balanced currents of r.m.s value I_a to the 3-phase windings in quadrature to the field, the torque is given by $T = kI_a$, where k is the machine constant, expressed in the SI units of volts per radian per second or the equivalent newton-meters per ampere. (In practice, manufacturers usually quote k in terms of volts per thousand revs per min.) These relationships are identical to those we discussed earlier for the d.c. machine, and once again they underline the unity of machines that operate on the '*BIl*' principle.

We have hinted previously that permanent magnet motors offer outstanding performance in terms of power density and performance in comparison with induction and d.c. machines, and in this section we look briefly at the underlying reasons. We then discuss the limitations that govern performance and finally give an example that illustrates the impressive results that can be obtained.

7.1 Advantages of permanent magnet motors

The stator windings of permanent magnet motors do not have to carry the excitation or magnetizing current required by the induction motor, so a given winding can carry a higher work current without generating more heat, thereby increasing the electric loading and the specific power output (as discussed in Chapter 1).

Cooling the rotor is difficult in any enclosed machine because ultimately the heat has to get to the stator, so the absence of current on the rotor not only improves efficiency by reducing the total copper loss, but also eases the cooling problem.

Historically, brushless permanent magnet motors only became practicable with the advent of power electronics, so it became normal for them to be supplied via power electronics with the associated expectation that they would operate in a speed-controlled drive. The majority were therefore not expected to operate directly off the utility supply, and as a result their designers had much greater freedom to produce bespoke designs, tailored to a specific purpose.

For example, suppose we require a motor that can accelerate very quickly, which implies that the ratio of torque to inertia should be maximized. We saw in Chapter 1 that with given values of the specific magnetic and electric loadings, the torque is broadly dependent on the volume of the rotor, so we are free to choose long and thin or short and fat. The inertia of a homogeneous rotor is proportional to the fourth power of its radius, so clearly for this application we want to minimize the rotor radius, so a long thin design is required. Fortunately, there is considerable flexibility in regard to the shape and size of the rotor magnets, so no serious constraint applies in relation to rotor diameter. Many so-called servo motors (see later) have this profile, as illustrated in Figure 9.12.

A different application with the same continuous torque and power requirements would require the same rotor volume, but if, for example, the application requires that unwanted changes in speed caused by step changes in load torque must be minimized, the inertia should clearly be maximized, with a short rotor of larger diameter.

A section through a high inertia permanent magnet motor typically rated up to 100 Nm and 3000 rev/min is shown in Figure 9.13 and Plate 9.1. It is interesting to

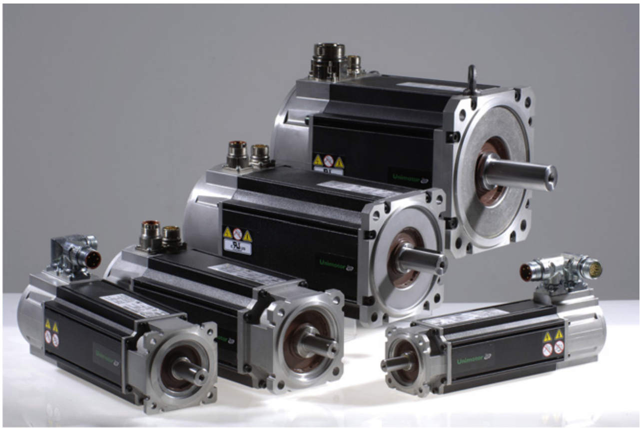

Figure 9.12 Permanent magnet servo motors. *(Courtesy of Emerson – Control Techniques Dynamics.)* (See Plate 9.12)

