Symplectic modeling of beam loading in electromagnetic cavities

Dan T. Abell,* Nathan M. Cook, and Stephen D. Webb

RadiaSoft, LLC, 1348 Redwood Ave., Boulder, Colorado 80304, USA

(Received 3 November 2016; published 22 May 2017)

Simulating beam loading in radio frequency accelerating structures is critical for understanding higher-order mode effects on beam dynamics, such as beam break-up instability in energy recovery linacs. Full wave simulations of beam loading in radio frequency structures are computationally expensive, while reduced models can ignore essential physics and can be difficult to generalize. We present a self-consistent algorithm derived from the least-action principle which can model an arbitrary number of cavity eigenmodes and with a generic beam distribution. It has been implemented in our new Open Library for Investigating Vacuum Electronics (OLIVE).

DOI: 10.1103/PhysRevAccelBeams.20.052002

I. INTRODUCTION

The interactions of charged-particle beams and electromagnetic fields drive a variety of applications and phenomena in accelerator systems. They range from klystrons to beam-loading in radio frequency (rf) cavities to high-order-mode instabilities in energy-recovery linacs. Accurate, self-consistent modeling of these interactions is an important and challenging problem for such applications.

Self-consistent simulations using approaches such as finite-difference time-domain (FDTD) particle-in-cell (PIC) algorithms [1] are computationally demanding and suffer from a variety of numerical artifacts—for example, the numerical Čerenkov instability [2], or dispersive errors induced by the mesh [3,4]—that can muddy the results. Detailed simulations of an electron beam passing through a rf cavity must resolve the current (i.e. the electron beam), and hence require cell dimensions that are small compared to the beam. Resolving the beam across the entire complex geometry can therefore require billions of cells, while the actual source terms are isolated to a mere few thousands of cells. As a consequence, load balancing becomes a serious concern for multi-core simulations.

The discrepancy of spatial scales is not the only challenge faced by FDTD-PIC algorithms. One must also initialize the fields on the numerical mesh in a manner that satisfies the appropriate numerical dispersion relation [3,5,6]. In addition, the presence of multiple bunches means that one must address the discrepancy of temporal scales that effectively prevents using such algorithms to study the effect of high-order modes on beam dynamics.

To bridge multiple scales, one often resorts to reduced models, as, for example, when studying beam loading (matbbu [7]) or beam instability (bi [8]). Such models typically treat the rf cavity as a thin lens, using reduced forms of the Shockley-Ramo theorem [9,10] to compute the energy transfer; and they can advance the field phases analytically once the beam passes. Reduced models are computationally much more efficient than full simulations. They also simplify the diagnostics, because the energy in each cavity mode is a dynamical variable and requires no additional computation to be extracted from the simulation. However, they can fail to demonstrate complicated phase-space structures that result from beam-beam disruption in a collider (as in the eRHIC energy recovery linac-ring design for an electron-ion collider [11,12]) or from a large energy spread (as results from advanced tapering schemes for free-electron lasers). The approximations made in reduced models frequently do not fail gracefully, and they are difficult to expand to next-leading-order, which can limit their versatility.

What we require is an algorithm that has the simplicity and physically intuitive feel of a reduced model, while being extensible for self-consistent simulations. In this paper, we present such an algorithm based on a symplectic map approach to electromagnetic particle-in-mode algorithms [13]. For this approach we use the cavity eigenmodes as our orthonormal basis for the electromagnetic field. We consider only the coupling of j_z to A_z , neglecting the transverse currents. And the modes themselves we compute using either an analytic model or interpolation of numerical data [14]. This approach leads to a fast, extensible model for beam loading with arbitrary numbers of modes and cavity geometries.

II. THE RATIONALE FOR A SPECTRAL TIME-BASED ALGORITHM

Self-consistent updates of the Maxwell equations require that the fields obey the boundary conditions. For FDTD-type algorithms—the canonical example being the

*dabell@radiasoft.net
www.radiasoft.net

Published by the American Physical Society under the terms of the *Creative Commons Attribution 4.0 International license*. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

staggered-Yee scheme [15]—this is not difficult: the basis used to decompose the fields is local in space, and it is easy to compute reflections off boundaries, transient effects, etc. FDTD algorithms are popular also because they can handle complex boundaries, and the electromagnetic part of the algorithm requires only local information to update the fields. However, these algorithms must resolve the smallest features in a simulation, typically the beam itself, and they have inherent dispersive errors that make accurate simulation of beam loading a challenge. Artificial instabilities such as numerical Čerenkov add to the challenge, thus overall making the FDTD approach unsuitable for studying beam loading in most structures.

One possible solution is to use a spectral algorithm that does a global decomposition of the fields as a superposition of the cavity eigenmodes. This is similar in spirit to the Condon method for computing wakefields [16–18]. Each mode frequency is known exactly,¹ meaning that in the absence of a source term it is possible to evolve the fields exactly. Adding the source term—the beam—is a convolution integral of the source currents and charges with each field eigenmode. The beam itself takes up a small volume of the simulation domain. Using the field eigenmodes removes the longest length scale from the simulation, dramatically reducing the computational resources required. It also greatly simplifies source deposition in an electromagnetic simulation, as compared to deposition for FDTD, which can encounter load balancing issues that impede parallel scaling.² This can make the spectral approach competitive for performance with FDTD algorithms. Because beam loading requires exact frequency information and field maps to compute properly, and because sources occupy only a small fraction of the simulation domain—a cause of severe load balancing issues for FDTD-PIC simulations—spectral algorithms are ideal for self-consistent beam loading simulations.

