

SOME UNSOLVED CHALLENGES IN RADIO-FREQUENCY HEATING AND CURRENT DRIVE

N. J. FISCH*

Princeton Plasma Physics Laboratory, Princeton, New Jersey 08543

Received April 24, 2013

Accepted for Publication May 28, 2013

doi:10.13182/FST13-682

Several unsolved challenges in radio-frequency heating and current drive are highlighted. These include current drive in magnetic geometries in which the toroidal magnetic field cannot be assumed to be dominant, current start-up with hyperresistivity, current drive with oscillating parameters, and synergistic effects between current drive and alpha channeling. These challenges are not necessarily straightforward to address, and it is possible that the challenges cannot even be met, but were they met,

at least in some cases, there is the potential of significant consequence.

KEYWORDS: *noninductive current drive, rf waves, rf current drive*

Note: Some figures in this lecture may be in color only in the electronic version.

I. INTRODUCTION

There are many methods by which radio-frequency (rf) waves drive the toroidal current in tokamak reactors.¹ A quick overview of these methods is given in another lecture in this issue.² However, it remains the case that the efficiency of these methods is marginal for the economical production of all the necessary current in modern designs for a reactor. By *necessary current*, we mean the current that is necessary after the plasma self-generated currents, like the bootstrap current, are taken into account. By *modern designs*, we mean a tokamak reactor that is not optimized for current drive efficiency, but is rather optimized in other ways.

A tokamak reactor optimized for current drive efficiency would tend to be large, hot, and not so dense. Methods of rf current drive tend to be less efficient in plasmas that are smaller, denser, and cooler. The original ideas on lower hybrid current drive³ (LHCD) anticipated that the reactor would carry a bootstrap current fraction of ~50% and that it would be very large and not very dense.

Such an optimized tokamak, hot, but not dense, might still have similar fusion reactivity, which varies as the pressure squared. But, it would be less collisional on both accounts, lower density and higher temperature, which allow currents more easily to persist. Also, a large reactor needs less current density. It also needs less fusion reactivity, at constant wall loading. Modern designs call for much smaller tokamak reactors, at higher pressures and at higher bootstrap fractions, but tend to be denser and colder. The modern designs are driven by the perception that fusion reactors must have low capital cost in order to be economically competitive with other sources of energy. This means that the current drive must be accomplished in an environment not optimized for current drive efficiency.

Nonetheless, there are, in principle, possibilities to achieve much more attractive reactors driven by rf waves. Here, we point out some of the speculative but high-payoff possibilities. These possibilities are extremely unlikely. If they were likely, then we would be presenting to you solutions rather than challenges. Instead, the approaches outlined are very challenging. Some of the ideas offered may even be absurd. But, they cover the space of possibilities with high upside potential.

*E-mail: fisch@princeton.edu

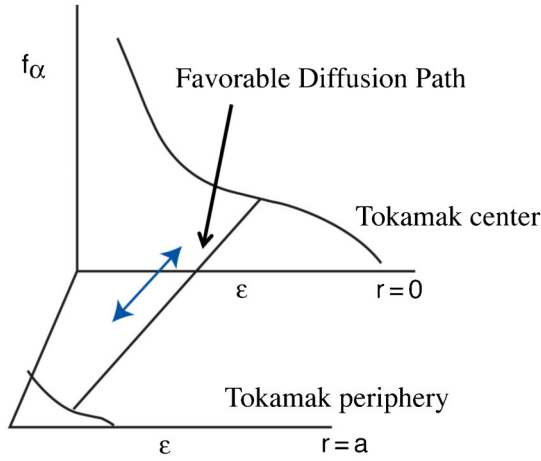


Fig. 1. Schematic distribution of alpha particles versus alpha-particle energy ϵ at $r=0$ (center of tokamak) and at $r=a$ (periphery of tokamak). At any radius, the alpha-particle distribution function is monotonically decreasing in energy. However, there is a population inversion in energy along the indicated favorable diffusion path.

II. CHALLENGES IN ALPHA CHANNELING

The accompanying lecture² points out that, on the one hand, the current drive in a reactor will have to be accomplished in the presence of alpha particles, which could damp any injected waves. However, on the other hand, there may be an opportunity to extract the alpha-particle energy, where there are several important unsolved challenges. Before we get into the unsolved issues, let us describe the alpha-channeling effect⁴ in greater detail.

The alpha-channeling effect relies upon a diffusive wave-particle interaction (see, for example, Ref. 5). Alpha

particles born with 3.5 MeV are born preferentially in the hot and dense plasma center. In the center, the alpha particles might assume a slowing-down distribution in more peripheral regions, and the alpha particles are less dense and less energetic. In Fig. 1, we show such a distribution of alpha particles monotonically decreasing in energy ϵ at various radii in a tokamak.

Thus, there is a population inversion in energy along the indicated favorable diffusion path. This diffusion path occurs in the joint energy-radius space. The alpha particles diffused along this diffusion path must leave the tokamak cold, because, absent collisions, the constraint to remain on the rf diffusion path is very strict. Leaving the tokamak with little energy is then the only way that the alpha particles can stop interacting with the wave. This contrasts with calculations that assumed diffusion paths in velocity space only,⁶⁻⁸ which described how alpha particles damp the wave. Here, the wave grows at the expense of the energy extracted from the alpha particles diffused along the favorable diffusion path. The energy that flows to the wave may then be used for other purposes such as achieving a hot-ion mode or current drive.