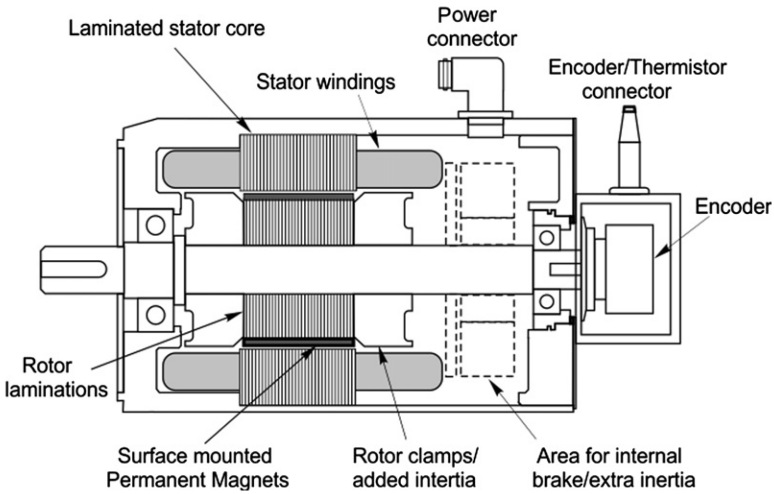

Figure 9.13 Typical high inertia brushless permanent magnet motor.

note how little of the volume of the motor is actually taken up with active material, namely the stator and rotor laminations and the surface-mounted magnets on the rotor. The end regions of the stator windings are seen to contribute significantly to the overall volume of the motor illustrated. (The move to using segmented stator windings reduces the impact of the end winding and can lead to significant reductions in total motor volume.) Most permanent magnet motors employ rare-earth magnets which have much higher energy product (in effect a measure of their magnetizing 'power') than traditional materials such as alnico, so they are very small, as shown in Figure 9.13.

As in other motors the heat dissipated in the stator diffuses into the air through the frame to the finned case and hence to the surrounding air. However, the design of some motors of this type is based on the requirement that a substantial proportion (perhaps 40%) of the loss is conducted through the mounting flange to a suitable heatsink, so this is an area where great care needs to be taken with the thermal properties of the mounting.

7.2 Industrial permanent magnet motors

In the preceding section we mentioned that bespoke design of permanent magnet motors has long been considered unexceptional, but recent years have seen the emergence of permanent magnet motors packaged in the same industrial motor (IEC or NEMA) housings that are being used for induction motors, as shown in Figure 9.14.

This type of motor is now marketed as a direct competitor of the induction motor in variable-speed applications. They are targeted at general applications

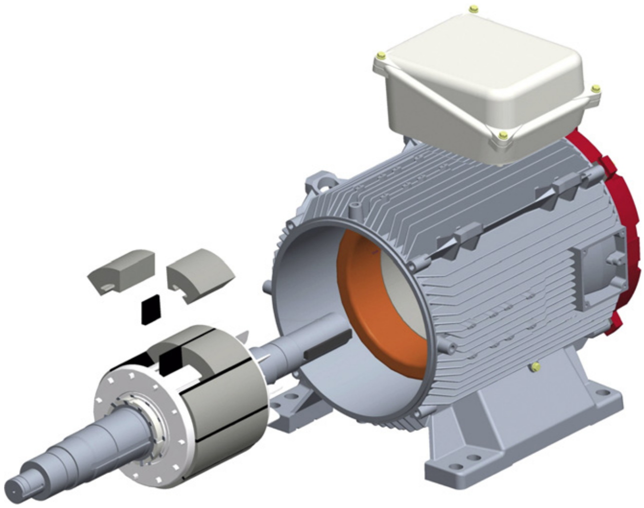

Figure 9.14 Permanent magnet industrial motor in standard IEC frame. *(Courtesy of Emerson Leroy Somer.)* (See Plate 9.14)

where the higher initial cost is offset by their higher efficiency and relatively high power density. The heat loss in the permanent magnet rotor is much less than in the corresponding induction motor, so the rotor runs cooler, which may also be an advantage in aspects such as bearing life.

