

Ion Acceleration with Beating Electrostatic Waves: *On-resonance and Off-resonance cases*

1. Introduction

In the 70's, an interesting phenomenon was observed in the upper ionosphere: O^+ and H^+ ions were accelerated to very high velocities—their velocities increasing by a factor of 30 (from around 0.3 eV to about 10 eV). Papers such as [1] described how this process may be explained as a result of interaction with multiple electrostatic waves which propagate along the geomagnetic field. Subsequent research in this area revealed that this process is most effective when the electrostatic waves differ in frequency by an integer times the cyclotron frequency. In other words: if we have two beating waves with frequencies ω_1 and ω_2 , they should satisfy the equation $\omega_2 - \omega_1 = n\omega_c$ [2]. The advantages of using two beating waves are (a) that their frequencies do not have to be close to the ion's thermal velocity, and (b) the ions can start off with arbitrarily low initial velocities [3]. This is in contrast to acceleration by a single wave, where the ions have to be at a certain velocity already in order to be subject to the acceleration effects of the electrostatic waves (see p3). Spektor and Choueiri [4] then determined that there were additional criteria for this effect to occur and determined the necessary and sufficient conditions required of the particles' trajectories in phase space. I studied Slava Spektor's review of the relevant research in [3, 4, 5] and then proceeded to explore the difference between on-resonance cases (where the normalized frequencies of the electrostatic waves are not integers – i.e. $\nu_1 = \frac{\omega_1}{\omega_c}, \nu_2 = \frac{\omega_2}{\omega_c} \in \mathbb{Z}$) and the off-resonance cases (where $\nu_1, \nu_2 \notin \mathbb{Z}$).

2. Mathematical prelude

The mathematical framework for our system is quite simple. Our system will consist of electrostatic waves propagating perpendicular to a uniform magnetic field. Without loss of generality, we will assume that the magnetic field is in the z direction and the electrostatic waves are in the x direction (Fig. 1), with our waves having amplitudes E_i , wave numbers k_i and frequencies ω_i . Our particle will be orbiting perpendicular to the magnetic field in the xy plane, with its Larmor radius ρ , and its phase angle θ .

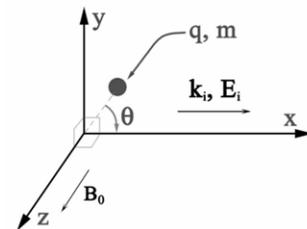


FIG. 1. Our system of a particle rotating in a magnetic field, interacting with electrostatic waves.

The only force that does work on the particle is the electrostatic waves, so our equation of motion will be

$$[F =]ma = qE, \quad \text{or} \quad \frac{d^2x}{dt^2} = \frac{q}{m} \sum_i E_i \sin(k_i x - \omega_i t + \phi_i).$$

From the equation of motion, some lengthy algebra manipulations and Legendre transformations enable us to derive

$$\text{the Hamiltonian, given by } H = \frac{\rho^2}{2} + \sum_i \frac{\varepsilon_i}{\kappa_i} \cos(\kappa_i \rho \sin \theta - \nu_i \tau + \phi_i) \quad [3]$$

Note that the value of ρ represents the Larmor radius, and is a direct measurement of the particles thermal velocity, since particles with higher speeds have larger orbits.

The dimensionless parameters in the Hamiltonian are defined as follows:

$\varepsilon_i = (k_i q E_i) / (m \omega_c^2)$ are the amplitudes of the electrostatic waves divided by some constants of the system to get a dimensionless measure of the wave amplitudes.

$\nu_i = \omega_i / \omega_c$ are the normalized frequencies of the waves. In terms of these normalized frequencies, the beating criterion may be expressed as $\nu_1 - \nu_2 = n$, where n is an integer. When referring to ‘‘on-resonance’’ cases, we mean that the normalized frequencies themselves are integers too, and not just the difference between them. The ‘‘off-resonance’’ case covers all cases where ν_i are not integers.

$\kappa_i = k_i / k_1$: so κ_i denotes the ratio of the i th wave number to the first. In my simulations, I assumed all the wave numbers were the same, so all κ_i are equal to 1.

3. Integration Algorithm

To perform the actual numerical integration, I used Candy & Rozmus’ symplectic integration algorithm for separable Hamiltonians [6], as it has a number of advantages in our case over the more popular Runge Kutta algorithm. As the name suggests, however, the algorithm requires that the Hamiltonian to be separable, and ours is not (there is a ‘ $\rho \sin \theta$ ’ term) and so a change of variables from $H(\rho, \theta)$ to Hamiltonian in a Cartesian xy frame of reference – $H(p, q)$ – was necessary to use this algorithm. The separable Hamiltonian was thus:

$$H = \frac{p^2}{2} + \frac{q^2}{2} + \sum_i \frac{\varepsilon_i}{\kappa_i} \cos(\kappa_i q - \nu_i \tau + \phi_i)$$

which is clearly separable in p and q .

When plotting our results, a convenient framework to use is a Poincaré surface of section, which effectively involves taking ‘snapshots’ of a particle’s position in phase space (we will use ρ - θ space) and accumulating its values of ρ vs. θ at regular intervals—the length of the intervals here being an integer multiple of the cyclotron period. The Poincaré section is a standard way of getting a good idea of the nature of the motion. Trajectories of particles undergoing regular motion will typically be neat and regular, whereas particles undergoing chaotic motion will leave a random, noisy signature on the Poincaré section.

4. Results for a single wave

In order to appreciate the advantages of using two beating waves, we should look briefly at the results for a *single* wave. In the plots below (Fig. 2), 60 particles were placed at random positions in phase space and their trajectories were traced over time. No collisions or interactions between particles were taken into account. As is evident from these plots, in both off-resonance (Fig. 2a) and on-resonance (Fig. 2b) cases there is a threshold at around $\rho = 20$,

below which particles do not undergo any acceleration, but rather drift along the Poincaré surface, maintaining the same value for ρ (some particles undergoing this type of motion have been colored different shades of green). Only once above this threshold do they undergo acceleration stochastically: if a particle starts off at, say, $\rho = 20$, it will