Conventional accelerator tracking codes almost universally use distance along the design orbit of a ring or linac as the independent variable (s -based tracking). When including collective effects, a code must address the issue of simultaneity in the Poisson equation: that leads to codes either adopting time as the independent variable (t -based tracking), or using various tricks to make the particle times simultaneous. Single particle integration through standing wave structures, such as radio frequency cavities, can be accomplished so long as the self-consistent evolution of the fields is neglected.

In a waveguide type structure, which has translational symmetry in the longitudinal direction, s - or t -based

tracking can be used. For systems with s -varying transverse geometries, such as rf cavities or traveling wave tubes, we must use t -based tracking. It is easiest to understand this by looking at how an ultrarelativistic beam ($\beta \approx 1$) would affect an s - versus t -based spectral representation. The fields from the bunch radiate purely transversely, before encountering the cavity surface. If that surface has any tilt to it, the fields reflect off the cavity surface and back into the cavity. In an s -based approach, this would produce a local-in- s increase in the individual cavity eigenmode strengths on the surface, creating surface currents. These surface currents are not self-consistent with any single eigenmode, so the fields would reradiate. An s -based eigenmode representation has a scattering term between multiple eigenmodes. Because the boundary conditions of a cavity are fixed in time, t -based tracking with spatial eigenmodes ensures that the self-consistent beam loading algorithm remains consistent with the boundary conditions.

Simply, an s -based algorithm with s -varying geometry has scattering terms between the eigenmodes, while a t -based algorithm does not. For the problem of a standing wave rf cavity, this means that a t -based algorithm is simpler. For a waveguide with no longitudinal variation in the boundaries, one can use either s - or t -based algorithms.

III. SELF-CONSISTENT BEAM LOADING ALGORITHM

A. Beam Low Lagrangian

We begin with the *Low Lagrangian* [19], devised by Low in 1958 to provide a variational formulation for describing a nonrelativistic ionized gas in an electromagnetic field. It has been used, for example, to study nonlinear waves in plasmas [20]. The extension to a relativistic form is straightforward, and has been used, in contexts related to the present work, to develop other algorithms for plasma simulations [21–23]. In Gaussian c.g.s. units, which we use throughout this work, it has the form

$$\begin{aligned} \mathcal{L} = \int d\mathbf{x}_0 d\mathbf{v}_0 & \left[-mc^2 \sqrt{1 - \left(\frac{d\mathbf{x}}{d\tau} \right)^2} - q\phi(\mathbf{x}, t) \right. \\ & \left. + q \frac{d\mathbf{x}}{d\tau} \cdot \mathbf{A}(\mathbf{x}, t) \right] \psi(\mathbf{x}_0, \mathbf{v}_0) \\ & + \frac{1}{8\pi} \int d\mathbf{x} \left[\left(-\frac{\partial \mathbf{A}}{\partial \tau} - \nabla \phi \right)^2 - (\nabla \times \mathbf{A})^2 \right]. \quad (1) \end{aligned}$$

Here the Boltzmann function $\psi(\mathbf{x}_0, \mathbf{v}_0)$ describes the phase-space distribution of particles at location $(\mathbf{x}_0, \mathbf{v}_0)$; ϕ and $\mathbf{A}$ denote the electromagnetic scalar and vector potentials; and $\tau = ct$ denotes our independent variable. The first integral describes the particles, including both their kinetic energy and their interaction with the

¹Or to within numerical tolerances, if the fields are solved numerically using a frequency domain solver.

²Because spectral algorithms are global instead of local algorithms, they do require an `MPI_AllReduce` to consolidate the source terms. This requirement, however, has not prevented properly-executed spectral codes from scaling well.

electromagnetic field. The second integral describes oscillations of the electromagnetic field.

If we neglect the beam space charge, and consider only the cavity eigenmodes, then $\phi = 0$. Furthermore, our beam has $|\mathbf{dx}_\perp/d\tau| \ll dz/d\tau \sim 1$, while $|\mathbf{A}_\perp| \sim A_z$. We therefore choose to neglect the transverse coupling terms. We do not *have* to do this: all the algorithms and computational results described in this paper can be (and in some cases have been) done without making this simplification; but doing so leaves us with what we call the *beam electromagnetic Low Lagrangian*:

$$\begin{aligned} \mathcal{L}\left(\mathbf{x}, \frac{d\mathbf{x}}{dt}, \mathbf{A}, \frac{\partial \mathbf{A}}{\partial t}\right) \\ = \int d\mathbf{x}_0 d\mathbf{v}_0 \left[-mc^2 \sqrt{1 - \left(\frac{d\mathbf{x}}{d\tau}\right)^2} + q \frac{dz}{d\tau} A_z(\mathbf{x}, t) \right] \\ \times \psi(\mathbf{x}_0, \mathbf{v}_0) + \frac{1}{8\pi} \int d\mathbf{x} \left[\left(\frac{\partial \mathbf{A}}{\partial \tau}\right)^2 - (\nabla \times \mathbf{A})^2 \right]. \end{aligned} \quad (2)$$

Our approximation—in essence neglecting $\mathbf{j}_\perp$ —is made in the interest of speed, as it dramatically reduces the number of required source depositions, which dominates the computation time for self-consistent simulations. For applications in which there is, for example, significant gyrotron motion, one may need to retain the transverse coupling terms.

