

Sigma Xi, The Scientific Research Society

Pushing Particles with Waves: New ways of using waves to generate electric current are being applied to the problem of controlled nuclear fusion

Author(s): Nathaniel J. Fisch

Source: *American Scientist*, Vol. 71, No. 1 (January-February 1983), pp. 27-35

Published by: [Sigma Xi, The Scientific Research Society](#)

Stable URL: <http://www.jstor.org/stable/27851815>

Accessed: 18/02/2015 01:16

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Sigma Xi, The Scientific Research Society is collaborating with JSTOR to digitize, preserve and extend access to *American Scientist*.

<http://www.jstor.org>

Pushing Particles with Waves

New ways of using waves to generate electric current are being applied to the problem of controlled nuclear fusion

Just as the ladder in Jacob's dream stretched to the heavens from the ground, so controlled thermonuclear fusion reaches to the sun for the clue to inexhaustible fuel reserves on earth. While the goal is lofty, recreating the energy of the sun in the laboratory seems to require advances in supporting technologies as numerous as the rungs on the celestial ladder. To promote controlled fusion, improvements have been made in materials, lasers, neutral beams, and superconducting magnets; new demands have been put on computers and computational methods; and fundamental theories in physics and mathematics, such as those relating to turbulence, collisions, and magnetic field structure, are being scrutinized from a new viewpoint. It is thought that another rung in the ladder might be a new way of creating an electric current, not by batteries or the standard sources of electromotive force, but by pushing particles with waves.

That waves can be employed to push particles is evident from the phenomenon of surfing. The surfer's knack is to stay exactly in phase with the wave, which involves traveling at the velocity of the wave crest—the wave phase velocity—so that he is always sliding down the crest. In so doing, the successful surfer is ex-

ploiting an important principle of wave-particle interactions. Energy is transferred to him from the wave in the push he receives as he slides down the crest. This push allows him to maintain his speed and, consequently, his favorable phase with respect to the wave. As long as the phase is unaltered, the same favorable push persists. Because this interaction relies on a perfect match between the velocities of the wave and the surfer in the direction of the wave, it is termed a resonant wave-particle interaction.

This kind of interaction is not limited to water waves and surfers: there are many other types of waves and particles. For example, there are sound waves and light waves that propagate in familiar solids, liquids, and gases. In most media, however, these waves travel so quickly that particles never manage both to get started in phase and then to move at the wave phase velocity, and thus the particles do not enjoy the subsequent benefits of resonance. Like the surfer who misses his wave, these particles are left bobbing up and down in a far less impressive nonresonant motion.

There is one medium besides surf, namely plasma, that is particularly supportive of resonant wave-particle interactions (1). A plasma is a gas of atoms which have been ionized, so that the constituents react to electromagnetic fields as well as the usual collisions between them. Hence the collective motion of these constituents, in the form of a wave, is far more complicated and varied than the motion of atoms in a nonionized gas.

A plasma supports not only sound waves and light waves but a host of other waves with no counterparts in other media. With new waves come new opportunities for

pushing particles by means of resonant wave-particle interactions. Pushing the constituents of a plasma with waves is particularly interesting, because the plasma consists of oppositely charged species, and relative motion between oppositely charged species is an electric current. Since waves can produce this motion, it follows that new ways of generating electric current may be at hand.

The plasma state of matter has received much attention recently, because, as a step toward controlled nuclear fusion, exceptionally hot plasmas are being produced in laboratories around the world. It turns out that new methods of producing electric current by pushing particles with waves can solve problems that arise in controlling nuclear fusion. Therefore, before focusing on the techniques for pushing particles, we should examine their major application, controlled nuclear fusion.

Nuclear fusion

An enormous amount of energy may be released in the fusing together of the nuclei of two light atoms, such as isotopes of hydrogen. This process, known as nuclear fusion, is the means by which the sun and other stars produce their light. It has also been accomplished by man, but not yet in a self-sustaining, controlled manner suitable for peacetime utilization of energy. The problem is that for fusion to occur the two nuclei must be banged against each other with enough force to overcome their mutual electrical repulsion. In the hydrogen bomb, this is done by exploding small atomic bombs of fissile uranium in order to implode a core of liquid hydrogen; the tremendous pressure fuses the hydrogen into helium, thereby creating an even larger explosion. For peacetime use,

Nathaniel J. Fisch is a staff member of Princeton Plasma Physics Laboratory. His B.S. (1972), M.S. (1975), and Ph.D. (1978) are all from the Massachusetts Institute of Technology. He has done research on two-way diffusion equations, plasma turbulence, and wave propagation in various media, but the principal focus of his work has been continuous operation of tokamak fusion reactors. He has proposed several methods for achieving this and has produced several patents relating to steady state tokamaks. Address: Plasma Physics Laboratory, Princeton University, Princeton, NJ 08544.

however, the hydrogen nuclear energy must be released slowly, in a manner resembling gradual burning rather than a sudden explosion.

