Chaotic Dynamics in Charged-Particle Beams: Possible Analogs of Galactic Evolution

COURTLANDT L. BOHN

Northern Illinois University, Department of Physics, DeKalb, Illinois, USA, and Fermi National Accelerator Laboratory, Batavia, Illinois, USA

ABSTRACT: During the last couple of years of his life, Henry Kandrup became intensely interested in using charged-particle beams as a tool for exploring the dynamics of evolving galaxies. He and I recognized that both galaxies and charged-particle beams can exhibit collisionless relaxation on surprisingly short time scales, and that this circumstance can be attributed to phase mixing of chaotic orbits. The chaos is often triggered by resonances caused by time dependence in the bulk potential, which acts almost identically for attractive gravitational forces as for repulsive electrostatic forces superposed on external focusing forces. Together we published several papers concerning evolving beams and galaxies, papers that relate to diverse topics such as the physics of chaotic mixing, the applicability of the Vlasov-Poisson formalism, and the production of diffuse halos. We also teamed with people from the University of Maryland to begin designing controlled experiments to be done at the University of Maryland Electron Ring. This paper highlights our collaborative findings as well as plans for future investigations that the findings have motivated.

KEYWORDS: Chaos, *N*-body problem, nonlinear dynamics, collisionless, halo

INTRODUCTION AND CONTEXT

Let us consider an example from the charged-particle-accelerator community for which collisionless evolutionary beam dynamics is a central concern. The example is a free-electron laser (FEL) for generating coherent x-rays.[1] This machine would generate the coherent radiation by passing an ultrarelativistic electron beam through a long, periodic array of magnets. The magnets deflect the beam, causing the electrons to "wiggle" and thereby radiate. Accordingly, the magnet array is called a "wiggler". Through a process of self-amplified spontaneous emission, the radiation gradually builds its intensity and becomes coherent as the electron beam traverses the wiggler.

The production of coherent x-rays hinges first on the production of a suitable electron beam. This beam consists of electron packets, called "bunches". Generating high-brightness x-rays requires the bunches to have high charge, at the nanocoulomb (nC) level. In addition, essentially all of the electrons must participate in the lasing process. This means the electrons in each bunch must fit within the optical mode, and in view of the small x-ray wavelength, the optical mode is correspondingly small. This sets a stringent limit on the size of the phase space that the electron bunch can span.

Voice: (815) 753-6473; fax: (815) 753-8565; clbohn@niu.edu

To produce the electron bunches, one uses an external laser to irradiate a photocathode. The profile of the individual bunches and the time structure of the train of bunches hinge on the profile and time structure of the laser radiation. As the bunches traverse the accelerator, various specially designed hardware components further shape their profiles. If all goes according to plan, the electron beam that enters the wiggler is of very high quality, meaning the electrons in each bunch together span a very small volume of phase space. From a practical perspective, "all goes according to plan" means collisionless processes associated with Coulomb self-forces within the bunch, a phenomenon called "space charge", have not seriously degraded the beam.

Space charge works against achieving the required quality of the electron beam. Nonlinear collective forces act rapidly to redistribute the electrons. Because the beam from the source is generally far from equilibrium, it carries free energy that, as it redistributes, irreversibly expands the occupied phase space. Space charge is thus a key concern regarding electron sources. In the laboratory frame, the space-charge force decreases inversely with the square of the beam energy. For the transverse component, this arises from the partial cancellation between the self-magnetic and self-electrostatic forces, while for the longitudinal component, it is due to Lorentz contraction.[2] Nonetheless, for nC bunch charges, space charge remains *the* important dynamic at beam energies up to ~100 MeV, i.e., through a substantial portion of the accelerator.

In addition to degrading the bulk properties of the beam's phase space, space charge also generates *beam halo*, a diffuse population of electrons that lie far from the bunch centroid. This is a special concern regarding accelerators that produce a beam with high average current.

