Fermilab FERMILAB-PUB-03-148-A

Chaos and collective relaxation in galaxies and charged-particle beams

Courtlandt L. Bohn,[1,4] Henry E. Kandrup,[2*] Rami A. Kishek[3]
Patrick G. O'Shea,[3] Martin Reiser,[3] Ioannis V. Sideris[2,4]

[1]Fermi National Accelerator Laboratory, Batavia, IL 60510
[2]Departments of Astronomy and Physics, University of Florida, Gainesville, FL 32611
[3]Institute for Research in Electronics and Applied Physics,
University of Maryland, College Park, MD 20742
[4] Department of Physics, Northern Illinois University, DeKalb, IL 60115

*To whom correspondence should be addressed; E-mail: kandrup@astro.ufl.edu.

Both galaxies and charged particle beams can exhibit collective relaxation on surprisingly short time scales. This can be attributed to the effects of chaos, often triggered by resonances caused by time-dependences in the bulk potential, which act almost identically for attractive gravitational and repulsive electrostatic forces. These similarities suggest that many physical processes at work in galaxies, albeit not subject to direct controlled experiments, can be tested indirectly using facilities such as the University of Maryland Electron Ring (*UMER*) currently nearing completion.

Introduction

Many-body systems interacting via long range $1/r^2$ Coulomb forces – both gravitational and electrostatic – can exhibit macroscopic relaxation and loss of coherence on time scales much shorter than might be expected on dimensional grounds. Observations and simulations agree that even a relatively gas-poor (and thus presumably nearly dissipation-free) elliptical galaxy displaced from an equilibrium as a result of an encounter with another galaxy can readjust itself towards a new equilibrium within a few hundred million years although the nominal relaxation time t_R is orders of magnitude longer than the age of the Universe. And similarly, charged particle beams traveling through an accelerator, which would be expected to maintain coherence for 10 km or more, can lose coherence and disperse within distances as short as 10 m.

This collective behaviour must reflect phase mixing, *i.e.*, the tendency of initially localised clumps of orbits to disperse. However, any phase mixing that can explain these phenomena must proceed far more efficiently than 'ordinary' phase mixing. The key point, then, is that these effects can be neatly explained if the bulk potential associated with the system admits large numbers of chaotic orbits. Phase mixing is *much* more efficient in a chaotic system than a system in which the bulk potential is integrable or near-integrable. An initially localised clump of regular, *i.e.*, non-chaotic, orbits will typically disperse as a power law in time; a clump of chaotic orbits will instead disperse exponentially.

Allowing for a bulk potential that is strongly chaotic and, hence, the possibility of chaotic phase mixing would enable one to understand how a galaxy can 'relax' towards an equilibrium or near-equilibrium state on a comparatively short time scale. The possibility of chaotic phase mixing in a charged particle beam can place strong constraints on emittance compensation, *i.e.*, the ability of the experimentalist to compensate for unwanted dispersal.

Theoretical considerations and detailed numerical simulations suggest that, in this setting, the origin of the chaos that drives the evolution is largely irrelevant. In particular, whether the two-body forces are attractive or repulsive should not be crucial. What is important is that the long-range scalings of gravitational and electrostatic forces are identical and that, in both cases, the early stages of evolution should be driven by long-range, collective interactions as opposed to short-range collisional encounters. All that seems to matter is whether the bulk potential associated with the many-body system admits a large measure of chaotic orbits (*1*).

That 'collisions' should be largely irrelevant in many settings relevant to galaxies and beams is easily seen. Viewing such effects as an incoherent sum of binary encounters, one computes, respectively, for galaxies and for charged particle beams (in cgs units) relaxation times

$$t_R \sim \frac{v^3}{G^2 m^2 n \log \Lambda} \qquad \text{and} \qquad t_R \sim \frac{v^3}{q^2 n \log \Lambda}.$$

Here v is a typical speed associated with random motions, G the gravitational constant, m and q typical stellar masses and charges, n a characteristic number density, and $\log \Lambda$ the so-called Coulomb logarithm (*2*) (*3*). In either case, the assumption that the bulk kinetic and potential energies are comparable in magnitude implies that $t_R \propto (N/\log \Lambda)t_D$, with N the number of 'particles' and $t_D \sim v/R$ a characteristic orbital time scale, defined in terms of the 'size' R of the system. Λ scales as a power of N. For large N the relaxation time t_R is clearly very long compared to the orbital time scale t_D. By contrast, chaotic phase mixing, *i.e.*, the phase mixing of chaotic orbits, can proceed extremely fast. The time scale associated with the exponential dispersal of an initially localised 'clump', given as the inverse of the largest positive Lyapunov exponent (*4*), is typically comparable in magnitude to the orbital time scale. This is, for example, the case for the systems illustrated in Figures 1 and 3 below.