We assume the cavity eigenmodes are known and form a complete orthonormal basis, so that we may decompose the vector potential in these modes:

$$\mathbf{A} = \sum_{\ell} a_{\ell}(\tau) \mathbf{f}_{\ell}(\mathbf{x}), \quad (3)$$

where the $\mathbf{f}_{\ell}$ denote spatial eigenmodes, and a_{ℓ} the corresponding mode amplitudes. Plugging this form into the Low Lagrangian, we obtain

$$\begin{aligned} \mathcal{L} = \int d\mathbf{x}_0 d\mathbf{v}_0 \left[-mc^2 \sqrt{1 - \dot{\mathbf{x}}^2} + q\dot{z} \sum_{\ell} a_{\ell} \mathbf{z} \cdot \mathbf{f}_{\ell}(\mathbf{x}) \right] \\ \times \psi(\mathbf{x}_0, \mathbf{v}_0) + \frac{1}{8\pi} \sum_{\ell} \left[\dot{a}_{\ell}^2 \int d\mathbf{x} |\mathbf{f}_{\ell}(\mathbf{x})|^2 \right. \\ \left. - a_{\ell}^2 \int d\mathbf{x} |\nabla \times \mathbf{f}_{\ell}(\mathbf{x})|^2 \right], \end{aligned} \quad (4)$$

where overdot denotes differentiation with respect to the proper time τ . In going from (2) to (4), we have used the orthogonality of electric and magnetic field eigenmodes to eliminate the cross-terms in the second integral. As a convenience, we also define the mode inductance $1/L_{\ell} = \frac{1}{4\pi} \int d\mathbf{x} |\nabla \times \mathbf{f}_{\ell}|^2$, and the mode capacitance $C_{\ell} = \frac{1}{4\pi} \int d\mathbf{x} |\mathbf{f}_{\ell}(\mathbf{x})|^2$.

To trace the particles, we introduce macroparticles by decomposing the phase-space density in discrete shapes in the usual manner [13,22–24]:

$$\psi(\mathbf{x}, \dot{\mathbf{x}}) = \sum_{j=1}^{N_{\text{macro}}} w_j \Lambda(\mathbf{x} - \mathbf{x}^{(j)}) \delta(\dot{\mathbf{x}} - \dot{\mathbf{x}}^{(j)}). \quad (5)$$

Here w_j denote the macroparticle weights, δ the Dirac delta function, and Λ the normalized particle shape functions (so that $\int d\mathbf{x} \Lambda = 1$). This decomposition transforms the Lagrangian (4) into a discrete set of coupled macroparticle-electromagnetic-mode Lagrangians:

$$\begin{aligned} \mathcal{L}(\mathbf{x}, \dot{\mathbf{x}}, a_{\ell}, \dot{a}_{\ell}) = \sum_{j=1}^{N_{\text{macro}}} \left[-w_j mc^2 \sqrt{1 - (\dot{\mathbf{x}}^{(j)})^2} \right. \\ \left. + w_j q \dot{z}^{(j)} \sum_{\ell} a_{\ell} F_{\ell}(\mathbf{x}^{(j)}) \right] \\ + \frac{1}{2} \sum_{\ell} \left[C_{\ell} \dot{a}_{\ell}^2 - \frac{1}{L_{\ell}} a_{\ell}^2 \right], \end{aligned} \quad (6)$$

where

$$F_{\ell}(\mathbf{x}^{(j)}) = \mathbf{z} \cdot \int d\mathbf{x} \mathbf{f}_{\ell}(\mathbf{x}) \Lambda(\mathbf{x} - \mathbf{x}^{(j)}) \quad (7)$$

represents the coupling of particle shape to field shape. Each of the individual $\mathbf{x}^{(j)}$ and a_{ℓ} are dynamical variables in this particle-in-mode discrete Lagrangian.

One approach to numerically integrating the resulting equations of motion revolves around using time-discrete Lagrangians [25]. This approach, however, typically yields an implicit algorithm, which is substantially slower than an explicit version. We therefore opt to develop explicit symplectic maps using a canonical formalism.

B. Particle-field Hamiltonian and a split operator approach

We can compute the Hamiltonian corresponding to the Lagrangian from the canonical coordinates and conjugate momenta for the particles and fields:

$$\mathbf{q}_j = \mathbf{x}_j, \quad (8a)$$

$$\mathbf{p}_{\perp}^{(j)} = w_j mc \frac{\dot{\mathbf{x}}_{\perp}^{(j)}}{\sqrt{1 - (\dot{\mathbf{x}}^{(j)})^2}}, \quad (8b)$$

$$p_z^{(j)} = w_j mc \frac{\dot{z}^{(j)}}{\sqrt{1 - (\dot{\mathbf{x}}^{(j)})^2}} - w_j q A_z, \quad (8c)$$

$$Q_{\ell} = a_{\ell}, \quad (8d)$$

$$P_{\ell} = C_{\ell} \dot{a}_{\ell}. \quad (8e)$$

The particle conjugate momenta are the usual, but with the $\mathbf{A}_\perp$ components neglected. The field conjugate momentum includes the mode capacitance, which acts like an effective mass for the mode.

