16.1 Introduction and synopsis

Migrating birds, some think, navigate using the earth's magnetic field. This may be debatable, but what is beyond question is that sailors, for centuries, have navigated in this way, using a natural magnet, *lodestone*, to track it. Lodestone is a mineral, magnetite (Fe_3O_4), that sometimes occurs naturally in a magnetised state. Today we know lodestone is a *ferri-magnetic* ceramic, or *ferrite*, examples of which can be found in every radio, TV set, tablet, and microwave oven. *Ferro-magnetic* materials, by contrast, are metals, typified by iron but including also nickel, cobalt, and alloys of all three. When placed in a magnetic field, these materials become magnetised, a phenomenon called *magnetic induction*; on removal of the field, some (the 'soft' magnets) lose their magnetisation, while others (the 'hard' magnets) retain it.

Magnetic fields are created by moving electric charge – electric currents in electro-magnets, electron spin in the atoms of magnetic materials. This chapter is about magnetic materials: how they are characterised, where their properties come from, and how they are selected and used. It starts with definitions of magnetic properties and the way they are measured. As in other chapters, charts display them well, separating the materials that are good for one sort of application from those that are good for others. The chapter continues by drilling down to the origins of magnetic behaviour and concludes with a discussion of applications and the materials that best fill them.

There are many different kinds of magnetic behaviour. Figure 16.1 gives an overview, with examples of materials and their applications.

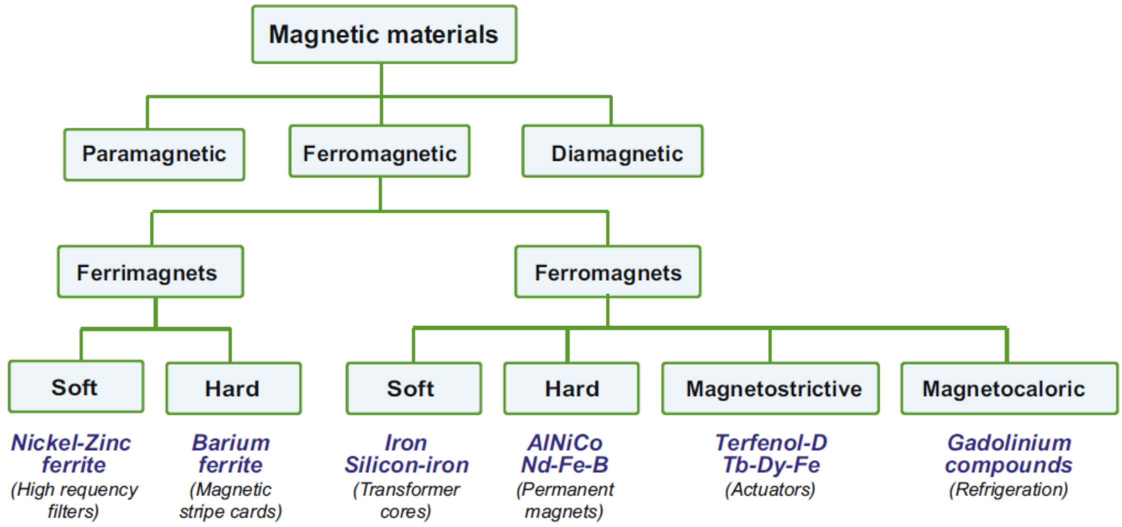

Figure 16.1 Magnetic behaviours, materials, and applications.

16.2 Magnetic properties: definition and measurement

Magnetic fields in vacuum First, some definitions. When a current i passes through a long, empty coil of n turns and length L as in Figure 16.2, a magnetic field is generated. The magnitude of the field, H, is given by Ampère's law as

$$H = \frac{ni}{L} \tag{16.1}$$

and thus has units of amps/metre (A/m). The field has both magnitude and direction – it is a vector field.

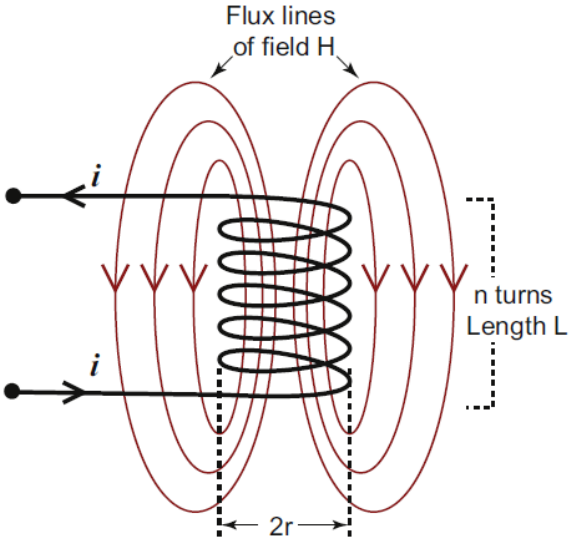

Figure 16.2 A solenoid creates a magnetic field, H.

Magnetic fields exert forces on a wire carrying an electric current. A current i flowing in a single loop of area S generates a dipole moment m where

$$m = iS \tag{16.2}$$

with units A.m^2, and it too is a vector with a direction normal to the loop S (Figure 16.3). If the loop is placed at right angles to the field H it feels a torque T (units: newton metre, or N.m) of

$$T = \mu_o mH \tag{16.3}$$

where μ_o is called the *permeability of vacuum*, $\mu_o = 4\pi \times 10^{-7}$ henry/metre (H/m). To link these we define a second measure of the magnetic field, one that relates directly to the torque it exerts on a unit magnetic moment. It is called the *magnetic induction* or *flux density*, B, and for vacuum or non-magnetic materials it is

$$B = \mu_o H \tag{16.4}$$

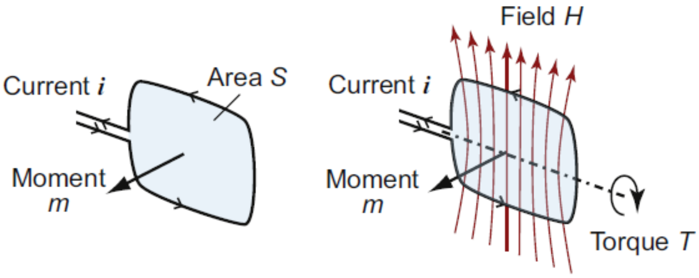

Figure 16.3 Magnetic moment and moment-field interaction.

Its units are *tesla*, T, so a tesla[1] is 1 HA/m^2. A magnetic induction B of 1 T exerts a torque of 1 N.m on a unit dipole at right angles to the field H.

Magnetic fields in materials If the space inside the coil of Figure 16.2 is filled with a material, as in Figure 16.4, the induction within it changes. This is because its atoms respond to the field by forming and aligning little magnetic dipoles in ways that are explained in Section 16.4. The material acquires a macroscopic dipole moment or *magnetisation*, M (its units are A/m, like H). The induction becomes

$$B = \mu_o (H + M) \tag{16.5}$$

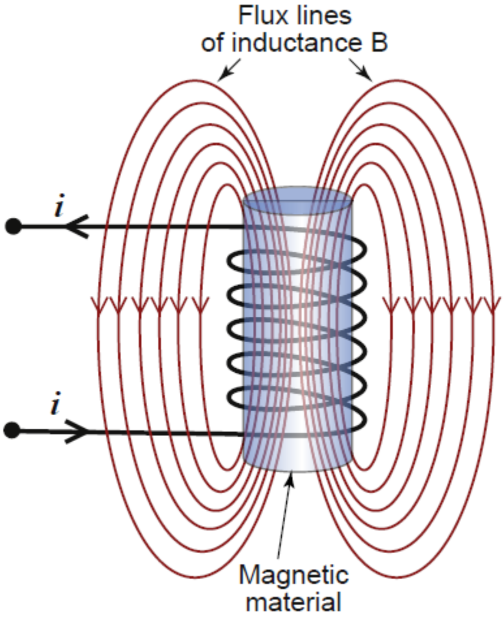

Figure 16.4 A magnetic material exposed to a field H becomes magnetised.