One approach is to heat hydrogen slowly to a high temperature by external means. As the temperature rises, individual hydrogen atoms collide with greater impact, and hence with greater probability of fusing, and eventually the hydrogen ignites—in other words, a self-sustaining reaction is achieved when the amount of heat generated by fusion reactions is large enough to replace the external heating. The temperature at the point of ignition is hundreds of millions of degrees, and at this temperature matter can exist only in the plasma state.

Where to put the plasma while it is being heated to ignition is quite a puzzle. It is far too hot to be in contact with material container walls. The approach that today is thought to be most promising is to confine the plasma by means of magnetic fields in a kind of magnetic bottle. There are several ingenious kinds of magnetic bottles, but the plasma, seemingly no less ingenious, manages to slip out of all of them. Luckily, it is necessary to confine the slippery plasma only long enough for fusion to occur, rather than indefinitely. Even so, the plasma is always finding ways to escape, and bottling it for even a short time is no mean task.

How does a magnetic bottle work? All magnetic confinement schemes exploit the fact that a plasma is a collection of charged particles. A charged particle cannot move in a straight line perpendicular to a magnetic field line—a line that points in the direction of the field. The magnetic field forces the particle into a circular orbit, and the larger the magnetic field, the tighter the orbit. A charged particle can, however, move parallel to a field line, so that the combination of circular motion perpendicular to the field and linear motion parallel to the field produces a helical orbit (Fig. 1).

A uniform magnetic field evidently achieves confinement in two dimensions. Particles are guided along the field lines, and their orbits are limited in extent perpendicular to these lines. The problem is then reduced to confinement in only one dimension, along the field line. One solution is to bend the magnetic field into a circle; particles following the

field lines would then be absolutely constrained. This is the operating principle of toroidal, or doughnut-shaped, confinement systems.

Unfortunately, however, bending the magnetic field introduces other forces on the plasma, which tend to destroy the confinement. The effects of these forces can be neutralized by twisting the magnetic lines of force while they are being bent into a doughnut shape. The resultant, twisted field is the sum of two fields, a toroidal field encircling the doughnut hole, and a poloidal field encircling the center of a minor cross section—what is exposed if you lay a doughnut flat on a table and slice it vertically (Fig. 2).

How are such magnetic fields generated? The toroidal field is easy—it is produced by a poloidal electric current in coils outside the plasma that encircle the minor cross section. The poloidal magnetic field is more difficult—it may be produced by a toroidal current, but this current must flow inside the plasma. This is the so-called transformer current, because the plasma is, in effect, the secondary in a transformer circuit. A current-carrying coil, the primary, in a plane parallel to the plane of the doughnut, is placed near the doughnut hole, but outside the plasma. If that current does not change with time, then no electric

field is produced. But if the current either increases or decreases with time, then a toroidal electric field, which may be steady, is produced inside the plasma. Since the plasma is composed of charged particles, it is conducting; thus, the electric field produces the required toroidal current, the transformer current. This is the tokamak approach to magnetic confinement, and it has emerged as the leading confinement scheme for controlled nuclear fusion.

The tokamak approach was formulated in 1950 by the Russian physicists and Nobel Prize winners Tamm and Sakharov (2). The word *tokamak*, coined by Golovin, is a Russian acronym for “toroidalnaya kamera i magnitnaya katushka,” meaning “toroidal chamber and magnetic coil,” after its salient features. A series of successful experiments on this concept, headed by Artsimovich of the Kurchatov Institute, was conducted in the 1960s. In 1968 both a very hot plasma and a long confinement time were reported on the Russian T-3 tokamak. This very persuasive argument for the tokamak approach led to immediate export of the idea. Following the T-3 experiment, a rash of experiments and calculations relating to tokamaks appeared in the US, Europe, and Japan.

Concomitant with the enthusiastic reception of the Russian achievement in the West came ideas for improvements. The tokamak was scaled up in size and optimized with respect to shape and other characteristics, and there were major advances in the means of heating and diagnosing the plasma (3).

It was also recognized that the tokamak was limited to pulsed operation, because the primary current cannot be increased indefinitely to maintain the transformer current. Even under optimal conditions, the tokamak fusion reactor could operate for only a few minutes at a time. Of course, the transformer current could be reinstated, the plasma could be reintroduced, and the fusion reactor could again engage in producing power. But this is a pulsed mode, and although many power production schemes—the gasoline engine found in today’s cars, for example—operate successfully in such a pulsed mode, the tokamak is expected to perform far more efficiently in a steady state mode.




Figure 1. Charged particles in a uniform magnetic field follow helical paths along field lines (*top*). In a circular magnetic field (*bottom*), such particles again follow helical paths along field lines, but because of the circular contours, they are confined.

Almost immediately upon the advent of the tokamak in the Western world, scientists began searching for ways to replace the transformer current by a continuous current. Among the leading replacement techniques were those that relied on pushing particles with waves.

Fast and slow electrons

That waves could be employed in generating toroidal currents was recognized as early as 1952, when Thonemann and his co-workers produced a current in a small cold plasma contained in a toroidal glass chamber (4). By today's standards, this current was miniscule, but it did establish the principle.