Such a diffusion path requires coupling diffusion in space to diffusion in energy. This can happen, for example, to an alpha particle in a uniform magnetic field B , as shown in Fig. 2. The magnetic field is into the paper, in the \hat{z} -direction, so that the alpha particles rotate in the counterclockwise direction with frequency $\Omega = qB/m$ and with gyroradius $\rho = v_{\perp}/\Omega$. Consider now an electrostatic wave with phase velocity ω/k_y , with a wavelength short compared to the gyroradius, so that $k_y v_{\perp}/\Omega \gg 1$. This wave interacts resonantly through a Landau resonance such that $\omega - k_y v_y = 0$. So long as the alpha particle is energetic enough, namely, $v_{\perp} > \omega/k_y$, there will be two points on the orbit in which this resonance is satisfied. When the resonance is satisfied, the alpha particle gets an instantaneous kick in the \hat{y} -direction, which, depending on

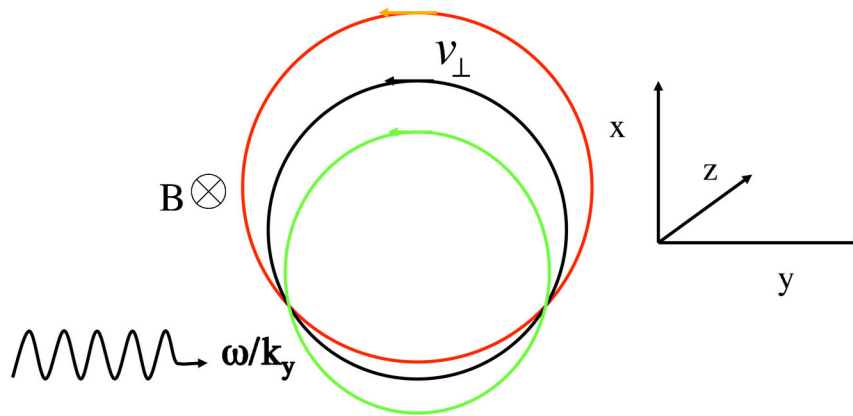


Fig. 2. Ion orbits in a homogeneous magnetic field in the presence of a resonant short-wavelength electrostatic wave traveling in the y -direction. If initially in the black orbit (middle circle) and the ion gains energy from the wave, the red orbit (upper circle) results; the green orbit (lower circle) results when the ion loses energy to the wave.

the phase of the wave, could be either to increase its energy or to slow it down. Thus, the velocity in the \hat{y} -direction changes instantaneously but randomly.

The random interaction with the wave produces an instantaneous velocity change in the \hat{y} -direction $v_y \rightarrow v_y + \Delta v_y$. The velocity change in the \hat{y} -direction moves the guiding center in the \hat{x} -direction, such that $x_{gc} \rightarrow x_{gc} - \Delta v_y / \Omega$. The perpendicular energy also changes instantaneously by $mv_y \Delta v_y$ (for Δv_y small). Thus, the change in the gyrocenter in the \hat{x} -direction, Δx_{gc} , is proportional to the energy absorbed $\Delta \varepsilon$ (Ref. 4), namely,

$$\frac{\Delta x_{gc}}{\Delta \varepsilon} = -\frac{1}{m\Omega v_y} = -\frac{k_y}{m\Omega\omega}, \quad (1)$$

where the last equality can be written since the interaction occurs instantaneously just when $v_y = \omega/k_y$.

This coupling of diffusion in space to diffusion in energy is depicted in Fig. 2, where the ion begins in the black orbit (central orbit). If the energy decreases as a result of the interaction, then the radius must become smaller (green orbit), with a guiding center lower in the \hat{x} -direction. If the energy increases, then the radius must become larger (red orbit), with a guiding center higher in the \hat{x} -direction. Upon repeated interactions with the wave, a particle will trace a line in $\varepsilon - x_{gc}$ space. With diffusion in energy thus coupled to diffusion in position, the favorable diffusion path depicted in Fig. 1 is enabled.

While the slab geometry derivation above captures the key features, the diffusion path in a tokamak occurs in toroidal geometry,⁹ rather than slab geometry. In toroidal geometry, particles interacting with one wave trace a line in $\varepsilon - \mu - P_\phi$ space, where $\mu = mv_\perp^2/2B$ is the magnetic moment; $\varepsilon = \mu B + mv_\parallel^2/2$ is the kinetic energy; and $P_\phi = R(mB_\phi v_\parallel/B - qA_\phi)$ is the canonical angular momentum, where A_ϕ is the vector potential.

There remain a number of challenges to realize or to assess the utility of this coupling, some of which have been solved in part and some of which have not yet been solved.