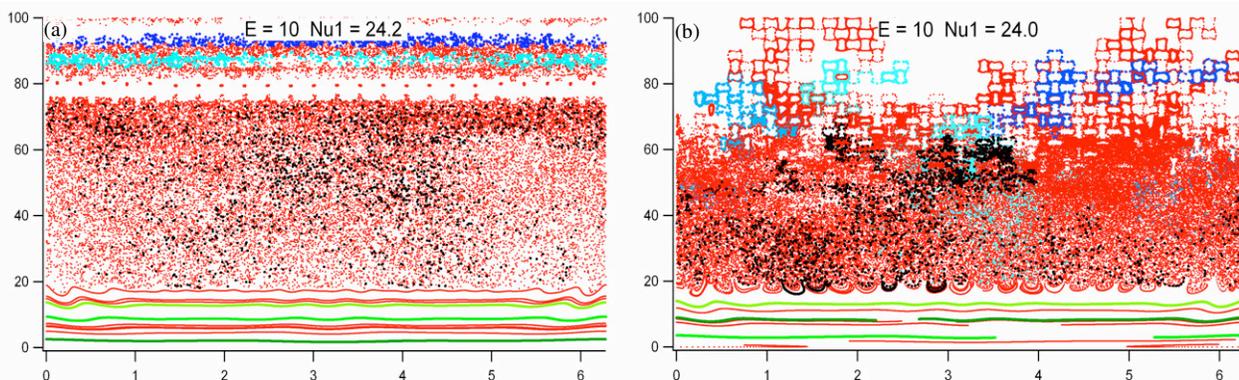


FIG. 2. Single wave results for the off-resonance (a) and on-resonance (b) cases. Between around $\rho = 20$ and $\rho = 75$ in both plots particles undergo chaotic motion and typically attain an average of about $\rho = 50$. In both plots, the trajectory of one particle in this region has been colored black to demonstrate its random nature. Below this region, the particles undergo regular motion. Above this region, particles follow a slightly chaotic pattern, but their average velocity does not change.

jump around the stochastic domain at random and attain an average ρ of around 50. The black points on each of the plots are the signature of one particle jumping around the stochastic domain. Here is where we see a slight difference between the two plots. Whereas the stochastic domain in the off-resonance case appears totally random, the motion of particles in the on-resonance case tends to remain in certain regions of the stochastic domain.

Above a certain value of ρ , though (around $\rho = 80$), the stochastic acceleration effects diminish rapidly. In the off-resonance case, they follow complex but regular stochastic webs, while in the on-resonance case the particles trace out squares in phase space. (some of these particles have been colored blue so their trajectories can be made out). In both cases there is no net gain in ρ , and thus no net increase in velocity.¹

5. Two waves: Off-resonance ($\nu \notin \mathbf{Z}$).

As we can immediately see from the Poincaré section on the right (Fig. 3), the ‘regular motion domain’, (which, in the single-wave case consisted of straight lines) now consists of trajectories which mostly move upwards towards the hyperbolic point (labeled H), and from there they escape to the stochastic domain, where they undergo acceleration chaotically. However, as Slava noted [3], not all of the particles make it to the stochastic domain: there is a trapped region in which particles never escape to the stochastic domain. The trapped region is defined by all those particles that are between the elliptical point (labeled E, and located at $\rho = \nu/2$) and the hyperbolic point (located at $\rho = \nu - \sqrt{\epsilon}$) when $\theta = \pi$. The

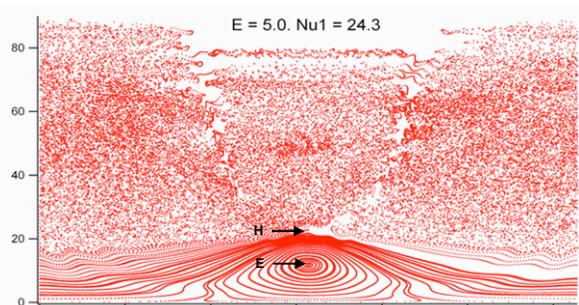


FIG. 3. A typical off-resonance plot. The normalized wave amplitude, ϵ , is 5.0, while ν_1 , the frequency of the lower-frequency wave is 24.3. The trapped region is centralized around an elliptical point (labeled E) and the stochastic domain begins from the vicinity of the hyperbolic point (labeled H)

¹ Note that in the on-resonance case, the ‘blue’ region is more more continuous with the stochastic domain than in the off-resonance case, and compare this with the beating-waves later.

particles in this trapped region orbit around the elliptical point and never escape to the stochastic domain. All particles not in this trapped region eventually make it to the stochastic domain. It is possible that a particle, moving chaotically in phase space may drop back into the regular motion domain. If it falls into the trapped region, then it will remain there, but if it does not, it will eventually escape back into the stochastic domain.

One or two hyperbolic points?

My plots also initially seemed to confirm that there was a single hyperbolic point located at $\nu - \sqrt{\epsilon}$. However, a closer and more careful inspection of the region around the hyperbolic point (involving tracing trajectories of individual particles – Fig. 4) revealed that there seem to be *two* hyperbolic points from which the particles enter the stochastic domain, one on each side of the central ‘hump’. All the particles that I traced that left the regular-motion domain only did so at these two points, and never from the center. This is rather surprising, as the analytical solution predicts only one hyperbolic point in this case..

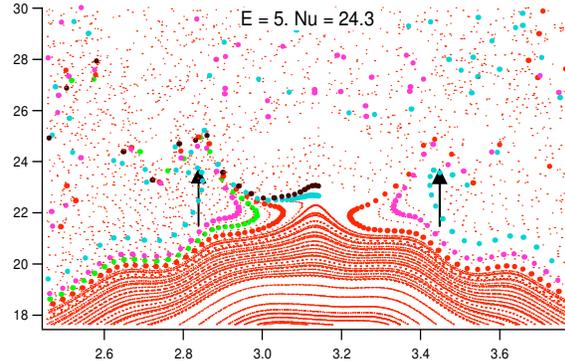


FIG. 4. Poincaré section for $\epsilon = 5.0$, $\nu_1 = 24.3$. Trajectories of various particles leaving the regular-motion domain are given different colors. The particles apparently do not leave from a single point, but rather from two points equidistant from $\theta = \pi$, each labeled by an arrow.

Higher values for ρ – stochastic domain

If we look at the off-resonance case for values of ρ above these hyperbolic point(s), we can see that there are two ‘streaks’ (indicated by the dashed arrows in Fig 3a.) which have a lower density of points, starting almost from the center ($\theta = \pi$), and moving out towards the sides.

Tracing the paths of individual particles on either side of these low-density streaks (Fig. 4b) shows that there is a strong tendency for the particles to stay on their ‘own’ side of them. Even though the motion of the particles is thus random, there are thus still certain bounds on its movement in the θ direction in the Poincaré section.

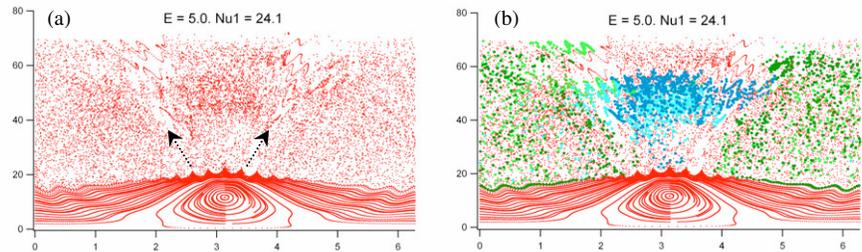


FIG. 5. Plots for $\nu = 24.1$. There are two ‘streaks’ indicated by arrows in (a) with a lower density of points. Color coding of individual particles in (b) indicates that there is very little exchange across these low density streaks. The ‘green’ particles tend to stay on the outside, while the ‘blue’ particles tend to stay on the inside.