It was recognized early on that mere application of Thonemann's results to a tokamak reactor would not suffice because of the intolerably high power requirement. The challenge in tokamaks is to generate toroidal currents while minimizing power dissipation. To be practical, only a small portion of the power output of a fusion reactor can be used to produce the toroidal current.

The first principle adduced in the search for greater efficiency was that it is easier to push a slow electron than a fast one. To see this, consider an electron of mass m and velocity v which is given a small velocity kick,

Δ , in the direction of v . The initial energy of the electron is $\frac{1}{2}mv^2$, and the final energy is $\frac{1}{2}m(v + \Delta)^2$. For a small kick, the energy imparted to the electron is then $E \approx mv\Delta$. The momentum imparted is $p = m\Delta$; therefore $E \approx vp$. Since the current an electron carries, like its momentum, is proportional to its velocity, it follows that to increase the current by imparting momentum to a slow electron (v small) requires far less energy than imparting the same amount of momentum to a fast electron (v large).

This principle led researchers to look for ways of pushing slow electrons. One method is to use waves with low phase velocities, such as Alfvén waves (5). Another is to use beams of neutral atoms (6). It is noteworthy that both these schemes were proposed within months of the Russian T-3 tokamak success.

There is a second efficiency principle, however, one that was ignored by early researchers in the field. According to this second principle, what matters is not so much the energy required to produce a given current but the power required to sustain it. In other words, directed momentum, or current, put into an electron persists only until that electron collides with an ion, at which point the current is destroyed and must be replenished. For high

efficiency, it is not enough that the energy to replace the current be small; the need to furnish this energy must, in addition, be rare. It turns out that fast electrons collide far less often than slow electrons, which means they may well be worth the larger, but longer-term, investment in energy (7).

This dependence of the collision frequency on speed, which makes the fast electrons worthwhile current carriers, is, at first glance, somewhat unexpected. We are used to the frequency of collisions increasing with speed. Consider, for example, a famous movie actor picking his way through Times Square on foot. Each time he is recognized by an admirer, he may be stopped and asked for an autograph. Inasmuch as they slow him down, these encounters may be referred to as collisions. And, as we all know, hurrying through a crowd is of no avail; you only meet up—or collide—with people more often. Thus, as the actor increases his speed, he finds that there is less and less time between autograph requests; in other words, his collision frequency increases with increasing speed—the usual phenomenon. But if he really hurries, he may move through the crowd so quickly that even his most ardent fans fail to recognize him. At that point, his collision frequency—the number of autograph requests he receives per unit time—begins to drop dramatically with increasing speed.

What slows an electron as it speeds through a plasma is deflection by the electric field of other charged particles. This electric field can only deflect the electron in a large way, however, if the electron is close to the deflecting particle long enough to recognize it, so to speak. The faster the electron zips by, the less time it spends in the vicinity of deflecting particles, and thus the smaller the deflection. In fact, it turns out that the amount of deflection is inversely proportional to the cube of the electron speed.

Now that we have identified the most efficient classes of electrons to push, we must come up with a method of pushing them. The plasma is quite cooperative in this regard, in that it supports many waves that may be easily and efficiently injected into it. To push slow electrons, waves of comparatively low frequency (Alfvén waves) may be used; these




Figure 2. Confinement of a plasma in three dimensions requires twisting magnetic field lines as they are being bent into a doughnut shape. The field that results has two components: a toroidal component encircling the doughnut hole and a poloidal component encircling the minor cross section. It is possible to inject waves from outside the doughnut; these waves then propagate in the plasma until they give up their energy and momentum to resonant particles.

may be excited by means of loop antennas. To push fast electrons, it is convenient to use a wave of higher frequency (around 1–10 GHz), the so-called lower-hybrid wave, which may be injected into the plasma by means of waveguides.

When it became known that the fast, relatively collisionless electrons were fair game for producing currents, experiments using some of the largest tokamak facilities were designed. Although these experiments are now the subject of ongoing debate and are probably destined to remain puzzling in many ways, there does seem to be a certain range of plasma parameters within which lower-hybrid waves generate current very efficiently. These experiments, which are being conducted at MIT, Princeton, and the Japan Atomic Energy Research Institute, among other places, all employ transformer coils in addition to lower-hybrid waves. A completely steady state tokamak is yet to be demonstrated in any parameter regime, but scientists are getting closer.

“Look Ma—no momentum!”

The principle that has been put forth is that wave momentum, whether residing in slow or fast waves, may be utilized to push electrons. The ions in the plasma remain relatively stationary, because the frictional forces on them due to collisions with drifting electrons are very small and are easily counterbalanced by other small forces not considered here. Current is thus produced by pushing the electrons. This principle, so intuitively satisfying, appears almost to dictate a physical law—that input of momentum is a prerequisite for generating a current.

Surprisingly, however, and in seeming defiance of our intuition, there is no such prerequisite. Before the introduction of waves, there is a thermal distribution of electrons, with exactly half traveling in one toroidal direction and exactly half traveling the other way. By virtue of collisions with ions, electrons going to the right try to push ions to the right, while electrons going to the left try to push ions to the left. Of course, since these forces are symmetric, they cancel each other out, and the ions are not persuaded to move in either direction.