II.A. Experimental Challenges

First, it remains to identify the wave that can accomplish the favorable diffusion path. In fact, it may be advantageous to employ more than one wave,¹⁰ which relaxes the hard constraint imposed by one diffusion path but accomplishes the channeling effect nonetheless. For tokamaks, the mode-converted ion Bernstein wave seems to be most appropriate for extracting most of the alpha-particle energy.¹¹ This wave grows at the expense of the alpha-particle energy and then, in a deuterium-tritium plasma, damps on the tritium fuel ions at the tritium resonance.¹²

To what extent has this identification been successful experimentally? On the one hand, there have been no

experiments to date that show that the mode-converted ion Bernstein wave extracts energy from alpha particles. On the other hand, these are hard experiments because to demonstrate such an effect, neither have enough alpha particles been produced nor have appropriate wave parameters been arranged. What has been demonstrated is that mode-converted ion Bernstein waves could produce diffusion paths in energy-position space. These experiments also confirmed the predicted wave characteristics of the mode-converted ion Bernstein wave, including a flip in the parallel wave number.¹³

However, the cooling effect could not be tested. With few fusion-produced alpha particles, the wave parameters were chosen so that the diffusion paths for 80-keV beams of deuterium ions connected cold in the center with hot on the periphery. These deuterium ions were then detected at 2.2 MeV at the periphery.¹⁴ This was, of course, not a cooling effect, but it did show that, in principle, the diffusion paths could operate as expected. Interestingly, the experimentally measured diffusion coefficient was a factor of 50 higher than expected. One possible explanation was that a high-Q cavity mode was excited by the mode-converted ion Bernstein wave,¹⁵ an explanation supported recently when related internal modes were observed on NSTX (Ref. 16).

However, while there is now evidence for the diffusion paths and the ion Bernstein wave characteristics, the experimental investigations demonstrate neither the cooling effect nor, given the cooling effect, how the energy that flows into the waves might then be used to accomplish the hot-ion mode or to generate current. Note that, if the current drive mechanism is via electrons, particularly through the tail electrons, say, by lower hybrid waves³ or electron cyclotron waves,¹⁷ then the electron tail heating may act in opposition to the hot-ion mode. But, even for electron-based methods that operate on the bulk electrons,^{18,19} the electron heating will make more challenging the diversion of substantive power to the ions. Thus, it will be a challenge to employ an ion-based method, such as minority species current drive.^{20,21}

Thus, while much has been accomplished experimentally, the following experimental challenges remain with respect to alpha channeling:

1. to experimentally demonstrate directly the alpha-cooling effect, with diffusion paths connecting hot in the center to cold on the periphery
2. to confirm and explain the anomalously large diffusion coefficient observed on TFTR
3. to demonstrate that waves amplified by the alpha particles can then be ultimately damped on ions, thus producing the hot-ion mode
4. to demonstrate that waves amplified by the alpha particles can then be ultimately damped on ions or electrons, producing a current drive effect.

II.B. Modeling Challenges

Second, in addition to the experimental challenges, there are theoretical or modeling challenges. In a reactor, simulations show that with two waves, it would be possible to divert more than half of the alpha-particle energy through waves.⁹ Moreover, under alpha channeling, as opposed to in the absence of alpha channeling, there is an optimal heat loss rate of energy from the electrons, which is actually finite.²² This is because large electron heat loss helps to make the electrons cooler relative to the ions, so that more pressure is available to the ions at the same confined plasma pressure.

The alpha-channeling effect remains very speculative, but it has such high upside potential that it merits serious investigation. The upside potential includes ~30% cheaper cost of electricity, compared to aggressively designed reactors (because of the increased reactivity at a given confined pressure and free current drive if the channeling can drive a current drive effect). Also, because of the transport associated with the channeling effect, there is automatic impurity removal and plasma fueling. The alpha-channeling effect may turn out to be even more important if ion heat transport is eventually tamed, but electron heat transport is not tamed, since it allows electron heat transport to be high. Also, since the top tokamak confinement and heating results were obtained under hot-ion mode operation using auxiliary heating, the present database of experiments is actually more supportive of the hot-ion mode than the equal-temperature mode, so extrapolation to a reactor of the tokamak transport properties can be made with greater confidence.

What has not been explored in any detail is how these advantages can be made more precise and optimized in reactor scenarios. Thus, in addition to supporting the experimental challenges in alpha channeling, there are a number of modeling challenges:

1. to assess how great can be the advantages in a practical reactor configuration
2. to assess how the alpha channeling could be accomplished synergistically with other advantageous effects, such as current recharge (see Sec. III).

III. TRANSFORMER RECHARGING

There may be significant advantages in terms of current drive efficiency to relax the constraint of *fully* steady-state operation. It turns out that there are significant cost savings if plasma parameters are allowed to vary, with the current almost steady state.

To see this, consider Fig. 3, which depicts a time during which the current is generated by, for example, LHCD. Thus, the lower hybrid power P_{LH} persists only

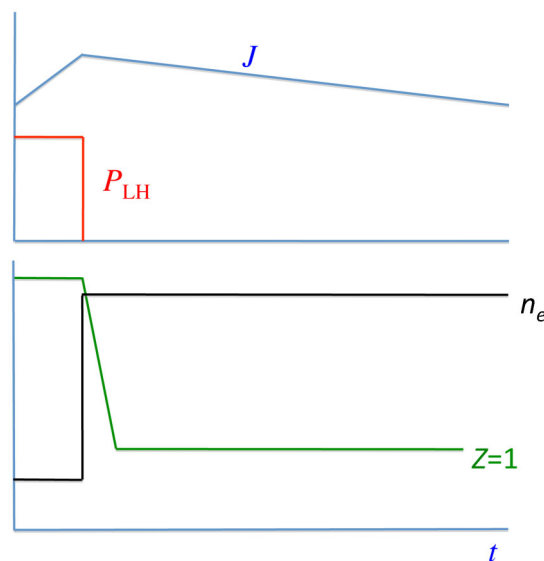


Fig. 3. Oscillating rf-driven current. The current-generation stage lasts while $P_{LH} \neq 0$. The current-relaxation stage lasts while $dJ/dt < 0$. When J relaxes to its pregeneration-stage value, the cycle of generation and relaxation repeats.

for duration T_g . This raises the current from its minimum value J_{min} to its maximum value J_{max} . The generation stage is followed by a relaxation stage of duration T_r , in which there is no current drive, so that the current decays in an L/R time, where L is the tokamak plasma inductance and R is the resistivity. The generation and relaxation cycles then repeat.