EXAMPLES OF RAPIDLY EVOLVING BEAMS

Rapid beam evolution due to space charge is an observational fact. Consider the example depicted in Fig. 1. This figure concerns beam measurements at the Fermilab photoinjector, a machine that accelerates nC-level bunches to about 15 MeV over a distance of a couple of meters. As Fig. 1 depicts, the quality of the output beam directly depends on its initial density profile. The beam tries to screen externally applied fields, and because its space charge is strong, the screening length is small compared to the beam size. This means the beam 'wants' to be near-uniform, and one thus wants to extract a near-uniform beam from the photocathode to minimize violent evolution. However, one is limited by the quality of the laser beam at the cathode, and this beam can be strongly nonuniform. The consequence of a strongly nonuniform initial density is much more serious beam degradation, which is the message of Fig. 1. This degradation takes place over just a few meters. What is the physical process that leads to such rapid degradation? Due to the large number of electrons in each bunch (some 10^{10}), evolutionary effects of two-body collisions would require kilometers of propagation distance to manifest themselves. The fact that degradation appears over just meters means the physics must necessarily be collisionless; it must be a process of phase mixing.

It was considerations like these that led to my collaboration with Henry Kandrup. Back in 2000 I applied the geometrodynamic theory of Pettini and collaborators[3] to estimate evolutionary time scales in both beams and galaxies associated with chaotic phase mixing. I submitted a paper[4] to Physical Review Letters that ultimately was rejected because one of the referees refused to believe that a beam can be chaotic. As a result, I contacted Henry, and thus began a sustained collaboration by which we, together with our colleagues, investigated many facets of the evolution of nonequilibrium *N*-body ‘Coulomb’ systems (those in which the interparticle force varies inversely with the square of their separation). Henry came to take a strong interest in the physics of beams over the last two years of his life, and he developed a special interest in doing laboratory experiments with beams that would clarify rapid dynamical processes in large stellar systems. He wanted to author papers on beams; unfortunately this desire came to fruition only posthumously.