But now imagine that electrons moving to the right collide less often with ions than do electrons moving to the left. Certainly, the situation is no longer symmetric, and the ions must feel a force pushing them to the left. On the other hand, no momentum has been injected into this system of particles; only the collision frequency of some electrons has been (at this point, let us say, mysteriously) altered. Hence, while the ions must, on the average, drift to the left, the total electron population must, on the average, drift to the right, since total momentum is conserved. Thus, slipping through a technical loophole that adheres to the letter but not the spirit of the physical laws, ions and electrons counterdrift and produce an electric current (8).

Imbalance in collision frequency, it turns out, can be achieved by exploiting the cyclotron motion of electrons—the orbiting motion in the direction perpendicular to the magnetic field lines to which we have already referred. The frequency of this motion is known as the cyclotron frequency and is typically about 100 GHz. Electromagnetic waves of this frequency propagate freely in the plasma until they encounter electrons orbiting with exactly the same frequency, at which point, the energy of the electromagnetic wave is transferred to the perpendicular motion of these electrons. Once the electrons absorb this energy, they travel faster. That means that they will collide with ions less often.

Happily, this transfer of energy is not entirely symmetric because of the phenomenon known as the Doppler shift. This phenomenon, which makes the toots of retreating trains sound lower in frequency than those of advancing trains, has a similar effect on the cyclotron frequency. Accordingly, if the cyclotron waves are traveling in one toroidal direction, their frequency will appear different to electrons traveling in opposite toroidal directions. We can thus pick the wave frequency so that only electrons traveling in one toroidal direction perceive themselves to be in resonance.

To see this effect in a different way, consider the following. If an electron moving with some parallel velocity v_0 is given a push in the direction perpendicular to the magnetic field lines, what will happen to

the total electron current, J_e , in the course of time? Since the initial push does not change the parallel momentum, there is no electron current initially. But electrons going to the left now collide more frequently than those going to the right, which implies that in about the time, t_c , that it takes for the pushed electron to collide with an ion, there will already be a noticeable effect, and the electron current will be at a maximum (Fig. 3). Further collisions will relax any asymmetry, so that, in the absence of repeated pushes, the current will eventually vanish. In sustaining a tokamak current, of course, electrons would be pushed repeatedly.

This unusual effect has opened up the possibility of generating current by means of waves that were hitherto considered inappropriate because of their small momentum. In particular, electron cyclotron waves, which are rather simple electromagnetic waves that could be transmitted even in the absence of the plasma, have now been recognized as eminently suitable candidates. Recommending them are their propagation characteristics and their ability to interact with very fast electrons. As with free space waves, the parallel phase velocity of these waves is greater than the velocity of light, which means that the ratio of their parallel momentum to their energy is very small. Nonetheless, by increasing the velocity of an electron in the perpendicular direction, these waves can generate current with almost the efficiency attained by pushing the same electron in the parallel direction, and, in accordance with the second efficiency principle, the faster the electron, the greater the efficiency. What this indicates, incidentally, is that even in the case of lower-hybrid waves, which have considerable parallel momentum, the dominant factor in driving the current is the decrease in the resonant electron's collision frequency (because of the injection of energy) and not the immediate increase in its parallel momentum.

Although electron cyclotron waves surpass lower-hybrid waves in their ability to interact with very fast electrons, there is a limit to how much benefit can be derived from increasing the speed of the electrons. Eventually, an interesting relativistic effect arises (9). A relativistic electron has the property that when its energy

is increased, its mass also registers a substantial increase. Thus, if a wave gives a relativistic electron perpendicular momentum without tampering with its parallel momentum, its parallel velocity must decrease if its parallel momentum is to remain unchanged. Since current is proportional to velocity, not momentum, this appears to be yet another mechanism for generating current without injecting momentum. Unfortunately, however, this current is always in the opposite direction to that produced by asymmetric collisions; because of this cancellation there is, in fact, an upper limit to how far efficiency can be increased by pushing very fast electrons.

Two species of ions

In view of the manifold possibilities for generating current by pushing electrons with waves, the question naturally arises as to whether the ions too might be pushed to advantage. Pushing ions, however, is quite different from pushing electrons. Not only are ions relatively heavy, which means that they must absorb much more momentum to generate the same amount of current, but, more important, the light, swift electrons will quickly collide with the ions, thereby canceling any ion current.

Can electrons be prevented from catching up with the ions? The answer is most easily understood if the current is calculated in the frame of reference in which the total ion current is zero. This is allowed because current in a neutral plasma is proportional to the difference of two velocities, namely, the ion and electron drift velocities, and is thus frame-independent, just as the relative velocity of two bicyclists viewed from a train is independent of the speed of the train.

Accordingly, consider two groups of ions in this frame of reference. The majority of the ions are to be found moving slowly to the left, while a minority comprises a faster beam moving to the right. How will the electrons respond? Each group of ions collides with electrons, pushing them in its direction, and the group that is most persuasive will be the group that collides most frequently. If the velocities of both groups of ions are less than the typical electron velocity, called the thermal velocity,

the typical impact speeds of electrons will be the same for the two groups. The collision frequencies, which depend on this impact speed, are then identical, and the tug-of-war will be a stalemate.