The advantage of doing so is that, even as the current does not deviate much from its average value, the average power dissipation can be much less if the plasma parameters in the two stages are different. In principle, these other parameters can be changed on timescales short compared to the L/R time, because the particle and heat confinement times are of the order of 1 s in a tokamak reactor, whereas the L/R time is approximately three orders of magnitude longer. Thus, the plasma can be exchanged many times, so that the current can be considered constant on the timescale on which the resistivity and other parameters might change; Fig. 3 depicts variations in the effective ion charge state Z and in the electron density n_e .

To see why the average power dissipation can be much less if the plasma parameters in the two stages are different, consider the role of the induced toroidal electric field. When the current is increasing, there is an induced electric field that opposes the increase in the current. However, this induced electric field operates mainly on the bulk thermal electrons rather than on the small number of electrons induced by the rf waves to carry the current. Hence, there would be an advantage to limiting this

opposing current by having higher plasma resistivity. On the other hand, during the current decay stage, the induced electric field supports the rf-generated current so as to prevent too rapid a decay of this current. Again, this current acts primarily on the thermal electrons, which are many. Therefore, there would be an advantage to assisting this supporting current by having lower plasma resistivity during this stage.

Suppose then that the L/R time in the generation stage is τ_g (with $\tau_g \gg T_g$) and that in the relaxation stage it is τ_r (with $\tau_r \gg T_r$). The maximum current density J_{max} occurs at the end of the generation stage, and the minimum current density J_{min} occurs at the end of the relaxation stage. The ratio of relaxation times that gives periodicity is $T_r/T_g = (J_{rf}\tau_r)/(J_{max}\tau_g)$. Since $\tau_g \gg T_g$, the plasma current itself can nonetheless remain approximately constant, so we also have that $J_{max} \simeq J_{min} \simeq J_0$, where J_0 is the average current density.

The rf-driven current density J_{rf} is defined as the plasma current were the rf on long enough for the current to achieve steady state. For simplicity, assume that the rf-driven current density J_{rf} is large compared to the actual current, so that $J_{rf} \gg J_0$. The average power dissipated can then be put in the form¹

$$\left\langle \frac{J}{P_d} \right\rangle_{avr} = \left\langle \frac{J}{P_d} \right\rangle_g \frac{\tau_r}{\tau_g}. \quad (2)$$

From Eq. (2), it can be seen that there are two ways to achieve large average efficiency.

First, one can maximize the ratio τ_r/τ_g . This produces the asymmetric resistivity effect, where the resistivity is high when the inductive field acts in opposition to the current. The high resistivity can be achieved by injecting higher ion charge state ions during the current-generation stage or by having lower electron temperature during this stage. There is opportunity to impede through higher ion charge state the ohmic counter-current, while not affecting to as large an extent the rf current, because of the relative insensitivity of the rf current drive efficiency to the ion charge state Z_i , at least for nearly pure hydrogen plasma. The ohmic resistivity is proportional to Z_i , whereas the rf current drive efficiency for methods that rely upon fast electrons [such as LHCD and electron cyclotron current drive (ECCD)] is proportional to $1/(5+Z_i)$ (Ref. 2). Thus, absent other effects, the higher ion charge state does more to impede the ohmic current than the rf current, as long as the charge state is less than ~ 6 .

Second, to achieve large efficiency, one can maximize $\langle J/P_d \rangle_g$. Here, we note that the rf current drive efficiency, as opposed to the ohmic efficiency, is inversely proportional to the density. Thus, it would be advantageous to capture the high rf current drive efficiency in the low-density current-generation stage. However, since the L/R time is independent of density, whereas the fusion

power varies as density squared, it makes sense to use a high-density current-relaxation stage.

These terms are multiplicative, so, by having both high resistivity and low density in the generation stage, the average current drive power can be very much reduced, perhaps by a factor of 10. The peak rf power is larger than the average rf power, but since the efficiency is calculated during the generation stage, where the density may be low, rather than the relaxation stage, this power may be less than the power that might have been needed were the power provided continuously with relaxation-stage parameters. Also, the lack of penetration of high-density plasma by lower hybrid waves reduces the parameter space in which the concepts can be applied (see, for example, Refs. 23 and 24). Thus, the LHCD effect is not only more efficient in low-density plasma, but it may better penetrate to the tokamak center, making the tokamak recharge approach even more attractive.