A complete understanding of these phenomena will require a synthesis of theory, simulations, and experiments. Performing experiments on self-gravitating systems like galaxies is impossible. However, controlled experiments *can* be performed with charged-particle beams; and combining the results of such experiments with simulations and theory should lead to a clear picture of the role of chaotic phase mixing in beams. Moreover, since the physics should not depend crucially on whether the force between particles is attractive or repulsive, one would expect that many results about beams should translate more or less directly into detailed predictions about the structure and evolution of galaxies.

Indeed, one can go one step further and argue that, in a real sense, carefully constructed experiments involving charged particle beams can be used as semi-direct probes of the physics of self-gravitating systems like galaxies. In this spirit, one aim here is to explain how the University of Maryland Electron Ring (*UMER*), currently nearing completion, can be used as a Laboratory for Galactic Dynamics.

Regular vs. Chaotic Orbits: A Tortured History

Chaos has been largely ignored until comparatively recently in both the galactic and accelerator dynamics communities. For example, although the famous Hénon-Heiles model (*5*) arose originally in attempts to understand meridional motions in axisymmetric galaxies, as recently as 15 years ago the potential role of chaos in galaxy structure and evolution was almost completely neglected (with the exception of a handful of groups in Europe). Only with the advent of high resolution photometry, facilitated in part by the Hubble Space Telescope, did many galactic astronomers begin to recognise that the bulk potentials associated with realistically shaped galaxies are likely to admit significant measures of chaotic orbits.

Chaos in galaxies. It has been long recognised that the dominant mechanism for relaxation in galaxies cannot be 'collisional'. For example, in the 1940's Chandrasekhar (*2*) showed that the relaxation time scale t_R on which binary encounters between individual stars could significantly alter the trajectories of stars in the Milky Way must be $\sim 10^{12}$ yr or more. Shorter-time relaxation must somehow involve collective effects. Two decades later, Lynden-Bell (*6*) proposed a theory of 'violent relaxation' which argued, *inter alia*, that phase mixing associated with a time-dependent potential might explain such collective effects. Substantial evidence for rapid relaxation accumulated over the next twenty years, derived both from numerical simulations of many body systems and from the interpretation of observations of galaxies that have been involved in collisions with other galaxies. Despite this, however, when subjected to closer scrutiny, it seemed that, at least in its simplest guise, violent relaxation could not explain why relaxation was as fast as it appears to be. Some ingredient seemed to be missing. Today there is good reason to think that the missing ingredient is chaos.

In the early 1990's Kandrup and Mahon (*7*) recognised a fact that seems obvious in retrospect, namely that, because of their exponentially sensitive dependence on initial conditions, chaotic orbits should phase mix far more rapidly than regular orbits, in fact exponentially fast. In the astronomical community, this phenomenon, now termed 'chaotic mixing' (*8*), led to speculations that chaos could play a critical role in violent relaxation. However, 'chaotic mixing' in itself does not constitute a complete and satisfactory explanation. Chaotic phase mixing cannot drive collective relaxation unless many/most of the orbits are chaotic, which seemed far from obvious. However, a few years later, motivated in part by the work of accelerator dynamicists (*9,10*), astronomers (*11*) recognised that time-dependent pulsations in the bulk potential of the form expected in a galaxy readjusting towards a (near-)equilibrium state could, via resonant couplings, make many/most of the orbits in a galaxy chaotic with large finite-time Lyapunov exponents (*12*), and that the resulting 'resonant phase mixing' might be sufficiently strong and ubiquitous to explain violent relaxation.

Chaos in charged particle beams. Concerns about collisionless relaxation in charged-particle beams have arisen with recent advances in technology for the production of high-brightness beams, where collective Coulomb self-forces, *i.e.*, space charge, becomes important. Examples include low-to-medium-energy hadron accelerators such as those that drive spallation-neutron sources or serve as boosters for high-energy machines, heavy-ion accelerators, and low-energy electron accelerators such as photoinjectors (*13*).

One of the earliest papers addressing space charge in beams was by Kapchinskij and Vladimirskij (*14*), who considered a continuous beam with uniform charge density and elliptical cross-section confined by linear external focusing forces, and derived the equations governing the motion of the beam envelope. The corresponding distribution function in the four-dimensional phase space of a single charge, commonly called the KV distribution, is a uniform-density hyperellipsoid. Some time later, Sacherer (*15*) noted that these results can be generalized easily to three-dimensional bunched beams so as to include the influence of space-charge on bunch length and energy spread. These two papers, regarded as classics in the accelerator community, set the stage for much of the subsequent investigations concerning space-charge, from which evolved now-conventional design strategies for high-brightness accelerators, strategies based on controlling root-mean-square (rms) properties of the beam.