7.3 Summary of performance characteristics

Permanent magnet motors with low inertia rotors are used in high-performance servo applications such as machine tools or pick and place applications where fast, precise movements are required, and motors with high inertia rotors (and high pole-numbers) suit low-speed applications such as gearless lift systems.

The performance characteristics of these drives are summarized below:

- They have excellent dynamic performance at speeds down to standstill when position feedback is used.
- For precision positioning the position feedback must define the absolute position uniquely within an electrical revolution of the motor. This can be provided with a position sensor or alternatively a sensorless scheme can be used. The performance of a sensorless scheme will be lower than when a position sensor is used.
- Field weakening of permanent magnet motors is possible to extend their speed range, but (as shown in section 5) this requires additional motor current, and so the motor becomes less efficient in the field weakening range. This form of control also increases the rotor losses and raises the temperature of the magnet material, thereby increasing the risk of demagnetization. Care is also needed in

such applications to avoid overvoltage at the motor and drive terminals in the event of a loss of control: at high speeds, the open-circuit voltage will exceed the rated value (see Figure 9.9).

- Permanent magnet motors exhibit an effect called cogging that results in torque ripple. It is caused by magnetic reluctance forces acting mainly in the teeth of the stator, and can be minimized by good motor design, but can still be a problem in sensitive applications.
- Permanent magnet motors can be very efficient as the rotor losses are very small.

To give an impression of the outstanding performance that can be achieved by a brushless permanent magnet motor, Figure 9.15 shows the results from a bench test in which the speed reference begins with a linear ramp from zero to 6000 rev/min in 0.06 s, followed shortly by a demand for the speed to reverse to 6000 rev/min, then back to full forward speed and finally to rest, the whole process lasting less than one second.

The motor was coupled to a high-inertia load of 78 times the rotor inertia, which makes the fact that the speed reversal is accomplished in only 120 ms even more remarkable. It takes only three revolutions to come to rest and a further three to accelerate in reverse. In common with many high-performance applications, the drive control is actually implemented in the form of position control, with the shaft angle being incremented at a rate equivalent to the required speed. The dotted trace in Figure 9.15 shows that the maximum position error of the motor shaft throughout the speed reversal is less than 0.05°, so by any standard this is truly impressive. It is no wonder that the brushless permanent magnet motor is frequently chosen for applications where closely coordinated motion control is called for.

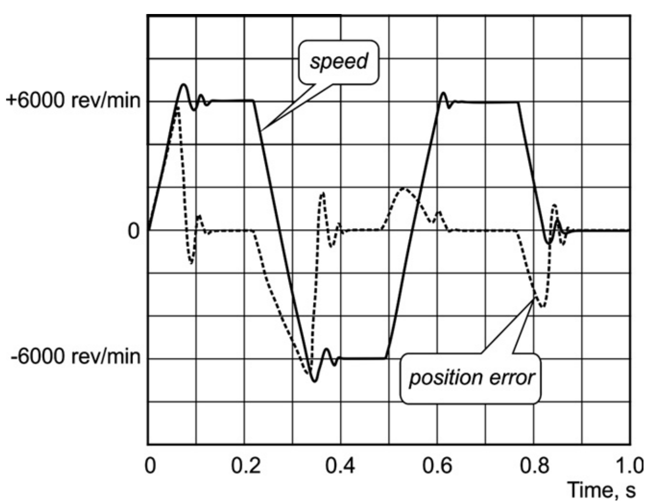

Figure 9.15 Permanent magnet motor drive performance under rapid reversal test.

Finally, on a matter of terminology, it is worth pointing out that brushless permanent magnet motors are sometimes referred to as 'servo' motors. The name 'servo' originates from 'servomechanism', defined as a mechanical or electrical system for control of speed or position. The term tends to be used loosely, but broadly speaking when it is applied to a motor it implies superior levels of performance.