There is considerable experimental and computational evidence supporting at least the underlying physics. At first, the high efficiency of the conversion of rf energy to magnetic field energy was puzzling, but it was explained by the phenomenon of tokamak recharging.²⁵ The first and most detailed studies of this phenomenon were conducted on PLT tokamak,²⁶ which produced such a detailed fit to the theory that the underlying equations can now be relied upon. The transformer recharging with LHCD has since been confirmed on many other tokamaks as well,²⁷⁻²⁹ including in a recent campaign of experiments on the tokamak EAST (Refs. 30, 31, and 32).

Note that transformer recharging can be implemented synergistically with alpha channeling.³³ The synergy occurs because, in the generation phase of the current, the plasma is at low density and therefore low reactivity. Because it is at low density, the electron and ion temperatures equilibrate more slowly, so it is easier to produce the hot-ion mode. Moreover, since, in the hot-ion mode, the fusion reactivity is greater, the fusion power production can be made more uniform in the generation and relaxation stages. The uniformity of the heat load reduces the thermal fatigue.

One of the issues in realizing the full benefit in transformer recharging is that it is not so easy to make the resistivity larger during the current-generation stage. Indeed, while much of the benefit might be realized by making the density lower, even more might be realized by simultaneously making the resistivity larger. However, there are effects that impede the increase in resistivity.

First, if fast electron-based methods are employed, such as LHCD or ECCD, then there is an additional or so-called *hot conductivity*, which is proportional to the rf power.³⁴ This added conductivity can be substantial. It can be avoided by employing ion-based methods of current drive, but those methods are less efficient and less sure, whereas LHCD in low-density plasma is a well-established technique.

Second, if the resistivity increase is sought through a higher ion charge state, then there tend to be effects that make the electrons hotter as well, thereby reducing the effective resistivity. Also, when the current is ramped-up quickly, there is production of the so-called *backward-runaways*, or electrons that reach high energy but carry current opposite to the desired current.³⁵

One speculative method to limit the electron conductivity might be through a transport mechanism that operates solely on the energetic backward-going electrons, possibly through some resonant mechanism, such as the stochastic instability suggested as being responsible for restraining energy in runaway electrons.³⁶ The idea is to remove electrons contributing to the high conductivity.

Another, even more speculative, method might to induce anomalous resistivity in the current-generation stage. Here, by an anomalous resistivity, we mean an effect that we have no idea really how to produce. But, it might involve, for example, magnetic turbulence. How would that work? Consider Fig. 4, which schematically, with utter speculation, shows how magnetic turbulence might increase the plasma resistivity. The idea is to lengthen field lines connecting an anode to a cathode; the upper frame depicts lower resistivity in the absence of turbulence. If electrons were confined to the field lines, then they would encounter more collisions on the way to the anode in the presence of the turbulence (lower frame), so that the resistivity would be higher. While this is not put forward really as a serious suggestion, it does indicate that if hyperresistivity or anomalously high resistivity could be made to occur through collisionless means, it would have high upside potential. Thus, a question of interest is the following: What power is necessary to tangle the magnetic field to make the resistivity higher? An alternative challenge is to prove that no power will do that.

Thus, the following unsolved challenges can be posed:

1. to assess how great can the advantages be in a practical reactor configuration, including the synergy with the alpha-channeling effect
2. to demonstrate experimentally the increase in average current drive efficiency

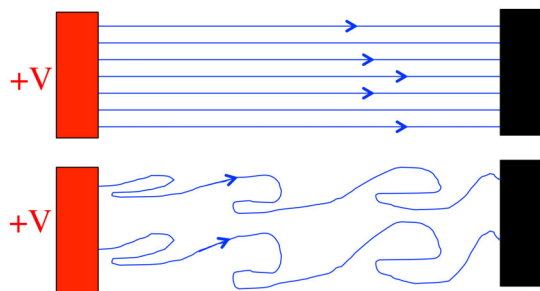


Fig. 4. A speculative picture for hyperresistivity.

3. to increase effectively the resistivity during the current-generation stage. A first goal would be to demonstrate the quickest possible ramp-up rate of the current.
4. to speculate how the resistivity might be increased through new methods, such as through inducing resonant transport or through generating turbulent fluctuations, or to prove that there are limits or costs in increasing the resistivity by these speculative methods.

IV. FREE ENERGY UNDER DIFFUSION

The following challenge is just for fun.

In the case of alpha channeling, energy is released from the alpha particles through diffusive effects only. The idea is that there are more particles in high energy states than there are in low energy states, so if only a diffusion path can be made to connect these two states, then energy will be released. We now pose the following question³⁷: Suppose that any desired rf diffusion path can be constructed. In such an ideal case, but under the constraint that the rf waves can only diffuse particles along these paths, how much energy exactly can be released from the plasma?

Note that this question is not the same as the free energy of a plasma subject to phase-space conservation. It is rather the question of the free energy in plasma under the diffusion constraint, which is not phase-space conserving, but certainly the case of more practical interest. After all, it is indeed the tendency of waves in plasma to be incoherent, and thus to diffuse particles.