However, the past decade has brought the realization that, albeit necessary, regulating rms properties is not sufficient. Perhaps the most prominent example concerns beam halos, *i.e.*, particles that reach large orbital amplitudes due to a time-dependent space-charge potential associated with transitions in the beam line that prevents the beam from reaching a long-lived equilibrium state. The concern is that impingement of beam particles on the beam line can generate sufficient radioactivation to preclude routine, hands-on maintenance. The radioactivation threshold is tiny: roughly 1 nA m^{-1} GeV^{-1} or less. For a 100 mA, 1 GeV light-ion beam such as that envisioned for driving high-yield spallation neutron sources, this criterion translates to just 1 in 10^8 particles lost per meter (*16*).

Early efforts to identify the fundamental mechanisms of halo formation involved generalizations of the work of Kapchinskij, Vladimirskij, and Sacherer (*9,17*). The basic recognition was that if a uniform-density beam 'core' is made to pulsate, particles resonating with the pulsating core can be driven to large amplitudes where they form a 'halo'. This led to the identification of parametric resonance as the essential mechanism of halo formation.

Albeit useful the stationary KV model suffers some serious deficiencies. The sharp boundary in phase space is unphysical and makes the distribution unstable to several classes of perturbations (*18*). Moreover, because the density is uniform, the net force on every particle is linear so that they all follow regular orbits. A real, inhomogeneous beam will impart nonlinear self-forces to the particles, forces that can induce globally chaotic motion (*19*). In a time-dependent potential such chaotic particles, as well as any other particles that become chaotic because of the time-dependence, can couple with the pulsations and experience a parametric resonance, phenomenology that has surfaced routinely in more comprehensive simulations that followed the particle-core model (*20*).

Efforts to understand space-charge-induced dynamics beyond the KV and particle-core models also brought the realization that early-time dynamics in fully self-consistent charged-particle systems is analogous to that of violent relaxation in stellar systems since both involve Coulomb interactions (*10*). However, that connection was explored no further – until now.

Evidence for Chaos and Chaotic Phase Mixing

Chaos in galaxies. High-resolution observations of galaxies over the past decade or so have provided compelling evidence that many galaxies are more irregularly shaped than had been assumed as recently as 15 years ago; and attempts to model such irregularly shaped objects have led many galactic dynamicists to conclude that the bulk potentials associated with realistic galaxies admit large measures of chaotic orbits. It has been argued (*21*) that nonaxisymmetric elliptical galaxies containing central density cusps of the form inferred from observations (*22*) are very likely to admit large numbers of chaotic orbits. And similarly, models of rotating barred spiral galaxies suggest (*23*) that breaking axisymmetry with even a comparatively weak bar can trigger large numbers of chaotic orbits, especially near certain resonances.

There is no guarantee that cuspy, nonaxisymmetric galaxies must be chaotic: contrived integrable models have been constructed (*24*) and it has proven possible, by careful preparation, to generate N-body realisations of nonaxisymmetric models which appear to contain few if any strongly chaotic orbits, with large positive Lyapunov exponents (*25*). However, there has emerged a general sense in much of the galactic dynamics community that 'generic' irregularly shaped galaxies might be expected to contain large numbers of chaotic orbits.

One intriguing possibility is that, possibly because of the presence of chaos, evolving galaxies will find it difficult, if not impossible, to approach a true equilibrium. Rather, it may well be (*21, 26*) that, at the time of formation, a galaxy settles down towards a quasi-equilibrium, rather than a true equilibrium; but that, in response, *e.g.*, to external irregularities associated with a high-density environment, it will subsequently exhibit a slow, secular evolution.

To the extent that this be true, a basic question is whether a galaxy originally in a non-axisymmetric near-equilibrium will evolve towards a more nearly axisymmetric state (*21*); or whether instead a galaxy originally containing large numbers of strongly chaotic orbits might evolve towards other near-equilibria, not necessarily more nearly axiymmetric, which contain smaller numbers of chaotic orbits (*27*). In any event, it is generally accepted that a robust, stable (near-)equilibrium must contain large measures of regular orbits to provide the 'skeleton' (*28*) of the interesting structures that generate chaotic orbits in the first place.

There is also emerging evidence that chaos should be even more ubiquitous in systems that feel a strongly time-dependent bulk potential, especially a time dependence involving roughly periodic oscillations. Nonlinear dynamicists argue that chaos typically arises via resonance overlaps (*4*), and this time-dependent chaos is simply another example thereof. When the time-dependence to which orbits in a galaxy are subjected has power at frequencies sufficiently close to (multiples of) the frequencies at which the orbits themselves have power, the orbits and the time-dependence can resonate with the result that the orbits become strongly chaotic.

If the time-dependence is weak, such resonances may require a near-perfect frequency match, but for a stronger time dependence it often suffices for the pulsation and orbital time scales to agree within an order of magnitude (*11*). However, in a nearly collisionless system like a galaxy, there is only one natural time scale, the dynamical time $t_D \sim 1/\sqrt{G\rho}$, with G the gravitational constant and ρ a characteristic density; so that the pulsation and orbital times are likely to be comparable in magnitude throughout much of the galaxy, thus rendering chaos extremely common. Simple models suggest that galaxies subjected to damped oscillations could (i) become almost completely mixed *and* (ii) settle down towards a nearly integrable equilibrium within a time as short as $\sim 10t_D$. Analogous effects can also be triggered by other nearly periodic phenomena, such as localised, nonstationary collective modes, or a supermassive black hole binary orbiting near the center of a galaxy, a mechanism of contemporary interest (*29*).