The Legendre transform to compute the coupled particle-field Hamiltonian is

$$\mathcal{H} = \sum_j c \mathbf{p}_j \cdot \dot{\mathbf{q}}_j + \sum_\ell c P_\ell \dot{Q}_\ell - \mathcal{L}. \quad (9)$$

Carrying out this Legendre transform for these variables yields the particle-field Hamiltonian

$$\mathcal{H} = \underbrace{\sum_j c \sqrt{(\mathbf{p}_\perp^{(j)})^2 + \left(p_z^{(j)} - w_j \frac{q}{c} \sum_\ell Q_\ell F_\ell(\mathbf{q}^{(j)}) \right)^2}}_{\mathcal{H}_{\text{pc}}} + \underbrace{w_j^2 m^2 c^2 + \frac{1}{2} \sum_\ell \left[\frac{P_\ell^2}{C_\ell} + \frac{1}{L_\ell} Q_\ell^2 \right]}_{\mathcal{H}_{\text{f}}}. \quad (10)$$

Here $\mathcal{H}_{\text{pc}}$ denotes the *particle-coupling Hamiltonian*, which describes both the particle dynamics and the particle-field coupling; and $\mathcal{H}_{\text{f}}$ denotes the *field Hamiltonian*, which describes the harmonic oscillation of the independent field eigenmodes.

The mode quantities C_ℓ and L_ℓ are determined only up to within a normalization constant. This means that the individual definitions of Q_ℓ and P_ℓ will depend on the normalization convention used for $\mathbf{f}_\ell$. It will therefore prove useful to make a canonical transformation that combines the two into a single scale-invariant quantity, which would be an intrinsic quantity for a given cavity eigenmode.

By defining the canonically conjugate variables

$$Q_\ell = \sqrt{C_\ell} \mathcal{Q}_\ell, \quad (11a)$$

$$P_\ell = \sqrt{\frac{1}{C_\ell}} \mathcal{P}_\ell, \quad (11b)$$

we transform the field and particle-coupling Hamiltonians into

$$\mathcal{H}_{\text{f}} = \frac{1}{2} \sum_\ell \left[\mathcal{P}_\ell^2 + \frac{1}{L_\ell C_\ell} \mathcal{Q}_\ell^2 \right] \quad (12a)$$

and

$$\mathcal{H}_{\text{pc}} = c \sum_j \left[(\mathbf{p}_\perp^{(j)})^2 + \left(p_z^{(j)} - w_j \frac{q}{c} \sum_\ell \mathcal{Q}_\ell \frac{1}{\sqrt{C_\ell}} F_\ell(\mathbf{q}^{(j)}) \right)^2 + w_j^2 m^2 c^2 \right]^{1/2}. \quad (12b)$$

All quantities in these Hamiltonians—including $\mathcal{Q}_\ell$, $\mathcal{P}_\ell$, and $F_\ell/\sqrt{C_\ell}$ —are independent of our choice of normalization for $\mathbf{f}_\ell$. The invariant combination of the mode capacitance and inductance is

$$\Omega_\ell^2 = \frac{1}{L_\ell C_\ell} = \frac{\int d\mathbf{x} |\nabla \times \mathbf{f}_\ell(\mathbf{x})|^2}{\int d\mathbf{x} |\mathbf{f}_\ell(\mathbf{x})|^2}, \quad (13)$$

which yields the eigenmode frequency, Ω_ℓ .

C. Splitting the Hamiltonian

The total particle-field Hamiltonian has no explicit time dependence, so the symplectic map for going from time τ to time $\tau + h$ is simply

$$\mathcal{M}(\tau \rightarrow \tau + h) = \exp(-h : \mathcal{H}_{\text{pc}} + \mathcal{H}_{\text{f}} :). \quad (14)$$

Here we use the colon notation introduced by Dragt [26–28] to denote that argument of the exponential is a Poisson-bracket Lie operator. The full map (14) is difficult to compute, but a symmetric splitting yields a map that is both straightforward to compute and second-order accurate in the step-size h :

$$\mathcal{M}(h) \approx \mathcal{M}_{\text{f}}(h/2) \mathcal{M}_{\text{pc}}(h) \mathcal{M}_{\text{f}}(h/2), \quad (15a)$$

where

$$\mathcal{M}_{\text{f}}(h/2) = \exp\left(-\frac{h}{2} : \mathcal{H}_{\text{f}} : \right), \quad (15b)$$

$$\mathcal{M}_{\text{pc}}(h) = \exp(-h : \mathcal{H}_{\text{pc}} :), \quad (15c)$$

and all maps are independent of the initial time τ . One can evaluate $\mathcal{M}_{\text{f}}$ exactly and $\mathcal{M}_{\text{pc}}$ to second order in h , leading to an overall second-order accurate symplectic integrator for both fields and particles.