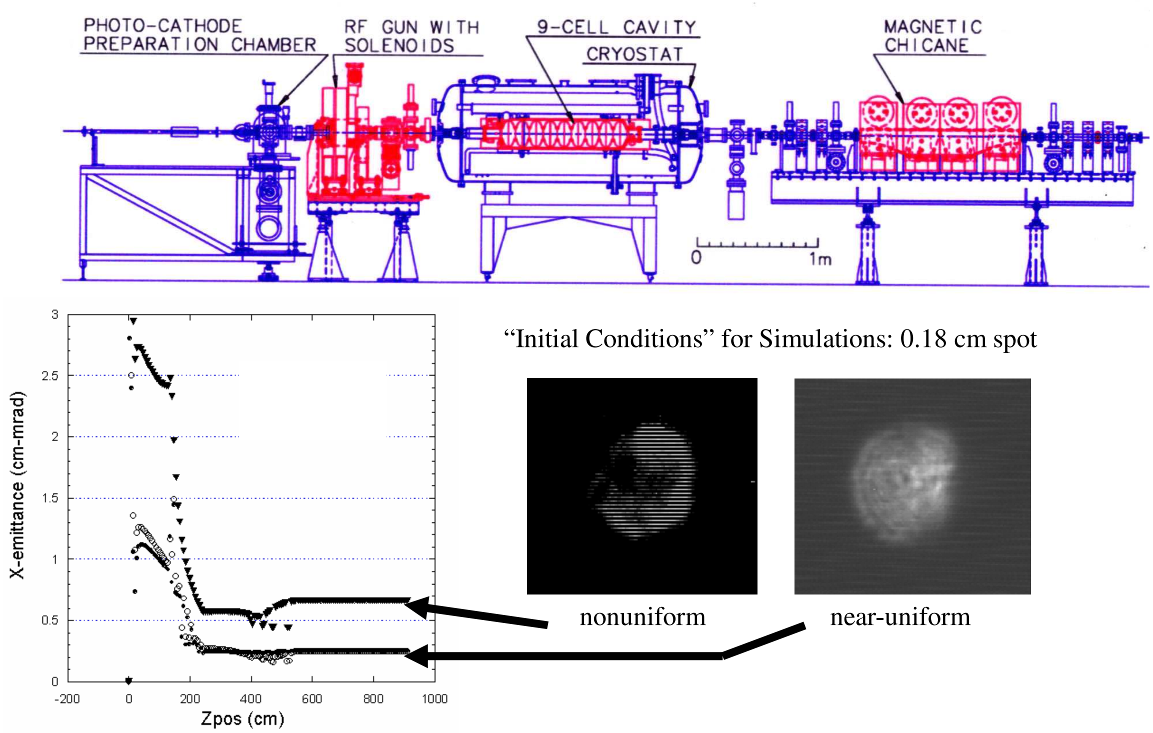

FIGURE 1. (Top) Schematic of Fermilab photoinjector. (Bottom right) Nonuniform and near-uniform laser spots on cathode. (Bottom left) Computational results for transverse *x*-emittance (a measure of phase-space area spanned by the beam for the *x*-axis degree of freedom). The nonuniform beam is seen to degrade twice as badly as the near-uniform beam. Measurements of the *x*-emittance of the output beams agree with the computed values. [Courtesy of Daniel Mihalcea]

A second example of a rapidly evolving beam, one that is perhaps more intuitive and is thoroughly documented, is an experiment conducted some 15 years ago at the University of Maryland that concerned the merging of five nonrelativistic electron beamlets initially constituting a quincunx pattern.[5] These beamlets reappeared only once, after a single dynamical time (orbital period), after which they phase-mixed and completely vanished. This phase mixing, again, took place over just a few meters, i.e., just a few dynamical times. For this beam the collisional relaxation time would correspond to some 100 km. Recent simulations of this experiment[6] have revealed that the underlying dynamics may be understood in terms of phase mixing of chaotic orbits, i.e., “chaotic mixing”.[7]

CHAOTIC MIXING: SOME SALIENT FEATURES

Our work has focused on chaotic orbital dynamics, both in charged-particle beams wherein space charge is important, and in self-gravitating stellar systems (i.e., galaxies). One topic of intensive study was the convergence of an N-body system to the continuum limit in complicated time-independent potentials. Another was the orbital dynamics and associated phase mixing in generic time-dependent potentials. Included in the latter is the formation of diffuse large-amplitude halos in the presence of parametric resonance and colored noise. Progress regarding each topic is summarized in what follows.

Time-Independent Potentials: The Continuum Limit

A standard, often tacit, assumption in theoretical investigations of charged particle beams (and galaxies as well) is that particle correlations are unimportant. With this assumption, one applies the Vlasov-Poisson equations to calculate the distribution function of the particles in the six-dimensional phase space of a single particle.[5, 8] If the system is in static equilibrium, then the distribution function can be expressed as a function of isolating integrals of particle motion in the mean potential. For example, any function of the Hamiltonian derived from the mean space-charge potential is a solution of the Vlasov-Poisson equations, though not all such functions correspond to stable equilibria.[9] A major justification typically provided for using Vlasov-Poisson is that the collisional relaxation time is long compared to, say, the transit time of a beam through any viable linear accelerator (or, for galaxies, long compared to a Hubble time).

As is evident in its derivation from the BBGKY hierarchy,[10] a Vlasov-Poisson system having total charge Q represents a 'continuum limit' in which the number of particles $N \to \infty$ while their individual charges $q \to 0$ such that $qN = Q$. By contrast, real systems contain a finite number of particles N, and real charges have nonzero magnitudes. Vlasov-Poisson is thereby unrealistic, yet for the reasons stated, it is used to describe systems in the real world. The obvious question, therefore, is (1) to what extent do predictions derived from Vlasov-Poisson adequately describe real finite-N systems? A related question is (2) to what extent do predictions derived from a simulation involving $N < \mathsf{N}$ macroparticles adequately describe the real N-body system?

