Chapter 6
The Remarkable Tokamak*

A Special Kind of Torus

The name *tokamak* comes from the Russian words ***to****roidalnaya* ***ka****mera* ***ma****gnitnaya* ***k****atushka* meaning "toroidal chamber magnetic coils," though it might have been appropriate to name it after the Russian word *tok*, meaning current. As mentioned in Chap. 4, this device was unveiled at the 1958 Geneva Conference. In those days, the Russians had the lead in space satellites, but their fusion research was done with poor equipment and considered primitive. The Americans and Britons, by contrast, had shiny, expensive, and well-engineered machines which they proudly displayed. The tokamak, however, turned out to be the one that worked the best and is the leading type of magnetic plasma container today. It was developed by a team led by Academician Lev Artsimovich on an idea of Andrei Sakharov and Igor Tamm and has been adopted by all nations working on magnetic fusion energy.

In Chap. 5, we showed that a magnetic bottle had to be a topological torus and that it had to have helically twisting magnetic field lines in order to compensate for the vertical particle drifts caused by the toroidal shape. The field lines also had to be sheared to stabilize the Rayleigh–Taylor interchange instability. In a stellarator, the proper magnetic field shape can be created with external helical windings carrying current. In a tokamak, this is simplified by driving a large amount of current through the plasma itself. The current flows in the toroidal direction (the long way around the torus), and it generates a poloidal magnetic field (the short way around the cross section). When this poloidal field is added to the main toroidal field from the large outside coils, the magnetic field inside the plasma is twisted into helices. Moreover, since the poloidal field is not the same on every magnetic surface, the helical field also has shear. This is illustrated in Fig. 6.1. A strong field in the toroidal direction is created by external coils, of which only three are shown for clarity. Inside the plasma, one magnetic surface is shown. A toroidal current is driven through the plasma inside this surface, and this creates a poloidal field, which adds to the toroidal field to form a twisted helical field. Depending on how much current

*Numbers in superscripts indicate Notes and square brackets [] indicate References at the end of this chapter.

F.F. Chen, *An Indispensable Truth: How Fusion Power Can Save the Planet*,
DOI 10.1007/978-1-4419-7820-2_6,

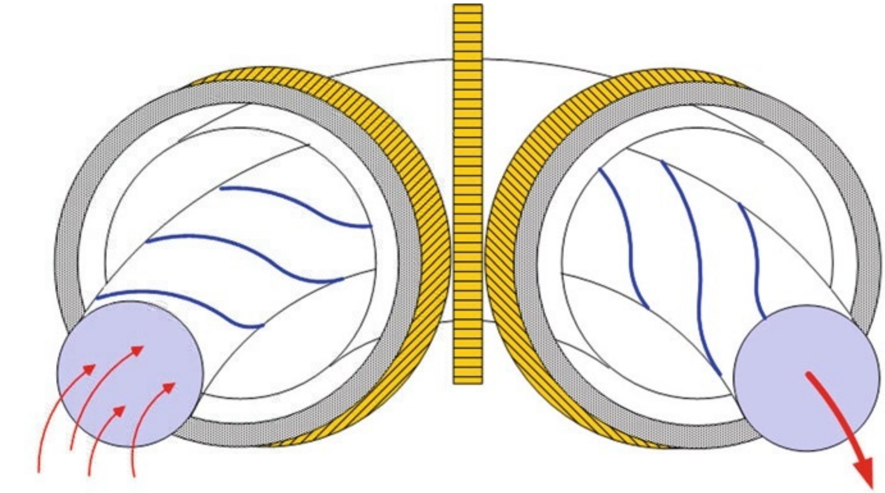

Fig. 6.1 Helical field lines created by external coils and a plasma current

there is inside each magnetic surface, the amount of twist differs from one surface to the next, and so the field is also sheared to prevent instabilities.

Using the plasma itself as a current-carrying coil to generate the twisting field would seem to be a great simplification, but we have not yet shown the hardware needed to drive this current. The advantage of the tokamak is more subtle. The current path for the poloidal field is not fixed by an external coil but can be varied by the plasma; and, fortuitously, the plasma has a self-curing property that distributes the current in a beneficial way. We explain this more fully later on.

Kink Instability and the Kruskal Limit

A toroidal plasma current serves two purposes: it generates the necessary twist in the magnetic field, and it can also raise the plasma temperature by ohmic heating. However, there is a limit to how much current can be driven because of yet another instability: the kink instability. Figure 6.2 shows an initially straight current path in the plasma that has bent itself into a kink. The circles show the field lines of the poloidal field that the current generates (the toroidal field is from left to right). Note that the lines are closer together on the inside of the kink than on the outside, indicating that the field is stronger on the inside. The magnetic pressure, therefore, is stronger at the bottom of this picture than at the top, and the kink is pushed further out. The bigger the kink, the larger the pressure difference; and the instability grows rapidly and disrupts the current. Remember that the poloidal field shown here is *not* the main (toroidal) field that supports the plasma pressure; it is the relatively small field that provides the twist. The toroidal field has a stabilizing influence, since it resists being pushed around by the plasma current. The onset of instability, therefore, depends on how strong the toroidal field is relative to the current. Conversely, onset of instability depends on how much current there is for a given toroidal field strength.

The limiting current for stable operation is called the Kruskal–Shafranov limit, and it is conveniently expressed in terms of the rotational transform, which is the number of times a field line goes around a torus the short way for each time it goes around the long way (Chap. 4). The critical rotational transform is exactly ONE! The critical current is that which creates a poloidal field large enough to twist the field lines just enough to give unity rotational transform, taking into account the strength of the main toroidal field. Transforms larger than 1 are unstable to kinks; transforms smaller than 1 are stable. The criterion for kink stability is actually quite complicated, since it depends on how the current varies across the plasma, but we can give a rough picture of why a rotational transform of 1 is a magic number.