1. Field map

The field map $\mathcal{M}_{\text{f}}$, is simply the harmonic oscillator, and it acts only on the field phase-space coordinates:

$$\begin{aligned} \begin{pmatrix} \mathcal{Q}_\ell \\ \mathcal{P}_\ell \end{pmatrix}_{\text{fin}} &= \mathcal{M}_{\text{f}}\left(\frac{h}{2}\right) \circ \begin{pmatrix} \mathcal{Q}_\ell \\ \mathcal{P}_\ell \end{pmatrix}_{\text{ini}} \\ &= \begin{pmatrix} \cos(\Omega_\ell \frac{h}{2}) & \Omega_\ell^{-1} \sin(\Omega_\ell \frac{h}{2}) \\ -\Omega_\ell \sin(\Omega_\ell \frac{h}{2}) & \cos(\Omega_\ell \frac{h}{2}) \end{pmatrix} \begin{pmatrix} \mathcal{Q}_\ell \\ \mathcal{P}_\ell \end{pmatrix}_{\text{ini}} \end{aligned} \quad (16)$$

As is usually the case with second-order splitting, the two half-step field maps can, for simplicity, be combined into a

single full-step map—so long as one retains the half-step maps at the *ends* of the simulation.

2. Particle-coupling map

The particle-coupling map $\mathcal{M}_{\text{pc}}$ is not immediately integrable, but it can be split using a method described by Wu, Forest, and Robin [29] and exploited by Webb *et al.* [13] for a cylindrical electromagnetic algorithm. This method involves tracking each particle with respect to its own proper time, which allows us to rewrite the

Hamiltonian as an effectively nonrelativistic Hamiltonian given by

$$\mathcal{H}_{\text{pc}} = \sum_j \frac{(\mathbf{p}_{\perp}^{(j)})^2 + (p_z^{(j)} - w_j \frac{q}{c} \sum_{\ell} \mathcal{Q}_{\ell} \frac{1}{\sqrt{C_{\ell}}} F_{\ell}(\mathbf{q}^{(j)}))^2}{2w_j m \gamma^{(j)}}, \quad (17)$$

where $\gamma^{(j)}$ denotes the Lorentz factor of the j th macroparticle:

$$\gamma^{(j)} w_j m c^2 = c \left[(\mathbf{p}_{\perp}^{(j)})^2 + \left(p_z^{(j)} - w_j \frac{q}{c} \sum_{\ell} \mathcal{Q}_{\ell} \frac{1}{\sqrt{C_{\ell}}} F_{\ell}(\mathbf{q}^{(j)}) \right)^2 + w_j^2 m^2 c^2 \right]^{1/2} \quad (18)$$

Because the vector potential has no *explicit* time dependence, $\gamma^{(j)}$ is a constant of the motion for this Hamiltonian, and the proper time of each macroparticle can be mapped to the lab time simply by multiplying by $\gamma^{(j)}$. Once written in the form (17), the Hamiltonian can be split and integrated using a pair of half-drift maps for the transverse coordinates, and another map for the longitudinal coordinate:

$$\mathcal{M}_{\text{pc}} \approx \exp \left[-\sum_j \frac{h}{2} : \frac{(\mathbf{p}_{\perp}^{(j)})^2}{2w_j m \gamma^{(j)}} : \right] \exp \left[-\sum_j h : \frac{(p_z^{(j)} - w_j \frac{q}{c} \sum_{\ell} \mathcal{Q}_{\ell} \frac{1}{\sqrt{C_{\ell}}} F_{\ell}(\mathbf{q}^{(j)}))^2}{2w_j m \gamma^{(j)}} : \right] \exp \left[-\sum_j \frac{h}{2} : \frac{(\mathbf{p}_{\perp}^{(j)})^2}{2w_j m \gamma^{(j)}} : \right]. \quad (19)$$

This result is again second-order accurate in h and hence preserves the overall order of the integrator. The leading and trailing half-drifts are particularly simple:

$$\exp \left[-\sum_j \frac{h}{2} : \frac{(\mathbf{p}_{\perp}^{(j)})^2}{2w_j m \gamma^{(j)}} : \right] \circ \left(\mathbf{p}_{\perp}^{(j)} \right) = \left(\mathbf{p}_{\perp}^{(j)} \right) \quad (20)$$

We can now apply the technique in [29] to deal with the vector potential. Lie transformations obey an important similarity transformation property:

$$e^{:f:} e^{:g:} e^{-:f:} = e^{:e^{:f:} g:}. \quad (21)$$

This property allows us to rewrite the central map in (19) as the product of three maps:

$$\begin{aligned} & \exp \left[-\sum_j h : \frac{(p_z^{(j)} - w_j \frac{q}{c} \sum_{\ell} \mathcal{Q}_{\ell} \frac{1}{\sqrt{C_{\ell}}} F_{\ell}(\mathbf{q}^{(j)}))^2}{2w_j m \gamma^{(j)}} : \right] \\ &= \underbrace{\exp \left[-\sum_j \sum_{\ell} w_j \frac{q}{c} : \mathcal{Q}_{\ell} \frac{1}{\sqrt{C_{\ell}}} \int dz F_{\ell}(\mathbf{q}^{(j)}) : \right]}_{\mathcal{A}_z} \underbrace{\exp \left[-\sum_j h : \frac{(p_z^{(j)})^2}{2w_j m \gamma^{(j)}} : \right]}_{\mathcal{D}_z} \\ & \times \underbrace{\exp \left[\sum_j \sum_{\ell} w_j \frac{q}{c} : \mathcal{Q}_{\ell} \frac{1}{\sqrt{C_{\ell}}} \int dz F_{\ell}(\mathbf{q}^{(j)}) : \right]}_{\mathcal{A}_z^{-1}}. \end{aligned} \quad (22)$$