The simplicity of this equation is misleading, since it suggests that M and H are independent; but in reality M is the response of the material to H, so the two are coupled. If the material of the core is *ferro-magnetic*, the response is a very strong one, and it is non-linear, as we shall see in a moment. It is usual to rewrite equation (16.5) in the form

$$B = \mu_R \, \mu_o \, H$$

where μ_R is called the *relative permeability*, and like the relative permittivity (the dielectric constant) of Chapter 15, it is dimensionless. The magnetisation, M, is thus

$$M = (\mu_R - 1) \, H = \chi H \tag{16.6}$$

where χ is the *magnetic susceptibility*. Neither μ_R nor χ are constants – they depend not only on the material but also on the magnitude of the field, H, for the reason just given.

---

[1] Nikola Tesla (1856–1943), Serbian-American inventor, discoverer of rotating magnetic fields, the basis of most alternating current machinery, inventor of the Tesla coil and of a rudimentary flying machine (never built).

The magnetic pressure P_{mag} (force per unit area) exerted by an electromagnet on a section of core material is

$$P_{mag} = \frac{1}{2}\mu_o H^2 = \frac{1}{2}\frac{B^2}{\mu_o} \tag{16.7}$$

provided that the core does not saturate. Iron has a high saturation magnetisation H_s of about 1.8×10^6 A/m, so the maximum pressure exerted by an electro-magnet with an iron core is

$$P_{mag} = \frac{1}{2} \times 4\pi \times 10^{-7} \times \left(1.8 \times 10^6\right)^2 = 2.9 \times 10^6 \ \mathrm{N/m^2} = 2 \ \mathrm{MPa}$$

The energy density stored in a magnetic field per unit volume, U/V, has the same form as the magnetic pressure:

$$\frac{U}{V} = \frac{1}{2}\frac{B^2}{\mu_o} = \frac{1}{2}\mu_o H^2 \ \mathrm{J/m^3} \tag{16.8}$$

Ferro-magnetic and ferri-magnetic materials Nearly all materials respond to a magnetic field by becoming magnetised, but most are paramagnetic with a response so faint that it is of no practical use. A few, however, contain atoms that have large dipole moments and the ability to spontaneously magnetise — to align their dipoles in parallel — as electric dipoles do in ferro-electric materials. These are called *ferro-magnetic* and *ferri-magnetic* materials (the second are called *ferrites* for short), and it is these that are of real practical use.

Magnetisation decreases with increasing temperature. Just as with ferro-electrics, there is a temperature, the Curie temperature T_c, above which it disappears, as in Figure 16.5. Its value for the materials we shall meet here is well above room temperature (typically 300–1000°C), but making magnets for use at really high temperatures is a problem.

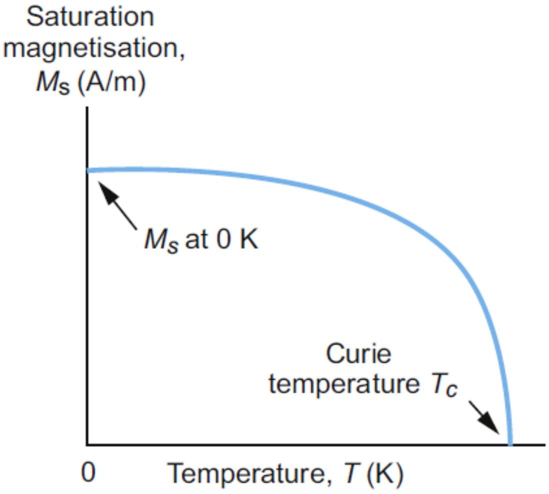

Figure 16.5 Saturation magnetisation decreases with increasing temperature.

Measuring magnetic properties The magnetic behaviour of a material is characterised by its $M-H$ curve (Figure 16.6). If an increasing field H is applied to a previously demagnetised sample, starting at A on the figure, its magnetisation increases, slowly at first and then faster, following the broken line, until it finally tails off to a maximum, the *saturation magnetisation* M_s at point B. If the field is now reduced, M does not retrace its original path; some magnetisation is retained when H reaches zero at point C. The residual magnetisation is called the *remanent magnetisation* or *remanence*, M_r, and is usually only a little less than M_s. To decrease M further, we must increase the field in the opposite direction until M finally passes through zero at point D when the field is $-H_c$, the *coercive field*, a measure of the resistance to demagnetisation. Beyond point D the magnetisation M starts to increase in the opposite direction, eventually reaching saturation again at point E. If the field is now decreased again, M follows the curve through F and G back to full forward magnetic saturation again at B to form a closed $M-H$ circuit, or *hysteresis loop*.

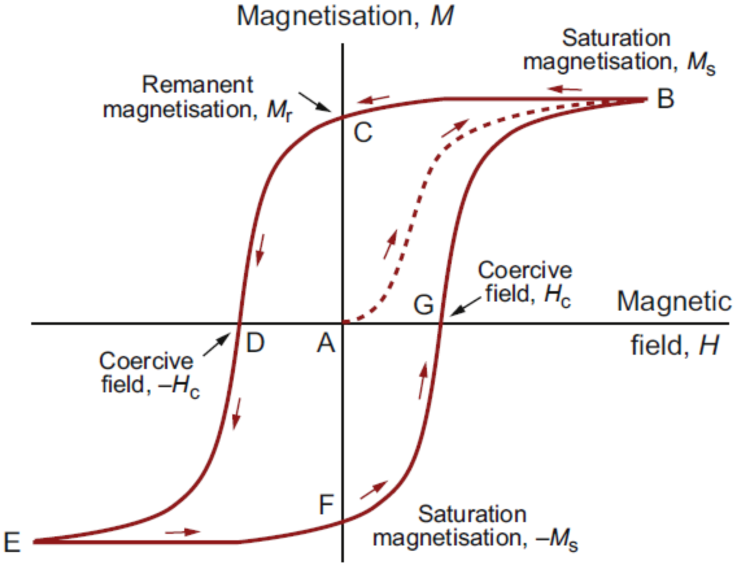

Figure 16.6 A magnetic hysteresis curve.

Magnetic materials are characterised by the size and shape of their hysteresis loops. The initial segment AB is called the *initial magnetisation curve*, and its average slope (or sometimes its steepest slope) is the magnetic susceptibility, χ. The other key properties – the saturation magnetisation M_s, the remanence M_r, and the coercive field H_c – have already been defined. Each full cycle of the hysteresis loop dissipates an energy per unit volume equal to the area of the loop multiplied by μ_o, the permeability of vacuum. This energy appears as heat; it is like magnetic friction. Many texts plot the curve of inductance B against H, rather than the $M-H$ curve. Equation (16.5) says that B is proportional to $(M + H)$. Since the value of M (for any magnetic materials worthy of the name) is very much larger than the H applied to it, $B \approx \mu_o M$ and the $B-H$ curve of a ferro-magnetic material looks very like its $M-H$ curve (it's just that the M axis has been scaled by μ_o).

There are several ways to measure the hysteresis curve, one of which is sketched in Figure 16.7. Here the material forms the core of what is in effect a transformer. The oscillating current through the primary coil creates a field H that induces magnetisation M in the material of the core, driving it round its hysteresis loop. The secondary coil picks up the inductance, from which the instantaneous state of the magnetisation can be calculated, mapping out the loop.

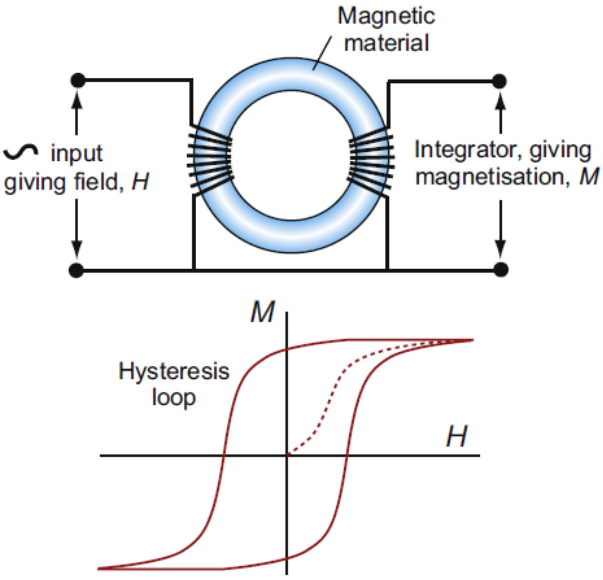

Figure 16.7 Measuring the hysteresis curve.