To appreciate the question, consider the classic bump-on-tail instability, as depicted in Fig. 5. This classic instability involves initially a *bump* on the tail of the electron distribution function f . However, through diffusion by electrostatic plasma waves, the bump relaxes time-asymptotically to a plateau in velocity space. The energy released from the particles flows into the waves. Thus, electrons diffuse in velocity space under the

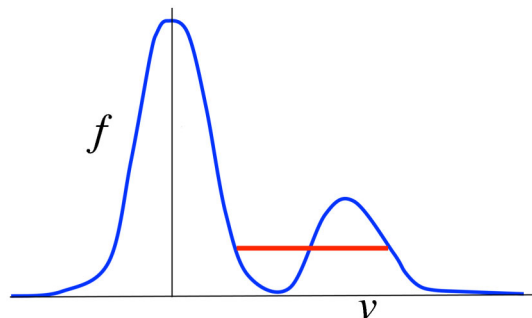


Fig. 5. The bump-on-tail instability and its saturated state.

influence of the unstable waves, until the distribution function is monotonically decreasing in energy.

How much energy is released? One can see, by construction, the well-known result that the maximum amount of energy release is when the plateau is at the unique height such that particles are conserved and the distribution function is monotonically decreasing. Where the plateau is lower than the initial distribution function, particles are removed, and transferred to fill in where the plateau is higher than the initial distribution function. If the plateau were drawn too high, then there would be too few energetic particles to remove and too large a space at lower energy to fill. If the plateau were drawn too low, then there would be too many particles to remove and too small a space to fit these particles. Thus, there is a unique height for the plateau.

However, the situation is considerably more complicated when there are two or more bumps on the tail of the electron distribution function. For example, consider the electron velocity distribution depicted in Fig. 6. Now it is not so clear how to draw the plateau. In fact, there will be more than one plateau, as one bump can be stabilized and then the other. In such a case, there are multiple time-asymptotic plateau solutions, and the amount of energy released differs with different plateau solutions.

The problem thus stated is actually more general than just the phenomenon of waves interacting with particles in plasma. For example, consider a ground-state atom and two excited states, as depicted schematically in Fig. 7. Suppose that laser frequency ν_{1g} can excite transitions between the ground state and the first excited state. Suppose further that the laser is incoherent, or its duration is hard to control. Then, the effect of this laser will be to equalize the population levels between the ground state g and the excited state 1. Similarly, other lasers can be tuned to equalize the population levels between other pairs of states. Note that the sequence in which these pulses are applied releases different amounts of energy.

In other words, for an inverted population in energy, it might be advantageous first to release energy by equalizing the populations in the two excited states or by first equalizing the populations between one of the excited

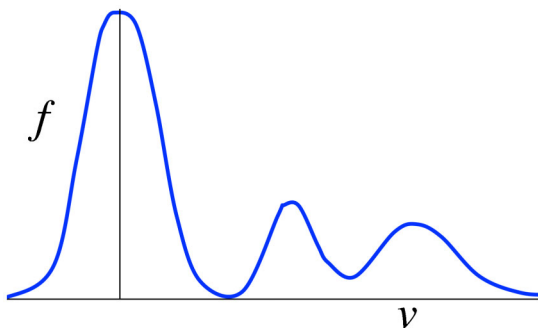


Fig. 6. Two-bump-on-tail velocity distribution function.

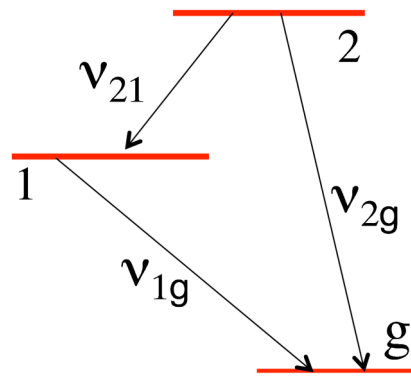


Fig. 7. Energy transfer between excited states 1 and 2 and ground state g via laser excitations ν_{ij} .

states and the ground state. A second laser pulse can release more energy by equalizing populations in two of the resulting states. This can be repeated until the population is no longer inverted. However, there are multiple final states that are not inverted.

Thus, we pose the following questions. First, and most fundamental, what sequence of pulses extracts the maximum amount of energy from atoms given an initially inverted energy population? Second, what is the complexity of this problem? In other words, as the number of distinct energy states grows, how does the time to calculate the optimum sequence of laser pulses grow?

V. CONCLUSIONS

The sampling of unsolved challenges presented here is by no means complete, but the challenges of how to accomplish alpha channeling, how to exploit quasi-steady-state operation, including how to combine with alpha channeling, offer significant upside potential to tokamak operation. The free energy available under the diffusion constraint is just a brain teaser; in practice, the full energy would never be released, but it is interesting to know in principle how much there is.

Among the other open areas, not touched upon here, is the maximum efficiency attainable under completely steady-state operation, which may involve manipulations in phase-space more complex than the ones attempted here.^{38,39} However, the upside potential in efficiency of these approaches is expected to be limited compared to approaches seeking quasi-steady-state operation. Another open area is to explain the plasma rotation in tokamaks that is apparently related to the LHCD (Ref. 40). Yet another area, not particularly tied to the tokamak, is the identification of current drive efficiencies in strange regimes, such as partially ionized plasma⁴¹ or Fermi degenerate plasma.⁴² Each of these areas is relatively unexplored and has the potential to surprise.

ACKNOWLEDGMENTS

This work was supported by the U.S. Department of Energy under contract DE-AC0209CH11466. The author acknowledges the hospitality of the Weizmann Institute of Science, where he held a Weston Visiting Professorship during much of the time over which this manuscript was prepared.

