



Integrating a ponderomotive guiding center algorithm into a quasi-static particle-in-cell code based on azimuthal mode decomposition

Fei Li ^{a,*}, Weiming An ^c, Frank S. Tsung ^b, Viktor K. Decyk ^b, Warren B. Mori ^{a,b,*}

^a Department of Electrical Engineering, University of California Los Angeles, Los Angeles, CA 90095, USA

^b Department of Physics and Astronomy, University of California Los Angeles, Los Angeles, CA 90095, USA

^c Department of Astronomy, Beijing Normal University, Beijing 100875, China



ARTICLE INFO

Article history:

Received 22 March 2022

Received in revised form 20 August 2022

Accepted 2 September 2022

Available online 7 September 2022

Keywords:

Laser wakefield accelerators

Ponderomotive guiding center model

Quasi-static approximation

Particle-in-cell algorithm

ABSTRACT

High fidelity modeling of plasma based acceleration (PBA) requires the use of three dimensional, fully nonlinear, and kinetic descriptions based on the particle-in-cell (PIC) method. In PBA an intense particle beam or laser (driver) propagates through a tenuous plasma whereby it excites a plasma wave wake. Three-dimensional PIC algorithms based on the quasi-static approximation (QSA) have been successfully applied to efficiently model the interaction between relativistic charged particle beams and plasma. In a QSA PIC algorithm, the plasma response to a charged particle beam or laser driver is calculated based on forces from the driver and self-consistent forces from the QSA form of Maxwell's equations. These fields are then used to advance the charged particle beam or laser forward by a large time step. Since the time step is not limited by the regular Courant-Friedrichs-Lowy (CFL) condition that constrains a standard 3D fully electromagnetic PIC code, a 3D QSA PIC code can achieve orders of magnitude speedup in performance. Recently, a new hybrid QSA PIC algorithm that combines another speedup technique known as an azimuthal Fourier decomposition has been proposed and implemented. This hybrid algorithm decomposes the electromagnetic fields, charge and current density into azimuthal harmonics and only the Fourier coefficients need to be updated, which can reduce the algorithmic complexity of a 3D code to that of a 2D code. Modeling the laser-plasma interaction in a full 3D electromagnetic PIC algorithm is very computationally expensive due the enormous disparity of physical scales to be resolved. In the QSA the laser is modeled using the ponderomotive guiding center (PGC) approach. We describe how to implement a PGC algorithm compatible for the QSA PIC algorithms based on the azimuthal mode expansion. This algorithm permits time steps orders of magnitude larger than the cell size and it can be asynchronously parallelized. Details on how this is implemented into the QSA PIC code that utilizes an azimuthal mode expansion, QPAD, are also described. Benchmarks and comparisons between a fully 3D explicit PIC code (OSIRIS), as well as a few examples related to laser wakefield acceleration, are presented.

© 2022 Elsevier Inc. All rights reserved.

* Corresponding authors.

E-mail addresses: lifei11@ucla.edu (F. Li), mori@physics.ucla.edu (W.B. Mori).

1. Introduction

Laser wakefield acceleration (LWFAs) [1,2] has attracted much interest and shown remarkable progress in the past four decades [3]. Due to the capability of producing large accelerating gradients and microscopic accelerating structures, LWFAs have immense promise for building compact and affordable accelerators for applications in high-energy physics and advanced radiation sources. In LWFA, when a relativistic laser pulse propagates through an under-dense plasma, the background electrons will be expelled sideways and forward by the ponderomotive force of laser pulse after which they are attracted back towards the axis by the Coulomb force from the immobile ions. They then execute a multi-dimensional plasma oscillation creating a wakefield. Charged particle beams riding at the correct phase of a wakefield (moving at close to the speed of light) will be continuously accelerated.