The kink shown in Fig. 6.2 is in a straight plasma, but the current channel actually flows around the torus and joins back on itself. Figure 6.3 shows the largest unstable kink, which is actually an off-center displacement of the plasma. The plasma has been made unrealistically thin in order to have room to show the effect. In the top view (a), the dashed lines indicate the cross sections viewed in panel (b). Let us assume that the rotational transform is exactly 1. On the right-hand side of either view, the plasma has been displaced toward the outer wall. On the left-hand side, half-way around the torus, the field lines have rotated half-way around the cross section, so the plasma is now close to the inside wall. If the transform is

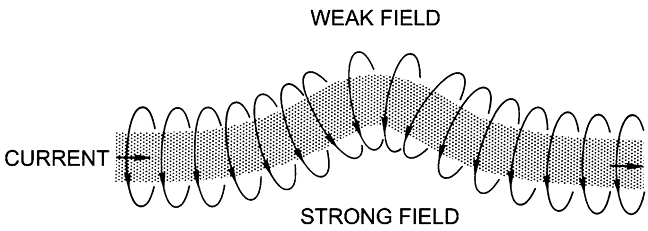

Fig. 6.2 A kink instability

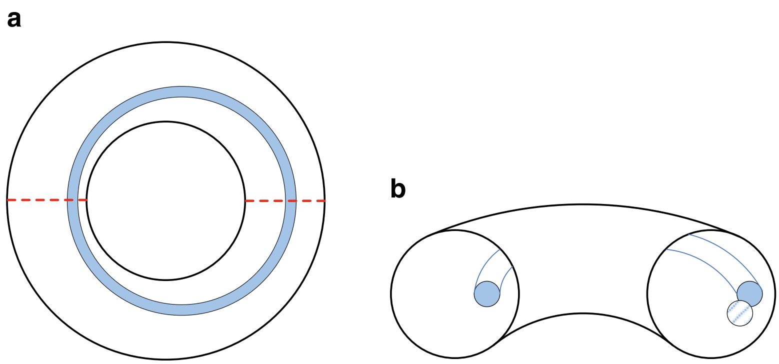

Fig. 6.3 A large kink distortion of the plasma in a torus: (**a**) top view and (**b**) cross-sectional view

exactly 1, when the field lines come back to the right-hand side, they will be in the same place where they started, so the current can flow in a closed path. Remember that the plasma is almost a superconductor; so, without collisions, the electrons carrying the current must stay on the same field line. Now let us assume that the rotational transform is less than 1. Then, upon coming back to the right-hand cross section, the current channel is in the position shown by the cross-hatched circle, which does not match up with its initial position. Since current must flow in a continuous path, this distortion of the current channel is not possible, and this kink cannot form. The plasma is stable for rotational transforms less than 1. In this simple picture, the plasma would also be stable if the transform is greater than 1, as long as it is not exactly 1. However, in that case, the current is strong enough to drive other shapes of kinks, and the plasma is kink-unstable in a way that is not easy to explain.

Since small rotational transform is good while large transform is bad, the *reciprocal* of the transform is used in tokamak lore. This is the *quality factor q* ("little q"), which is high when the plasma is kink-stable and low when it is kink-unstable. If the rotational transform is larger than 1, q is less than 1, and the plasma iskink-unstable. If the rotational transform is smaller than 1, q is larger than 1, and the plasma is kink-stable. What if q is a rational fraction so that the current channel joins up to itself after several trips around the torus? Then very interesting things happen, which we will get to.

Mirrors, Bananas, and Neoclassicism

Walking past Harold Furth's office one day, I saw this huge Chiquita Banana balloon hanging down from the ceiling. "What's going on?" I asked. "Welcome to banana theory," he replied, "the *fruitful* approach to fusion!" This was the beginning of a new understanding of how particles move in a torus. We knew that bending a cylinder into a torus would induce vertical drifts, and we knew how to counteract those by twisting the field lines into helices. But there were more subtle toroidal effects that we did not know about for the first 15 years. To explain banana orbits, we first have to describe magnetic mirrors.

If a magnetic field is not uniform – that is, if its strength changes as you move along a field line – it can reflect a charged particle and cause it to go backwards. This is the same effect that makes two permanent magnets repel each other when you turn one around so that their polarities don't match. There are toys that use this repulsion effect to suspend a magnetic object in midair. In Fig. 4.3b in Chap. 4, we showed that an electromagnet can create a magnetic field with coils of wire carrying a current. The ions and electrons gyrating in their circular orbits in a magnetic field are like electromagnets, since they are like one-turn coils carrying a current, even if the current is lumped into one charged particle. Figure 6.4 shows the field of a gyrating ion immersed in the nonuniform field of a normal electromagnet. The ion's magnetic field is always in the opposite direction to that of the field it's immersed in. Why? Because a physical system always tries to fall into the

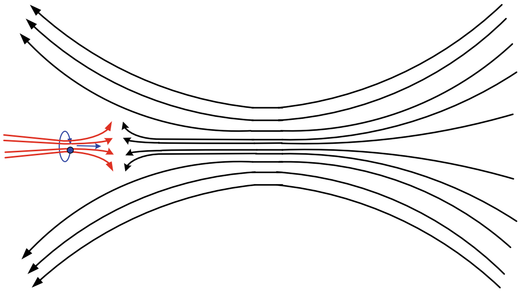

Fig. 6.4 Reflection of an ion heading into a stronger magnetic field. The field generated by the ion's gyration is shown in *red*

lowest energy state. By canceling part of the background magnetic field, the ions can lower the total magnetic energy. Electrons will do the same even though they have negative charge. They rotate in the opposite direction, but being negative, they carry current in the same direction as the ion do.

In Fig. 6.4, an ion, carrying the magnetic field that it generates, moves to the right. The field lines on the right are of a background magnetic field generated by large coils outside the plasma. The field lines generated by the current of the gyrating ion are shown in red. The opposing fields push the ion backwards, like two permanent magnets with opposite polarity. The ion's motion to the right is slowed up. The ion is moving into a stronger field, since the black lines are getting closer together. When the external field gets too strong, the ion cannot go any farther and is reflected back. How far the ion goes depends on how fast it was moving from left to right. However, not all ions will get reflected because the background field has a maximum strength. If the ion comes in with enough energy to go through the maximum, it gets slowed up there, but it is able to go through and regain its velocity on the other side. A converging magnetic field is a magnetic mirror that can reflect all but the fastest ions. This mechanism of magnetic mirroring was used by Enrico Fermi to explain the origin of cosmic rays. There, the interstellar magnetic fields are moving very rapidly, and they can push ions up to very high energies. Why can't we use magnetic mirrors to trap and hold a plasma? Indeed, we can, but magnetic mirror systems have not worked out as well as tokamaks. Mirrors will be described in Chap. 10.

Now we can get to the bananas. Tokamaks also have magnetic mirrors, but they hinder rather than help the confinement. Recall from Fig. 4.14 in Chap. 4 that the magnetic field is always stronger on the inside of a torus, near the hole, than on the outside because the coils are closer together in the hole, and therefore the field near one coil also gets contributions from the neighboring coils. That means that there is a nonuniform magnetic field, and particles going from a weak field to a strong field might get reflected. Ideally, particles travel along helical field lines on a magnetic surface and never leave it. However, magnetic mirroring prevents this, as shown in Fig. 6.5.