Magnetic materials differ greatly in the shape and area of their hysteresis loop, the greatest difference being that between *soft magnets*, which have thin loops, and *hard magnets*, which have fat ones, as sketched in Figure 16.8. In fact, the differences are much greater than this figure suggests: the coercive field H_c (which determines the width of the loop) of hard magnetic materials like alnico is greater by a factor of about 10^5 than that of soft magnetic materials like silicon–iron.

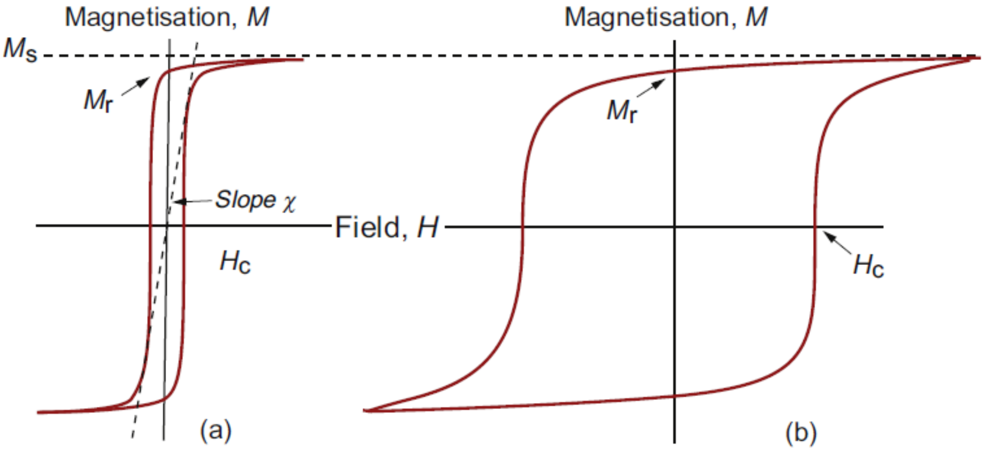

Figure 16.8 Magnetisation hysteresis loops: (a) a thin, square loop typical of soft magnets, (b) a fat loop typical of hard magnets.

16.3 The big picture: charts for magnetic properties

The remanence–coercive field chart The differences between families of soft and hard magnetic materials are brought out by the chart in Figure 16.9. The axes are remanent magnetisation M_r and coercive field H_c. The saturation magnetisation M_s, more relevant for soft magnets, is only slightly larger than M_r, so an M_s–H_c chart looks almost the same as this one. There are 12 distinct families of magnetic materials, each enclosed in a coloured envelope. The members of each family, shown as smaller ellipses, have unhelpful tradenames (such as 'H Ferrite YBM-1A') so they are not individually labelled. Soft magnets require high M_s and low H_c; they are the ones on the left, with the best near the top. Hard magnets must hold their magnetism, requiring a high H_c, and to be powerful they need a large M_r; they are the ones on the right, with the best again at the top.

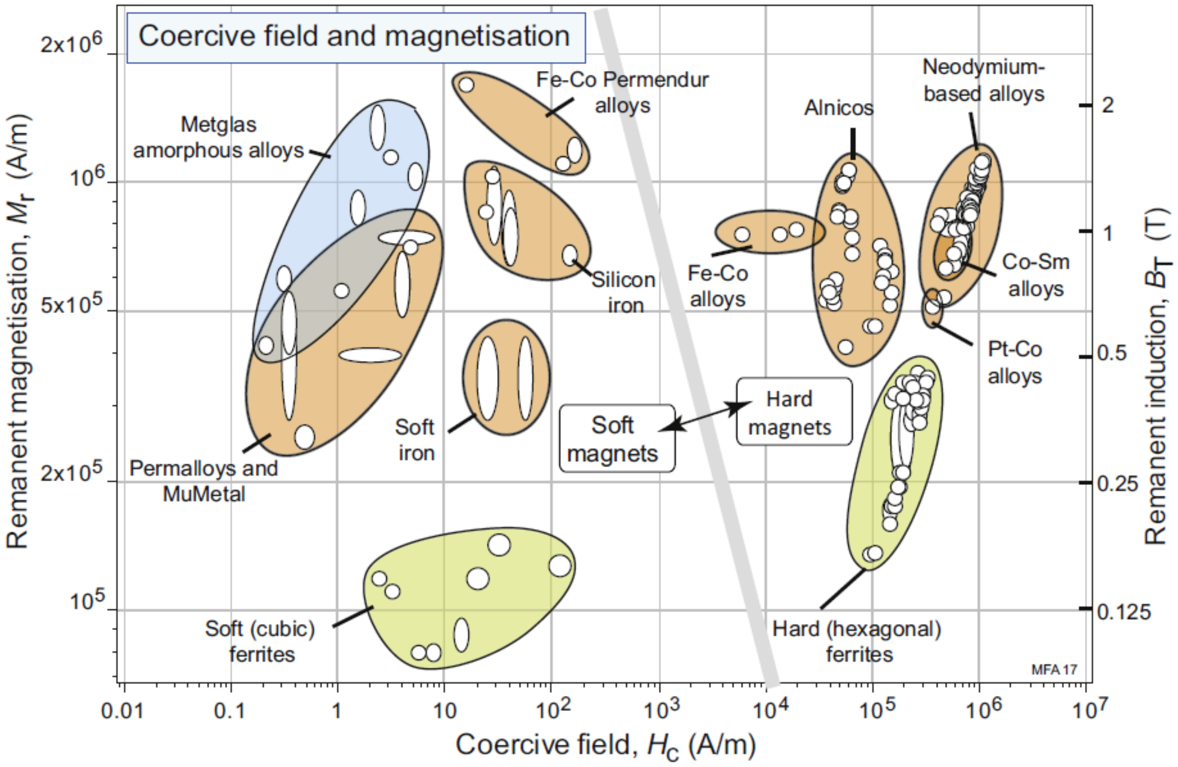

Figure 16.9 Remanent magnetisation and coercive field. Soft magnetic materials lie on the left, hard magnetic materials on the right. The orange and blue envelopes enclose electrically conducting materials, the green ones enclose insulators.

In many applications the electrical resistivity is also important because of eddy-current losses, which will be described later. Changing magnetic fields induce eddy currents in conductors but not in insulators. The orange and blue envelopes on the chart enclose metallic ferro-magnetic materials, which are good electrical conductors. The green ones describe ferrites, which are electrical insulators.

The saturation magnetisation–susceptibility chart This is the chart for selecting soft magnetic materials (Figure 16.10). It shows the saturation magnetisation, M_s – the ultimate degree to which a material can concentrate magnetic flux – plotted against its susceptibility, χ, which measures the ease with which it can be magnetised. Many texts use, instead, the saturation inductance B_s and maximum relative permeability, μ_R. They are shown on the other axes of the chart. As in the first chart, materials that are electrical conductors are enclosed in orange and blue envelopes; those that are insulators have green ones.

