home

Physics 2415 Lecture 18: Sources of Magnetic Field II

Michael Fowler UVa

Magnetostatics: the Biot-Savart Law

Finding the magnetic field from a steady current distribution is called *magnetostatics*, in analogy with electrostatics, which is finding the electric field from a stationary charge distribution. But there we had a very definite prescription: we knew the inverse-square field from a point charge, and we used the principle of superposition, adding together the fields from all the charges, to find the total field.

When Ampère was doing his experiments, in Paris in the 1820's, his colleagues included some of the world's best mathematicians, and they set about doing for magnetostatics what had already been accomplished in electrostatics: they looked for a formula for the magnetic field from a little bit of current, so that using superposition they'd be able to find the field from any distribution of currents by adding all the elements, just like the electric field from many point charges.

In fact, two of them succeeded in finding a formula that worked, but it's a very strange formula. It's called the *Biot-Savart* law, and here it is:

The magnetic field at $\vec{r}$ from an infinitesimal length $d\vec{\ell}$ of wire carrying current I at the origin is

$$d\vec{B} = \frac{\mu_0}{4\pi}\frac{I d\vec{\ell} \times \hat{r}}{r^2}$$

Notice it is inverse-square, like electrostatics (remember $\hat{r}$ is a unit vector).

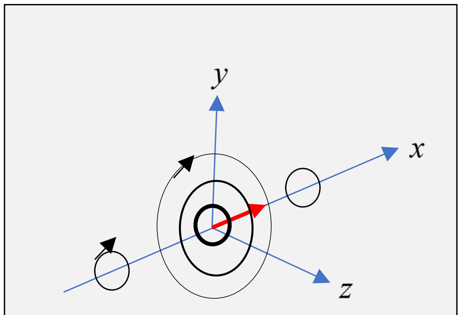

Field from current element $I d\vec{\ell}$

It's worth thinking about this field a little. Take $d\vec{\ell}$ at the origin and pointing in the x-direction. What are the field lines in the plane $x = 0$? (This plane includes the y and z axes.) They are circles, curling around as given by the right-hand rule. But they're not like the field from a wire along the x-axis: the field strength from this little current element goes down as the inverse square, evidently Ampere's law doesn't work for this current. Actually on the x-axis, anywhere, the field is zero, and everywhere else it's perpendicular to the x-axis and circling around it.

In one way at least, this formula is just a mathematical trick: you can't physically have a little element of steady current, opposite sign charges would be piling up at the two ends. Such an element only has meaning as part of a complete circuit.

The forces between two current elements ***aren't even equal and opposite***. Consider this by taking a second current element, at $\begin{pmatrix}1 & 0 & 0\end{pmatrix}$ pointing in the y-direction. It will feel no force from the first current element, but the field from the second current element is certainly nonzero at the origin.

(*Another point*: suppose these two "current elements" are just nonrelativistic moving charged particles. If we assume the Biot-Savart law is still good, apparently Newton's Third Law doesn't work for the magnetic interaction? The answer is that the electric and magnetic *fields* carry energy and momentum—not just the particles—so *total* momentum can be conserved even if the forces between *particles* are not equal and opposite. But this is too complicated to analyze here.)

So this is a strange formula, but it works. Consider a straight finite stretch of wire, part of some circuit. What's the field at distance R ?

The Biot-Savart rule tells us

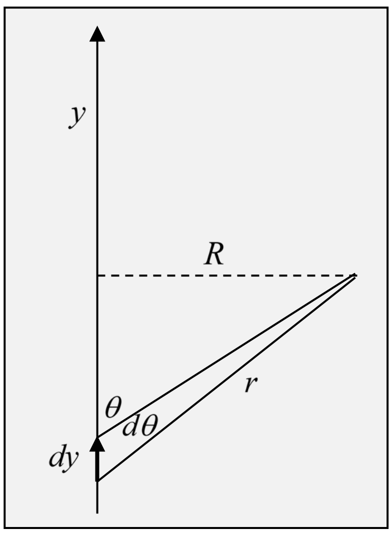

$$dB = \frac{\mu_0}{4\pi}\frac{Idy\sin\theta}{r^2}$$

and from the discussion above all the dB's point into the screen, so we just integrate over the (finite) length of wire we're considering. It's easiest to switch variables from y to θ, notice $\cot\theta = y/R$, so $dy = R\cos ec^2\theta d\theta = r^2 d\theta / R$, the integral just becomes

$$B = \frac{\mu_0 I}{4\pi}\int_{\theta_1}^{\theta_2}\sin\theta d\theta / R = \frac{\mu_0 I}{4\pi}(\cos\theta_1 - \cos\theta_2)/R.$$

This is done in the book for an infinite wire—but it works for a finite length too, provided it's part of a circuit, so charge isn't piling up. For example, you could find the field from a square coil.

For a round coil, though, the integral isn't that easy, *except* for the important case of the field on the axis, let's work it out.