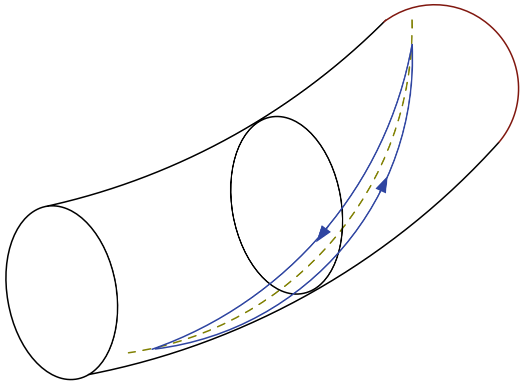

Fig. 6.5 A banana orbit in a tokamak. In reality, this orbit drifts around the torus

In this figure, the dashed line is a helical field line. An ion does not actually follow this line exactly unless its Larmor radius is zero. When it gyrates in a finite-sized circle, it will drift slowly from one line to another, as shown in Fig. 4.10, if the magnetic field strength is not the same on every side of its orbit.[1] The helical twisting cancels out the vertical drift on the average, but the averaging is disrupted by the mirror effect. The actual ion orbit is like the one shown by the solid line in Fig. 6.5. This ion starts out on the outside of the torus, where the field is weak, and it loops around toward the inside, where the field is strong. If it is not moving fast enough, it will be reflected by the magnetic mirror effect and come back on a slightly different path. Only ions with enough energy parallel to the field line will make it around to the inside of the torus and sample all parts of a magnetic surface as we envisioned in our earlier naïve picture of magnetic bottles. If we project the path of the ion in Fig. 6.5 onto the cross section of the torus shown there, it will look something like Fig. 6.6.

These are the so-called banana orbits. In each case, the outside of the torus is on the right side of the cross section, and the strong field near the hole in the doughnut is on the left. The small banana in panel (a) is for a particle with small velocity parallel to the magnetic field; it gets reflected before it gets very far toward the inside. The *dashed* line is the path of a passing particle, one that gets through the mirror and can come all the way around. In panel (b), the particle has larger parallel velocity and goes farther to the left, describing a larger banana. The limiting case is shown in panel (c), where the particle nearly makes it through the mirror. Tom Stix whimsically dubbed this the WFB, the *World's Fattest Banana*.

Banana orbits were discovered theoretically. They have never been seen in experiment because it would be very hard to track the path of a single ion or electron in a plasma with more than a trillion particles per cubic centimeter. However, theory predicts the consequences of banana orbits, and these unfavorable effects are well established by experiment. It's easy to understand why these bananas bear

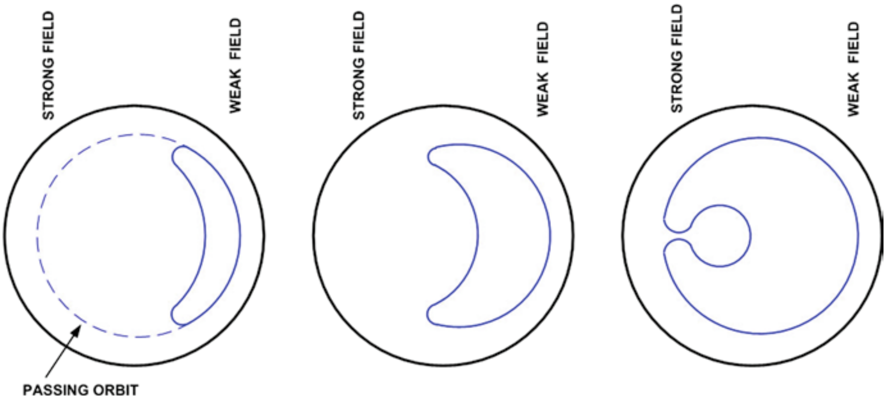

Fig. 6.6 Banana orbits of particles with increasing parallel velocities

bitter fruit. When an ion makes a collision, instead of jumping from one Larmor orbit to an adjacent one, it jumps from one banana orbit to the next; and banana orbits are much wider.[2] Instead of the very slow rate of "classical" diffusion that we described in Chap. 5, the rate of plasma transport across the magnetic field is much faster in a torus than in a straight cylinder. The rate of banana diffusion can be calculated easily and is called *neoclassical diffusion*. It is a characteristic of toruses that was not initially foreseen. The good news is that it is still a classical effect; that is, it can be calculated using a known theory. Figure 6.7 shows how banana diffusion differs from classical diffusion. At the left-hand side, the collision rate between ions and electrons is very small, so small that an ion can traverse one or many banana orbits before making a collision. In the middle, flat part of the curve, the trapped ions (those making banana orbits) make collisions during a banana orbit, but the passing particles, being faster, do not. In the right-hand part, the collision rate is high enough that all particles make collisions in traversing the torus. Under fusion conditions, the plasma is so hot and so nearly collisionless that it is well into the banana regime, at the extreme left of the graph. Therefore, it is clear that the banana diffusion rate is much higher than the classical one, shown by the straight line at the bottom.

One might think that the closer a torus is to a cylinder, the smaller the banana effects will be. The aspect ratio A of a torus is the major radius R divided by the minor radius a, as shown in Fig. 6.8. A fat torus would have small A and a skinny one, large A. One would think that large A would have smaller banana diffusion, but this is not always true. It depends on many subtle effects which can cancel one another. The Kruskal–Shafranov limit states that q (the inverse rotational transform) has to be larger than 1. For a given value of q, banana diffusion is actually larger for large A. This is primarily because the ion has to go a long way around the torus before it turns around, and it is drifting vertically the whole time.

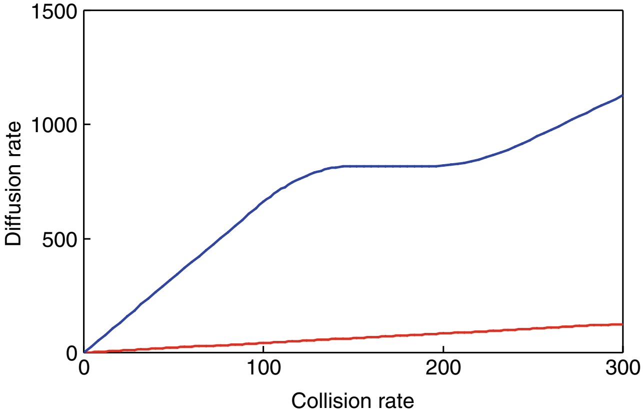

Fig. 6.7 Neoclassical (*top curve*) and classical (*bottom curve*) diffusion rate for ions as a function of collision frequency[3]