One of our recent investigations[11] concerned how discreteness effects, i.e., granularity, influence the answers to both questions. The investigation invoked three qualitatively different time-independent space charge potentials. The first was a uniform-density ellipsoid such that, in the corresponding smooth potential, the particles execute simple-harmonic-oscillator orbits; the smooth potential is thus completely regular. The second was the same as the first, except a spike of charge was added at the center making the smooth potential entirely chaotic. The third was a configuration of thermal equilibrium[12] for which the smooth potential supports populations of both regular and chaotic orbits. One finding is that, in all three potentials viewed macroscopically, there is a precise sense in which, as the number of macroparticles N increases, trajectories in frozen-N systems converge toward their counterparts in the corresponding smooth potential. For very small N, $<10^4$ or so, the notion of an average bulk potential fails and orbits in frozen-N systems are very different from smooth-potential characteristics. In particular, the usual distinctions between regularity and chaos that exist in a smooth potential seem completely lost. However, for larger N the distinctions become clearer, and they are manifest in the evolution of 'regular' versus 'chaotic' clumps of initially localized particles. Just as for clumps evolved in a smooth potential, the 'emittance' of a regular clump evolved in a frozen-N potential, i.e., the volume of coarse-grained phase space it occupies, was seen to grow as a power law in time, whereas for a chaotic clump it grows roughly exponentially. However, in both cases the growth is more rapid than in the smooth potential. Discreteness effects accelerate emittance growth for *both* regular and chaotic clumps.

In terms of both the statistics of collections of orbits and the complexities of individual orbits, Gaussian white noise was found to mimic well discreteness effects in the context of a Fokker-Planck/Langevin description. This appears true even when considering the short-time behavior of individual orbits. These findings suggest strongly that Langevin simulations are useful for assessing the importance of discreteness effects in real beams for which the constituent number of particles $N \to \mathsf{N}$ is too large to allow honest direct-summation integrations. Accordingly, Langevin simulations were applied to the three potentials considered, and the results indicated that discreteness effects can remain important even if N is very large. Such is the case for the physical beam corresponding to the thermal-equilibrium configuration that was considered, for which $\mathsf{N} > 10^9$. This is especially true when chaotic orbits are present.

Discreteness effects were also seen to trigger transitions between regular and chaotic behavior. The larger the value of N, the longer it takes for these transitions to become important, and they become impossible in the continuum limit. However, for any finite N there appears to be a finite time beyond which it is unsafe to ignore these discreteness-induced transitions. Even if discreteness effects were too weak to facilitate frequent transitions between regularity and chaos, they may nonetheless play an important role in accelerating diffusion through a complex chaotic phase space. Generic smooth potentials admitting both regular and chaotic orbits have phase spaces in which chaotic regions are partitioned by complex structures associated with cantori in two dimensions and the Arnol'd web in three dimensions. Although they are not absolute obstructions, they serve as 'entropy' barriers that impede phase-space transport.[13] However, even very-low-amplitude Gaussian white noise has been shown dramatically to accelerate diffusion through such barriers.[14] In view of the finding that discreteness effects can be modeled as Gaussian white noise, they should likewise present a significant source of accelerated phase-space transport.

The meaning of 'chaos' in regard to beams is somewhat subtle, just as it is in regard to galaxies. Two distinct sources of chaos were found to exist, associated with physics on different scales. Chaos associated with close encounters between individual charges is always present. They cause nearby orbits to diverge exponentially until a distance comparable to the interparticle spacing separates them. In that context beams are *always* 'locally chaotic'. However, if in the continuum limit the bulk potential admits global stochasticity, then the orbits will continue to separate and exponentially fill global regions of phase space. These two distinct epochs of phase mixing are characterized separately by different sets of Lyapunov exponents. Close encounters trigger an exponential separation of nearby trajectories at a rate χ_N, but the separation saturates at only microscopic scales. The bulk potential triggers an exponential separation at a rate χ_S typically much smaller than χ_N (though much larger than the rate of collisional relaxation), but it saturates on macroscopic scales. Hence, when global stochasticity is present in beams, it leads to macroscopic, operationally irreversible evolution.