The tug-of-war can be won by the majority ions—although it turns out to be a Pyrrhic victory—if the minority ions have a velocity that is large compared to the electron thermal velocity. In that case, the majority ions still collide with electrons with an impact speed equal to the electron thermal velocity, but the minority ions collide with a much greater impact speed, on the average, and, consequently, with a much lower frequency. Electrons, because they collide more often with majority ions, will therefore be urged to drift to the left. Unfortunately, however, production of current by this mechanism is not feasible in a tokamak reactor because of the tremendous energy required to create such ions in the first place.

A practical scheme employing ions to drive current, suggested by Ohkawa (6), exploits the disproportionate deflection of electrons by ions with a high charge state (i.e., a high atomic number). In the presence of counterstreaming ions with disparate charge states but zero total current, electrons collide more often with ions of higher charge state and thus tend to follow them. Since this effect does not rely on the impact speed, both groups of ions may be drifting with speeds considerably less than the electron thermal speed.




Figure 3. If an electron is given a push in the direction perpendicular to the magnetic field lines, without there being any change in the parallel direction, the current, J_e , will increase from zero at the time of the push to a maximum value by the time the pushed electron collides with an ion, t_c . In the absence of any further push, the current will subsequently decline to zero again.

Ohkawa reasoned that if a beam of atoms with high atomic number were aimed tangentially into a tokamak, the atoms, being neutral, could slip through the magnetic fields confining the plasma. Once inside the plasma, the atoms would quickly ionize, thereby producing the streaming ion beam required. The attractiveness of this scheme is related to the first efficiency principle, since the injected beam could be slow enough to have a large ratio of momentum to energy.

This is all quite nice, except that it would be better to do it with waves, since wave energy is considerably cheaper, and the powerful neutral beams required are cumbersome and intrusive. Moreover, the apparatus for delivering wave energy to the tokamak can survive more easily than the neutral beam apparatus in the harsh reactor environment. But the plasma ions are going to have to be tricked into counterstreaming.

In thermal equilibrium, half the minority ions stream to the right and half to the left. The frequency of collisions between ions and electrons is roughly independent of the ion velocity (which is, in any case, far smaller than that of the electrons), but the frequency of collisions between pairs of ions is very sensitive to the speed of each. Suppose that minority ions going to the right are given a push in energy in the direction perpendicular to that in which they are traveling. Although they collide with electrons as frequently as before, they now collide less often with majority ions. So, whereas the electrons feel no immediate effect, the majority ions, which collide more frequently with minority ions going to the left, are pushed to the left. For momentum to be conserved, minority ions must, therefore, drift to the right on the average. Thus, counterstreaming populations of ions are produced, and if their charge states differ, an electron current develops (10).

There are waves that have already been shown to interact primarily with minority ions, which adds to the attractiveness of this proposal. In the Princeton Large Torus (PLT) experiment, low-frequency waves have heated minority ions both with very high efficiency and without directly heating either majority ions or electrons. This selective heating identified minority

ions on the basis of their charge/mass ratio, which determines their gyro-frequency—the frequency with which they encircle the magnetic field lines.

The PLT experiment was designed to heat plasma, not to drive current, and waves were launched symmetrically in both toroidal directions. Nonetheless, success in selective heating of minority ions by symmetric launching of waves naturally bodes well for the generation of current by asymmetric launching.

Hairy doughnuts and groundhogs

Having discussed a variety of methods that have been proposed for driving steady state electric current, methods that would allow tokamaks to operate continuously, I shall now say more about why continuous operation is important. The quest for such operation is motivated by considerations that impinge upon very different aspects of what is boldly envisioned as the economical tokamak reactor.

First, a continuously operating tokamak would be less susceptible to the metal fatigue that arises in a pulsed device from the heat stress to which the structural components are brutally subjected. This stress is caused by the fact that the device is hot when on and cold when off. Only continuous operation would provide the uniform heat load needed to avoid such stress; other ways of avoiding it are, in principle, possible, but they have been discounted as technologically undesirable. Moreover, a truly steady state device would be less vulnerable to the disruptions—sudden and unpredictable losses of control of the plasma—to which a pulsed device is particularly prone, and which are damaging to its structural components.

Second, although no one is yet sure what a tokamak reactor will look like, it is expected that the toroidal magnetic field will be provided by superconducting coils, which must be kept at extremely low temperatures. In a pulsed reactor, these coils would be heated by the time-dependent poloidal magnetic field; the cooling requirements, and hence the refrigeration costs, would thus be

greater than they would be in a continuous reactor.

The third objection to the pulsed tokamak relates to the premium location in the doughnut center that the cumbersome transformer coils must occupy. Equipment required for confinement of the plasma or extraction of the fusion energy also competes for this valuable space. Further, it is more economical to construct a tokamak as a fat doughnut—that is, a doughnut with a very small hole. If the toroidal current is driven by waves, which need not be introduced from the doughnut hole, it is possible to construct such a geometrically superior reactor.