Modeling the complex physics in a LWFA requires self-consistently solving Maxwell's equations and the relativistic equations of motion for plasma and trailing beam particles. Kinetic simulations based on the particle-in-cell (PIC) algorithm [4–6] have and continue to play a critical role in understanding and utilizing the LWFA processes. However, carrying out full 3D explicit start-to-end PIC simulations for LWFAs is very challenging mainly due to the need to resolve the shortest spatial and temporal scales (typically the laser wavelength and period) and the huge disparity between these smallest physical scales and the acceleration length L_{acc} (\sim mm to \sim m). We note that for some beam loading scenarios [7,8] the trailing beam can be very narrow and this transverse scale is the smallest spatial scale needed to be resolved. For simplicity when discussing computational needs we assume that the smallest transverse scale is much larger than the laser wavelength. Furthermore, there can be subtle numerical issues which can affect the fidelity of these simulations, e.g., numerical Cherenkov instability [9–12], numerical dispersion, and spurious forces [13,14]. To mitigate or eliminate these deleterious effects often requires finer spatial-temporal resolution or more sophisticated numerical schemes, which inevitably increases the computational cost. For example, even with the today's most powerful supercomputers [15], simulating a LWFA driven by a peta-watt laser pulse with ten-centimeter- to meter-class acceleration distance needs $\sim 10^8$ CPU hours (assuming the use of the moving window [16]) and therefore it still lacks feasibility with 3D explicit PIC codes. To achieve highly efficient modeling of LWFA processes, various speedup techniques based on reduced models have been developed, such as a Lorentz boosted frame technique [17,18], quasi-static approximation [19], ponderomotive guiding center model [19–21] and azimuthal Fourier decomposition [22–24]. In some cases these methods have been combined.

The ponderomotive guiding center (PGC) approximation is based on the idea of averaging the motion of a particle in the laser and wakefields over the laser frequency. This permits solving for the envelope of the laser and then using the laser envelope to calculate an average or ponderomotive force on the plasma electrons. This approximation works when there is no back-scattered light and there are no relativistic particles that reside within the laser pulse. With this time-averaging, the smallest physical scales needed to be resolved now become the spot size and wavelength of the wake which are typically the laser pulse spot size and duration. Therefore, the PGC approximation can provide speedups on the order of the factor $\sim (\omega_0/\omega_p)^2$ against the standard PIC simulations, where ω_0 and ω_p are the laser and plasma frequency respectively. We typically have $\omega_0/\omega_p \gg 1$ for LWFA problems of interests and thus the speedup provided by the PGC approximation is considerable. The PGC makes the computational needs of LWFA simulations similar to those of a PWFA per c/ω_p of driver propagation distance. The Lorentz boosted frame method can achieve speedups on the order of a factor of $\sim \gamma_f^2 = (1 - \beta_f^2)^{-2}$ by modeling the physics in a boosted frame moving at a speed of $\beta_f c$. In the boosted frame, the laser frequency is Doppler-shifted from ω_0 to $\omega'_0 = \omega_0/\gamma_f(1 + \beta_f)$ and the plasma frequency $\omega_p \equiv \sqrt{\frac{4\pi e^2 n_p}{\gamma_f m_e}}$ is a Lorentz invariant. Thus, the cell size and time step used to resolve the laser wavelength are increased (the number of cells remains the same as the number of laser cycles is a Lorentz invariant). In this frame the plasma length (acceleration length) is also Lorentz contracted by a factor γ_f^{-1} . The maximum speedup is achieved when $\omega'_0 \sim \omega'_p$, which leads to a speedup that scales as $\gamma_f^2 \sim (\omega_0/2\omega_p)^2$. The Lorentz boosted frame technique also requires that there is no back-scattered radiation and all forward moving modes have similar phase velocities. For both the PGC and Lorentz boosted frame methods it is assumed that the number of particles per cell is kept fixed when estimating the speedups. Therefore, the two approaches can provide similar speedups. Although the PGC can reduce the disparity between the minimum and maximum scales to be resolved, the use of a standard explicit PIC simulation is still subject to the limitation of time step size, which is a necessary condition to guarantee the numerical stability [known as the Courant–Friedrichs–Lowy (CFL) condition]. We note that even without the CFL constraint the time steps need to be small enough that particles do not move a distance approaching the smallest physical scale. The quasi-static approximation (QSA), initially presented as an analytical approach to study short-pulse laser interacting with plasma [19], was developed to treat the disparity in the physical scale between the evolution of lasers/charged particle beams and the plasma response. The QSA is essentially a multi-scale method that can separate the fast-varying plasma response, characterized by ω_p^{-1} , from the slow-varying evolution of lasers/charged particle beams, characterized by Rayleigh length, $z_R \equiv \frac{\pi w_0^2}{\lambda_0}$ or betatron length, $\beta^* \equiv \frac{\sigma_0^2}{\epsilon}$, where w_0 and σ_0 are the spot sizes of the laser of beam respectively, λ_0 is the wavelength of the laser, and ϵ is the geometrical emittance of the beam. The QSA thus leads to speedup of $\beta^* \omega_p/c$ for PWFA simulations and an additional $z_R \omega_p/c$ speedup over the PGC method in standard PIC codes for LWFA simulations. We note that the PGC was first developed within the QSA [19] and was then implemented into the fully electromagnetic PIC description [20].