We show the ring of current as copper-colored, and focus on an increment $I\vec{d\ell}$ at the top of the loop (this little vector is perpendicular to the screen/page). At a point P on the ring's axis distance $\vec{r}$ from the current increment, the magnetic field contribution from this increment is perpendicular to $\vec{r}$, and has a component parallel to the x-axis

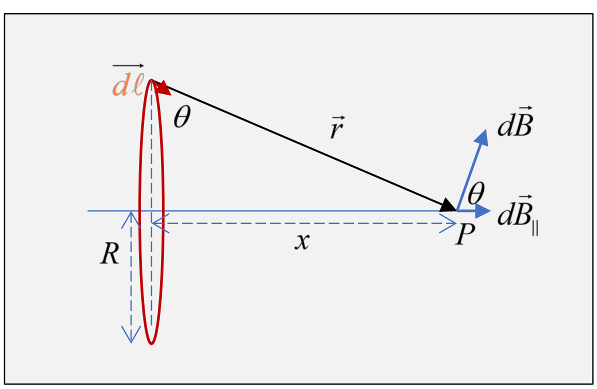

$$dB_{\parallel} = \frac{\mu_0 I d\ell}{4\pi r^2}\cos\theta$$

Adding the contributions from all around the circle, the components perpendicular to the x-axis cancel out by symmetry, those along the x-axis add to give

$$B = B_{\parallel} = \frac{\mu_0 I 2\pi R}{4\pi r^2}\frac{R}{r} = \frac{\mu_0 I}{2}\frac{R^2}{\left(R^2 + x^2\right)^{3/2}}.$$

At distances $x \gg R,\ B \cong \dfrac{\mu_0}{2\pi}\dfrac{\pi R^2 I}{x^3}$, which is the field from a dipole of strength $\pi R^2 I = AI$, with A the area of the loop, a formula that turns out to be good in this limit for any shaped loop.

Helmholtz Coils

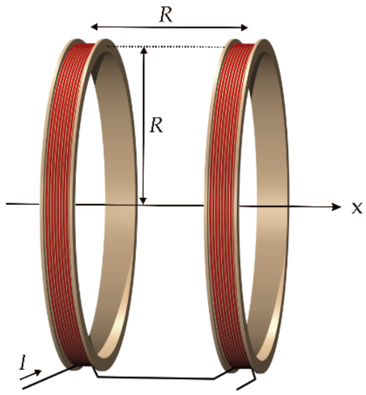

Two identical circular coils of radius R are distance R apart along their common axis, as shown. This provides an experimentally useful very uniform field at the midpoint of the setup, the field only increases 7% on going from that center point to the plane of a coil. This configuration is sometimes used to cancel the Earth's field.

If the currents in the two coils are opposite, the magnetic field at the center point is zero and linear in x to a very good approximation, this is useful in magneto-optical traps.

Exercise: check these field facts.

Magnetic Field inside a Long Solenoid

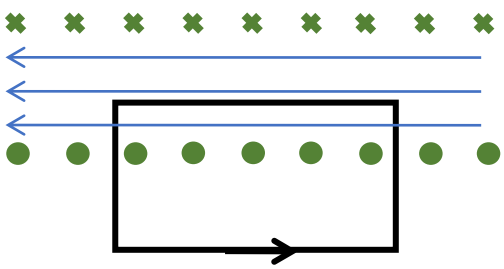

The dots and crosses are the loops of current going in and out of the paper. The blue represents the fairly uniform magnetic field inside the solenoid, compared with which the field outside is negligible. We take the current I, and the number of coil turns per unit length n.

Then from Ampère's Law integrated around the black rectangular contour, of length L,

$$\oint \vec{B} \cdot \overrightarrow{d\ell} = BL = \mu_0 nIL,$$

so

$$B = \mu_0 nI$$

and is uniform across the cross section, provided we can neglect end effects (so a very long solenoid).

Diamagnetism and Paramagnetism: Permeability

Diamagnetism is a molecular version of Lenz' law. Orbiting electrons in atoms and molecule are little currents, and when the magnetic field going through a current loop changes, the current itself changes in a way to minimize the total change in magnetic field. In other words, if the magnetic field through a diamagnetic solid changes, the solid generates its own field to lessen the change.

A quantitative measure of diamagnetic response is the permeability, denoted by μ. This is a measure of the material's response to a magnetic field, so if the long solenoid discussed just above is filled with material the field inside will be

$$B = \mu n I.$$

The constant μ_0 is often called the *permeability of the vacuum*.

For many substances μ is very close to μ_0, and a convenient parameter is the magnetic susceptibility

$$\chi_m = \frac{\mu}{\mu_0} - 1.$$

For diamagnets, the effect is usually small, $\chi \sim -10^{-5}$ except bismuth, $\chi \approx -1.66 \times 10^{-4}$, and one big exception: superconductors, which exhibit the *Meissner effect*: on putting one into a field, surface currents appear and the field (below a certain strength) cannot penetrate the superconductor, so $\chi = -1.$

All solids have some diamagnetic response, but if some of the atoms also have nonzero magnetic moments, there is also a (usually stronger) *paramagnetic* response, as the molecular magnetic moment tends to align with the applied field. For paramagnetic materials $\chi \sim 10^{-5\pm1}$.