All of these considerations have practical implications for beams. In particular, discreteness effects can be important in real beams and over real acceleration time scales, thereby vitiating at least to some extent the Vlasov-Poisson methodology. In turn, simulations that correctly account for the full scale of evolutionary mechanisms may require a huge number of macroparticles, possibly comparable to the number of particles in the real beam bunch itself. Inasmuch as the same phenomenology is reflected in stellar systems,[15] the same conclusion would seem to pertain in that context as well, though application of Langevin techniques to infer discreteness effects in, e.g., giant elliptical galaxies with very large N has yet to be done.

Time-Dependent Potentials: Large, Diffuse Halos

Beam loss from impingement of halo particles on accelerator hardware is a major concern for, e.g., high-current light-ion accelerators. Just a tiny impingement, ~1 W/m, could generate radioactivation that would preclude routine, hands-on maintenance.[16] Given a 1 mA, 1 GeV light-ion beam, e.g., for baseline beam parameters of the Spallation Neutron Source (SNS) presently under construction at Oak Ridge National Laboratory, this criterion translates to just 1 in 10^6 particles lost per meter, a quantity that scales linearly with average beam current. Accordingly, a comprehensive understanding of beam-halo formation is imperative.

Early efforts to identify fundamental mechanisms of beam-halo formation centered on using a 'particle-core' model.[17] The basic recognition was that if a uniform-density core is made to pulsate, particles that initially lay outside the core and that resonate with its pulsations could reach large amplitudes and form a 'halo'. This led to the identification of parametric resonance as the essential mechanism of halo formation. A key feature of parametric resonance in the context of the particle-core model is a hard upper bound to the amplitude that a halo particle can reach.[18] Because the particle's orbital frequency is a function of its amplitude, at sufficiently large amplitude the particle falls out of resonance with the core and thereby its amplitude ceases from growing further. The prospect that the beam halo is 'self-collimating' had led to hope that aperture requirements for beamline components might be modest.

However, one feature that is unavoidable in real accelerators, but is commonly overlooked in simulations, is the presence of noise. The noise manifests itself by way of imperfections in the electromagnetic fields external to the beam, which then self-consistently influence the beam's evolving space-charge potential. A charged particle will experience all of the noise inherent to the total potential. Moreover, the noise will generally comprise a superposition of 'colored' noise, i.e., noise with nonzero autocorrelation time. For example, the autocorrelation time of noise in the collective space-charge potential could be short, say of the order of a plasma period, whereas for hardware irregularities/misalignments it could be long, say several betatron (orbital) periods. By generalizing simple particle-core models to include this noise, we showed that the presence of colored noise can boost statistically rare particles to ever-growing amplitudes by keeping them more in phase with the core oscillation. This leads to rapid formation (within ~3-4 global pulsations of the beam) of a large halo and removes the fundamental limit on the halo amplitude predicted by the particle-core model.[19, 20]

We also pointed out the importance of this mechanism in the context of a self-gravitating stellar system for which environmental noise from surrounding galaxies self-consistently influences the dynamics.[19] By considering a perturbed Plummer sphere, a configuration for which the unperturbed collective potential scales as $(1+r^2/3)^{-1/2}$, we found that, notwithstanding it is a restoring force, gravity is so weak that, when combined with the noise, only a relatively tiny oscillatory perturbation suffices to pump stars to very large amplitudes. This suggests that the phenomenology applies generically to nonequilibrium Coulomb systems: colored noise combined with parametric resonance will drive a statistically small number of particles to much larger amplitudes than parametric resonance can do on its own, leaving an extended halo as a byproduct.

The results for the Plummer sphere were generated only as a quick 'existence proof' of the phenomenology in the context of galactic dynamics. We alerted Henry Kandrup of our findings, at which time Henry teamed with Sideris for a more in-depth, systematic study of noise-induced formation of galactic halos.[21] That study was based on spherically symmetric, time-dependent Dehnen potentials[22] subject to low-amplitude, strictly periodic perturbations. Pseudo-random variations in the pulsation frequency were added, these being modeled as colored noise.