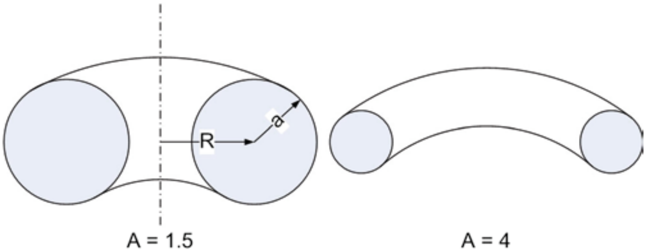

Fig. 6.8 Toruses with small and large aspect ratios

An even stranger, counter-intuitive effect has to do with the width of a banana orbit. It turns out that this width depends only on the strength of the poloidal field B_p and not on the toroidal field B_t. Remember that B_p is only the small field generated by the plasma current that gives the field lines a small twist. The banana width is approximately the Larmor radius of an ion calculated with B_p instead of B_t. This is much larger than the real Larmor radius, calculated with B_t. Since banana diffusion goes by steps of the size of a banana width, which depends only on the relatively weak B_p, does this mean that the much stronger toroidal field is useless? No! The toroidal field is needed to make the real Larmor radius small so that we can consider only the movement of the guiding centers, not the actual particles. If the toroidal field were eliminated,[4] the gyration orbits would be so large that magnetic confinement would be no good at all, and furthermore there would be nothing to hold the plasma pressure.[5]

Turbulence and Bohm Diffusion

A picture of David Bohm was taped to the wall of Bob Motley's office, and our group of experimentalists at Princeton's Plasma Physics Laboratory took turns throwing darts at it. The frustration came from an unexplained phenomenon called "Bohm diffusion," which caused plasmas in toruses to escape much faster than any classical or neoclassical theory would predict.[6] In spite of all efforts to suppress the known instabilities, the plasma was always unstable, vibrating, rippling, and spitting itself out, like the foam on violently breaking surf. In Chap. 5, classical diffusion was described. This is a process in which collisions between ions and electrons cause them to jump from one field line to another one about one Larmor radius away. The classical confinement time is long, of the order of minutes. In this chapter, we described neoclassical diffusion, in which particles jump from one banana orbit to the next. The neoclassical confinement time is still of the order of seconds, longer than needed. Bohm diffusion caused the plasma to be lost in milliseconds. Major instabilities like the Rayleigh–Taylor or kink were no longer there, else the confinement time would have been microseconds. There were obviously some other instabilities that the theorists had not foreseen.

Bohm diffusion was first reported by physicist David Bohm when he was working on the Manhattan project and, in particular, on a plasma device for separating uranium isotopes. From measurements of the plasma's escape rate, he formulated a scaling law for this new kind of diffusion. It reads as follows. The diffusion rate across the magnetic field, given by the coefficient $D_\perp$ (pronounced *D*-perp), is proportional to 1/16 of the electron temperature divided by the magnetic field:

$$D_\perp \propto \frac{1}{16}\frac{T_e}{B}.$$

The 1/16 makes no sense here because I have not said what units T_e and B are in, but that number has a historical significance. Whenever Bohm diffusion is observed, there are always randomly fluctuating electric fields in the plasma. Regardless of what is causing these fluctuations, the plasma particles will respond by drifting with their $\mathbf{E}\times\mathbf{B}$ drifts (Chap. 5). Since the size of the noise is related to T_e, which supplies the energy for it, and the drift speed is inversely proportional to B, it is not hard to show that the T_e/B part is to be expected [1]. But how did Bohm come up with the number 1/16? Bohm had disappeared from sight after he was exiled to Brazil for un-American activities. In the 1960s, Lyman Spitzer tracked him down and asked him where the 1/16 came from. He didn't remember! So we'll never know. It turns out that the Bohm coefficient depends on the size and type of turbulence and can have different values, but always in the same ballpark.

Plasma turbulence is the operative term here. Any time there was unexplained noise, it was called "turbulence." Doctors do the same thing with "syndrome" or "dermatitis." Figure 6.9 is an example of turbulence; it is simply a wave breaking on a beach. As the wave approaches the beach, it has a regular, predictable up and down motion. But as the water gets shallower, the wave breaks and even foams. The

Fig. 6.9 Turbulence at the beach

motion of the water is no longer predictable, and every case is different. That's the turbulent part. The regular part is called the *linear* regime; this is a scientific term that has to do with the equations that govern a physical system's behavior. Linear equations can be solved, so the linear behavior is predictable. The turbulent part, in the *nonlinear* regime, can be treated only in a statistical sense, since each case is different. Nonlinear generally means that the output is not proportional to the input. For instance, taxes are not proportional to income, since the rate changes with income. Compound interest is not proportional to the initial investment, even if the interest rate does not change, so the value increases nonlinearly. Population growth is nonlinear even with constant birth rate, in exact analogy with compound interest. Waves, when they are small and linear, will have sizes proportional to the force that drives them. But they cannot grow indefinitely; they will saturate and take on different forms when the driving force is too large. What a wave will look like after it reaches saturation can be predicted with computers, but the detailed shape will be different each time because of small differences in the conditions. Then you have turbulence. The smoke rising from a cigarette in still air will always start the same way, but after a few feet each case will look different.

The turbulence in every fusion device in early experiments was always fully developed; we could never see the linear part, so we could not tell what caused the fluctuations to start in the first place. An example of plasma turbulence in a stellarator is shown in Fig. 6.10. This is what "foam" looks like in a plasma. These are fluctuations in electric field inside the plasma. These noisy fields make the particles do a random walk, reaching the wall faster than classical diffusion would take them.

Turbulence is well understood in hydrodynamics. If you try to push water through a pipe too fast, the flow breaks up into swirling eddies, slowing down the flow. Hydrodynamicist A.N. Kolmogoroff once gave an elegant proof, using only dimensional analysis, that the sizes of eddies generally follows a certain law; namely, that the number of eddies of a given size is proportional to the power 5/3 of the size.[7] Attempts to do this for plasmas yielded a power of 5 rather than 5/3 [1],

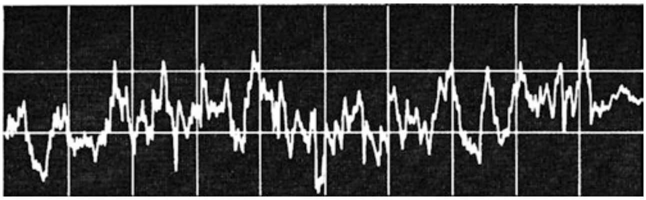

Fig. 6.10 Fluctuations in a toroidal plasma

and this has been observed in several experiments. However, plasmas are so complex (because they are charged) that no such simple relation holds in all cases.