Higher values for ρ – beyond the stochastic domain

As we get to higher values of ρ we reach the end of the stochastic domain (defined² as $\rho_{ub} = \left(\frac{2}{\pi}\right)^{1/3} (4\epsilon\nu)^{2/3}$) [1], which, for $\nu = 24.1$ and $\epsilon = 5$ is about 52, and so from

² This formula is for the upper bound for the stochastic domain for a *single* wave, but some experimentation has shown that it applies to some degree of accuracy to the beating waves case as well. For ν , I used ν_1 .

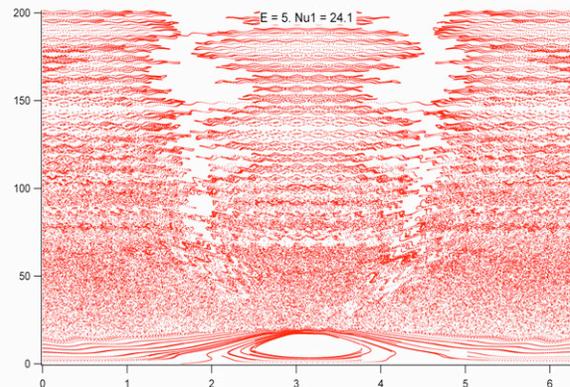


FIG. 6. Plot for $\nu_1 = 24.1$, illustrating the flat webs at high values of ρ , and the low-density ‘streaks’ which seem to converge to $\theta = \pi/2$ and $\theta = 3\pi/2$.

around this point upward the behavior of the particles is not due to stochastic effects, but rather to the nature of phase space (as defined by the parameters of the system, such as ν).

At these higher values of ρ , we see (Fig. 6) that the Poincaré plot resolves into flat, wavy webs, but the webs never seem to cross the low-density ‘streaks’. We also notice that these ‘streaks’ converge towards $\theta = \pi/2$ and $\theta = 3\pi/2$. Thus, for high values of ρ , we can divide phase space into two domains: the ‘inner’ domain $[\pi/2, 3\pi/2]$ and the outer domain $[0, \pi] \cup [3\pi/2, 2\pi]$ (these last two are obviously connected at $\theta = 0$).

A closer look at a single web

The unusually large gaps between the two domains at the top of the plot in Fig. 6 is not because of any intrinsic nature of the motion at higher values of ρ , but because many of the trajectories at the center are ‘trapped’ in smaller orbits of their domain. This becomes clearer if we zoom in to take a closer look at just one of some of the horizontal webs, at a very high value for ρ , say $\rho = 300$.

This close-up view (Fig. 7) reveals that there two main types of trajectories: Most of the particles follow a simple (rather wavy) web and are confined to their own domain. However, the particles at the edge of the domains follow a rather different path, and are able to jump from one domain to the other. It will be interesting to compare these two types of paths to those of the on-resonance case (see

p7). These horizontal webs are not discernible in the region of $\rho = 30$ or so, since there stochastic effects can easily cause particles to jump to higher or lower webs – something they can not do above the stochastic domain.

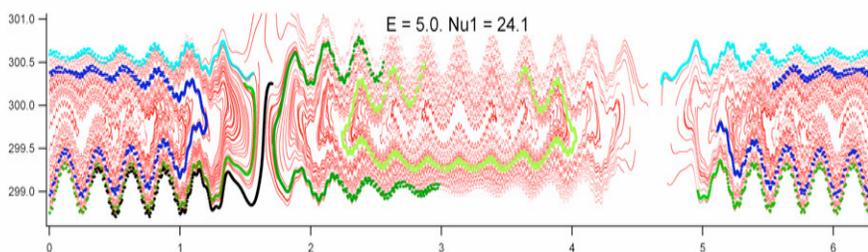


FIG. 7. A close up view of one of the horizontal webs in Fig. 4. Most of the particles’ trajectories (such as light blue, dark blue, light green and dark green) remain in their own ‘domain’, bounded by $\theta = \pi/2$ and $\theta = 3\pi/2$. Only particles which are on the edge of their domain (such as black) can pass to the other.

Thus, we may understand the stochastic domain in the off-resonance case (between $\rho = [20, 50]$ in Fig. 6) as a combination of two effects. First, there are the regular horizontal webs (whose boundaries grow from around $\theta = \pi$ to $\theta = \pi/2$ and $3\pi/2$). Second, this effect is complicated by a ‘veil’ of stochasticity, which obscures most of the webs, except for the low-density streaks in between the two domains.

Back to the hyperbolic point

Even though the low-density streaks converge towards $\theta = \pi/2$ and $3\pi/2$, they start out close to $\theta = \pi$ at the bottom of the stochastic domain. Their exact path is difficult to tell because of the stochastic effects, but if we take a low ε plot³ (Fig. 8), we can make out that these ‘streaks’ seem to originate from the region of the main hyperbolic point(s), and in fact, it might have something to do with the fact that there are two hyperbolic points instead of the one as predicted by the analytical solution. Depending on the mathematical explanation behind these streaks, it may or may not provide us with some insight into this anomaly.

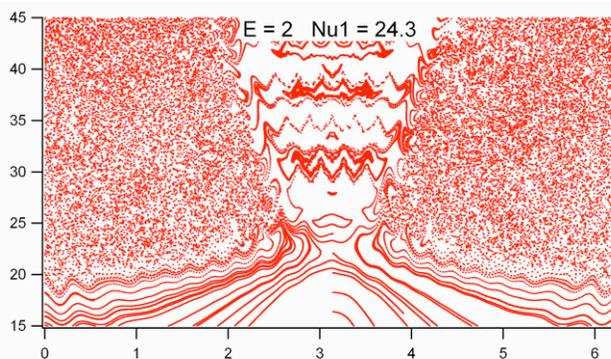


FIG. 8. A low ε plot for $\nu = 24.3$, illustrating how the low-density streaks seem to originate from the dual hyperbolic point noted above.

³ There is a threshold value for ε below which stochastic effects are reduced drastically [2], given by $\varepsilon_{th} \cong \nu^{3/4} / 4$ (which is ≈ 2.7 when $\nu = 24.3$). As we can see from Fig. 6, this threshold is not as simple for the beating wave case, but by dropping ε to 2, we can see what happens in the central domain much clearer than before.