As is true for beams, periodic driving generically tends to pump energy into the stars in a galaxy, thus displacing them toward higher energies and larger radii, an effect resulting from a resonant coupling between the driving frequency and the frequencies of the orbits.[23] Even relatively weak driving can have significant effects within just a few tens of dynamical times (orbital periods), and larger amplitudes can account for violent relaxation (at least in principle). Variations in the driving frequency were found not to vitiate the effects of such resonant couplings. In fact, modest variations in frequency tended on average to increase the maximum radii to which orbits are displaced. However, this does not imply that, on average, such variations result in more energy being pumped into the orbits. To the extent that the orbits have all become 'wildly' chaotic, i.e., that they have one or more positive Lyapunov exponents and they are not locked to the (near-constant) driving frequency, allowing for different random frequency variations leads to a Gaussian distribution of maximum energies centered about the maximum energy attained by an orbit subjected to strictly periodic driving. The systematic increase in the average orbital amplitude arises because realistic near-equilibrium systems have phase-space distributions that are monotonically decreasing functions of energy, which means a symmetric spread in energies occasions an increase in the number of larger-amplitude orbits at the expense of smaller-amplitude orbits. More importantly, a random frequency can have a very large impact on at least a small number of stars. In particular, it is statistically probable that a few orbits will experience a noisy variable frequency that will continually keep them more in phase with the global oscillation, thereby displacing them to very large radii.

IMPLICATIONS FOR FUTURE RESEARCH

Accounting for intricacies of space charge requires, in principle, an N-body simulation code. For computations that accurately reproduce details of the beam's phase space, N needs to be large and the computing time correspondingly long. To explore the parameter space of, e.g., photoinjectors, fast codes are needed that also encompass sufficiently accurate models of the beam physics to enable them to be used with confidence. For detailed production runs, inasmuch as a nC bunch charge corresponds to 6.25×10^9 electrons, an "exact" simulation is impractical, and accurate alternatives must be found. Analogous statements apply in the context of galactic dynamics; large galaxies comprise some 10^{11}-10^{12} stars.

Preserving a hierarchy of scales in the time-dependent space-charge potential is dynamically important. To reiterate, recent research has revealed that nonlinear, time-dependent forces commonly establish large populations of globally chaotic orbits in beams that are out of equilibrium, and such orbits can even be present in thermal-equilibrium beams.[12, 24] When present, these chaotic orbits mix exponentially throughout their accessible phase space with a time scale of only a few orbital periods, i.e., very much faster than collisional relaxation. The presence of colored noise due to space-charge fluctuations and/or machine imperfections can, when combined with parametric resonance associated with low-order oscillatory modes, generate much larger halos than would be inferred from parametric resonance alone. Thus, all scales are potentially important to the dynamics.

Wavelets constitute a mathematical tool that is inherently designed to represent a hierarchy of scales. Accordingly, we have recently embarked on the development of a space-charge algorithm based on innovative use of wavelets, as is reported in the accompanying paper by B. Terzić.[25] To start, we are developing a wavelet-based solution of Poisson's equation on a grid. One advantage is simultaneous denoising of macroparticle *N*-body simulations: because artificial noise arising from the use of macroparticles is present on all scales, using wavelets removes most of the noise without altering the beam's inherent structure.

In addition to wavelet denoising, we also plan to develop an algorithm for solving the Vlasov-Poisson system using wavelet decomposition. The idea involves the use of a continuous wavelet transform to decompose the six-dimensional phase space of a single electron in terms of 'mother wavelets' multiplied by time-dependent coefficients. Doing so then reduces the Vlasov-Poisson equations to a set of coupled equations for the time-dependent coefficients. Solving these equations numerically thereby yields the solution for the distribution function, from which any desired beam property can then be calculated. This method may prove to be much faster than large *N*-body simulations.