REFERENCES

1. N. J. FISCH, "Theory of Current Drive in Plasmas," *Rev. Mod. Phys.*, **59**, 1, 175 (1987); <http://dx.doi.org/10.1103/RevModPhys.59.175>.
2. N. J. FISCH, "Methods of Radio-Frequency Current Drive," *Fusion Sci. Technol.*, **65**, 1 (2014); <http://dx.doi.org/10.13182/FST13-670>.
3. N. J. FISCH, "Confining a Tokamak Plasma with rf-Driven Currents," *Phys. Rev. Lett.*, **41**, 13, 873 (1978); <http://dx.doi.org/10.1103/PhysRevLett.41.873>.
4. N. J. FISCH and J. M. RAX, "Interaction of Energetic Alpha-Particles with Intense Lower Hybrid Waves," *Phys. Rev. Lett.*, **69**, 612 (1992); <http://dx.doi.org/10.1103/PhysRevLett.69.612>.
5. J. M. RAX, "Physics of Landau Resonances, Cyclotron Resonances, and Current Generation," *Fusion Sci. Technol.*, **65**, 10 (2014); <http://dx.doi.org/10.13182/FST13-634>.
6. K. L. WONG and M. ONO, "Effects of Ion-Cyclotron Harmonic Damping on Current Drive in the Lower Hybrid Frequency-Range," *Nucl. Fusion*, **24**, 5, 615 (1984); <http://dx.doi.org/10.1088/0029-5515/24/5/008>.
7. E. BARBATO and F. SANTINI, "Quasi-Linear Absorption of Lower Hybrid Waves by Fusion Generated Alpha-Particles," *Nucl. Fusion*, **31**, 4, 673 (1991); <http://dx.doi.org/10.1088/0029-5515/31/4/005>.
8. N. J. FISCH and J. M. RAX, "Current Drive by Lower Hybrid Waves in the Presence of Energetic Alpha-Particles," *Nucl. Fusion*, **32**, 4, 549 (1992); <http://dx.doi.org/10.1088/0029-5515/32/4/102>.
9. M. C. HERRMANN and N. J. FISCH, "Cooling Energetic α Particles in a Tokamak with Waves," *Phys. Rev. Lett.*, **79**, 1495 (1997); <http://dx.doi.org/10.1103/PhysRevLett.79.1495>.
10. N. J. FISCH and M. C. HERRMANN, "Alpha Power Channelling with Two Waves," *Nucl. Fusion*, **35**, 12, 1753 (1995); <http://dx.doi.org/10.1088/0029-5515/35/12/140>.
11. N. J. FISCH, "Alpha Power Channeling using Ion-Bernstein Waves," *Phys. Plasmas*, **2**, 2375 (1995); <http://dx.doi.org/10.1063/1.871454>.
12. E. J. VALEO and N. J. FISCH, "Excitation of Large- k_{θ} Ion-Bernstein Waves in Tokamaks," *Phys. Rev. Lett.*, **73**, 3536 (1994); <http://dx.doi.org/10.1103/PhysRevLett.73.3536>.
13. N. J. FISCH and M. C. HERRMANN, "A Tutorial on Alpha-Channelling," *Plasma Phys. Controlled Fusion*, **41**, 3A, A221 (1999); <http://dx.doi.org/10.1088/0741-3335/41/3A/015>.
14. N. J. FISCH, "Physics of Alpha Channelling and Related TFTR Experiments," *Nucl. Fusion*, **40**, 6, 1095 (2000); <http://dx.doi.org/10.1088/0029-5515/40/6/307>.
15. D. S. CLARK and N. J. FISCH, "The Possibility of High Amplitude Driven Contained Modes During Ion Bernstein Wave Experiments in the Tokamak Fusion Test Reactor," *Phys. Plasmas*, **7**, 2923 (2000); <http://dx.doi.org/10.1063/1.874143>.
16. N. N. GORELENKOV, N. J. FISCH, and E. FREDRICKSON, "On the Anomalous Fast Ion Energy Diffusion in Toroidal Plasmas Due to Cavity Modes," *Plasma Phys. Control. Fusion*, **52**, 055014 (2010); <http://dx.doi.org/10.1088/0741-3335/52/5/055014>.
17. N. J. FISCH and A. H. BOOZER, "Creating an Anisotropic Plasma Resistivity with Waves," *Phys. Rev. Lett.*, **45**, 720 (1980); <http://dx.doi.org/10.1103/PhysRevLett.45.720>.
18. D. J. H. WORT, "Peristaltic Tokamak," *Plasma Phys.*, **13**, 258 (1971); <http://dx.doi.org/10.1088/0032-1028/13/3/008>.
19. N. J. FISCH and C. F. F. KARNEY, "Current Generation with Low-Frequency Waves," *Phys. Fluids*, **24**, 27 (1981); <http://dx.doi.org/10.1063/1.863243>.
20. N. J. FISCH, "Current Generation by Minority-Species Heating," *Nucl. Fusion*, **21**, 15 (1981); <http://dx.doi.org/10.1088/0029-5515/21/1/002>.
21. M. J. MANTSINEN et al, "Application of ICRF Waves in Tokamaks Beyond Heating," *Plasma Phys. Control. Fusion*, **45**, A445 (2003); <http://dx.doi.org/10.1088/0741-3335/45/12A/028>.
22. N. J. FISCH and M. C. HERRMANN, "Utility of Extracting Power from Alpha Particles by Waves," *Nucl. Fusion*, **34**, 1541 (1994); <http://dx.doi.org/10.1088/0029-5515/34/12/101>.
23. A. A. TUCILLO et al, "Progress in LHCD: A Tool for Advanced Regimes on ITER," *Plasma Phys. Control. Fusion*, **47**, B363 (2005); <http://dx.doi.org/10.1088/0741-3335/47/12B/S26>.
24. C. GORMEZANO et al, "Chapter 6: Steady State Operation," *Nucl. Fusion*, **47**, S285 (2007); <http://dx.doi.org/10.1088/0029-5515/47/6/S06>.
25. N. J. FISCH and C. F. F. KARNEY, "Conversion of Wave Energy to Magnetic Field Energy in a Plasma Torus," *Phys. Rev. Lett.*, **54**, 897 (1985); <http://dx.doi.org/10.1103/PhysRevLett.54.897>.
26. C. F. F. KARNEY, F. C. JOBES, and N. J. FISCH, "Comparison of the Theory and the Practice of Lower Hybrid Current Drive," *Phys. Rev. A*, **32**, 2554 (1985); <http://dx.doi.org/10.1103/PhysRevA.32.2554>.
27. F. LEUTERER et al, "Recharging of the Ohmic-Heating Transformer by Means of Lower-Hybrid Current Drive in the ASDEX Tokamak," *Phys. Rev. Lett.*, **55**, 75 (1985); <http://dx.doi.org/10.1103/PhysRevLett.55.75>.
28. Y. TAKASE et al, "Plasma Current Ramp-Up and Ohmic-Heating Transformer Recharging Experiments Using Lower-Hybrid Waves on a Tokamak," *Phys. Fluids*, **30**, 1169 (1987); <http://dx.doi.org/10.1063/1.866316>.