6. On-resonance ($\nu \in \mathbb{Z}$).

When we look at a typical Poincaré section of the on-resonance case (Fig. 9), we find a very different picture from the off-resonance case. There are many differences in the particles' paths, but two things stick out in particular:

- (1) the presence of numerous hyperbolic points on the boundary of the stochastic domain, from which particles leave the regular-motion domain.
- (2) the diagonal, V-shaped webs, which center around $\theta = \pi$ and extend upwards towards the sides of the graph.

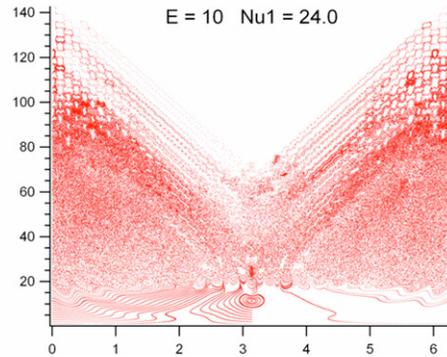


FIG. 9. A typical on-resonance case. Here, $\varepsilon = 10$, $\nu_1 = 24.0$

Hyperbolic and elliptical points on the border of the stochastic domain

If we look closely at the border of the stochastic domain, we note a discrete number of hyperbolic points. We should expect, then, that there should be a corresponding number of elliptical points in between them (otherwise, a particle could leave the regular-motion domain at any point, and not only at a hyperbolic point). But it is difficult to see these elliptical points in Fig. 9, for two reasons: (1) All the particles being simulated start off from very far away from these elliptical points,

and so do not enter the 'trapped' region around them. (2) At our value of ε , the stochastic domain completely overlaps these points, and so particles can stochastically enter and exit these 'trapped' regions, making them invisible on a Poincaré section. However, if we remove these two factors by (1) setting our starting conditions precisely in the vicinity of these elliptical points, and (2) lowering the value of ε below the stochastic

threshold (mentioned above – it was about 2.7), then we can see these elliptical points quite clearly (Fig. 10a). Taking similar plots for different values of ν reveals a direct relationship between the number of elliptical points and the value of ν (see Fig. 10b and 10c). There are, in fact, precisely 2ν elliptical points in each plot. There are, then, also 2ν hyperbolic points in between them, through which the particles escape into the stochastic domain (although at this low value of ε , there *is* no stochastic domain). This is reminiscent of the single wave case, where the boundary between the regular and stochastic domain was not a simple straight line, but was made of what looked like hyperbolic and elliptical points, the number of which was directly proportional to ν (Fig. 11).

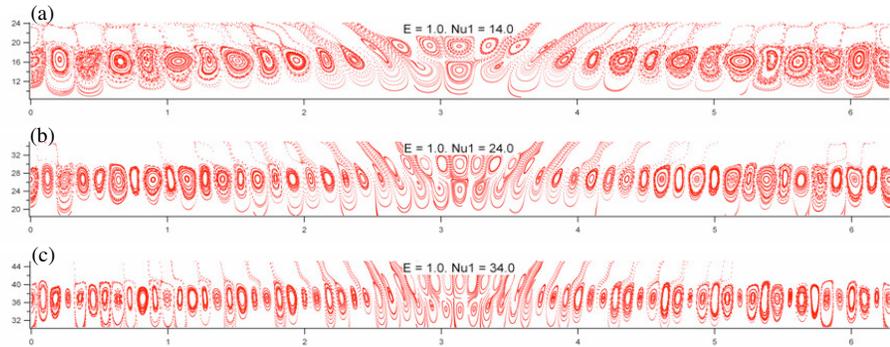


FIG. 10. Plots of the boundary between the regular and stochastic domain for $\nu = 14.0$ (a), $\nu = 24.0$ (b) and $\nu = 34.0$ (c). These plots at very low values of ε , so that these elliptical points are visible on the Poincaré plot. Although counting the elliptical points can become slightly ambiguous in the region of $\theta = \pi$, a careful count of the *bottom* row of elliptical points shows that there are exactly 2ν of them, and thus 2ν hyperbolic points in between them as well.

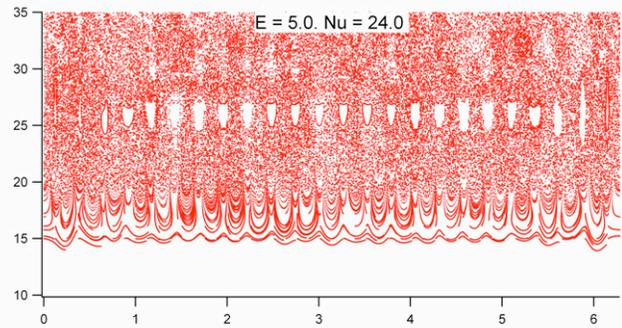


FIG. 11. Plot for a single wave, $\varepsilon = 5$, $\nu = 24.0$. Just as with two waves (with low ε) as above, the number of elliptical points at the border of the stochastic domain is related to ν . Here, in the single wave case it is, in fact, exactly equal to ν .

The V-shaped webs

Another characteristic feature of the on-resonance case is the V-shaped webs which dominate the plots. The plot in Fig. 12 is taken at a very high value of ρ , where stochastic effects are negligible, and we can thus see the nature of the motion with the stochastic elements removed. As in the case of off-resonance, there are two main types of trajectories: There are particles which remain ‘trapped’ around elliptical points (colored in various shades of green), and there are particles which manage to weave their way between all the elliptical points (colored black and grey), and are not bound to any particular values of θ . Most of the particles, though, actually lie somewhere in between these two extremes, and weave their way in a fixed web which encompasses a number of elliptical points (colored various shades of blue).

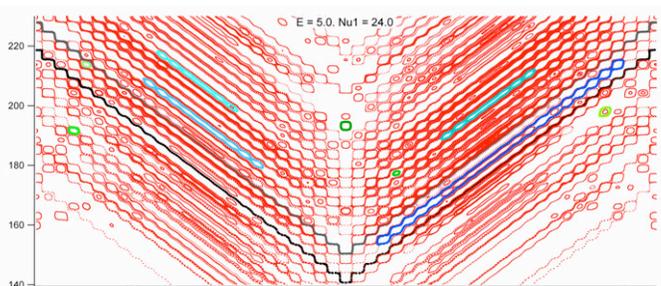


FIG. 12. On-resonance plot illustrating the V-shaped webs, and the various types of trajectories possible for a particle in them: fixed webs around a single elliptical point (green), unbounded paths weaving in between the elliptical points (black), and paths that are a combination of these (blue).

Height of the V-shaped webs and ν

The height of the V-shaped webs (I have defined the height of a web as $\rho_{\theta=0} - \rho_{\theta=\pi}$, i.e. the difference between the highest and lowest possible points of the ‘black’ particle in Fig. 12) is determined by how many elliptical points a particle has to ‘weave’ its way around to get to the top of the web. Thus, we expect that the height should increase with higher values of ν , since we have seen above that the number of elliptical points (i.e. in a horizontal direction – between $\theta = 0$ and $\theta = \pi$) is 2ν .