A key advantage of wavelet-based solutions is the representation of the potential at each time step in terms of a modest number (something like ~100) wavelet coefficients. This *greatly* eases the computational storage of the history of the system: one now does *not* need to store the coordinates and velocities of the (very many!) macroparticles at each time step. The net result from a simulation therefore becomes a realistic time-dependent potential stored compactly. One can then use this potential to integrate efficiently orbits of very many test particles. Doing so enables in-depth studies of mixing in realistic time-dependent systems (an area that has heretofore largely been inaccessible) to include, as just one example, realistic studies of halo dynamics. This is true not only for beams with space charge, but of course also for galaxies.

We have also made significant progress in formulating a rapidly computed measure of orbital chaos, as is reported in the accompanying paper by I. Sideris.[26] It is based on the existence of morphological patterns in the time series associated with the Poincaré section of the orbit. This gives the technique the important advantage of being applicable, without extra computational effort, to any number of degrees of freedom. For every orbit, and over relatively localized time intervals (a few dynamical times), the code correctly identifies the existing patterns and decides if the orbit is regular or chaotic. If it is chaotic, the code distinguishes between orbit segments that are sticky and segments that are wildly chaotic.

The first major application will be to identify transitions of orbits between regular and chaotic behavior in time-dependent potentials. This way we will be able to build a picture of the evolving phase space and thereby increase dramatically our understanding of the underlying dynamics. The ultimate application is to the self-consistent evolution of *N*-body systems. Then, not only will we have a new, efficient method for quantifying the dynamics, but also we will be able to compare the performance of different codes based on how their respective outputs, i.e., phase spaces, evolve.

EPILOGUE

October 18, 2003, the date of Henry Kandrup's death, was a shocking day for all of us who knew and appreciated him. I personally regard his work and innovative insights to be foundational to the understanding of nonequilibrium *N*-body systems. I have tried to document here how these ideas have enabled a much-improved understanding of space charge, with considerable interplay back into the field of galactic dynamics. In learning of his death, I resolved to keep Henry's line of investigation alive and growing. This endeavor would seem to be thus far successful. His former student, Ioannis Sideris, has been my postdoc since Fall 2002, and his former postdoc, Balša Terzić, joined us in April 2004. Henry had a graduate student when he died, Ileana Vass, who has relocated to my group to pursue her Ph.D. on the topic of galactic halos. We are collaborating with Lawrence Berkeley Lab in regard to the pursuit of wavelet techniques; Ilya Pogorelov, another former Kandrup doctoral student, being our principal collaborator. The need to preserve hierarchies of scale, something we learned as we worked with Kandrup, is what motivates our pursuit of wavelet techniques for improved simulation codes. Developing a rapidly computed measure of chaos will permit us to quantify the efficacy of the wavelet code vis-à-vis 'conventional' codes, like those based on Green functions and fast Fourier transforms, for modeling time-dependent *N*-body systems. Although we launched these new pursuits subsequent to Henry's death, we are confident he would have been intensely interested in them.

There is an additional collaboration formed some two years ago, i.e., while Henry was alive, and it is a collaboration he valued highly. He and we teamed with the University of Maryland to begin devising controlled experiments with intense nonrelativistic electron beams as laboratory analogs of galactic systems. Plans remain in place to do these experiments using the University of Maryland Electron Ring, a facility described in the accompanying paper by P. O'Shea.[27] We have documented this collaboration in a paper published in these Annals,[6] a paper that was started with Henry prior to his death, and that we subsequently completed, retaining him as coauthor. We will fondly remember Henry and his influence as we pursue the new experimental and theoretical investigations mentioned herein; we sorely miss him.

ACKNOWLEDGEMENTS

This work was supported by the U.S. Departments of Education under Grant No. P116Z010035, Energy under Grant No. DE-FG02-04ER41323, and Defense under Air Force Contract No. FA9451-04-C-0199.