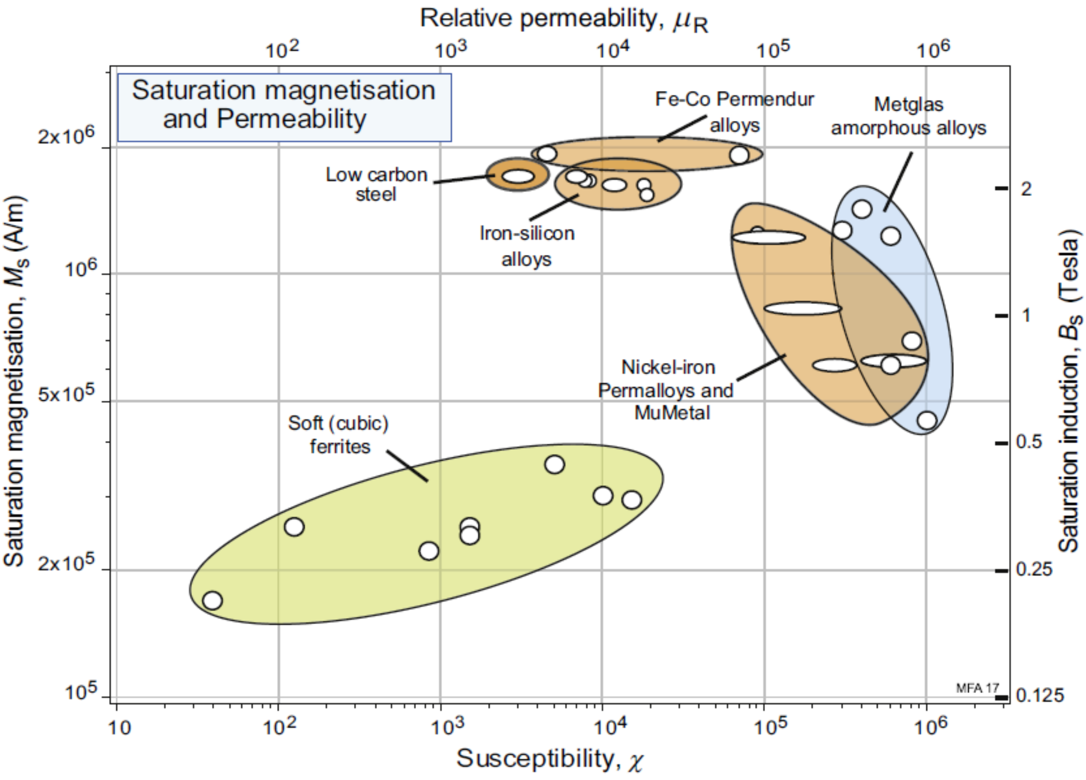

Figure 16.10 Saturation magnetisation, susceptibility, and relative permeability for soft magnetic materials. The orange and blue envelopes enclose electrically conducting materials, the green ones enclose insulators.

The coercive field–Curie temperature chart Two things can disrupt ferro-magnetic behaviour: temperature and demagnetising fields. The ferro-magnetic response disappears above the Curie temperature, which, as the chart of Figure 16.11 shows, is as low as 100°C for some materials and as high as 1000°C for others. Remanent magnetisation is cancelled or reversed by fields exceeding the coercive field, a risk when permanent magnets must operate in environments with stray fields. The chart is useful for selecting magnetic materials for use at elevated temperatures – such as actuators in control equipment of engines, for instance. As before, the colours distinguish conducting ferro-magnetic metals from insulating ferrites.

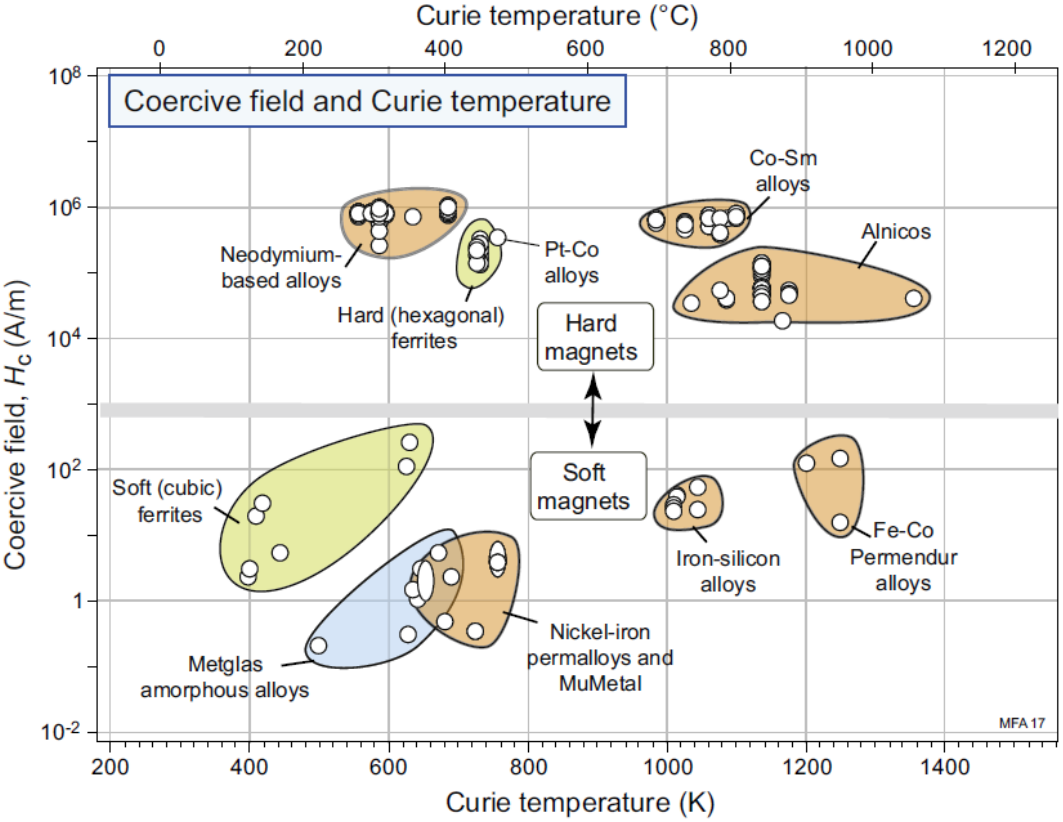

Figure 16.11 Coercive field and Curie temperature. The orange and blue envelopes enclose electrically conducting materials, the green ones enclose insulators.

16.4 Drilling down: the physics and manipulation of magnetic properties

The classical picture of an atom is that of a nucleus around which swing electrons, as in Figure 16.12. Moving charge implies an electric current, and an electric current flowing in a loop creates a magnetic moment, as in Figure 16.3. Thus each orbiting electron creates a magnetic dipole. That is not all. Each electron has an additional moment of its own: its spin moment. A proper explanation of this requires quantum mechanics, but a way of envisaging its origin is to think of an electron not as a point charge but as slightly spread out and spinning on its own axis, again creating a rotating charge and a dipole moment – and this turns out to be large. The total moment of the atom is the sum of the whole lot.

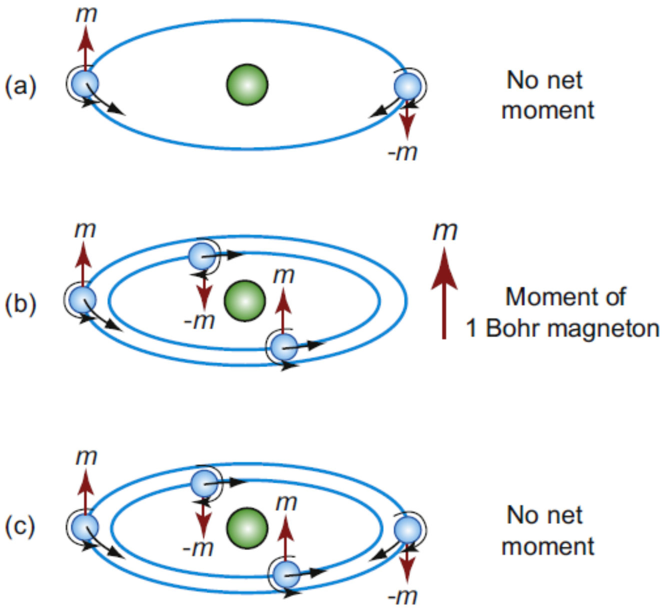

Figure 16.12 Orbital and electron spins create a magnetic dipole. Even numbers of electrons filling energy levels in pairs have moments that cancel, as in (a) and (c). An unpaired electron gives the atom a permanent magnetic moment, as in (b).

A simple atom like helium has two electrons per orbit, and they configure themselves such that the moment of one exactly cancels the moment of the other, leaving no net moment – as in Figure 16.12(a), which shows helium, or in Figure 16.12(c). But now think of an atom with three electrons as in Figure 16.12(b). The moments of two may cancel, but there remains the third, leaving the atom with a net moment represented by the red arrow at the right. Thus atoms with electron moments that cancel are non-magnetic; those with electron moments that don't cancel carry a magnetic dipole. Simplifying a little, one unpaired electron gives a magnetic moment of 9.27×10^{-24} A.m^2, called a Bohr[2] magneton; two unpaired electrons give two Bohr magnetons, three gives three, and so on.