Indeed, if we look at a couple of on-resonance plots with different values of ν (Fig. 13) this is exactly what we find: We see that as ν increases, the height of the webs grows, proportional to the increase in ν .

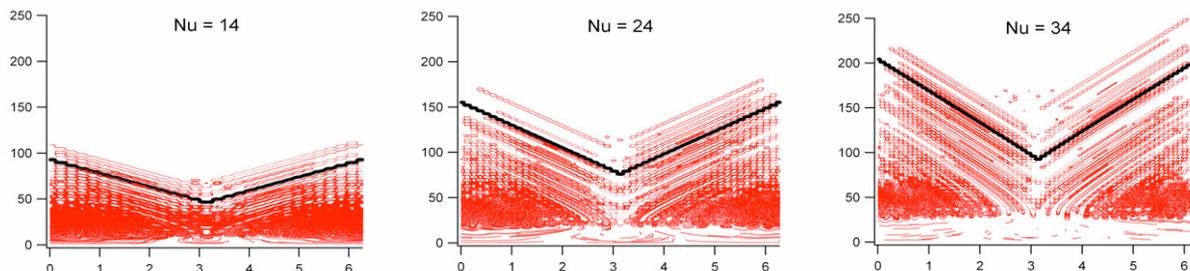


FIG. 13. Plots for different values of ν ($\varepsilon = 5.0$). In each plot, one of the webs has been colored black, so that the difference between the heights is more apparent. As expected, the height of the V-shaped webs increases in direct proportion to ν .

In other words, since the structure of phase space is such that there are more elliptical points (in the θ direction) with a higher ν , a particle will ‘step over’ more elliptical points and will thus travel further up or down for higher values of ν . More precisely, since the particle starting from $\theta = \pi$ passes precisely ν elliptical points as it works its way to $\theta = 0$, we should expect the height of the webs to be $H = \nu h$, where h is height of a single elliptical point.⁴ Extensive numerical exploration reveals, interestingly enough, that the value of h does not seem to have any dependency on ν or ε (or ρ), and remains a constant height of approximately 3.3 ± 0.05 . The total height of the web from top (at $\theta = 0$) to bottom (at $\theta = \pi$) is thus 3.3ν . Numerical exploration also confirms that the shape of these V-shaped webs does not change at all for higher values of ρ .

⁴ It is easiest to measure this near $\theta = 0$ or $\theta = \pi$ where the height of the ‘steps’ is more distinct.

Upper boundary of the stochastic domain and ε

As is the case for a single wave, the value of ε affects the upper bound of the stochastic domain. As mentioned above, the predicted upper bound for the *single* wave case is $\rho_{ub} = (\frac{2}{\pi})^{1/3} (4\varepsilon\nu)^{2/3}$ [2]. For beating waves, the upper boundary of stochastic domain also increases with increasing ε , and it does follow the predicted upper bound in a general sense, but if we look carefully we see it is a little more complex (Fig. 14): Particles may rise above this boundary around $\theta = 0$ (and $\theta = 2\pi$), but this is presumably due to them following the V-shaped webs mentioned earlier. (This is supported by the observation that open circles—indicating the presence of elliptical points—appear prominently only above ρ_{ub} in the region of $\theta = 0$. This is a good indication that the stochastic effects have dropped significantly, since in the stochastic domain particles can randomly jump in out of elliptical points, making them invisible on a Poincaré section.) However, around $\theta = \pi$, the border of stochastic domain drops significantly below ρ_{ub} . In this respect, the boundary of the stochastic domain deviates somewhat from that in a single wave.

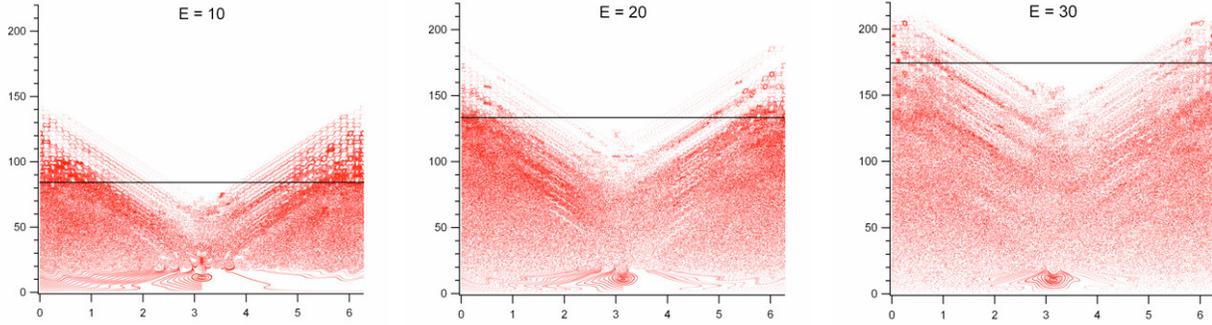


FIG. 14. On-resonance plots for different values of ε . The predicted value of the upper bound of the stochastic domain (for a single wave) is marked on each line by a horizontal line. The stochastic domain appears to end below this line at $\theta = \pi$, and go above it at $\theta = 0$.

“Half-Resonance”

I will note briefly here the interesting ‘half-resonance’ case, i.e. where $\nu_1 = 24.5$ (Fig. 15). Even though it is not an on-resonance case, it displays—to a lesser extent—some of the characteristic features of the on-resonance case: (1) It has some diagonal webs (colored blue), although they are not as long and clear as in the on-resonance case. (2) There are numerous elliptical points (colored green) which fit together, curiously, in a tessellating diamond-shape. (3) It has a number of hyperbolic points along the bottom border of the stochastic domain.

On the other hand, the ‘half’-resonance case also bears many similarities to the off-resonance case – in particular, for high values of ρ , we see horizontal webs in addition to the V-shaped webs. It appears that the “half-resonance” case is an interesting hybrid of the off and on-resonance cases.

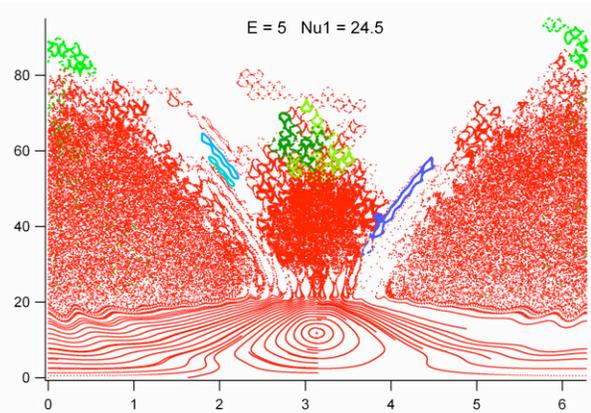


FIG. 15. “Half-resonance” case of $\nu = 24.5$. This plot has numerous similarities to the on-resonance case: some short diagonal webs (blue) reminiscent of the V-shaped webs; many elliptical points (green) which here are diamond-shaped instead of circular; and numerous hyperbolic points can be found on the border of the stochastic domain.