An example of such resonant phase mixing is illustrated in Figure 1, which tracks three initially localised clumps evolved in a toy galactic potential subjected to periodic driving that damps as a power law in time. The left and center panels exhibit the x- and y-coordinates at several different times; the right panel exhibits the exponential growth of components of the *emittance* ϵ_i ($i = x, y, z$), which is a measure of the phase space area of the occupied phase-space planes corresponding to the coordinate r_i. Here, *e.g.*, $\epsilon_x^2 = \langle x^2 \rangle \langle v_x^2 \rangle - \langle x v_x \rangle^2$, where $\langle \ \rangle$ denotes an ensemble average.

There can of course be no direct experimental evidence for chaos in galaxies. However, careful analysis of observable velocity fields provides compelling evidence that the gas flows in such spiral galaxies as *NGC* 3632 could be chaotic, especially near various resonances (*30*).

Chaos in charged particle beams. Anisotropy in a beam can have a number of causes, including anisotropic focusing and the acceleration process itself. However, a recent computational study has provided strong evidence that chaotic mixing due to nonlinear forces from space-charge waves is intimately connected with equipartitioning in beams (*31*). These computations were done using the '$2\frac{1}{2}$-dimensional' version of the particle-in-cell code *WARP* (*32*), which tracks randomly distributed macroparticles in pre-specified external electric and magnetic fields, along with the self-consistently calculated self-fields. The work focused on a highly space-charge-dominated, direct-current, cylindrical beam in which the initial momentum space reflected an anisotropic pressure with $p_{xx} = 2p_{yy}$. As the beam evolved, the pressure isotropized on a rapid time scale, fully equipartitioning in ~ 5 m and exhibiting anisotropic oscillations that largely damped by ~ 50 m. The equipartitioning time scales were found to correlate with the evolution of initially localized clumps of globally chaotic particles, which dispersed exponentially with an e-folding 'time' ~ 2 m (roughly two plasma periods) and filled their accessible phase spaces in ~ 50 m.

This first study considered a form of 'symmetry breaking', with the broken symmetry appearing in momentum space rather than configuration space. The beam began in a nonequilibrium state and evolved toward a meta-equilibrium in which the particle orbits filled an invariant measure of phase space. The transient dynamics reflects an intricate, evolving network of space-charge waves that set up a complicated potential in which a substantial population of particle orbits becomes globally chaotic. However, equipartitioning does not lead to a halo, as long as the rms properties of the beam are 'matched' (*33*) to the strength of the focusing forces, thus minimising large scale time-dependent oscillations. By contrast, a symmetric, isotropic system evolved analogously establishes a potential that is integrable, in which the orbits are regular. The presence of chaotic orbits is evident in Fig. 2, which shows the trajectories of 20 test particles randomly selected from a given initially localized clump in both isotropic and anisotropic systems. Progressively reducing the area of the phase space spanned initially by the clump reveals that the test-particle orbits are regular in the isotropic beam. However, the orbits are clearly chaotic in the anisotropic beam, this reflecting the complicated network of space-charge waves that arise in the presence of anisotropy.

A unique experiment concerning violent relaxation in beams, conducted in the early 1990's, involved the propagation and merging of five beamlets in a periodic solenoidal transport channel (*34*). The beamlets were initially oriented in a quincunx pattern and were close enough to each other that mutual interactions were important. The beam was nonrelativistic and subject to considerable space-charge forces. *A priori* the beamlets would be expected to 'dissolve' and reappear periodically. However, irrespective of how well the rms beam properties were matched to the transport channel, the beamlets were seen to reappear only once, at a point ~ 1 m from the source. Moreover, mismatched beams led to the formation of an extended halo, with the density of the halo increasing with the degree of mismatch. As discussed in (*34*), detailed simulations with a particle-in-cell code were successful in reproducing the measurements.

The failure of the beamlets to reappear again would seem to reflect a collisionless process that, in effect, causes the particle orbits to lose memory of their initial conditions. To explore how chaotic mixing influences the dynamics of such a manifestly nonequilibrium beam, we redid the simulations using *WARP*. Our simulations differed from the experiment in that we took the transport channel to impart a constant, linear external focusing force, whereas in the experiment the channel comprised a periodic solenoidal focusing lattice. However, our results correlated well with the measurements.

One might expect the strongly time-dependent space-charge potential to drive a large population of globally chaotic orbits. That this is the case is illustrated in Figure 3, which shows that clumps of representative test particles initially localised in phase space grew exponentially to fill much of their accessible phase-space regions. In each case, an initial extremely fast growth rate subsequently gave way to a slower rate, the transition occuring after a distance ~ 5 m, when the beamlets had lost their identity and the phase-space distribution had become rounder.