---

[2] Niels Henrik David Bohr (1885–1962), Danish theoretical physicist, elucidator of the structure of the atom, contributor to the Manhattan Project, and campaigner for peace.

Think now of the magnetic atoms assembled into a crystal. In most materials the atomic moments interact so weakly that thermal energy is enough to randomise their directions, as in Figure 16.13(a) and the structure as a whole has no magnetic moment; these materials are *paramagnetic*. In a few materials, though, something quite different happens. The fields of neighbouring atoms interact such that their energy is reduced if their magnetic moments line up. This drop in energy is called the *exchange energy*, and it is strong enough that it beats the randomising effect of thermal energy so long as the temperature is not too high. The shape of the Curie curve of Figure 16.5 showed how thermal energy overwhelms the exchange energy as the Curie temperature is approached. They may line up anti-parallel, head to tail, so to speak, as in Figure 16.13(b), and we still have no net magnetic moment; such materials are called *anti-ferro-magnets*. But in a few elements, notably iron, cobalt and nickel, exactly the opposite happens: the moments spontaneously align so that – if all are parallel – the structure has a net moment that is the sum of those of all the atoms it contains (Figure 16.13(c)). These materials are *ferro-magnetic*. Iron has three unpaired electrons per atom, cobalt has two, and nickel just one, so the net moment if all the spins are aligned (the saturation magnetisation M_s) is greatest for iron, less for cobalt, and still less for nickel.

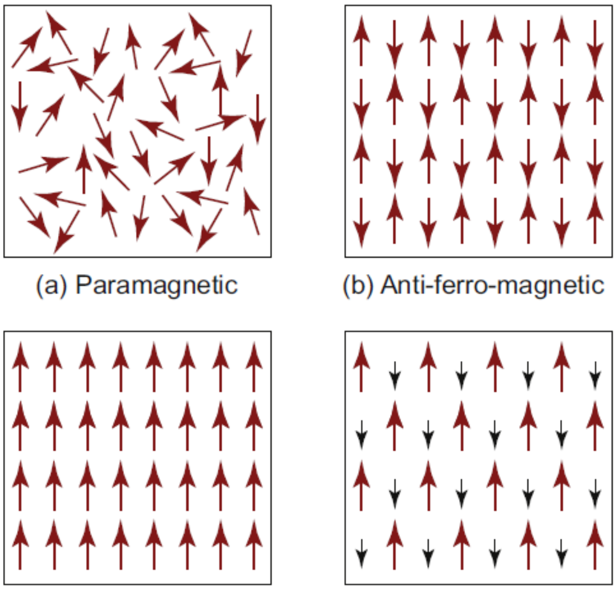
(a) Paramagnetic

(b) Anti-ferro-magnetic

(c) Ferro-magnetic

(d) Ferri-magnetic

Figure 16.13 Types of magnetic behaviour.

Compounds give a fourth possibility. The materials we have referred to as *ferrites* are oxides; one class of them has the formula MFe_2O_4 where M is also a magnetic atom, such as cobalt, Co. Both the Fe and the Co atoms have dipoles but they differ in strength. They line up in the anti-parallel configuration, but because the moments differ, the cancellation is incomplete, leaving a net moment M; these are *ferri-magnets*, or ferrites for short (Figure 16.13(d)). The partial cancellation and the smaller number of magnetic atoms per unit volume mean they have lower saturation magnetisation than, say, iron. But as they are oxides they have other advantages, notably that they are electrical insulators.

Domains If atomic moments line up, shouldn't every piece of iron, nickel or cobalt be a permanent magnet? Although they are magnetic materials they do not necessarily have a magnetic moment. Why not?

A uniformly magnetised rod creates a magnetic field, H, like that of a solenoid. The field has an energy

$$U = \frac{1}{2}\mu_o \int_V H^2 dV \tag{16.9}$$

where the integral is carried out over the volume V within which the field exists. Working out this integral is not simple, but we don't need to do it. All we need to note is that the smaller is H, and the smaller the volume V that it invades, the smaller is the energy. If the structure can arrange its moments to minimise its overall H or get rid of it entirely (remembering that the exchange energy wants neighbouring atom moments to stay parallel) it will try to do so.

Figure 16.14 illustrates how this can be done. The material splits up into *domains* within which the atomic moments are parallel, but with a switch of direction between mating domains to minimise the external field. The domains meet at *domain walls*, regions a few atoms thick in which the moments swing from the orientation of one domain to that of the other. Splitting into parallel domains of opposite magnetisation, as at Figure 16.14(b) and (c), reduces the field substantially; adding end regions magnetised perpendicular to both as at (d) kills it almost completely. The result is that most magnetic materials, unless manipulated in some way, adopt a domain structure with minimal external field – which is the same as saying that, while magnetic, they are not magnets.

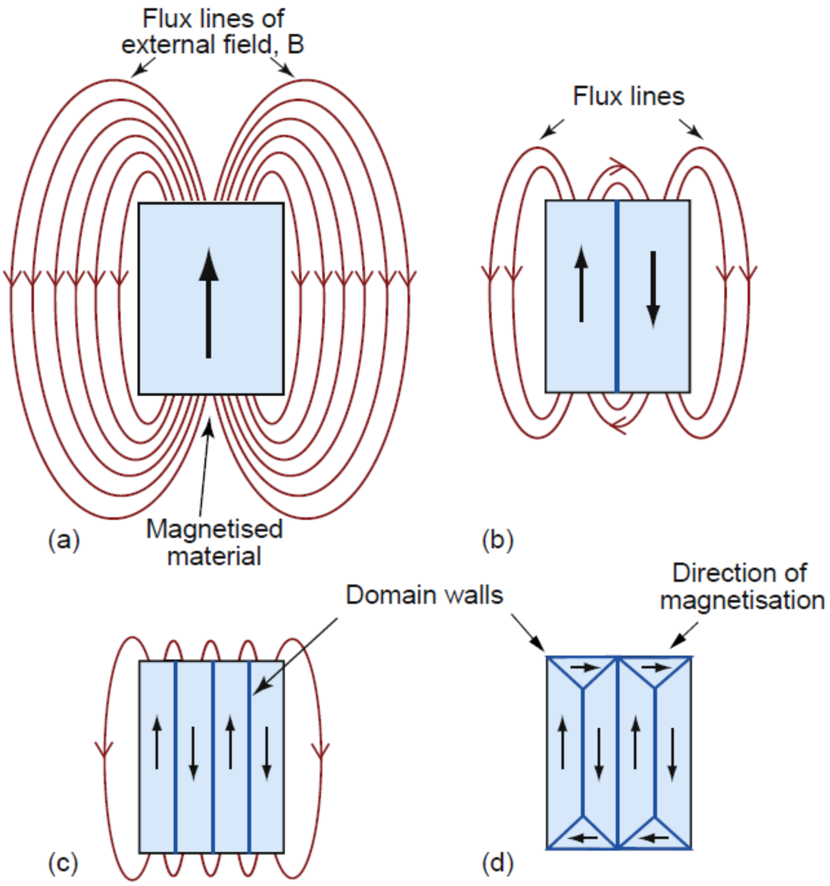

Figure 16.14 Domains cancel the external field while retaining magnetisation of the material itself.

How can they be magnetised? Figure 16.15, left, shows the starting domain wall structure. When placed in a magnetic field, created, say, with a solenoid, the domains already aligned with the field have lower energy than those aligned against it. The domain wall separating them feels a force, the *Lorentz*[3] *force*, pushing it in a direction to make the favourably oriented domains grow at the expense of the others. As they grow, the magnetisation of the material increases, moving up the $M-H$ curve, finally saturating at M_s when the whole sample is one domain, oriented parallel to the field, as in the right of Figure 16.15.