In fact, when push comes to shove, we must admit that it would be desirable to eliminate the doughnut hole entirely. Fuel is always burned, it seems, in much simpler places than toruses: coal for generating steam is burned in a furnace, and gasoline to drive a piston is exploded in an engine chamber, for example. Such places, which have geometries that mathematicians call “simply connected” (as opposed to rings or doughnuts, which are not simply connected), are always conducive to the containment and combustion of fuel and the recovery of the energy. Some variants on the tokamak, such as the so-called spheromak, do try to do away with the hole, but there are problems, which are related, strangely enough, to hairy objects such as groundhogs and doughnuts.

The difference between tokamaks and other fuel-burning devices lies in their walls. There is no problem, in principle, in fashioning the walls of a furnace in any shape whatsoever. But a plasma must be contained in a magnetic bottle, and the magnetic fields that comprise the walls have certain geometrical properties that preclude them from covering the simply connected spaces in which fuel is most easily handled (11).

To see this, consider the hairy groundhog, the simply connected North American rodent that resembles a ball covered completely by hair (Fig. 4). It might be imagined that the appearance of this animal would be improved by combing. Its resistance to combing is far more serious than that of unkempt and recalcitrant children, however; these animals

resist combing mathematically! Imagine making horizontal strokes with a comb, beginning at the groundhog’s “equator” with wide circular sweeps and then progressing to smaller circles at “higher latitudes.” As we approach his “north and south poles,” the circles become very small indeed, and when we reach the “poles,” the hair has no place to go: if flattened in any direction, it will be perpendicular to other hair. Even the best-groomed groundhog is still left with two cowlicks.

On the other hand, the hairy doughnut, happily a variety not yet on the market, can be combed. It is possible to comb it either poloidally, in small circles that intersect the doughnut hole, or toroidally, in large circles that are concentric with the hole. Indeed, the doughnut may be combed in an infinity of combinations of these two directions, and each and every hair may be placed parallel to its neighboring hairs so that there are no cowlicks.

The magnetic field is somewhat like hair that must be combed. The containment of charged particles in a tokamak relies on the presence of bounding surfaces, to which the magnetic field is always parallel. Since the particles try to follow the field lines, they remain on the surface. That the field lines at one point on the surface always point to another point on the surface can be represented mathematically as a mapping of the surface, and the inevitable presence of cowlicks on some objects is a consequence of what is known as the fixed point theorem.

To understand this elegant theorem, take an image of a sheet of paper, deform it any way you like by stretching and shrinking, but not ripping, and then place the deformed image in any orientation on the original sheet. The fixed point theorem states that at least one point on the original will coincide with its image. If, instead of a sheet of paper, you take a spherical surface, there will be at least two fixed points. With a torus, however, there need not be any fixed points, which means that in a doughnut-shaped tokamak, there need not be any places at which the magnetic field has nowhere to point—that is, at which it would vanish—and at which particles could leak out. Because fixed points must be




Figure 4. If a naturally unkempt hairy groundhog is combed, using parallel strokes in the horizontal (or any other) direction, he will be left with two cowlicks, even in the best-groomed case. If a hairy doughnut is combed, on the other hand, it is possible to end up with no cowlicks, since there is an infinity of directions in which every hair can be made to lie parallel to its neighboring hairs.

avoided, the doughnut hole cannot be eliminated entirely; pushing particles with waves can only shrink it.

The versatility of waves

Considerable attention has been paid to many different ways of using waves, partly because at this point there is no demonstration or guarantee that any one technique will perform as advertised, and partly because our conception of what precise form a reactor will take is incomplete. On the drawing board, tokamak reactors have already assumed a wide range of sizes, temperatures, plasma densities, magnetic field strengths, and so on. Certain combinations of values for these parameters are more hospitable to some techniques of producing current than to others. How a reactor is designed will depend on which techniques pass experimental tests in the coming years. Given all the unknowns, there is safety in exploring a large number of techniques. Comfort may also be taken in their reliance, as a group, on three different frequency regimes and on a variety of mechanisms of interaction of waves and particles to produce current.

It would be too much to describe all the techniques here or to determine which is to be preferred in a particular type of reactor and what trade-offs must be made with other design considerations. But it is impossible not to reflect on the sheer variety and remarkable versatility of waves in pushing particles. The option provided by waves of pinpointing exactly which particles are to be pushed and in what direction is itself cause for comment, let alone its potential for employment to great ad-

vantage in maximizing the efficiency with which current is generated.

The capability of pinpointing particles is largely a consequence of the resonance condition, which selects with precision particles of a given parallel velocity. The component of velocity perpendicular to the magnetic field is not relevant as far as resonance is concerned, because particles are constrained in small orbits in that direction; thus, to stay in phase with the wave crests, the velocity component of the particle parallel to the field must match the phase velocity of the wave.