The importance of turbulence and Bohm diffusion is not only that it is much faster than classical diffusion, but also that it depends on $1/B$ instead of $1/B^2$. In classical diffusion, doubling the magnetic field B would slow the diffusion down by a factor of 2^2 or 4. In Bohm diffusion, it would take *four* times larger B to get the same reduction in loss rate. It was this unforeseen problem of "anomalous diffusion" that held up progress in fusion for at least two decades. Only through the persistence of the community of dedicated plasma physicists, was the understanding and control of anomalous diffusion achieved. Modern tokamaks have confinement times approaching those required for a D–T reactor.

The Culprit: Microinstabilities

If the plasma in a torus always thrashes around violently, there must be an energy source that drives the thrashing. An obvious source is the electric field applied to drive the current in ohmic heating. In the 1960s, a new method was devised for heating without a large DC current in the plasma. This was Ion Cyclotron Resonance Heating or ICRH. A radiofrequency (RF) power generator was hooked up to an antenna around the plasma, the way an FM station is hooked up to its antenna on a tower. The frequency was tuned to the gyration frequency of the ions in their cyclotron orbits. As the ions moved in circles, the RF field would change its direction so as to be pushing the ions all the time, just as in a real cyclotron. This could heat up the plasma without having to drive a DC current in it.[8] Would this kill the turbulence and make the plasma nice and quiet, without Bohm diffusion? A case of champagne was bet on it. It didn't work. The thrashing was as bad as ever. The darts in Bohm's picture stayed there.

The problem was a failure of magnetohydrodynamics, MHD for short. MHD theory treats a plasma as a pure superconductor, with zero resistivity, and neglects the cyclotron orbits of the particles, treating them as points moving at the speed of their guiding centers. Though this simplified theory served us well in the design of toroidal confinement devices and in the suppression of the gravitational and kink instabilities, it did not treat a plasma in sufficient detail. First of all, there have to be *some* collisions in a fusion plasma or else there wouldn't be any fusion at all! These infrequent collisions cause the plasma's resistivity to be not exactly zero, and that has dire consequences on stability. The fact that the Larmor orbits of the ions are not mathematical points gives rise to the finite Larmor radius (FLR) effect. In some cases, even the very small inertia of the electrons has to be taken into account. Finally, instabilities could even be caused by distortions of the particles' velocity distributions away from a pure Maxwellian, an effect called Landau Damping. These small deviations from ideal MHD turned out to be important, making the theorists' task much more difficult.

The first inkling of what can happen was presented by Furth, Killeen, and Rosenbluth in their classic paper on the tearing mode [2]. If a current is driven along the field lines in a plasma with nonzero resistivity, the current will break up into filaments; and the initial smooth plasma will tear itself up into pieces! So "tearing" rhymes with "bearing," and not "fearing," though the latter interpretation may have been more appropriate. The tearing mode is too complicated to explain here, but we describe other instabilities which caused even more tears.

One of the tenets of ideal MHD is that plasma particles are "frozen" to the field lines, as shown in Fig. 4.10. Without collisions or one of the other microeffects named above, ions and electrons would always gyrate around the same field line, even if the field line moved. Bill Newcomb once proved a neat theorem about this [3], saying that *plasma cannot move from one field line to another as long as* $E_{\|}$ *(E-parallel) is equal to zero*. Here, $E_{\|}$ is the electric field along a magnetic field line, and it has to be zero in a superconductor, since in the absence of resistance even an infinitesimal voltage can drive an infinite current. But if there are collisions, the resistivity is not zero, $E_{\|}$ can exist, and plasma is freed from one of its constraints.

So it was back to the drawing board. While the theorists enjoyed a new challenge and a new reason for their employment, the experimentalists pondered what to do. In previous chapters, we showed that (1) a magnetic bottle had to be shaped like a torus, (2) bending a cylinder into a torus caused vertical drifts of ions and electrons, (3) these drifts could be canceled by twisting the field lines into helices, (4) this twist could be produced by driving a current in the plasma, and (5) this current could cause other instabilities, even in ideal MHD, but that those could be controlled by obeying the Kruskal–Shafranov limit. In spite of these precautions, the plasma is always turbulent, even when the current is removed by using a stellarator rather than a tokamak. How can we get a plasma so smooth and quiet that we can see a wave grow bigger and bigger until it breaks into turbulence, as in Fig. 6.9? Obviously, if one could straighten the torus back into a cylinder, much of the original cause of all the trouble would be removed. But how can one hold the plasma long enough just to do an experiment? The plasma will simply flow along the straight magnetic field into the endplates that seal off the cylinder so that it can hold a vacuum. The solution came with the invention of the Q-machine (Q for Quiescent). Developed by Nathan Rynn [4] and Motley [5], this is a plasma created in a straight cylinder with a straight magnetic field. Inside each end of the vacuum chamber is a circular tungsten plate heated to a red-hot temperature. A beam of cesium, potassium, or lithium atoms is aimed at each plate. It turns out that the outermost electron in these atoms is so loosely bound that it gets sucked into the tungsten plate upon contact. The electron is then lost, and the atom comes off as a positively charged ion.

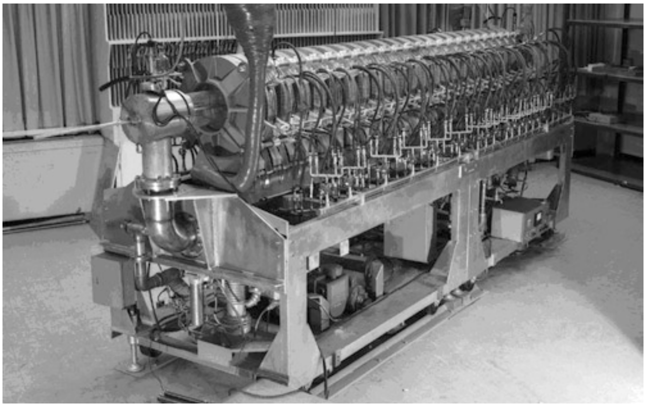

Fig. 6.11 Example of a Q-machine

Of course, a plasma has to be quasineutral, so the tungsten has to be hot enough to emit electrons thermionically, the way the filament in light bulb does. So both ions and electrons are emitted from the tungsten plates to form a neutral plasma. *No electric field has to be applied!* Only tungsten or molybdenum, in combination with the three elements above, can perform this kind of thermal ionization. In this clever device, all sources of energy to drive an instability have been removed or so we thought. Figure 6.11 shows a typical Q-machine, covered with the coils that create the steady, straight, and uniform magnetic field.