7. Approaching Resonance

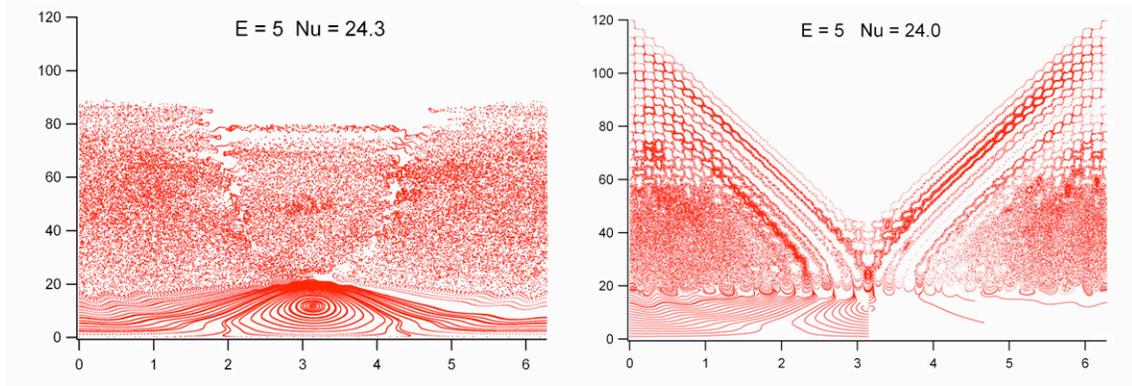


FIG 16. The sharp contrast between the off-resonance (left) and on-resonance (right) plots.

As is evident from Fig. 16 above, there are many stark differences between the off-resonance and on-resonance cases. For example, the numerous hyperbolic points on the border of the stochastic domain in the on-resonance case, versus the one (or two) in the off-resonance case. Also apparent is the V-shaped webs in the on-resonance case, versus the horizontal webs in the off-resonance case, both of which we have explored above. However, despite appearances, these differences are not a sudden, discrete jump, but rather a continuous process, which we can watch in ‘slow motion’ if we look at the cases very close to resonance, such as 24.1 and 24.01. By doing so, we will hopefully gain an intuitive understanding of how, as we approach resonance, the two cases merge into a beautiful continuum.

(a) The appearance of hyperbolic points

If we watch the border of the stochastic domain, near the hyperbolic point(s) of the off-resonance case, we find that even as we approach 24.2 and 24.1, we notice the border of the regular-motion domain becoming more jagged (Fig. 17), hinting to us where the hyperbolic points will appear. In fact, already at $\nu_1 = 24.1$, tracing individual particles shows that particles do not only leave from the center, but also from near the sides.

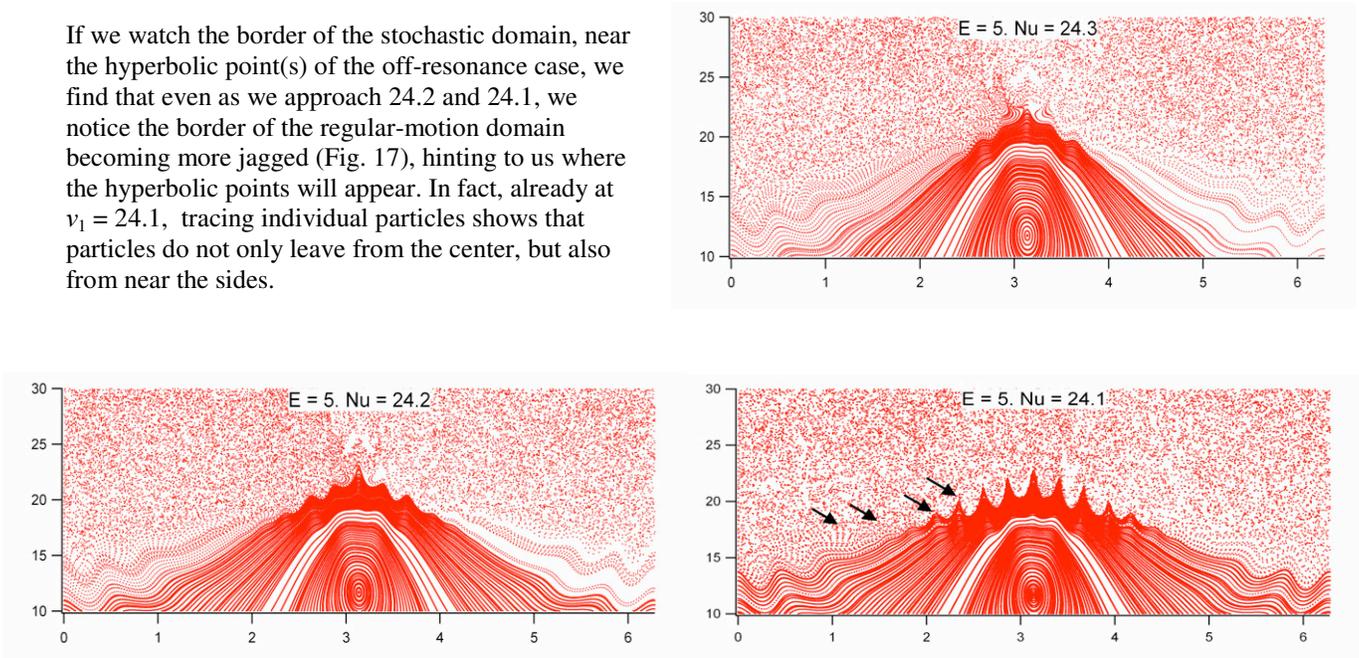


FIG. 17. Plots of $\nu_1 = 24.3$, 24.2 and 24.1. As we get closer to resonance, the boundary gets more ‘jagged’, giving us a clue as to where the hyperbolic points will be. In the 24.1 case, various particles have been traced, and the place where they have left the regular motion domain has been marked with an arrow.

(b) The change in size of the trapped domain.

One of the consequences of the appearance of many hyperbolic points is that the trapped domain begins to shrink in size. Trajectories which otherwise might have remained in the trapped domain can now escape via one of the new hyperbolic points. In Fig. 18, we can see how the ‘blue’ particles, which might otherwise have remained in the trapped domain, are now able to escape through a new hyperbolic point.