29. Z. Y. CHEN et al, "Ohmic Radio-Frequency Synergy Current Drive and Transformer Recharging Experiments in the HT-7 Tokamak," *Chin. Phys. Lett.*, **22**, 1721 (2005); <http://dx.doi.org/10.1088/0256-307X/22/7/044>.
30. B. J. DING et al, "Current Ramp-Up with LHCD in EAST," *Phys. Plasmas*, **19**, 122507 (2012); <http://dx.doi.org/10.1063/1.4773049>.
31. H. W. LU et al, "Investigation of Runaway Electrons in the Current Ramp-up by a Fully Non-Inductive Lower Hybrid Current Drive on the EAST Tokamak," *Phys. Scr.*, **87**, 055504 (2013); <http://dx.doi.org/10.1088/0031-8949/87/05/055504>.
32. M. LI et al, "Investigation of LHCD Efficiency and Transformer Recharging in the EAST Tokamak," *Plasma Sci. Technol.*, **14**, 201 (2012); <http://dx.doi.org/10.1088/1009-0630/14/3/04>.
33. N. J. FISCH, "Transformer Recharging with Alpha Channeling in Tokamaks," *J. Plasma Phys.*, **76**, 3-4, 627 (2010); <http://dx.doi.org/10.1017/S0022377809990857>.
34. N. J. FISCH, "Conductivity of RF-Heated Plasma," *Phys. Fluids*, **28**, 245 (1985); <http://dx.doi.org/10.1063/1.865186>.
35. C. F. F. KARNEY and N. J. FISCH, "Current in Wave Driven Plasmas," *Phys. Fluids*, **29**, 180 (1986); <http://dx.doi.org/10.1063/1.865975>.
36. L. LAURENT and J. M. RAX, "Stochastic-Instability of Runaway Electrons in Tokamaks," *Europhys. Lett.*, **11**, 219 (1990); <http://dx.doi.org/10.1209/0295-5075/11/3/006>.
37. N. J. FISCH and J. M. RAX, "Free-Energy in Plasmas Under Wave-Induced Diffusion," *Phys. Fluids B*, **5**, 1754 (1993).
38. N. J. FISCH, J. M. RAX, and I. Y. DODIN, "Current Drive in a Ponderomotive Potential with Sign Reversal," *Phys. Rev. Lett.*, **91**, 205004 (2003); <http://dx.doi.org/10.1103/PhysRevLett.91.205004>.
39. I. Y. DODIN, N. J. FISCH, and J. M. RAX, "Ponderomotive Barrier as a Maxwell Demon," *Phys. Plasma*, **11**, 5046 (2004); <http://dx.doi.org/10.1063/1.1787771>.
40. A. INCE-CUSHMAN et al, "Observation of Self-Generated Flows in Tokamak Plasmas with Lower-Hybrid-Driven Current," *Phys. Rev. Lett.*, **102**, 035002 (2009); <http://dx.doi.org/10.1103/PhysRevLett.102.035002>.
41. D. FARINA, M. LONTANO, and R. POZZOLI, "Radiofrequency Current Drive in a Weakly Ionized Plasma," *Nucl. Fusion*, **27**, 155 (1987); <http://dx.doi.org/10.1088/0029-5515/27/1/015>.
42. S. SON and N. J. FISCH, "Current Drive Efficiency in Degenerate Plasma," *Phys. Rev. Lett.*, **95**, 225002 (2005); <http://dx.doi.org/10.1103/PhysRevLett.95.225002>.