Moreover, each of these maps—the drift $\mathcal{D}_z$ and the transformation $\mathcal{A}_z$ —can be evaluated exactly:

$$\mathcal{D}_z \circ \begin{pmatrix} p_z^{(j)} \\ z^{(j)} \end{pmatrix} = \begin{pmatrix} p_z^{(j)} \\ z^{(j)} + h \frac{p_z^{(j)}}{w_j m \gamma^{(j)}} \end{pmatrix}, \quad (23a)$$

$$\mathcal{A}_z \circ \begin{pmatrix} \mathcal{P}_\ell \\ \mathcal{Q}_\ell \end{pmatrix} = \begin{pmatrix} \mathcal{P}_\ell + \sum_j w_j \frac{q}{c} (\sqrt{C_\ell})^{-1} \int dz F_\ell(\mathbf{q}^{(j)}) \\ \mathcal{Q}_\ell \end{pmatrix} \quad (23b)$$

$$\mathcal{A}_z \circ \begin{pmatrix} \mathbf{p}^{(j)} \\ \mathbf{q}^{(j)} \end{pmatrix} = \begin{pmatrix} \mathbf{p}^{(j)} + \nabla_{\mathbf{q}} \sum_\ell w_j \frac{q}{c} (\sqrt{C_\ell})^{-1} \mathcal{Q}_\ell \int dz F_\ell(\mathbf{q}^{(j)}) \\ \mathbf{q}^{(j)} \end{pmatrix}. \quad (23c)$$

Because the algorithm is based on explicit symplectic maps, it is possible to implement each of these maps as a single function, and the update sequence is just a series of function calls. Once each of these maps is implemented, it is straightforward to begin simulations.

Before going on to discuss our numerical results, we should clarify a point that may seem mysterious: namely, how does the Lorentz γ factor change? The method given by Wu, Forest, and Robin is described in the context of magnetostatic systems. Moreover, the absence of (explicit) time dependence in the vector potential that appears in the Hamiltonian $\mathcal{H}_{\text{pc}}$ of (17) is consistent with a magnetostatic system, which cannot change the Lorentz factor. And indeed the map $\mathcal{M}_{\text{pc}}$ derived from the Hamiltonian $\mathcal{H}_{\text{pc}}$ does not do so.

To resolve the apparent mystery, recall that $\mathcal{H}_{\text{pc}}$ denotes just the particle-coupling term in the full Hamiltonian: the remaining part is the field Hamiltonian, given in (12a). When stepping through the full map (15) to update the system for a single time step, one must apply the field map $\mathcal{M}_{\text{f}}$ of (16), which rotates the field coordinates, $\mathcal{Q}_\ell$, and momenta, $\mathcal{P}_\ell$; and it is effectively *this* operation that updates—see (18)—the Lorentz γ factor.

IV. NUMERICAL RESULTS

To test the algorithm described in this paper, we considered a number of configurations. In all tests, we considered a rectangular pillbox cavity with a handful of test macroparticles. We list the specific parameters in

TABLE I. Simulation parameters for testing the algorithm.

Parameter	Quantity
Bunch Charge	192 nC
e -beam Energy	25.5 MeV
Total Bunch Energy	4.8×10^7 ergs
Accelerating Gradient	5.1 MV/m
Cavity Length	10 cm
Cavity Width $\times$ Cavity Height	50 cm $\times$ 40 cm
Fundamental Frequency	$2\pi \times 959$ MHz

Table I. In all simulations, we considered four cavity modes: an accelerating mode, two transverse dipole modes, and a single transverse quadrupole mode.

An artificially high bunch charge was selected to demonstrate the field energy loss—otherwise the total energy of the bunch is so much smaller than the total energy in the fields that the variations are not easy to discern.

FIG. 1. Energy gain and eigenmode energy decomposition for the case of one macroparticle on axis.

In the first test, we considered a configuration with a single macroparticle representing the entire bunch, accelerated from $\gamma = 50$ to $\gamma = 51$. The particle is on-axis, so no higher order modes should be excited. The result is shown in Fig. 1. As can be seen, the energy conservation is excellent, and none of the higher order modes are excited.

To test the excitation of higher order modes, we split the beam into four macroparticles—one on-axis; one offset in x and another offset in y to excite dipole modes; and one offset along the x - y diagonal to excite both dipole and quadrupole modes. As shown in Fig. 2, the total energy is conserved to high precision, and all the modes are excited, although the fundamental mode sees the most variation.

Because this is a symplectic integrator, we expect the total energy to be well-behaved for arbitrarily long time. To test this, we considered a much longer rectangular pillbox, 125 cm long, with the same transverse eigenmodes. As shown in Fig. 3, the algorithm retains stable energy behavior for longer time scales. This captures oscillatory behavior as the beam accelerates and decelerates, and the particles slip in phase for the higher order modes.

When simulating considerably longer cavities, ~ 180000 cm in length, we detected a slow spurious decay

FIG. 2. Total particle-field energy and individual mode energies for the four-particle configuration.