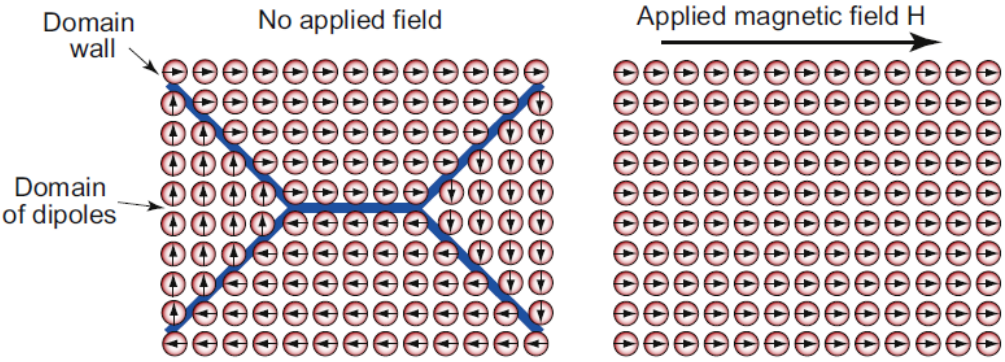

Figure 16.15 An applied field causes domain boundaries to move. At saturation the sample has become a single domain, as on the right. Domain walls move easily in soft magnetic materials. Hard magnetic materials are alloyed to prevent their motion.

The saturation magnetisation, then, is just the sum of all the atomic moments contained in a unit volume of material when they are all aligned in the same direction. If the magnetic dipole per atom is $n_m m_B$ (where n_m is the number of unpaired electrons per atom, and m_B is the magnetic moment of a Bohr magneton) then the saturation magnetisation is

$$M_s = \frac{n_m m_B}{\Omega} \tag{16.10}$$

where Ω is the atomic volume. Iron has the highest M_s because its atoms have the most unpaired electrons and they are packed close together. Anything that reduces either one reduces the saturation magnetisation. Nickel and cobalt have lower saturation because their atomic moments are less; alloys of iron tend to be lower because the non-magnetic alloying atoms dilute the magnetic iron. Ferrites, which are oxides, have much lower saturation, both because of dilution by oxygen and because of the partial cancellation of atomic moments sketched in Figure 16.13(d).

There are, nonetheless, good reasons for alloying even if it dilutes the dipoles. A good permanent magnet retains its magnetisation even when placed in a reverse field. That's achieved by pinning the domain walls to stop them from moving. Impurities, precipitates, and porosity all interact with domain walls, tending to pin them in place. They act as obstacles to dislocation motion too (Chapter 6), so magnetically hard materials are mechanically hard as well.

---

[3] Hendrik Antoon Lorentz (1853–1928), Dutch mathematical physicist, friend and collaborator with Raleigh and Einstein, originator of key concepts in optics, electricity, relativity, and hydraulics (he modelled and predicted the movement of water caused by the building of the Zuyderzee).

And there are subtler barriers to domain wall motion. One is *magnetic anisotropy*, arising because certain directions of magnetisation in the crystal are more favourable than others. Another relates to the shape and size of the sample itself: small, single-domain magnetic particles are particularly resistant to having their magnetism reversed, which is why they can be used for information storage.

The performance of a hard magnet is measured by its *energy product*, roughly proportional to the product $B_r H_c$, the axes of Figure 16.9. The higher the energy product, the more difficult it is to demagnetise it.

16.5 Materials selection for magnetic design

Selection of materials for their magnetic properties follows the same process of translating design requirements followed by screening and ranking. In this case, however, the first screening step — 'must be magnetic' — immediately isolates the special set of materials shown in the property charts of this chapter. The key selection issues are then:

- whether a soft or hard magnetic response is required;
- whether a fat or a thin $M-H$ hysteresis loop is wanted;
- what frequency of alternating field the material must carry;
- whether the material must operate above room temperature in a large external field.

Soft magnetic devices Electromagnets, transformers, electron lenses, and the like have magnetic cores that must magnetise easily when the field switches on, be capable of creating a high flux density, yet lose their magnetism when the field is switched off. They do this by using soft, low H_c, magnetic materials, shown on the left of the chart of Figure 16.8. Soft magnetic materials are used to 'conduct' and focus magnetic flux (Figure 16.16). A current-carrying coil induces a magnetic field in the core. The field, instead of spreading in the way shown in Figure 16.4, is trapped in the extended ferro-magnetic circuit. The higher the susceptibility χ, the greater is the magnetisation and flux density induced by a given current loop, up to a limit set by its saturation magnetisation M_s. The magnetic circuit conducts the flux from one coil to another, or from a coil to an air gap where it is further enhanced by reducing the section of the core. The field in the gap may be used, for example, to actuate or to focus an electron beam. Thus the first requirement of a good soft magnet is a high χ and a high M_s. These are the two properties plotted in Figure 16.10. Permendur (Fe—Co—V alloys) and Metglas amorphous alloys are particularly good, but being expensive they are used only in small devices. Silicon—iron (Fe 1—4% Si) is much cheaper and easier to make in large sections; it is the staple material for large transformers.

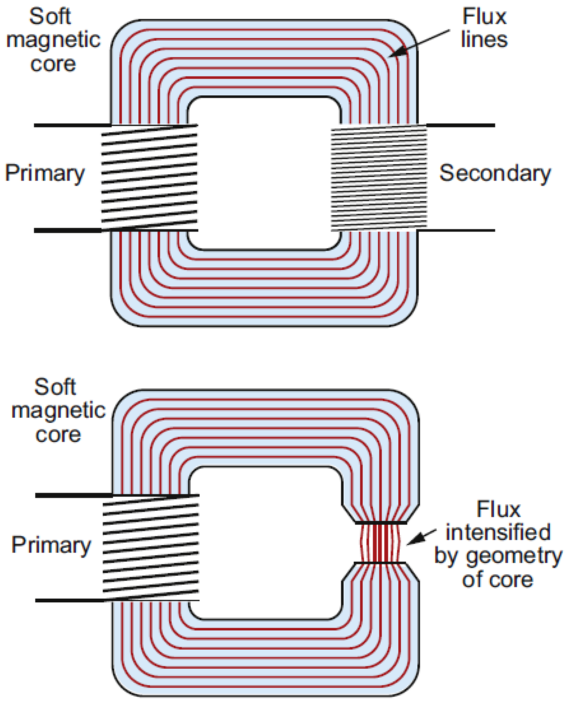

Figure 16.16 Soft magnets 'conduct' magnetic flux and concentrate it if the cross-section of the magnet is reduced.

Most soft magnets are used in AC devices, and then energy loss, proportional to the area of the hysteresis loop, becomes a consideration. Energy loss is of two types: the hysteresis loss already mentioned, proportional to switching frequency f, and eddy current losses caused by the currents induced in the core itself by the AC field, proportional to f^2. In cores of transformers operating at low frequencies (meaning f up to 100 Hz or so) the hysteresis loss dominates, so for these we seek high M_s and a thin loop with a small area. Eddy current losses are present but are suppressed when necessary by *laminating*: making the core from a stack of thin sheets interleaved with thin insulating layers that interrupt the eddy currents. At audio frequencies ($f < 50$ kHz) losses of both types become greater because they occur on every cycle and now there are more of them, requiring an even narrower loop and thinner laminates; here the more expensive permalloys and Permendurs are used. At higher frequencies still (f in the MHz range) eddy current loss, rising with f^2, becomes very large and the only way to stop it is to use magnetic materials that are electrical insulators. The best choice is ferrites, shown in green envelopes on all three charts, even though they have a lower M_s. Above this ($f > 100$ MHz) only the most exotic ceramic magnets will work. Table 16.1 lists the choices, guided by the charts.