REFERENCES

1. O'SHEA, P.G. & H.P. FREUND, 2001. Free-electron lasers: status and applications, Science **292**, 1853.
2. CHAO, A.W., 1993. Physics of Collective Beam Instabilities in High Energy Accelerators, Wiley, NY, p. 26.
3. CASETTI, L., C. CLEMENTI, & M. PETTINI, 1996. Riemannian theory of Hamiltonian chaos and Lyapunov exponents, Phys. Rev. E **54**, 5969.
4. BOHN, C.L., 2000. Rapid irreversible mixing in charged-particle beams, FERMILAB-Pub-00/052-T.
5. REISER, M., 1994. Theory and Design of Charged Particle Beams, Wiley, NY, cf. Chapter 6.
6. BOHN, C.L., H.E. KANDRUP, R.A. KISHEK, P.G. O'SHEA, M. REISER, & I.V. SIDERIS, Chaotic collisionless evolution in charged-particle beams and galaxies, these Annals.
7. MERRITT, D. & M. VALLURI, 1996. Chaos and mixing in triaxial stellar systems, Astrophys. J. **471**, 82.
8. DAVIDSON, R.C., 1990. Physics of Nonneutral Plasmas, Addison-Wesley, Redwood City, CA.
9. DAVIDSON, R.C., 1998. Nonlinear stability theorem for high-intensity charged particle beams, Phys. Rev. Lett. **81**, 991.
10. BALESCU, R., 1997. Statistical Dynamics: Matter Out of Equilibrium, Imperial College Press, London.
11. KANDRUP, H.E., I.V. SIDERIS, & C.L. BOHN, 2004. Chaos and the continuum limit in charged particle beams, Phys. Rev. ST Accel. Beams **7**, 014202.
12. BOHN, C.L. & I.V. SIDERIS, 2003. Chaotic orbits in thermal-equilibrium beams: existence and dynamical implications, Phys. Rev. ST Accel. Beams **6**, 034203.
13. LICHTENBERG, A.J. & M. A. LIEBERMAN, 1992. Regular and Chaotic Dynamics, Springer, Berlin.
14. POGORELOV, I.V. & H.E. KANDRUP, 1999. Noise-induced phase space transport in two-dimensional Hamiltonian systems, Phys. Rev. E **60**, 1567.
15. KANDRUP, H.E., 2002. Should elliptical galaxies be idealised as collisionless equilibria?, Space Sci. Rev. **102**, 101.
16. JAMESON, R., 1996. Beam losses and beam halos in accelerators for new energy sources, Fusion Eng. and Des. **32-33**, 149.
17. GLUCKSTERN, R.L., 1994. Analytic model for halo formation in high current ion linacs, Phys. Rev. Lett. **73**, 1247.
18. WANGLER, T.P., K.R. CRANDALL, R. RYNE, & T.S. WANG, 1998. Particle-core model for transverse dynamics of beam halo, Phys. Rev. ST Accel. Beams **1**, 084201.
19. BOHN, C.L. & I.V. SIDERIS, 2003. Fluctuations do matter: large noise-enhanced halos in charged-particle beams, Phys. Rev. Lett. **91**, 264801.
20. SIDERIS, I.V. & C.L. BOHN, 2004. Production of enhanced beam halos via collective modes and colored noise, Phys. Rev. ST Accel. Beams **7**, 104202.
21. SIDERIS, I.V. & H.E. KANDRUP, 2004. Noise-enhanced parametric resonance in perturbed galaxies, Astrophys. J. **602**, 678.
22. TERZIĆ, B. & H.E. KANDRUP, 2004. Orbital structure in oscillating galactic potentials, Mon. Not. R. Astron. Soc. **347**, 957.
23. KANDRUP, H.E., I.V. SIDERIS, B. TERZIĆ, & C.L. BOHN, 2003. Supermassive black hole binaries as galactic blenders, Astrophys. J. **597**, 111.
24. KISHEK, R.A., S. BERNAL, C.L. BOHN, D. GROTE, I. HABER, H. LI, P.G. O'SHEA, M. REISER & M. WALTER, 2003. Simulations and experiments with space-charge-dominated beams, Phys. Plasmas **10**, 2016.

25. TERZIĆ, B. Wavelet-based Poisson solver for use in particle-in-cell simulations (these Annals).
26. SIDERIS, I.V., Characterization of the stochasticity of the phase space: a new fast method (these Annals).
27. O'SHEA, P.G., Nonlinear dynamics experiments with electron beams (these Annals).