A recently completed Los Alamos experiment involved the production and measurement of a halo generated in a proton beam that was purposely mismatched to a periodic focusing channel comprised of quadrupole magnets (*35*). The beam energy and current were 6.7 MeV and 75 mA, respectively, which means that the beam was nonrelativistic and space charge was strong. The length of the focusing channel spanned ~ 10 mismatch oscillations. The principal conclusions from this experiment and its accompanying simulations (*36*) were (i) that the phase-space volume of the beam grew in conjunction with the conversion of free energy from mismatch into 'thermal' energy of the beam, and (ii) that parametric resonance was the principal mechanism driving halo formation. However, the quantitative data appears to be sensitive to the exact distribution of the input beam which could not be measured with precision, and the finite sensitivity of the halo diagnostics precluded characterization of the outlying wings of the halo profile. Moreover, the theoretical model provides no prediction of growth rates; the simulations were used to extract this information for comparison with the experiment.

In summary, it would seem evident that physical effects observed in experiments involving charged-particle beams, many of which can be reproduced numerically, bear striking similarities to effects that have been predicted to act in galaxies. In particular, the obvious similarities between Figures 1 and 3 reinforce the expectation that the physics of collective relaxation in galaxies and charged-particle beams is very similar, if not identical.

The smooth potential approximation. The discussion hitherto has assumed implicitly that, viewed 'on large', many-body systems of stars or charges can be approximated by a continuous distribution and a smooth bulk potential. However, this assumption has been questioned in both the galactic (*37*) and accelerator (*38*) communities. To what extent is it really true that there is a smooth continuum limit; and, even assuming that there is such a limit, how many 'particles' must there be before discreteness effects can be ignored? Can one, *e.g.*, treat a realistic accelerator bunch, comprised of $\sim 10^{10}$ charged particles, as a continuous charge distribution?

Numerical computations performed over the last several years, for both gravitating (*39*) and electrostatically (*40*) interacting systems, suggest strongly that, viewed macroscopically, there is a well-defined continuum limit; and that discreteness effects, which exist for finite particle number, can be extremely well modeled, even for individual orbits over comparatively short times, by Gaussian white noise in the context of a Fokker-Planck description. Indeed, one can extract estimates of smooth potential Lyapunov exponents from N-body simulations (*41*).

That a Fokker-Planck description can be justified is nontrivial since the standard derivations (*3*) and most experimental tests focus on the long-time behaviour of orbit ensembles. Even more interesting, however, is the fact that, when applied to chaotic systems, a Fokker-Planck description implies that discreteness effects can have important effects on time scales much shorter than the collisional relaxation time t_R! Discreteness effects can dramatically accelerate diffusion through a complex phase space, both by facilitating transport along the Arnold web (*4*) and, in some cases, by transforming regular orbits into chaotic or vice versa. Indeed, under certain circumstances, *e.g.*, for systems with 'lumps' or large density gradients, discreteness effects could be important on very short time scales even for N as large as 10^{10} (*40*)! However, it is likely that those effects can be adequately modeled by a Fokker-Planck description.

Proposed Beam Experiments with UMER

As discussed above, laboratory experiments conducted to date have not explored explicitly the role of chaotic mixing via Coulomb forces on the evolution of nonequilibrium beams. Yet the combined effects of transient chaos and resonances are the keys toward a full understanding of violent relaxation in both beams and galaxies. Accordingly, we plan a series of experiments to study phase mixing and attendant collisionless relaxation using the University of Maryland Electron Ring (*UMER*), a facility slated to be completed and commissioned during Summer 2003 (*42*). The ring is 11 m in circumference and transports an electron beam with 10 keV kinetic energy, 100 mA current, and 50 μm effective emittance. The nominal bunch contains $\sim 10^{10}$ electrons spanning a volume 1 cm in radius and 3 m in length. The bunch charge, and hence the collective space-charge force, is adjustable over a wide range.

Two beam sources are available, a thermionic cathode and a laser-driven photocathode system. A localized modulation of the thermionic electron current can be applied using a 5 ns laser pulse, a technique that was recently demonstrated by the *UMER* group (*43*). The level of modulation far exceeds what is achievable by grid pulsing alone, and the technique enables the formation of initially localized particle clumps of desired strength and position in the beam. In addition, any desired multibeamlet distribution can easily be created by masking the source beam. The beam is then injected into the ring by means of a magnetic kicker system. The ring confines and steers the beam by means of a magnet system comprising alternating-gradient quadrupoles for transverse confinement, dipoles for bending, and inductive modules for longitudinal confinement. The system is designed so that the beam can be transported over 1 km, a distance that spans some 500 - 1000 plasma periods. Beam diagnostics presently installed on *UMER* include phosphor screens, fast beam position monitors, fast energy analyzers, and both fast and integrated transverse phase space monitors.