Table 16.1 Materials for soft magnetic applications

Application	Frequency f	Material requirements	Material choice
Electro-magnets	<1 Hz	High M_s, high χ	Silicon iron Fe–Co alloys (Permendur)
Motors, low-frequency transformers	<100 Hz	High M_s, high χ	Silicon–iron AMA (amorphous Ni–Fe–P alloys)
Audio amplifiers, speakers, microphones	<100 kHz	High M_s, very low H_c, and high χ	Ni–Fe (permalloy, MuMetal)
Microwave and UHF applications	$<$MHz	High M_s, very low H_c, electrical insulator	Cubic ferrites: $M\mathrm{Fe_2O_4}$ with $M = \mathrm{Cu/Mn/Ni}$
Gigahertz devices	>100 MHz	Ultra-low hysteresis, excellent insulator	Garnets

Hard magnetic devices Many devices use permanent magnets: fridge door seals, magnetic clutches, loudspeakers and earphones, DC motors (as in hybrid and electric cars), and power generation (as in wind turbines). For these, the key property is the remanence, M_r, since this is the maximum magnetisation the material can offer without an imposed external field to keep it magnetised. High M_r, however, is not all. When a material is magnetised and the magnetising field is removed, the field of the magnet itself tries to demagnetise it. It is this demagnetising field that makes the material take up a domain structure that reduces or removes the net magnetisation of the sample. A permanent magnet must be able to resist being demagnetised by its own or other fields. A measure of its resistance is its coercive field, H_c, and this too must be large. A high M_r and a high H_c mean a fat hysteresis loop as on the right of Figure 16.8. Ordinary iron and steel can become weakly magnetised with the irritating result that your screwdriver sticks to screws and your scissors pick up paper clips when you wish they wouldn't, but neither are 'hard' in a magnetic sense. Here 'hard' means 'hard to demagnetise', and that is where high H_c comes in; it protects the magnet from itself.

Hard magnetic materials lie on the right of the M_r–H_c chart of Figure 16.9 and the upper half of Figure 16.11. The workhorses of the permanent magnet business are the Alnicos and the hexagonal ferrites. An Alnico (there are many) is an alloy of aluminium, nickel, cobalt and iron – as mentioned earlier – with values of M_r as high as 10^6 A/m and a huge coercive field: 5×10^4 A/m. The hexagonal ferrites, too, have high M_r and H_c and, like the cubic soft magnet ferrites, they are insulators. Being metallic, the Alnicos can be shaped by casting or by compacting and sintering powders; the ferrites can only be shaped by powder methods. Powders, of course, don't have to be sintered – you can bond them together in other ways, and these offer opportunities. Magnetic particles mixed with a resin can be moulded to complex shapes. Bonding them with an elastomer creates magnets that are flexible, as discussed later.

The charts of Figures 16.9 and 16.11 show three families of hard magnets with coercive fields that surpass those of Alnicos and hexagonal ferrites. These, the products of research over the last 30 years, provide the intense permanent magnet fields that have allowed the miniaturisation of earphones, microphones, and motors. The oldest of these is the precious metal Pt–Co family but they are so expensive that they have now all but disappeared. They are outperformed by the cobalt–samarium family based on the inter-metallic compound $SmCo_5$, and they, in turn, are out-performed by the neodymium–iron–boron family. None of these are cheap, but in the applications they best suit, the quantity of material is sufficiently small, or the value of the product sufficiently high, that this is not an obstacle. Table 16.2 lists the choices.

Table 16.2 Materials for hard magnetic applications

Application	Material requirements	Material choice
Large permanent magnets (electric propulsion, wind turbine generators)	High M_r, high H_c, and high-energy product $(MH)_{max}$	Alnicos Hexagonal ferrites Neodymium–iron–boron
Small high-performance magnets (acoustic, electronic, and miniature mechanical applications)		Cobalt–samarium Neodymium–iron–boron
Information storage	Thin, 'square' hysteresis loop	Elongated particles of Fe_2O_3, Cr_2O_3, or hexagonal ferrite

Flexible magnets Flexible magnets are made by mixing ferrite powders into an elastomeric resin that is then extruded, moulded or calendered (rolled). Flexible magnets can be bent, punched or cut, without loss of magnetic properties. They are relatively cheap and because of this they are used in signs, promotional magnets, and door seals.

Typical of these is the range of magnets with strontium ferrous oxide particles embedded in styrene butadiene rubber (SBR) or vinyl (flexible PVC). Table 16.3 lists typical magnetic properties.

Table 16.3 Properties of flexible magnets at room temperature

Property	Value range
Remanent induction B_r	$0.15 - 0.25$ T
Coercive force H_c	$1 \times 10^5 - 2 \times 10^5$ A/m
Maximum energy product	$4.5 - 15$ kJ/m^3

Magnetic storage of information Magnetic information storage requires hard magnets, but those with an unusual loop shape: they are rectangular (called ‘square’). The squareness gives binary storage: a unit of the material can flip in magnetisation from $+M_r$ to $-M_r$ and back when exposed to fields above $+H_c$ or $-H_c$. The information is read by moving the magnetised unit past a read-head where it induces an electric pulse. The choice of the word ‘unit’ was deliberately vague because it can take several forms: discrete particles embedded in a polymer tape or disk, or a thin magnetic film that has the ability to carry small stable domains (called ‘bubbles’) that are magnetised in one direction, embedded in a background that is magnetised in another direction.

Figure 16.17 shows one sort of information storage system. A signal creates a magnetic flux in a soft magnetic core, which is often an amorphous metal alloy (AMA, the blue envelopes on the charts) because of its large susceptibility. The direction of the flux depends on the direction of the signal current; switching this reverses the direction of the flux. An air gap in the core allows flux to escape, penetrating a tape or disk in which sub-micron particles of a hard magnetic material are embedded. When the signal is in one direction, all particles in the band of tape or disk passing under the write-head are magnetised in one direction; when the signal is reversed, the direction of magnetisation is reversed. The same head is used for reading. When the tape or disk is swept under the air gap, the stray field of the magnetised bands induces a signal in the coil, which is amplified and processed to read it.

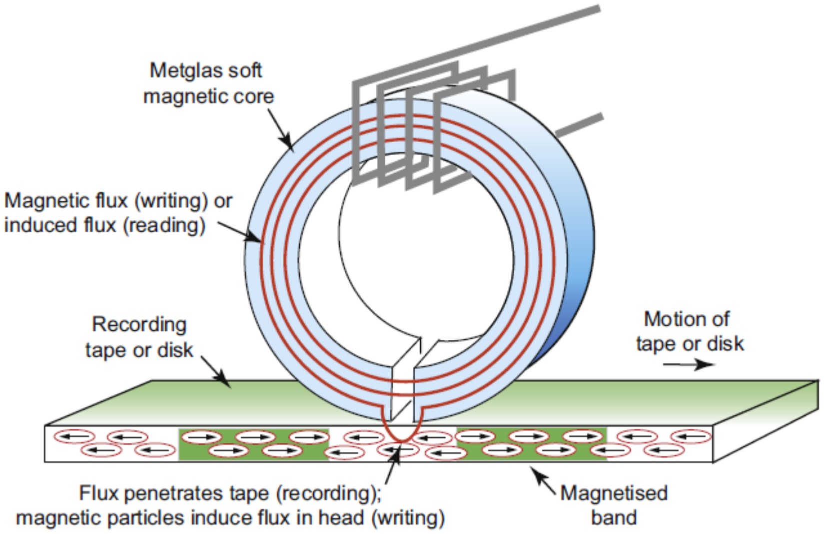

Figure 16.17 A magnetic read-write head. The gap in the soft magnetic head allows flux to escape, penetrating the tape and realigning the field in the particles.

Why use fine particles? It is because domain walls are not stable in very small particles: each particle is a single domain. If the particle is elongated, it prefers to be magnetised parallel to its long axis. Each particle then behaves like a little bar magnet with a north (N) and a south (S) pole at either end. The field from the write-head is enough to flip the direction. Rewritable tapes and disks use particles of Fe_2O_3 or CrO_2, typically 0.1 microns long and with an aspect ratio of 10:1. Particles of hexagonal ferrites $(MO)(Fe_2O_3)_6$ have a higher coercive field, so a more powerful field is required to change their magnetisation – but having done so, it is hard to erase it. This makes them the best choice for read-only applications like the identification strip of credit and swipe cards.