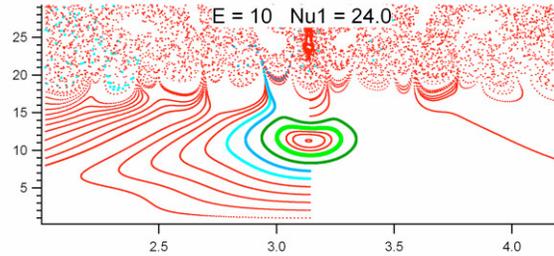


FIG. 18. Plot for $\nu = 24.0$, demonstrating why the trapped region diminishes in size in the on-resonance case. The ‘blue’ particles can now escape through one of the new hyperbolic points, leaving only the ‘green’ particles in the trapped region.

In the plots in Fig. 19 (below), we can see the way that the trapped domain changes as we get closer to resonance. All the trajectories of particles that eventually escaped to the stochastic domain have been removed, and so the trapped domain is clear and unambiguous. There are separate graphs for different values of ε , but each graph has plots for various values of ν superimposed, ranging from 24.1 (light blue) to 24.0 (black). For relatively low values of ε ($\varepsilon = 5$ and $\varepsilon = 10$) we can see—for each plot—how the trapped domain gets smaller as we get closer to resonance (as we go from light blue to dark blue to black).

As Slava showed in [1], increasing the value of ε also causes the trapped domain to diminish in size, but we see here that this is only part of the picture. Firstly, that result was taken for off-resonance plots: in contrast, we can see how the on-resonance plots (all the black traces from Fig. 19a-d) get consistently larger for larger ε , and at some point they start growing sideways in weird and wonderful ways. But even for the off-resonance plots, the trapped domain

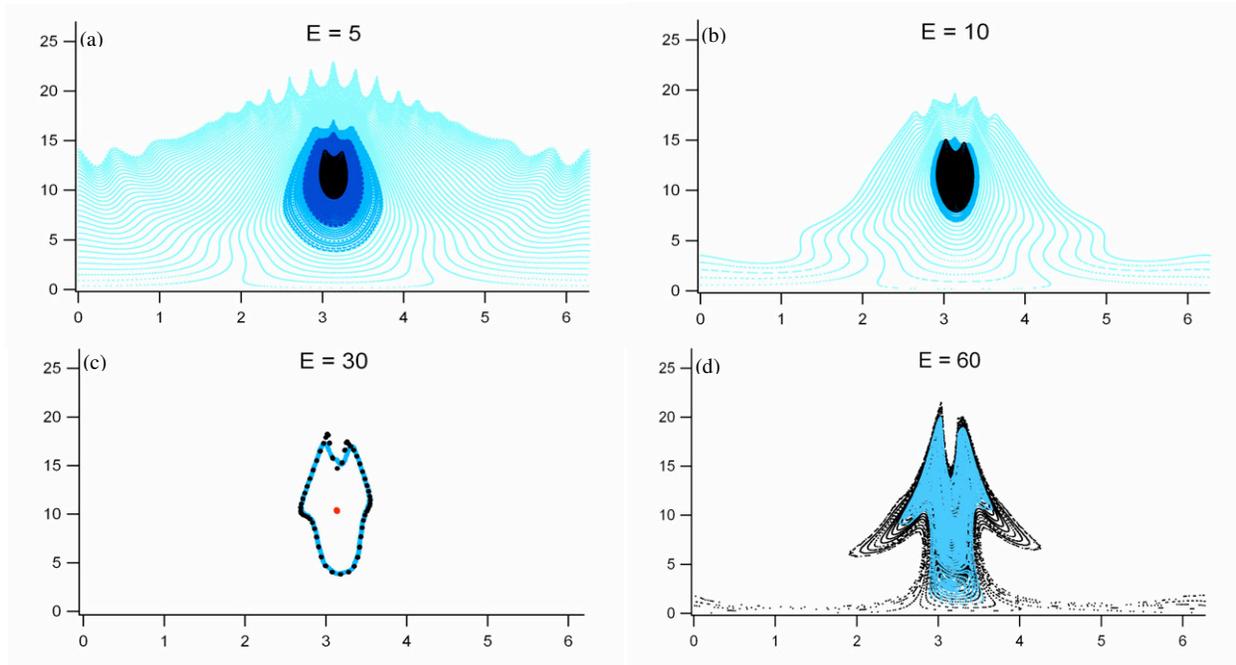


FIG. 19. Superpositions of the trapped regions in different Poincaré plots. In these plots, I excluded all paths of particles which eventually escape to the stochastic domain.

- (a) Poincaré plots for $\varepsilon = 5$, for $\nu_1 = 24.1$ (lightest blue), $\nu_1 = 24.01$ (sky blue), $\nu_1 = 24.001$ (dark blue) and the on-resonance case of $\nu_1 = 24.0$ (black).
- (b) Poincaré plots for $\varepsilon = 10$, for $\nu_1 = 24.1$ (lightest blue), $\nu_1 = 24.01$ (sky blue), and the on-resonance case of $\nu_1 = 24.0$ (black). Note that the trapped region of the on-resonance case is slightly larger for $\varepsilon = 10$ than for $\varepsilon = 5$.
- (c) Poincaré plots for $\varepsilon = 30$, for $\nu_1 = 24.1$ (blue), and $\nu_1 = 24.0$ (black). Since the trapped regions for these two cases overlap so much, I included just the outlines of these regions for the sake of clarity.
- (d) Poincaré plots for $\varepsilon = 60$, for $\nu_1 = 24.1$ (blue), and $\nu_1 = 24.0$ (black). Here, we see that the trapped region is larger for the on-resonance case.

gets smaller only up to a point (somewhere between $\varepsilon = 10$ and $\varepsilon = 60$), but then it starts growing again, eventually looking taking on the blue dressing-gown shape in Fig. 19d. As these changes occur, the upper boundary of the trapped domain at $\theta = \pi$ seems to drop lower and lower, and conversely, almost as if in compensation for this, the trapped domain starts to expand at the sides. For $\varepsilon = 30$, we have the interesting case that the size of the trapped domain in the off and on-resonance cases is about the same.

Thus, if we are looking for as small a trapped domain as possible, we would be interested either in low values for ε with resonance, or at $\varepsilon \approx 30$ with off-resonance.

(c) Horizontal \rightarrow Diagonal webs

Another critical way in which the on-resonance case differs from the off-resonance case is in the behavior of the particle at high values of ρ . Whereas the particles in the off-resonance case travel along paths of more or less constant ρ (Fig. 20a), in the on-resonance case the particles travel along V-shaped stochastic webs (Fig. 20b), and can have significant increase (or decrease) in their value of ρ .