Collectively the diagnostics are capable of detailed, time-resolved measurement of the distribution function in the six-dimensional phase space of a single beam particle. The diagnostics measure the same macroscopic observables as generated in the simulations and thereby provide the means for direct comparison. Accordingly, *UMER* serves as a platform for a virtually unlimited range of experiments to explore nonlinear, transient dynamics of Coulomb systems, and our overarching plan is to exploit this capability to access the physics of collisionless relaxation that large charged-particle and self-gravitating systems share in common.

Conclusions

It is clear that, in principle, chaotic phase mixing can serve as a trigger for rapid macroscopic dynamics, including collective relaxation. Moreover, there is substantial numerical evidence that such mixing could play an important role in the evolution of galaxies and charged particle beams. It thus seems natural to look for evidence of chaos and chaotic phase mixing in real laboratory experiments. Unfortunately, it is impossible to perform controlled experiments on self-gravitating systems like galaxies. However, there are strong indications, both theoretical and numerical, that the relevant physics is virtually identical in galaxies and charged particle beams; and, for this reason, it would seem possible – and highly desireable – to use accelerators like *UMER* as laboratories in which to perform indirect tests of the predictions of galactic dynamics.

References and Notes

1. H. E. Kandrup, *Springer Lecture Notes in Physics*, in press (2003) [astro-ph/0212031].
2. S. Chandrasekhar, *Principles of Stellar Dynamics* (University of Chicago, Chicago, 1943).
3. M. N. Rosenbluth, W. M. MacDonald, D. L. Judd, *Phys. Rev.* **107**, 1 (1957).
4. A. J. Lichtenberg, M. A. Lieberman, *Regular and Chaotic Dynamics* (Springer, New York, 1992).
5. M. Hénon, C. Heiles, *Astron. J.* **69**, 73 (1964).
6. D. Lynden-Bell, *Mon. Not. R. Astron. Soc.* **136**, 101 (1967).
7. H. E. Kandrup, M. E. Mahon, *Phys. Rev.* **E 49**, 3735 (1994).
8. D. Merritt, M. Valluri, *Astrophys. J.* **471**, 82 (1996).
9. R. L. Gluckstern, *Phys. Rev. Lett.* **73**, 1247 (1994).
10. C. L. Bohn, *Phys. Rev. Lett.* **70**, 932 (1993); C. L. Bohn, J. R. Delayen, *Phys. Rev.* **E50**, 1516 (1994).
11. H. E. Kandrup, I. M. Vass, I. V. Sideris, *Mon. Not. R. Astron. Soc.* **341**, 927 (2003).
12. See, *e.g.*, P. Grassberger, R. Badii, A. Politi, *J. Stat. Phys.* **51**, 135 (1988); T. M. Antonsen, Z. Fan, E. Ott, E. Garcia-Lopez, *Phys. Fluids* **8**, 3094 (1996).

13. See, *e.g.*, *Space Charge Dominated Beams and Applications of High-Brightness Beams*, ed. S. Y. Lee, AIP Conf. Proc. No. 377 (AIP, New York, 1996).

14. I. M. Kapchinskij, V. V. Vladimirskij, in *Proceedings of the International Conference on High-Energy Accelerators and Instrumentation, CERN 1959*, pp. 274 - 288.

15. F. J. Sacherer, *IEEE Trans. Nucl. Sci.* **NS-18**, 1105 (1971).

16. R. A. Jameson, *Fus. Eng. Design* **32 - 33**, 149 (1996).

17. J. S. O'Connell, T. P. Wangler, R. S. Mills, K. R. Crandall, in *Proceedings of the 1993 Particle Accelerator Conference*, IEEE Catalog No. **93CH3279-7**, pp. 3666 - 3668 (1993).

18. I. Haber, A.W. Maschke, *Phys. Rev. Lett.* **42**, 1479 (1979).

19. C. L. Bohn, I. V. Sideris, *Phys. Rev.* **ST-AB 6**, 034203 (2003).

20. See, *e.g.,* Q. Qian, R. C. Davidson, C. Chen, *Phys. Plas.* **2**, 2674 (1995).

21. D. Merritt, *Pub. Astron. Soc. Pacific* **111**, 129 (1999)

22. See, *e.g.*, T. Lauer *et al*, *Astron. J.* **110**, 2622 (1995).

23. See, *e.g.,* A. El-Zant, I. Shlosman, *Astrophys. J.* **577**, 626 (2002).

24. S. Sridhar, *Mon. Not. R. Astron. Soc.* **238**, 1159 (1989).

25. K. Holley-Bockelmann, J. C. Mihos, S. Siggurdsson, L. Hernquist, *Astrophys. J.* **549**, 862 (2001). See also M. Y. Poon, D. Merritt, *Astrophys. J. Lett.* **568**, 89 (2002).