7.4 Limits of operation of a brushless permanent magnet motor

We have previously talked about the limits of operation that determine the rating and operating envelope of other types of electrical machine, so we will conclude this section by taking a closer look at the limits of operation for a brushless permanent magnet servo motor (which usually has no external cooling fins or fan). A typical torque–speed characteristic is shown in Figure 9.16.

The individual limits shown in Figure 9.16 are discussed below, but the most striking feature is clearly the very large area where operation above rated torque is possible (albeit on an intermittent basis). This provision clearly reflects the potential application areas, such as rapid positioning systems, where high acceleration is needed for relatively short periods.

The continuously rated region is, as usual, limited by the allowable temperature rise of the motor. At standstill the predominant source of loss is the stator copper loss (shown as I^2R limit in Figure 9.16), but at higher speeds the iron loss becomes significant and the full-load torque at rated speed is therefore less than the standstill torque.

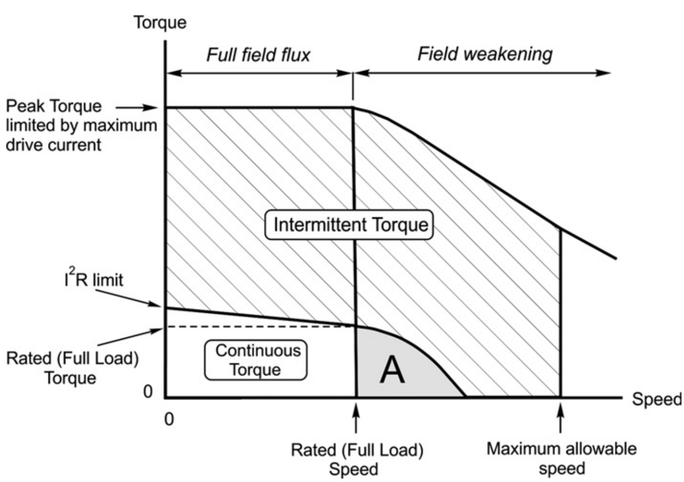

Figure 9.16 Limits of operation of a typical brushless permanent magnet servo motor.

The upper boundary on the intermittent torque region is usually determined by the maximum current that the drive converter can supply: only intermittent operation is possible otherwise the motor overheats.

Continuous operation is not possible when field weakening either, because, as explained in section 5, the stator current is large in order to reduce the magnet flux. For this thermal reason region A shown in Figure 9.16 can be somewhat truncated compared with the characteristic field weakening region of the induction motor.

Operation of permanent magnet motors at excessive temperatures can lead to demagnetization of the rare earth magnets in the motor, as well as the hazards common to all other electrical machines such as loss of robustness of insulation. Good thermal protection is therefore necessary. Below base speed, for applications involving few excursions outside the continuous operating region, a relatively simple motor thermal model in the drive control scheme may be adequate. For applications involving significant operation in the intermittent torque region, and certainly where field weakening operation is used, more complex thermal models would be needed and these are usually supplemented by thermistors embedded in the stator windings.

One final note of warning in relation to operating permanent magnet motors in the field weakening region: the maximum motor speed must be selected with due consideration of what may happen if the drive were to trip. In such a case the stator voltage could rise to a dangerously high level, as discussed previously, causing damage to the drive, the motor itself or indeed the cables or connectors. Crowbar circuits can be designed to provide protection against overvoltages, but the motor and drive system must be designed to withstand the short circuit currents that result.

7.5 Brushless permanent magnet generators

As with almost all forms of electrical machines the permanent magnet motor can be operated as a generator, with mechanical energy supplied to the shaft being converted to electrical energy. The advantages of high power density and efficiency offered by the permanent magnet motor of course apply equally when the machine is being used as a generator.

The majority of commercial wind generators up to 75 kW use permanent magnet synchronous machines. Much larger wind generators are also in service with some multi-pole, multi-MW motors being applied in utility-scale turbines with direct drive systems, i.e. systems which do not employ a gearbox between the wind turbine and the generator.