For paramagnets, the response to an external field can be found by analyzing a single spin. The analysis breaks down at very low temperatures when enough atoms are aligned with the field to contribute a sufficient extra overall magnetic field. This was at first thought to explain ferromagnetism, but, on running the numbers, in iron crystals this magnetic moment lining-up doesn't occur above a few degrees absolute, evidently something else is lining up the moments in iron. Read on.

Ferromagnetism: Domains

In a f*erromagnet*, the individual atoms are little magnets, but in contrast to a paramagnet, there are powerful *quantum mechanical* forces causing nearest neighbors to align magnetically. (So this alignment does *not* come from the much weaker magnetic dipole-dipole interaction.)

In particular, the atoms of Fe, Co and Ni (and rare earths) are little magnets: in the incompletely filled shell of electrons, the electron spins line up—and electrons are themselves magnets. (If you're familiar with quantum mechanics, the spins line up to give a symmetrical spin wave function, which means the spatial wave function must be antisymmetric, and that keeps the electrons from getting too close, so minimizing the repulsive electrostatic energy.)

So why isn't every piece of iron magnetic?

What actually happens is best illustrated by considering a very pure small crystal of iron called a whisker. On looking closely, it turns out that almost all atoms *are* aligned with their neighbors, but the crystal as a whole is divided into fully aligned regions, called *domains*, as shown below.

The reason is that this is the state of lowest energy. To create a magnetic field costs energy, and if all the atoms were aligned, there would be a strong dipole field in the surrounding space. The arrangement shown here generates very little external field. Of course, the domain walls cost some energy, so

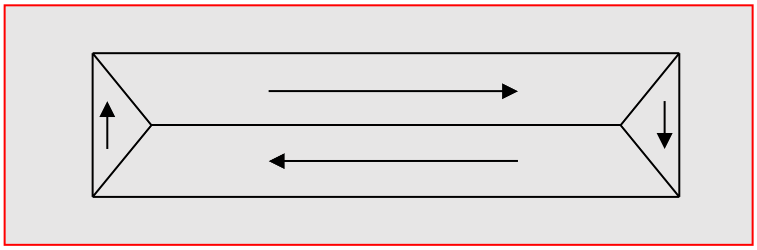

calculations are needed to find the optimum configuration.

If this whisker is placed in an external magnetic field, the domain closest to parallel with the external field will grow at the expense of the others. This lowers the energy of the whisker in this new environment, just as a compass needle will swing around to align itself with an external field.

Actually this whisker is very pure iron, a single crystal, and so very soft, magnetically speaking. That means it readily responds to a change in external field, which in turn means that the boundaries between domains move easily. This is good for iron in the core of an electromagnet, but not desirable in a permanent magnet. There you need to be able to line up domains and make it hard for them to readjust. This will be the case if the single crystal is replaced by many small crystals having different axes, and also certain non-iron atoms, such as in an alloy, can pin the domain walls and make movement difficult. Obviously, to make a permanent magnet takes a more intense external field than is necessary for magnetizing soft iron temporarily.

Using Soft Iron to Make a Strong Electromagnet

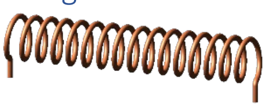

We've already seen that a way to concentrate the magnetic field from a current in a wire is to form a solenoid, which then resembles a bar magnet. A much stronger field can be achieved by filling the core of the solenoid with a soft ferromagnetic material. "Soft" in this context means a material very responsive to the field, the domain walls move readily to ensure a full magnetic response, that is, the atomic magnets in the material fully line up with the prevailing field. A soft iron core can increase the magnetic field in a solenoid by a factor

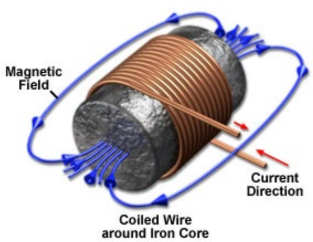

of thousands (the factor μ/μ_0). Mumetal goes to hundreds of thousands, and there are alloys at one million. Mumetal is used to shield a hard disk from the field of the motor.

(*Note*: in some books, you might see the notation $B = \mu H$ where B is the physical magnetic field and H is the field that would be produced by the currents in the wires if no magnetic material were present. This is perhaps useful for engineering design, but we won't be using it.)

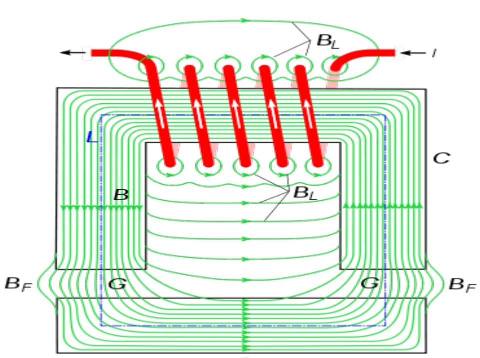

Here is a basic design for the type of magnet used to pick up pieces of wrecked cars, etc. Notice how the soft iron concentrates the field lines. The Π shape is the body of the magnet.

The bar along the bottom is the piece of car being picked up.

Click the picture for more details.