26. H. E. Kandrup, *Space Sci. Rev.* **102**, 101 (2002).

27. C. Siopis, H. E. Kandrup, *Mon. Not. R. Astron. Soc.* **319**, 43 (2000).

28. J. Binney, *Comments Astrophys.* **8**, 27 (1978).

29. H. E. Kandrup, I. V. Sideris, B. Terzić, C. L. Bohn, *Astrophys. J.*, submitted [astro-ph/0303173].

30. A. M. Fridman, O. V. Khoruzhii, E. V. Polyachenko, *Space Sci. Rev.* **102**, 51 (2002).

31. R. A. Kishek, P. G. O'Shea, M. Reiser, *Phys. Rev. Lett.* **85**, 4514 (2000); R. A. Kishek, C. L. Bohn, I. Haber, P. G. O'Shea, M. Reiser, H. Kandrup, in *Proceedings of the 2001 Particle Accelerator Conference*, IEEE Catalog No. 01CH37268, pp. 151 - 154 (2001).

32. D. P. Grote *et al.*, *Fus. Eng. Design* **32 - 33**, 193 (1996). Distance down the accelerator is viewed as a 'time' coordinate – hence the appellation '$2\frac{1}{2}$-dimensional'.

33. As described in § 4.3 of (*34*), 'matching' involves tuning the bulk properties of a beam to the external focusing forces so as to minimise large scale density readjustments.

34. M. Reiser, *Theory and Design of Charged Particle Beams* (Wiley, NY, 1994), § 6.2.2, *cf.* Fig. 6.10.

35. C.K. Allen *et al.*, *Phys. Rev. Lett.* **89**, 214802 (2002).

36. J. Qiang, *et al.*, *Phys. Rev.* **ST-AB 5**, 124201 (2002).

37. D. Heggie, in *Predictability, Stability and Chaos in N-Body Dynamical Systems*, A. Roy, Ed. (Plenum, New York, 1991), pp. 47 - 62.

38. J. Struckmeier, *Phys. Rev.* **E 54**, 830 (1996).

39. I. V. Sideris, H. E. Kandrup, *Phys. Rev.* **E 61**, 5753 (2002).

40. H. E. Kandrup, I. V. Sideris, C. L. Bohn, *Phys. Rev.* **ST-AB**, submitted [physics/0303022].

41. H. E. Kandrup, I. V. Sideris, *Astrophys. J.* **585**, 244 (2003).

42. R. A. Kishek, *et al.*, AIP Conference Proceedings **642**, 319 (2002).

43. D. W. Feldman, *et al.*, in *Proceedings of the 2001 Particle Accelerator Conference*, IEEE Catalog No. 01CH37268, pp. 2132 - 2135 (2001).

44. This research was supported by NSF AST-0070809, Department of Education Grant G1A62056, and DOE contracts DE-FG02-94ER40855, DEFG02-92ER54178, and DE-AC02-76CH00300.