FIG. 3. Total particle-field energy and individual mode energies for the four-particle configuration in a longer cavity.

in the total energy. Over this length, the error was $\sim 10^{-6}\%$ total, as shown in figure 4. We have identified the origin of this spurious decay as numerical roundoff error in the field rotation matrix: for our implementation, the floating-point

FIG. 4. Total system energy evaluated for 1000 particles and 6 modes in a cavity of length ~ 1.8 km. 1 million steps of length $\frac{\lambda}{30}$ produced relative energy $(\frac{\Delta E}{E})$ losses on the order of 10^{-8} .

value of the determinant was very slightly less than 1. Over many time steps this can lead, depending on the roundoff error, to spurious cooling or heating of the system; but for a single bunch pass this is almost undetectable. Furthermore, when advancing the fields in between bunch passes, it suffices to perform a single matrix multiplication to advance the eigenmodes. Doing this reduces the round-off-induced energy change in multibunch and multipass simulations performed when studying, for example, energy-recovery linacs.

V. ALGORITHM PERFORMANCE

To demonstrate the second-order behavior of the algorithm, we performed a simple test with a rectangular $10\text{ cm} \times 10\text{ cm} \times 20\text{ cm}$ cavity instantiated with three modes. We propagated four particles through the cavity for a single step and evaluated the change in the system Hamiltonian as the step size was varied. The resulting behavior is shown in Fig. 5. For decreasing step size, h , the error falls as h^3 , indicative of a second-order algorithm. For step sizes below 1 mm, the single-step error quickly approaches machine precision, on the order of 10^{-14} relative error.

Initial benchmarks for the algorithm's performance demonstrate the scalability to many hundreds of modes and thousands of particles. We performed 100 steps of fixed step size on a single processor with varying macroparticle number and mode number to illustrate the single-step evaluation time per-macroparticle-per-mode; the results are displayed in Fig. 6. Although Python-induced overhead associated with array creation hinders performance for small samples, the step times settle near $\Delta t \sim 0.5\text{ }\mu\text{s}$ per-particle-per-mode. When particle and mode numbers cause array size to exceed cache size, a

FIG. 5. The change in Hamiltonian evaluated after a single step reveals 2nd order scaling with step size h .

FIG. 6. Timings for a single step weighted per particle and per mode are shown for upwards of 256 modes and 12800 particles. For large systems, we find $\Delta t \sim 0.5\text{ }\mu\text{s}$ per-particle-per-mode.

small drop in performance occurs, as can be seen in the top right of the figure.

VI. CONCLUSIONS

We have presented a new approach to simulating beam loading in electromagnetic cavities. The approach is based on a particle-in-mode symplectic map approach to electromagnetic charged-particle simulations. The algorithm is suitable for modeling a variety of vacuum electronic devices, including radio frequency cavities, traveling wave tubes, and klystrons.

For applications in energy recovery linacs, the beam remains relativistic at all times, effectively suppressing space-charge forces. We therefore chose to ignore it in our current implementation. But it is possible to include space-charge within the framework of our algorithm, and we mention two possible approaches: One possibility is to stick with the Weyl gauge ($\phi = 0$) and introduce additional high-order modes to represent the space-charge fields. When the bunch is very short, the required modes will correspond to very-high-frequency, essentially free-space, modes that are broadly separated from the cavity modes. This approach will work well if the space-charge variation is predominantly longitudinal, or if one includes the $\mathbf{j}_\perp$ terms we neglected in transforming from the Lagrangian of (1) to that of (2). The other possibility is to use a spectral algorithm described in a paper by one of the present authors [24].

Another issue we did not address in our current implementation is the use of realistic rf cavities. To model realistic cavity eigenmodes, one may use the technology of *generalized gradients* described in Ref. [14]. There the fields are computed on the basis of generalized gradients, which one computes as integrals over a surface enclosing the relevant volume. This surface integration acts to damp any imprecision in the simulated fields, and it means that the field

evaluations—including the derivatives—needed for the maps in (23) can be computed to high accuracy.

The algorithm presented here avoids the various pitfalls of both the finite-difference time-domain electromagnetic particle-in-cell and the reduced-model approaches that are typically applied to this problem.

Because the algorithm is directly spectral, there are no grid aliasing artifacts that can result in the numerical instabilities and dispersive errors that tend to infect the full electromagnetic PIC simulations. By using a modal decomposition, it is straightforward to identify, directly from the dynamical data, which cavity eigenmodes are excited and how they evolve—there is no need for post-processing of the simulation data to extract the eigenmodes. The approach also solves the load balancing problem, as the field data is global, and the particles can be evenly distributed across multiple processors. The approach is also much faster than the fully self-consistent approach: our early simulations of the single particle traversing the pillbox cavity, including the time to generate the figures, required less than a second using an implementation that performed symbolic mathematical computations and is far from optimized. See Fig. 6 for performance data based on a more sophisticated implementation.

The symplectic particle-in-mode (SymPIM) algorithm described in this paper in many ways resembles the reduced models, but it is distinct in a number of important ways. Because the cavity is not treated as a thin element, it is possible to model the complex phase-space of realistic beams, such as those with large energy spreads, or those disrupted by beam-beam collisions or various wakefield instabilities. The SymPIM algorithm thus lies between the fully self-consistent approach—which includes all resolvable cavity modes at great computational expense—and the reduced-model approach—which neglects many details of the beam distribution and how it evolves within the cavity. This makes the SymPIM algorithm suitable for determining the validity of reduced models, as well as for modeling those cases where the reduced models do not apply.