Some of the resonant electrons may be pushed harder than others on account of their perpendicular velocity. A particle with a large perpendicular velocity undergoes large

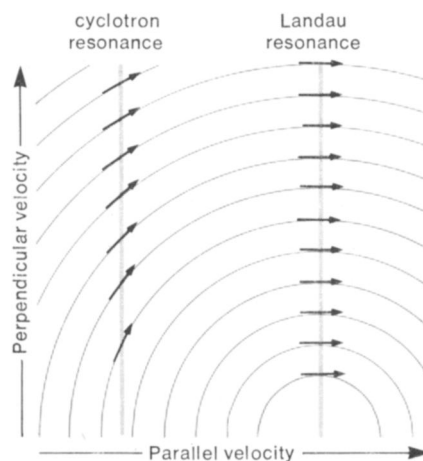


Figure 5. Particles remain with constant energy in the frame of reference of a wave when exchanging energy and momentum with that wave. Particles responding by virtue of a cyclotron resonance experience an increase in energy predominantly in the perpendicular direction. Particles whose velocity equals the parallel wave phase velocity satisfy the Landau resonance; these particles can get only a parallel push from the wave.

perpendicular excursions as it glides along the magnetic field lines. It thus samples wave fields at neighboring locations and can be subjected to new forces, should the wave amplitude vary in the perpendicular direction. This creates the opportunity of distinguishing it from particles with small perpendicular velocities. An example of a wave that discriminates in this way is the Alfvén wave, which preferentially pushes electrons with large perpendicular velocities in much the same way that magnetic mirror machines, an alternate means of confinement, preferentially reflect such electrons.

The gyrating particle motion opens the door to even more varied behavior. When the wave frequency matches the cyclotron frequency, particles can be pushed substantially, just as a child on a swing is responsive to pushing at the pendulum frequency (which entails a push only when he reaches that point of his orbit where his parent awaits him—a steady force, such as a wind, will not keep the child swinging). By means of this cyclotron resonance, particles can also be pushed in directions other than the parallel direction. To see this, consider that the energy of a wave is proportional to its frequency, whereas the momentum is proportional to its wave number. In the frame of reference of the wave—i.e., the frame moving along with the wave crests—the wave frequency is zero, which implies that the wave has no energy to impart to particles. The wave can still exchange momentum with particles, however, as long as the energy of the particles is constant in the wave frame. In other words, particles interacting with the wave are constrained to move along semicircular paths, as depicted in Figure 5.

This constraint applies also to those wave-particle interactions where the resonance does not depend on the cyclotron frequency—those involving lower-hybrid waves or surfers. At such a resonance, known as the Landau resonance after a pioneer in wave-particle interactions and other fields, the paths lie in the parallel direction, which means that particles can get only a parallel push from the wave. At a cyclotron resonance, the push can be in some other direction; indeed, if either the wave number or the velocity in the

parallel direction is small, the push will be almost entirely in the perpendicular direction. By choosing the wave parameters carefully, it is possible to push particles from highly populated to sparsely populated states by means of diffusion; in general, this makes it possible to push particles in all directions that point toward higher energy.

Waves can not only finely select particles on the basis of their velocity in both parallel and perpendicular directions, they can also discriminate on the basis of species of charged particles, since different species have different cyclotron frequencies. And now, as we have seen, the particles can be pushed in any direction. All these possibilities for driving current with waves have been explored. Waves can be chosen such that one type of resonance is dominant and the corresponding resonant particles are preferentially pushed. In every case, particles are diffused by waves from a highly populated, low-energy state to a sparsely populated, high-energy state.

For completeness, it should be remarked that a fascinating twist to the traveling wave, the focus of consideration here, is the accelerating wave, a wave whose phase velocity changes in time or space. This kind of wave can be tailored to remain in resonance with accelerating or decelerating particles. As might be imagined, new opportunities arise from this additional degree of freedom. Accelerating waves, which push particles coherently rather than diffusively, are useful for free electron lasers and intense charged-particle beams but are not really applicable to tokamaks because of the difficulty of maintaining coherence.

Laws of resistivity

All the proposed schemes for driving current with waves would consume very little of the output power of large reactors of the type under consideration in the early 1970s. At that time, tokamak reactors were envisioned to grow to about 26 m in major diameter and about 10 m in cross-sectional, or minor, diameter, and the power required to generate the current with waves could have been only a few percent of the several gigawatts these reactors were designed to generate. More recently, however,

tokamak reactors have been thought of as much smaller objects, perhaps only 15 m in major diameter and 6 m in minor diameter. The power required to generate current by waves in the smallest and coldest of these would comprise a large fraction of the fusion output, too much to be worth the trouble, perhaps. If waves are to be used in smaller and intermediate-size reactors, it is important that the current be generated with maximum efficiency by carefully choosing among the wave candidates and other remaining free parameters.