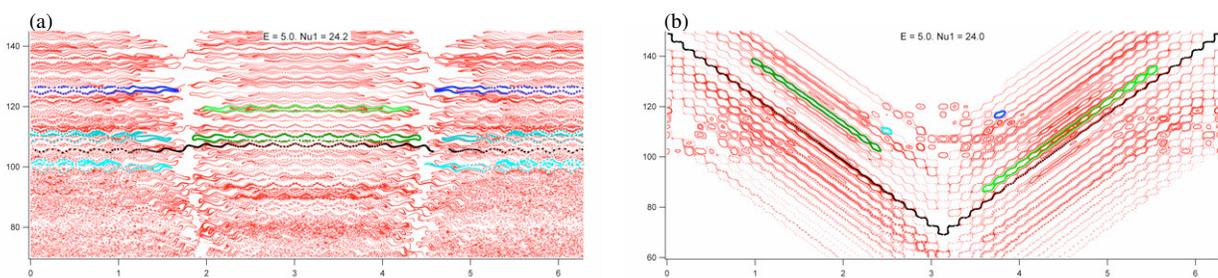


FIG. 20. Plots at high ρ for the off-resonance (a) and on-resonance (b) cases reveal major differences between the two cases. A particle in the on-resonance case travels along V-shaped webs, while the off-resonance case sees particles moving in horizontal webs.

This abrupt change between the on- and off-resonance cases may initially be surprising, but we can find an intuitive approach to understanding it by looking closely at how the horizontal (off-resonance) webs change as we get closer to resonance.

With $\nu = 24.1$, we note that the horizontal webs are completely independent of one another (Fig. 21a). Each web is quite ‘wavy’, but although the webs almost seem to overlap, there is no crossing over between one web and the ones just above or below it. However, with $\nu = 24.01$ (Fig. 21b), we notice that the trajectories of many of the webs seem to cross over what were before (in the 24.1 case) different, separate webs. This effect is even more pronounced in the case of $\nu = 24.005$ (Fig 21c). It seems that as we get closer and closer to resonance, there is more and more crossing over between the different horizontal webs, until eventually we get a continuous, diagonal chain which forms the characteristic V-shape of the on-resonance case.

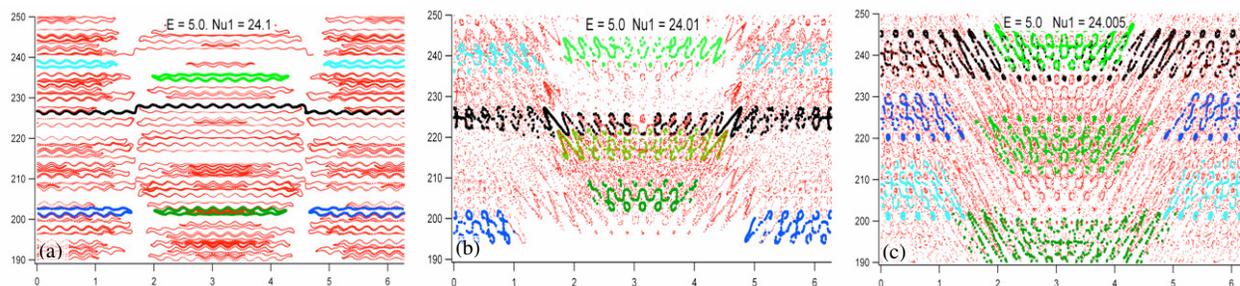


FIG. 21. Plots for values of ν closer and closer to resonance: $\nu = 24.1$ (a), $\nu = 24.01$ (b) and $\nu = 24.005$ (c). As we get closer and closer to resonance, adjacent horizontal webs connect more and more connected. These plots are taken at very high values for ρ , so stochastic effects do not play any role here.

8. Conclusion: Applications for Plasma Acceleration

After looking at some of the basics of the beating wave phenomenon, it is useful to stop and assess what we have gleaned so far from the plots, particularly with regard to the problem of acceleration particles in a plasma. When considering how to relate the theory to practice, a number of thoughts come to mind:

1) When searching for the region in parameter-space which gives the most favorable results for acceleration, we should keep in mind that ideally, we would like the trapped domain to be as small as possible. As Slava showed in his paper [3], collisions between particles may actually reduce the effect of the trapped domain, since particles may be knocked out of it via collisions in the plasma. However, we have to consider plasmas that may be hot enough to be collisionless, and thus the size of the trapped domain is a factor must take into consideration. The relationship between ε , ν , and the size of the trapped domain is quite complex, as noted above, and might warrant further investigation.

2) The average speed is determined by the upper bound of the stochastic domain. A particle starting out at arbitrarily low velocity can reach up to 50 times its original speed by stochastic acceleration. However, the efficiency of this process is, again, governed by our parameters. For high values of ν , even though they may have a higher average velocity, the stochastic domain begins much higher, and the particles may move much slower in the regular-motion domain, such that it may take up to 100 times longer to get to the stochastic domain.

3) Although it has been commonly agreed upon [2] that a particle interacting with electrostatic waves can only accelerate stochastically, it is very tempting to try utilize the V-shaped webs of the on-resonance case as a ‘loophole’ to this rule. A single particle that starts out at $\theta = \pi$ will on average attain a higher speed than it’s initial speed (climbing up the diagonal web). However, a particle that starts out at $\theta = 0$ will experience the reverse effect, and its average speed will decrease. Thus, if we place a whole set of particles randomly in phase space (above the stochastic domain) and watch the energy vs. time, on average we would see no net energy change. However, if we could find a process which treats the $[\pi/2, 3\pi/2]$ and $[0, \pi/2] \cup [3\pi/2, 2\pi]$ domains asymmetrically, such that it results in more particles at $\theta = \pi$ than at $\theta = 0$, then it might be possible to use this process in conjunction with the on-resonance case to produce an acceleration affect without stochasticity.

References:

- [1] K. Ram, A. Bers, and D. Benisti. *Ionospheric ion acceleration by multiple electrostatic waves*. Geophysical Research Letters, 103:9431–9440, 1998.
- [2] D. Benisti, A.K. Ram, and A. Bers. *Ion dynamics in multiple electrostatic waves in a magnetized plasma, I. Coherent acceleration*. Physics of Plasmas, 5(9):3224–3232, 1998.
- [3] E.Y. Choueiri and S. Spektor. *Coherent ion acceleration using two electrostatic waves*. In 36th AIAA Joint Propulsion Conference, Huntsville, AL, USA, 2000. AIAA-2000-3759.
- [4] R. Spektor and E. Y. Choueiri. *Ion acceleration by beating electrostatic waves: Domain of allowed acceleration*. Physical Review, E, 69(4):Article No. 046402, 2004
- [5] R. Spektor and E.Y. Choueiri. *Effects of ion collisions on ion acceleration by beating electrostatic waves*. In 28th International Electric Propulsion Conference, Toulouse, France, 2003. IEPC-03-65.
- [6] J. Candy and W. Rozmus. *A Symplectic integration algorithm for separable Hamiltonian functions*. Journal of Computational Physics, 92:230–256, 1991.