ACKNOWLEDGMENTS

This work was supported in part by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics under Award No. DE-SC0015211. In addition, S. D. W. thanks the sponsorship of the Air Force Office of Scientific Research, Young Investigator Program, under Contract No. FA9550-15-C-0031.

- [1] C. K. Birdsall and A. B. Langdon, *Plasma Physics via Computer Simulation* (McGraw-Hill, New York, 1985).
- [2] B. B. Godfrey, Numerical Cherenkov instabilities in electromagnetic particle codes, *J. Comput. Phys.* **15**, 504 (1974).
- [3] L. N. Trefethen, Group velocity in finite difference schemes, *SIAM Rev.* **24**, 113 (1982).

- [4] A. Taflov and S. C. Hagness, *Computational Electrodynamics*, 2nd ed. (Artech House, Boston, 2000).
- [5] J.-L. Vay, C. G. R. Geddes, E. Cormier-Michel, and D. P. Grote, Modeling laser wakefield accelerators in a Lorentz boosted frame, [arXiv:1009.2727](#).
- [6] B. B. Godfrey and J.-L. Vay, Numerical stability of relativistic beam multidimensional PIC simulations employing the Esirkepov algorithm, *J. Comput. Phys.* **248**, 33 (2013).
- [7] K. B. Beard, L. Merminga, and B. Yunn, TJNAF Report No. JLAB-TN-02-044, 2003.
- [8] Ivan V. Bazarov, bi—Beam instability code, Technical Report (Cornell University, New York, 2003).
- [9] W. Shockley, Currents to conductors induced by a moving point charge, *J. Appl. Phys.* **9**, 635 (1938).
- [10] S. Ramo, Currents induced by electron motion, in *Proc. IRE* **27**, 584 (1939).
- [11] Y. Hao, V. N. Litvinenko, and V. Ptitsyn, Mitigation of the kink instability in the energy recovery linac based electron-ion collider, *Phys. Rev. ST Accel. Beams* **16**, 101001 (2013).
- [12] Y. Hao, V. N. Litvinenko, and V. Ptitsyn, Beam-beam study of ERL-based eRHIC, [arXiv:1410.6000](#).
- [13] S. D. Webb, D. T. Abell, N. M. Cook, and D. L. Bruhwiler, A spectral symplectic algorithm for cylindrical electromagnetic plasma simulations, [arXiv:1609.05095](#).
- [14] D. T. Abell, Numerical computation of high-order transfer maps for rf cavities, *Phys. Rev. ST Accel. Beams* **9**, 052001 (2006).
- [15] K. S. Yee, Numerical solution of initial boundary value problems involving Maxwell's equations in isotropic media, *IEEE Trans. Antennas Propag.* **14**, 302 (1966).
- [16] E. U. Condon, Forced oscillations in cavity resonators, *J. Appl. Phys.* **12**, 129 (1941).
- [17] P. B. Wilson, Stanford Linear Accelerator Center Technical Report No. SLAC-pub-4547, 1989.
- [18] A. W. Chao, *Physics of Collective Beam Instabilities in High Energy Accelerators*, Wiley Series in Beam Physics and Accelerator Technology (John Wiley & Sons, New York, NY, 1993).
- [19] F. E. Low, A Lagrangian formulation of the Boltzmann-Vlasov equation for plasmas, *Proc. R. Soc. A* **248**, 282 (1958).
- [20] R. L. Dewar, A Lagrangian theory for nonlinear wave packets in a collisionless plasma, *J. Plasma Phys.*, **7**, 267 (1972).
- [21] E. G. Evstatiev and B. A. Shadwick, Variational formulation of particle algorithms for kinetic plasma simulations, *J. Comput. Phys.* **245**, 376 (2013).
- [22] A. B. Stamm, B. A. Shadwick, and E. G. Evstatiev, Variational formulation of macroparticle models for electromagnetic plasma simulations, *IEEE Trans. Plasma Sci.* **42**, 1747 (2014).
- [23] B. A. Shadwick, A. B. Stamm, and E. G. Evstatiev, Variational formulation of macro-particle plasma simulation algorithms, *Phys. Plasmas* **21**, 055708 (2014).
- [24] S. D. Webb, A spectral canonical electrostatic algorithm, *Plasma Phys. Controlled Fusion* **58**, 034007 (2016).

- [25] J.E. Marsden and M. West, Discrete mechanics and variational integrators, *Acta Numer.* **10**, 357 (2001).
- [26] A.J. Dragt and J.M. Finn, Lie series and invariant functions for analytic symplectic maps, *J. Math. Phys. (N.Y.)* **17**, 2215 (1976).
- [27] A.J. Dragt, Lectures on nonlinear orbit dynamics, *AIP Conf. Proc.* **87**, 147 (1982).
- [28] A.J. Dragt, Lie methods for nonlinear dynamics with applications to accelerator physics (2016), latest version available at <http://www.physics.umd.edu/dsat/>.
- [29] Y.K. Wu, É. Forest, and D. S. Robin, Explicit symplectic integrator for s -dependent static magnetic fields, *Phys. Rev. E* **68**, 046502 (2003).