Numerically, a QSA PIC code [19,25–31] calculates the plasma response using QSA based field equations by assuming the laser or particle beam is static and then the resultant plasma-induced fields are used to advance the laser or particle beam by a large time step. The QSA field equations exclude radiative fields thus permitting time steps in excess of the CFL condition while still resolving the evolution of the driver. For laser drivers, the QSA based PIC codes require a time averaging approach to model the laser and its associated ponderomotive force. This was implemented into a 2D r-z QSA based PIC code as described by [19] and a 3D QSA based code as described by [26,30].

Another technique to boost the computational performance is known as the azimuthal Fourier decomposition [22–24] (also referred to as a quasi-3D or cylindrical mode expansion in the literature). In this technique, all the electromagnetic fields, and the charge and current density are expanded into a truncated Fourier series in ϕ (the azimuthal coordinate). The maximum number of Fourier harmonics that are kept is determined by the degree of asymmetry of the physical problem. This method turns the original 3D field updates into a series of 2D updates by only solving the complex Fourier coefficients defined on the 2D r-z grid, and thus its computational complexity is greatly reduced. Thus it is a hybrid approach that uses a PIC description on an $r - z$ grid and a gridless description in ϕ . This approach has been successfully implemented into an explicit PIC framework [22–24].

Very recently, it was described how to merge the quasi-3D approach and the QSA. The result is the PIC code QPAD [29] which is based on the numerical workflow of QuickPIC [26,27]. It can enhance the computation efficiency by orders of magnitude compared to the full 3D standard and 3D QSA PIC codes. Implementing the PGC algorithm into a quasi-3D QSA based PIC code such as QPAD will enable simulating meter long LWFA simulations (that require peta-watt scale lasers) on small scale parallel computers (including desktop) while maintaining high fidelity so long as self-injection is not occurring and the laser does not overlap with the relativistic electron beam. The QSA and its neglect of radiative fields also has the advantage that it is more straightforward to handle small cell sizes in the transverse direction as can be needed for some nonlinear beam loading scenarios.

In this paper, based on the existing PGC theoretical model [19], we extend the PGC algorithm from regular QSA [19,26] and standard electromagnetic PIC codes [20,21] to a Fourier decomposed QSA code. In Ref. [21] the PGC was implemented into an $r-\xi$ ($\xi = ct - z$) PIC code and the envelope solver included all derivatives. The implementation into QPAD will be described in detail. This algorithm consists of two components – a Fourier decomposed finite difference solver and a QSA based particle pusher that includes the ponderomotive force. The proposed PGC laser solver is not constrained by the conventional CFL condition and is unconditionally stable for modeling lasers in vacuum. In the presence of plasma a large time step is still allowed. Moreover, the proposed algorithm is compatible with the pipelining technique in QPAD (based on ideas from QuickPIC [32]) and thus can be executed asynchronously in parallel.

The paper is organized as follows: In Section 2, we first briefly review the PGC theory, including the envelope equation for laser and the ponderomotive force for plasma particles. In Section 3, we describe the finite difference laser solver and provide a proof of unconditional numerical stability. This is followed by the detailed description of the azimuthal Fourier decomposition form for the laser solver. Next, the plasma susceptibility deposition scheme and the plasma particle pusher are presented in Section 4 and 5. In Section 6, comparisons between QPAD [29] and OSIRIS [33,34] for a freely-drifting laser pulses in vacuum, standard LWFA driven by a regular Gaussian beam, and LWFA driven by a higher order Laguerre-Gaussian beam are presented.