The plasma in a Q-machine *has to be quiescent*, right? To everyone's surprise, it was still turbulent! The trace shown in Fig. 6.10 actually came from a Q-machine. Fortunately, it was possible to stabilize the plasma by applying shear, as shown in Fig. 5.9, or by applying a small voltage to the radial boundary of the plasma. A quiescent plasma in a magnetic field was finally achieved. Then, by adjusting the voltage, one could see a small, sinusoidal wave start to grow in the plasma; and, with further adjustment, one could see the wave get bigger and bigger until it broke into the turbulence seen in Fig. 6.10. With a regular, repetitive wave like a wave in open water, one could measure its frequency, its velocity, its direction, and how it changed with magnetic field strength. These were enough clues to figure out what kind of wave it was, what caused it to be unstable, and, eventually, to give it a name: a *resistive drift wave*.

As its name implies, the wave depends on the finite resistivity of the plasma. It also depends on microeffects: the finite size of the ion Larmor orbits. Before showing how a drift instability grows, let's find the source of energy that drives it. In a Q-machine, we have eliminated all toroidal effects and all electric fields normally needed to ionize and heat the plasma. In fact, the plasma is quite cold, as plasmas go. It is the same temperature as the hot tungsten plates, about 2,300 K, so that the plasma temperature is only about 0.2 eV. You can heat a kiln up to that temperature, and it would stay perfectly quiescent. A magnetically confined plasma, however, has one subtle source of energy: its pressure gradient. When everything is at the same temperature and there are no energy sources such as currents, voltages, or drifts, there is still one source of energy when the plasma is *confined*. And confinement is the name of the game. Since ions and electrons recombine into neutral atoms when they strike the wall, plasma is lost at the walls. The plasma will be denser at the center than at the outside, and this causes a pressure that pushes against the magnetic field. By Newcomb's theorem, the plasma would remain attached to the field lines, and nothing can happen; but once there is resistivity, all bets are off. The plasma is then able to set up electric fields that allow it to move across the magnetic field in the direction that the pressure pushes it. Even if there are no collisions, other microeffects like electron inertia or Landau damping can cause the drift instability to grow. For this reason, the resistive drift instability and others in the same family are called *universal instabilities*. They are fortunately weak instabilities because the energy source is weak, and they can be stabilized with the proper precautions.

The Drift Instability Mechanism

There are many microinstabilities, but they all share the same types of plasma motion. As an example, we shall try to explain how a resistive drift wave goes unstable. This instability has stood the test of time as other theoretical predictions have come and gone. In general, it is easier to derive an instability mathematically than to figure out exactly what the plasma is doing. If this part is difficult to follow, you can skip to the next section without losing essential information. Let's start with a plasma in a straight cylinder with a straight magnetic field, as shown in Fig. 6.12. The plasma is necessarily denser at the center than on the outside. The white arrows show a density ripple, like a wave, propagating in the azimuthal direction. We shall focus on the plasma's behavior inside the small rectangle at the bottom. This rectangle is shown enlarged in Fig. 6.13. On the left, we see Larmor orbits of ions whose guiding centers may be outside the rectangle. If the magnetic field is

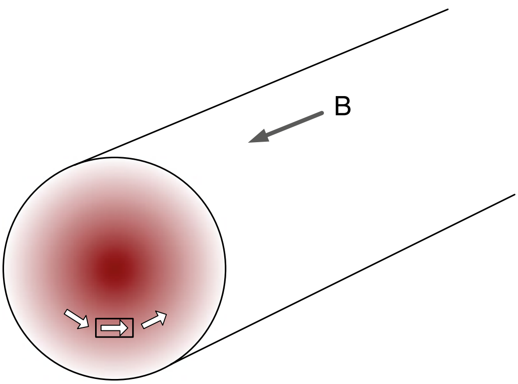

Fig. 6.12 A drift wave in an inhomogeneous plasma

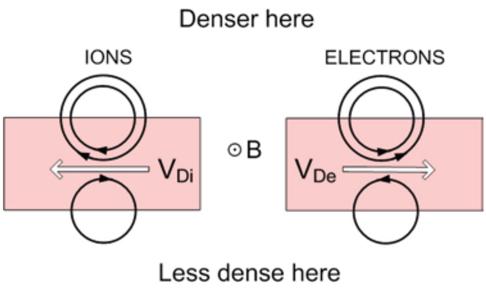

Fig. 6.13 Definition of the diamagnetic drift. The electron orbits are actually much smaller than those of the ions

out of the page, as shown, the ions will be rotating clockwise. Remember that the plasma density is higher at the top than at the bottom because the top is closer to the center of the plasma. To show this, we have drawn two orbits at the top and only one at the bottom. There are obviously more ions going left than going right. The ion fluid in this small volume therefore has an average flow toward the left. This effect is called the *ion diamagnetic drift*, and the drift velocity is called v_{Di}. Note that this drift is perpendicular both to the magnetic field and to the direction in which the density is changing. The diagram on the right is the same thing for electrons. With their negative charge, electrons gyrate counterclockwise. Their diamagnetic drift velocity, v_{De}, is therefore in the opposite direction, to the right. This motion of the ions and electrons, considered as fluids occupying the same space, is there even if the guiding centers are not moving. The existence of the diamagnetic drift depends on the gradient in density and would be zero if the density were uniform everywhere. If you have a problem with two fluids occupying the same space, just think of the vermouth and vodka in a martini.

Now we can proceed with the wave. Our little rectangle is shown three times in Fig. 6.14. At the bottom of the first diagram, (a), a density ripple is shown. A slice of the rectangle near the peak of the wave, where the density is high, is shown in a darker shade. The *background* density is high at the top and low at the bottom, as seen in Fig. 6.12. The diamagnetic drift of the ions in the *background* density gradient is to the left for ions and to the right for electrons, as shown in Fig. 6.13. Because the *wave density* is high near its peak, the diamagnetic drifts bring an excess of positive charge to the left side of the small slice and an excess of negative charge to the right side. These electric charges create the electric field **E** shown in panel (b). Recall from Chap. 5, Fig. 5.4, that an electric field causes an **E** × **B** drift, v_E, perpendicular to both **E** and **B**. In this case, the drift is downwards, as shown in panel (c). Since the *background density* is high at the top, v_E brings more density into the slice, and the wave gets more density where the *wave density* was already high. Therefore, the wave grows; it is unstable. Figure 6.15 shows what happens at a wave trough. There, the density is less than average, so the diamagnetic drifts bring *less* density to the edges of the slice, causing the buildup of charges of the opposite sign. The resulting electric field, shown in panel (b), is in the opposite direction from before. This causes the **E** × **B** drift in panel (c) to be upwards instead