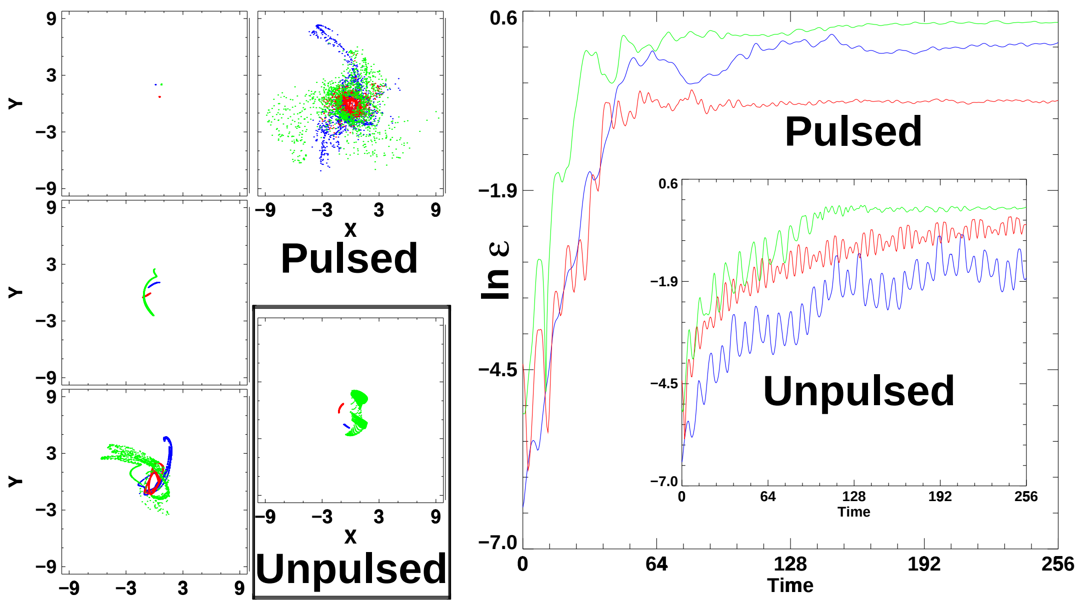

Figure 1: Left hand panels: x and y coordinates for three different clumps of 1600 stars evolved in the time-dependent galactic potential $V(t) = -A/(1 + x^2 + y^2 + z^2)^{1/2}$, with $A(t) = 1 + a \sin \omega t\,/(1 + t/t_0)^2$, for $a = 0.5$, $\omega = 1.25$, and $t_0 = 100$. From top to bottom, the snapshots are at times $t = 0$, 32, 64. The dynamical time $t_D \sim 20$. The clumps had initial size $\delta x = \delta y = 0.04$. Top center panel: The same ensemble at $t = 128$. Bottom center panel: A snapshot at $t = 128$ for the same clumps evolved in a time-independent potential with $A \equiv 1$. Right panel: Natural logarithm of the quantity $\epsilon = (\epsilon_x^2 + \epsilon_y^2)^{1/2}$, defined in terms of the components of emittance ϵ_x and ϵ_y, for the clumps evolved in the time-dependent potential.

Figure 2: Trajectories of 20 test particles in x-y space for the isotropic beam (top); and the anisotropic beam (bottom). The initial clump "emittance" decreases progressively by factors of 10 from left to right.

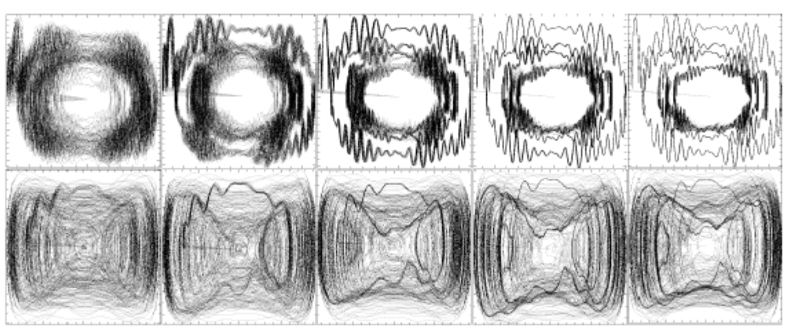

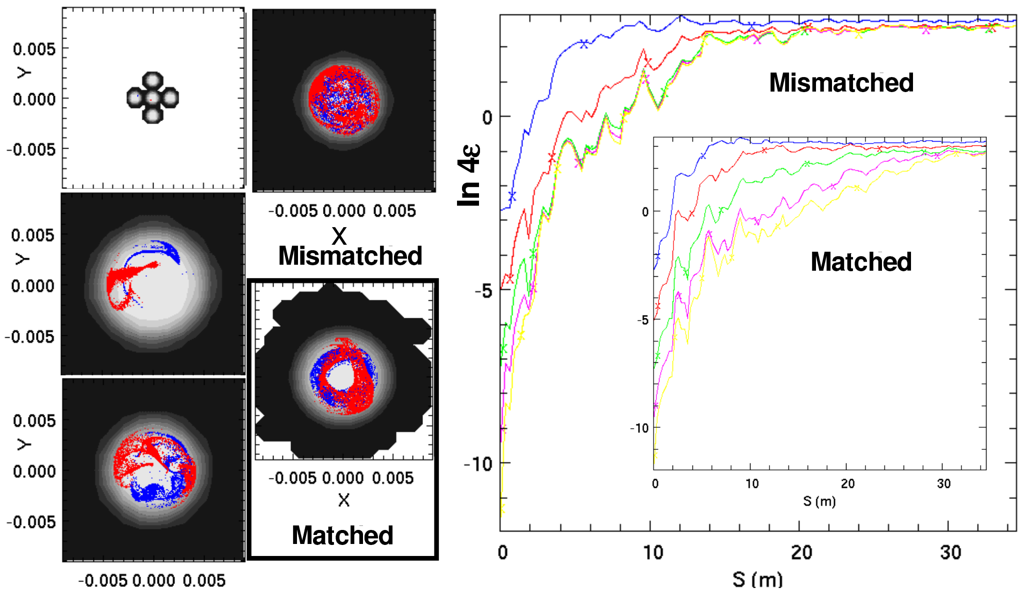

Figure 3: Left hand panels: x and y coordinates, expressed in meters, for two different clumps of 20,000 electrons evolved in the self-consistent potential calculated using *WARP*. The snapshots are taken from a simulation with a mismatched beam, at locations $s = 0.0$, 10.08, 14.4, and 20.16 m along the beam, with the exception of the bottom right snapshot, which is at $s = 31.68$ m from the simulation of the matched beam. The plasma wavelength is 0.47 m and the betatron wavelength is 2.0 m. The initial emittance of each clump is 10^{-5} the full beam emittance ($\epsilon_x = \epsilon_y = 6.48 \times 10^{-10}$ μm.). Right hand panel: Natural logarithm of $4\epsilon_x$ (as defined in the text) for 5 clumps, each sampling a portion of the 'red' clump on the left, but with progressively smaller emittances.