8. RELUCTANCE AND HYSTERESIS MOTORS

There are a number of synchronous motors which have niche applications and so a brief introduction is appropriate.

8.1 Reluctance motors

The reluctance motor is arguably the simplest synchronous motor of all, the rotor consisting simply of a set of laminations shaped so that it tends to align itself with the field produced by the stator. This 'reluctance torque' action is discussed in Chapter 10, which deals with stepping and switched reluctance motors.

Here we are concerned with utility-frequency reluctance motors, which differ from steppers in that they only have saliency on the rotor, the stator being identical with that of a 3-phase induction motor. In fact, since induction motor action is required in order to get the rotor up to synchronous speed, a reluctance-type rotor resembles a cage induction motor, with parts of the periphery cut away in order to force the flux from the stator to enter the rotor in the remaining regions where the air-gap is small, as shown in Figure 9.17(a). Alternatively, the 'preferred flux paths' can be imposed by removing iron inside the rotor so that the flux is guided along the desired path, as shown in Figure 9.17(b).

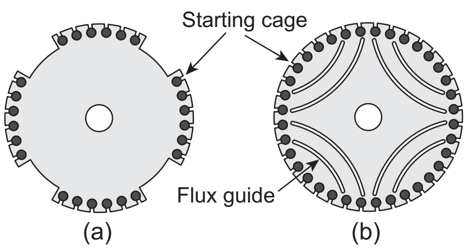

Figure 9.17 Reluctance motor rotors (4-pole) – (a) salient type, (b) flux-guided type.

The rotor will tend to align itself with the field, and hence is able to remain synchronized with the traveling field set up by the 3-phase winding on the stator in much the same way as a permanent magnet rotor. Early reluctance motors were invariably one or two frame sizes bigger than an induction motor for a given power and speed, and had low power-factor and poor pull-in performance. As a result they fell from favor except for some special applications such as textile machinery where cheap constant speed motors were required. Understanding of reluctance motors is now much more advanced, though their fundamental performance still lags the induction motor as regards power output, power-factor and efficiency.

Historically, reluctance motors were frequently used in large numbers connected in parallel to large 'bulk' power electronic inverters. The fact that the motors were locked in synchronism was a great advantage over the induction motor for synchronized/coordinated applications, notably in the textile industry. However, the falling cost of power electronics and advances in control saw such applications being addressed by means of individual inverters with dedicated induction motors and overarching control. Recently, however, some drives companies are once again promoting reluctance motor drives utilizing off-the-shelf converter and control hardware, perhaps stretched to the limit, to obtain performance comparable to that of an induction motor.

8.2 Hysteresis motors

Whereas most motors can be readily identified by inspection when they are dismantled, the hysteresis motor is likely to baffle anyone who has not come across it before. The rotor consists simply of a thin-walled cylinder of what looks like steel, while the stator has a conventional single-phase or 3-phase winding. Evidence of very weak magnetism may just be detectable on the rotor, but there is no hint of any hidden magnets as such, and certainly no sign of a cage. Yet the motor runs up to speed very sweetly and settles at exactly synchronous speed with no sign of a sudden transition from induction to synchronous operation.

These motors (the operation of which is quite complex) rely mainly on the special properties of the rotor sleeve, which is made from a hard steel which exhibits pronounced magnetic hysteresis. Normally in machines we aim to minimize hysteresis in the magnetic materials, but in these motors the effect (which arises from the fact that the magnetic flux density B depends on the previous 'history' of the m.m.f.) is deliberately accentuated to produce torque. There is actually also some induction motor action during the run-up phase, and the net result is that the torque remains roughly constant at all speeds.

Small hysteresis motors were once used extensively in office equipment, fans, etc. The near constant torque during run-up and the very modest starting current (of perhaps 1.5 times rated current) mean that they are also suited to high inertia loads such as gyro compasses and small centrifuges.