The most efficient way to generate current is by the transformer direct current toroidal electric field, which consumes very little power. This field exerts an equal force on all electrons, regardless of velocity, and an equal but opposite force on all ions. The amount of momentum transferred to particles does not depend on their collision frequency. What makes the current generation so efficient is that particles are pushed directly and not by means of diffusion, which always pushes particles to higher energy. An electron traveling with parallel velocity v and its symmetric counterpart traveling in the opposite direction with parallel velocity $-v$ are given the same velocity increment, Δv . The first electron gains energy from the push, the second loses it, and it may be verified that the total energy gained by the two is proportional to the square of Δv . Since the current is proportional to the average velocity increment Δv , the power dissipated will be proportional to the square of the current. The constant of proportionality is a measure of the resistivity and depends on the average collision frequency, which in turn is some function of the plasma temperature. This is Ohm's law, which serves well for both copper wires and tokamak plasmas.

Pushing particles with waves gives rise to a completely different law of resistivity, according to which the power dissipated is proportional to the current rather than to its square. This law implies further differences between the two mechanisms for generating current. In denser plasmas, for example, more ions try to destroy a wave-generated current, so that more power is required to sustain it. The same ions, of course, try to scatter current-carrying

electrons in the case of transformer-generated currents, but with the important difference that more electrons are available to carry the current, so that their average drift velocity, Δv , can be smaller. Since the power required to maintain these currents decreases with decreasing Δv , the two effects cancel, rendering plasma resistance to transformer-generated currents independent of plasma density.

In a completely steady state reactor, only the law of resistivity for wave-generated currents applies, and this law may result in a high price for generating current, at least in smaller and denser designs. Not to be overlooked, however, is the possibility of confusing the plasma terribly by periodically shutting off the waves. This would cause the plasma current to decay, but very lethargically because of inductive effects, which are nature's way of resisting sudden change. The decaying magnetic field accompanying the decaying current would produce an electric field that drives current in just the direction to reinforce the current. The resistivity law obeyed by this supportive, or additive, current would, however, be Ohm's law. Subsequent employment of waves to return the sagging current to a respectable value would give rise to magnetic fields that vary with time, which would again generate inductive currents obeying Ohm's law; this time, however, the currents would be subtractive—that is, in the opposite direction to the wave-driven current.

The additive and subtractive inductive currents must cancel on the average, if all the plasma parameters are fixed. There is now room in town, however, for a new game. Many of the advantages of completely steady state operation would be retained were the tokamak to operate as a partly steady state device, in which most, but not all, of the plasma parameters are held steady. The current might be generated by waves and held very nearly constant, for example, but the balance between the wave-generated and inductive contributions could vary synchronously with other plasma parameters such as density, temperature, or impurity content. The additive and subtractive inductive currents would not then cancel on the average; juggling the two laws of resistivity in this manner could indeed lead to a sharp reduc-

tion in the average power required to sustain the current.

The completely steady state reactor may be the ultimate preference, but the milestones leading to it, probably smaller and colder varieties of reactor, may be partly steady state devices. The possibility of oscillating some plasma parameters clearly introduces many new levers to pull and buttons to push, as well as many new rules by which to play the game. The possibilities and ramifications are fascinating (for a summary, see ref. 12), but to expound on them here would be to begin another story. What started out as a struggle between wave and particle has developed into a contest between man and plasma, the tricky versus the slip-

pery, with surfers and groundhogs mere spectators at what promises to be an interesting sideshow to center-stage fusion.

References

1. T. H. Stix. 1962. *The Theory of Plasma Waves*. McGraw-Hill.
2. I. E. Tamm and A. D. Sakharov. 1961. In *Plasma Physics and the Problem of Controlled Thermonuclear Reactions*, ed. M. A. Leontovich, pp. 1-47. Pergamon.
3. H. P. Furth. 1979. Progress toward a tokamak fusion reactor. *Sci. Am.* 241:50.
4. P. C. Thonemann, W. T. Cowhig, and P. A. Davenport. 1952. Interaction of traveling magnetic fields with ionized gases. *Nature* 5:34.
5. D. J. H. Wort. 1971. The peristaltic tokamak. *Plasma Phys.* 13:258.
6. T. Ohkawa. 1970. New methods of driving plasma current in fusion devices. *Nuclear Fusion* 10:185.
7. N. J. Fisch. 1978. Confining a tokamak plasma with rf-driven currents. *Phys. Rev. Lett.* 41:873.
8. N. J. Fisch and A. H. Boozer. 1980. Creating an asymmetric plasma resistivity with waves. *Phys. Rev. Lett.* 45:820.
9. N. J. Fisch. 1981. Current generation in a relativistic plasma. *Phys. Rev. A* 24:3245.
10. N. J. Fisch. 1981. Current generation by minority-species heating. *Nuclear Fusion* 21:15.
11. M. D. Kruskal and R. M. Kulsrud. 1958. Equilibrium of magnetically confined plasma in a toroid. *Phys. Fluids* 1:265.
12. N. J. Fisch. 1982. Current generation in toroidal plasma. In *Heating in Toroidal Plasmas: Proceedings of the Third Joint Varenna-Grenoble International Symposium*, vol. 3, ed. C. Gormezano, C. G. Leotta, and E. Sindoni, p. 841. Brussels: Commission of the European Communities.



“Wake up, Dr. Erskine—you’re being transferred to low energy physics.”