Magnetostriction The electrical 'hum' of transformers is due to magnetostriction of the core. Magnetostriction is shape change caused by a magnetic field. Positive magnetostriction means that the length increases parallel to the field; negative magnetostriction means the opposite. The saturation magnetostriction is the maximum strain that can be induced in a material by applying a magnetic field. For most materials this is small; the exceptions are the 'giant magnetostrictive' materials – those with extremely high saturation magnetostriction. Of these, Terfenol-D (an alloy of terbium, dysprosium, and iron) is the most commonly used. An alternative is Galfenol (an alloy of gallium and iron), which is less brittle and requires a lower coercive field to saturate it.

This ability to convert magnetic energy into kinetic energy and vice versa enables magnetic sensing, an effect exploited in sonar detection. It also allows magnetic actuation (Figure 16.18), and this is particularly attractive because it allows positioning with nanometer resolution and millisecond response times. The effect is used in fuel injectors, atomic-force microscopes, microcircuit fabrication, and vibration damping.

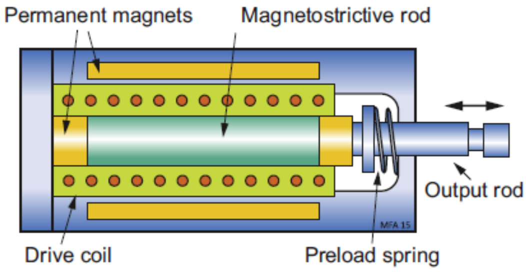

Figure 16.18 A magnetostrictive actuator.

Magneto-caloric materials To melt ice to water, you have to supply the latent heat of fusion. In part this latent heat reflects the entropy change associated with the change in order of the material (from an ordered crystal to a disordered liquid). In the old days, fridges worked using the melting of ice: you put a block of ice in the bottom; as it melted it drew its latent heat from the contents of the fridge. When a magnetic material is put into a magnetic field, its magnetic moments tend to align with the applied field. This too represents an increase in the order in the material, and a decrease in entropy. The entropy change occurs extremely quickly, with a release of latent heat of magnetisation, causing heating. Removing the magnetic field adiabatically allows disordering of the magnetic moments, increasing the magnetic entropy, absorbing heat, and causing a temperature decrease (Figure 16.19).

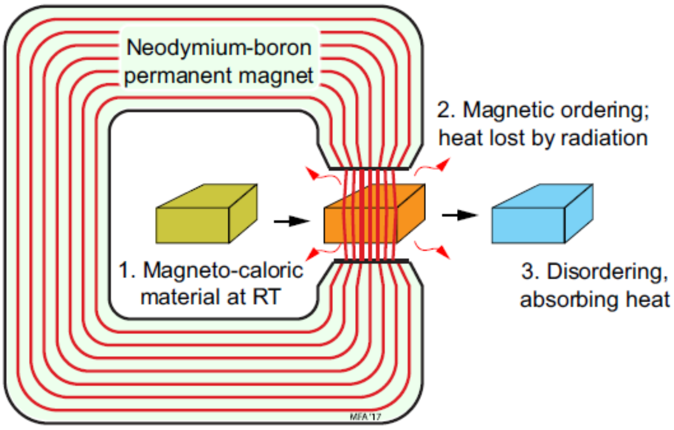

Figure 16.19 Magneto-caloric cooling.

The temperature change per unit applied field generated by the magneto-caloric effect is typically about 2°C per tesla. This on its own is not large, but when combined with a heat pump it is large enough to change the temperature by 50°C, more than enough for household refrigeration. Solid-state magnetic refrigerators are about 30% more efficient than conventional vapour-compression refrigerators and are more reliable because they have fewer moving parts and with no ecologically harmful coolants. Magnetic refrigeration has the potential to displace conventional refrigeration in everything from domestic fridge-freezers to industrial air conditioning.

Superconducting magnets Superconducting magnets allow high-field solenoids with no steady-state power consumption, but they do need very low temperatures. Early superconducting materials had such a low critical temperature, T_c, that they had to be cooled using liquid helium (boiling point 4.2K), an expensive undertaking. More recently, superconducting oxides have been developed with critical temperatures above the boiling point of liquid nitrogen (77.4K), making superconducting magnets much more practical (see Figure 15.3). The high-T_c superconductors also have a high critical field – the magnetic field strength that quenches superconduction – that sets an upper limit to the performance of a superconductor. The critical field depends on temperature in the way shown in Figure 16.20 for four of them. Ordinary electro-magnets get hot because of the resistance of the core, limiting them to flux densities of about 2 T. Superconducting magnets dissipate no energy in the core, allowing them to reach 80 T. They are used in MRI scanners, NMR equipment, mass spectrometers, magnetic separation processes, and particle accelerators.

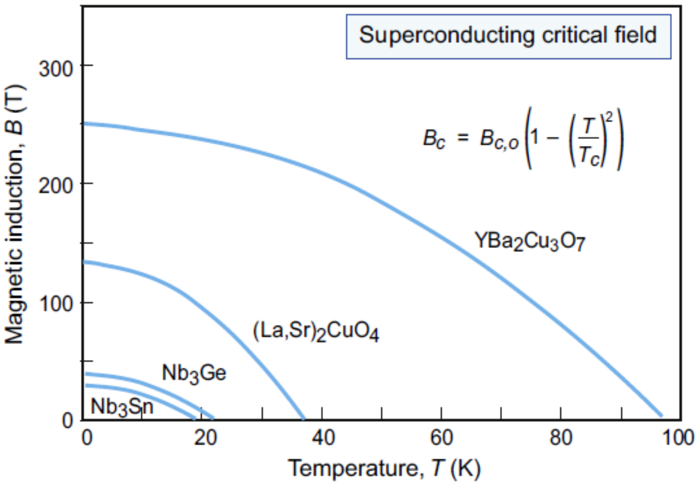

Figure 16.20 The dependence of the critical field on temperature, falling from a peak of $B_{c,o}$ at 0 K to zero at the superconducting critical temperature T_c.

The Meissner effect and magnetic suspension When a superconductor is placed in a magnetic field it behaves like any other conductor in a moving field – a current is induced in it. But because it has no resistance, the current continues to flow, preventing the flux lines of the field from entering. The result is that the field exerts a repulsive pressure (given by equation (16.7)) on the superconductor in the same way that two bar magnets repel when two north or south poles are face to face. It sounds as if this so-called Meissner effect might be used to make magnetic bearings and levitation suspensions, but in practice most magnetic bearings rely on magnetic attraction rather than repulsion. The shaft or track is surrounded by a symmetrical array of electro-magnets that pull it toward them. This is intrinsically unstable because the closer the shaft or track gets to one of the magnets, the greater is the attractive force it experiences. The instability is overcome by gap sensors that monitor the position of the shaft (or of the vehicle in magnetic levitation systems), adjusting the power in the individual electromagnets continuously to correct the position.

16.6 Summary and conclusions

The classical picture of an atom is that of a nucleus around which swing electrons in discrete orbits, each electron spinning at the same time on its own axis. Both spins create magnetic moments that, if parallel, add up, but if opposed, cancel to a greater or lesser degree. Most materials achieve near-perfect cancellation either within the atomic orbits or, if not, by stacking the atomic moments head to tail or randomising them so that, when added, they cancel. A very few, most of them based on iron, cobalt or nickel, have atoms with residual moments and an inter-atomic interaction that causes them to line up to give a net magnetic moment or magnetisation. Even these materials can find a way to screen their magnetisation by segmenting themselves into domains: a ghetto-like arrangement in which atomic moments segregate into colonies or domains, each with a magnetisation that is oriented such that it tends to cancel that of its neighbours. A strong magnetic field can override the segregation, creating a single unified domain in which all the atomic moments are parallel, and if the coercive field is large enough, they remain parallel even when the driving field is removed, giving a ‘permanent’ magnetisation.

There are two sorts of characters in the world of magnetic materials. There are those that magnetise readily, requiring only slight urging from an applied field to do so. They transmit magnetic flux, and require only a small reversal of the applied field to realign themselves with it. And there are those that, once magnetised, resist realignment; they give us permanent magnets. The charts of this chapter introduced the two, displaying the properties that most directly determine their choice for a given application.

Aligning and randomising atomic magnetic moments gives other functionality. The microscopic shape change induced by magnetisation (magnetostriction) allows ultra-precise actuation. And the entropy change of randomising aligned moments allows solid-state cooling devices.