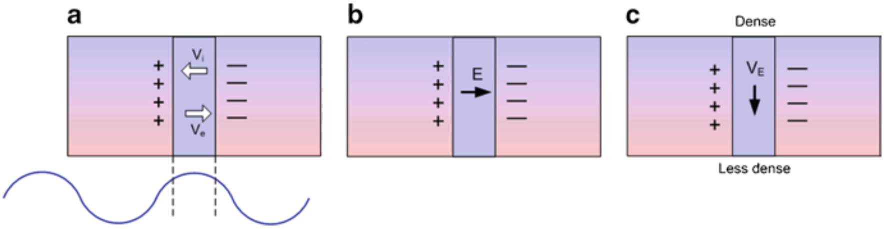

Fig. 6.14 The charges, fields, and velocities at the peak of a drift wave

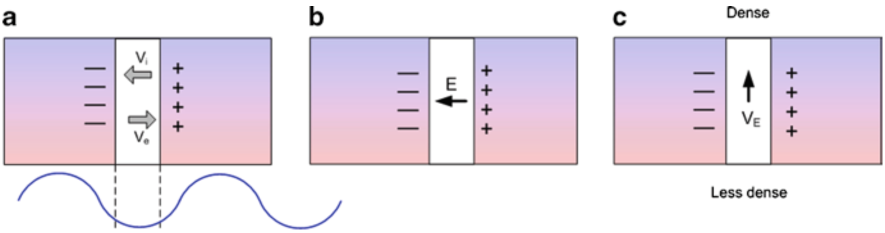

Fig. 6.15 The charges, fields, and velocities in the trough of a drift wave

of downwards. But an upward motion brings lower *background* density into the slice where the *wave* density is already low. This adds to the growth of the wave. We can now give it its rightful name: a *drift wave*. If we average over the cycles of the drift wave, more density is moved downwards at the peaks of the wave than is lost at the troughs, and consequently the wave causes plasma to move outwards, away from the center, toward the wall. Another insidious, cunning way the plasma finds to escape from its magnetic trap.

However, we are not quite finished; there is a three-dimensional part, shown in Fig. 6.16. The rectangular slices at the peaks and troughs of the wave in the last two figures are shown together at two cross sections of the cylinder, now considered as part of a torus. There are four slices: peak, trough, peak, trough. Between the slices are the electric charges shown in Figs. 6.14 and 6.15. Recall that a toroidal confinement requires a poloidal field to twist the magnetic field. This twist causes the field line going through a positive charge to connect to a negative charge in a cross section further downstream. Electrons, being very light and mobile, almost instantaneously move along the field line to cancel the charges. The electric field of the drift wave is nullified, and the wave can never grow. Ah, but if there are collisions, the electrons are slowed down, and they cannot cancel the charges fast enough. This is another example of Newcomb's theorem: if $E_{\|}$ is not zero, all bets are off! The growth of the drift instability depends on the existence of resistivity, one of our microeffects. Even without collisions, electron inertia or Landau damping can slow down the electrons and allow the instability to grow. Hence, it is a universal instability which can occur whenever there is a density gradient in a magnetic confined plasma.

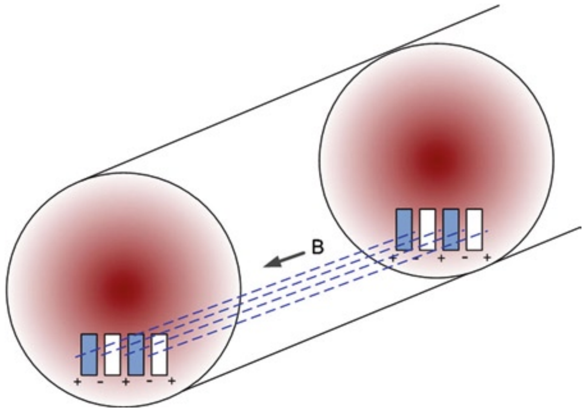

Fig. 6.16 A drift wave in three dimensions

The obvious question is, “What if the plasma density is uniform all the way out to the wall?” That can’t happen, since the density has to be essentially zero at a cold wall. If the density gradient occurs in a thin layer near the wall, the sharp gradient there will make the instability grow even faster. It then eats away the plasma so that the thickness of the gradient layer gets larger and larger. Drift instabilities can be stabilized by shortening the *connection length* between the cross sections shown in Fig. 6.16, so that the electrons can move between them fast enough. This requires a larger helical twist of the magnetic field. Fortunately, this can be done without violating the Kruskal–Shafranov limit.

There are many other possible microinstabilities. The ion-temperature-gradient instability is another one that is worrisome. This example of the resistive drift instability serves to give an idea of how complicated plasma behavior is and how Bohm diffusion was solved. What happens when an unstable wave breaks and becomes turbulent? It is no longer possible to identify which instability started the turbulence, but one can apply known stabilization methods to see if the fluctuations can be suppressed. There are turbulence theories that purport to predict how the turbulence will look and how much anomalous diffusion it will lead to. A powerful modern method is to do a computer simulation. A computer does not care whether an equation is nonlinear or not. It does not even need to solve an equation; it just follows the particles around to see where they go. There will be some examples later; it’s not as simple as it sounds. Or, one can use physical intuition to make a guess. Figure 6.17 shows a guess on what a resistive drift instability might become when it goes nonlinear. The waves break up into blobs of density which are drifted out to the boundary by their internal electric fields. Thus, plasma is lost in bunches. This guess was made in 1967, before diagnostic techniques were available to detect such blobs. However, in 2003, physicists at M.I.T. (Massachusetts Institute of Technology) developed a special technique which allowed them to take pictures of blobs as they carried plasma radially outward. One such picture of two simultaneous blobs is shown in Fig. 6.18 [6]. This behavior is not accidental, since it was observed also in several other tokamak machines. However, this is just a example

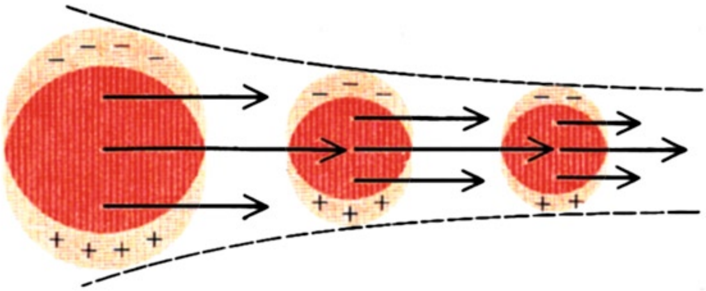

Fig. 6.17 Anomalous transport of plasma in blobs (adopted from Chen [7]). These are not spheres but long tubes of plasma curving with the magnetic field lines

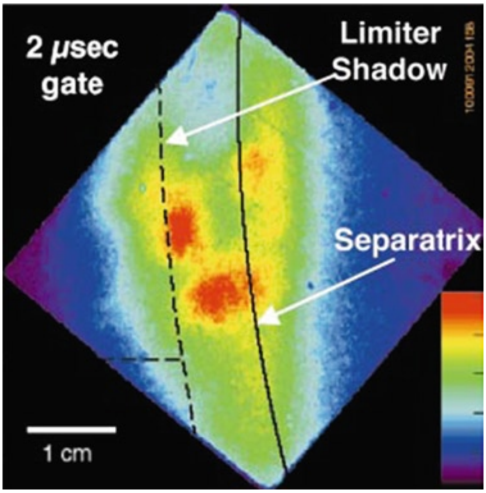

Fig. 6.18 A picture of blobs leaving a tokamak, taken at a shutter speed of 2 millionths of a second. The outside of the torus is on the *left* [6]

of how an instability, starting as a simple wave, can grow and carry plasma outwards. Other instabilities have been found to develop into other shapes as they do their dirty work.

In Chap. 4, we showed why a torus was chosen as a possible shape of a magnetic bottle used to hold a plasma hot enough to produce fusion power. In Chap. 5, we discussed the general features that had to be built into toruses in order to hold plasmas. In this chapter, we described the unexpected difficulties that were encountered in tokamaks and how these were overcome. These are the concepts which guided our work in the early days of fusion. In the four decades since that time, experiments on dozens, or even hundreds, of tokamaks, stellarators, and other magnetic devices throughout the world have led to improvements in design and advances in theoretical understanding. Tokamaks no longer look like simple circular toruses. The next chapter will tell why.

Vertical Fields

Before leaving this basic description of a tokamak, there is one more essential part that needs to be described: the vertical field. A ring of hot plasma will try to expand. Its internal pressure will push outwards so as to make the cross section fatter, and we have countered this force with a strong toroidal magnetic field. However, the plasma pressure will also tend to make the entire ring expand in radius, as shown in Fig. 6.19. The toroidal field is not good at restraining this motion because it is weaker on the outside of the ring than on the inside. Furthermore, the toroidal current in a tokamak creates a hoop force that also pushes on the ring to expand its major radius. This force arises from the magnetic field that the plasma current generates. This field is also stronger inside the hole of the torus than outside, so that its magnetic pressure is outward. Fortunately, these hoop forces are easily balanced by applying a small magnetic field in the vertical direction. Remember that in a tokamak there is always a current in the toroidal direction to give the field lines a twist. This current is mostly carried by the electrons. The Lorentz force on a moving charge, described in Chap. 4, is perpendicular to both the velocity of the particle and the magnetic field. By superposing a magnetic field in the vertical direction, either up or down, depending on the direction of the current, the tokamak current creates a Lorentz force that pushes the plasma ring inwards, toward the center of the torus.

Note that this effect is different from all the plasma drifts that we discussed previously. Those concerned individual particle motions; it did not matter how many particles there were. Here, we are considering the immense pressure of a hot gas.

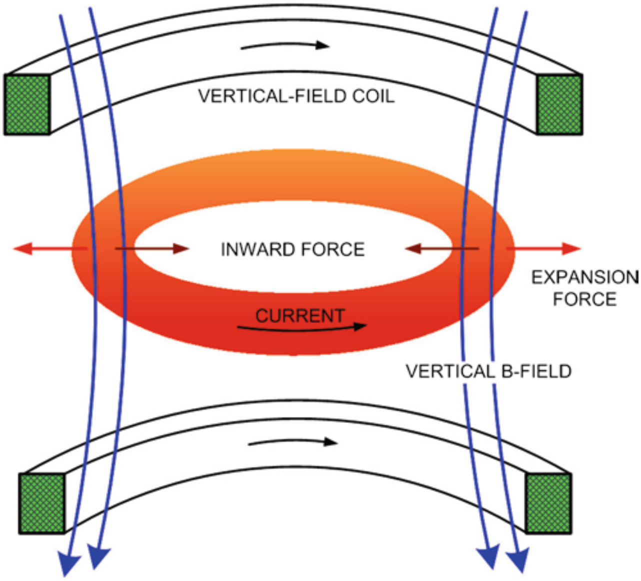

Fig. 6.19 Application of a vertical B-field to keep a ring of plasma from expanding

Thus, there are three main types of fields in a tokamak: the toroidal field generated by poloidal coils; the poloidal field generated by a plasma current; and a vertical field generated by large toroidal coils above and below the torus. These vertical-field coils can be combined with the coils that drive the plasma current, as will be described in the next chapter, so they do not always appear as a separate set.

Notes

1. In addition to the vertical drift due to the gradient of the toroidal field, there is also a smaller vertical drift due to the centrifugal force of particles whizzing around the torus the long way.
2. More likely, a collision takes a particle from a banana orbit to a passing orbit, and a second collision takes it from the passing orbit into another banana orbit.
3. (What is plotted here is perpendicular diffusion coefficient in m^2/s against Spitzer collision frequency in kHz. We have assumed 10 keV ions, 1 T magnetic field, aspect ratio $A=2.5$, quality factor $q=2$, and major radius $R=1$ m. The densities required to trace this whole curve would be unreasonably high. Fusion conditions have the very low diffusion rates in the extreme lower left corner.)
4. There are other devices, called reversed-field pinches, that have a very large toroidal current and only a small toroidal B-field. These depend on other stabilization mechanisms such as wall currents. But we are concentrating on tokamaks here because their development is further along.
5. The fact that banana diffusion does not depend on B_t comes from a cancelation between the vertical drift velocity, which varies as $1/B_t$, and the time a particle spends drifting in one direction, which varies as B_t. This is because increasing B_t for fixed B_p decreases the twist of the field lines.
6. For historical accuracy, neoclassical diffusion was discovered *after* Bohm diffusion was.
7. This holds only for an intermediate range of sizes.
8. This was not in a tokamak but in a stellarator. In tokamaks, a DC current is needed to create the rotational transform; in stellarators external coils are used to do this.

References

1. F.F. Chen, *Spectrum of low-β plasma turbulence*. Phys. Rev. Lett. **15**, 381 (1965)
2. H.P. Furth, J. Kileen, M.N. Rosenbluth, Phys. Fluids **6**, 459 (1963)
3. W.A. Newcomb, Ann. Phys. **3**, 347 (1958)
4. N. Rynn, N. D'Angelo, Rev. Sci. Instrum. **31**, 1326 (1960)
5. R.W. Motley, *Q-Machines* (Academic, New York, 1975)
6. S.J. Zweben et al., Phys. Plasmas **9**, 1981 (2002)
7. F.F. Chen, *The leakage problem in fusion reactors*. Sci. Am. **217**, 76 (1967)