2. Mathematical description of PGC model

This section will briefly review the main components of the PGC theoretical model [19,20]. The basic idea of PGC model is separating the physics into the rapidly varying components caused by the laser oscillation ($\tilde{\mathbf{a}}$) and the slowly varying components of the wake arising from averaging over the laser period (\mathbf{A}). Assuming the Coulomb gauge $\nabla \cdot (\mathbf{A} + \tilde{\mathbf{a}}) = 0$, the wave equation under the QSA for the vector potential is given by [19]

$$\left(\nabla^2 - \frac{\partial^2}{\partial t^2} \right) (\mathbf{A} + \tilde{\mathbf{a}}) = -\mathbf{J} - \tilde{\mathbf{j}} + \frac{\partial}{\partial t} \nabla \Phi \quad (1)$$

where the uppercase symbols \mathbf{A} , Φ and \mathbf{J} are the time-averaged vector potential, scalar potential and plasma current, while the lowercase symbols $\tilde{\mathbf{a}}$ and $\tilde{\mathbf{j}}$ are the vector potential and the current density associated with the fundamental frequency of the laser field. Note that the high frequency contribution to the scalar potential is dropped [19]. Here and in the remainder of the paper, normalized units are used. Length, time, mass and charge are normalized to $c\omega_p^{-1}$, ω_p^{-1} , electron rest mass m_e , and electron charge e respectively. The rapidly varying current density can be expressed in terms of the laser vector potential and relativistic factor [Eq. (A34) in Ref. [19]] as

$$\tilde{\mathbf{j}} = -\tilde{\mathbf{a}} \sum_i \frac{q_i \rho_i}{\bar{\gamma}_i m_{0i}} \quad (2)$$

where the ρ_i , q_i , m_{0i} and $\bar{\gamma}_i$ are the charge density, charge, rest mass and averaged Lorentz factor of the i -th particle species. The derivation of Eq. (2) uses the relationship between the rapidly varying current and plasma transverse momentum $\tilde{\mathbf{p}}_{\perp i}$, i.e., $\tilde{\mathbf{j}} = \sum_i \rho_i \tilde{\mathbf{p}}_{\perp i} / (\bar{\gamma}_i m_{0i})$, and the canonical momentum conservation $\tilde{\mathbf{p}}_{\perp i} = -q_i \tilde{\mathbf{a}}$. A more rigorous derivation can be found

in the appendix of Ref. [19]. Separating the rapidly varying component of Eq. (1) out and combining with Eq. (2), the evolution of laser fields satisfies the following PDE

$$\left(\nabla^2 - \frac{\partial}{\partial t^2} - \sum_i \frac{q_i \rho_i}{\bar{\gamma}_i m_{0i}} \right) \tilde{\mathbf{a}} = 0. \quad (3)$$

By assuming $\tilde{\mathbf{a}}$ can be expressed as the product of a rapidly varying phase term and a slowly varying complex envelope, i.e.,

$$\tilde{\mathbf{a}} = \frac{a(\mathbf{x}_\perp, \xi, s)}{2} \exp(i k_0 \xi) \hat{\mathbf{e}}_\perp + \text{c.c.} \quad (4)$$

the evolution of the laser envelope a satisfies

$$2 \frac{\partial}{\partial s} \left(i k_0 + \frac{\partial}{\partial \xi} \right) a = (\nabla_\perp^2 + \chi) a \quad (5)$$

where the Galilean transformation $\xi \equiv t - z$ (z is the laser propagation direction), $s = t$ has been carried out, k_0 is the wavevector corresponding to ω_0 and χ , which describes the plasma response, is defined as

$$\chi = - \sum_i \frac{q_i \rho_i}{\bar{\gamma}_i m_{0i}}. \quad (6)$$

Mora and Antonsen [19] showed that the time-averaged Lorentz factor is given by

$$\bar{\gamma} = \left[1 + \bar{u}^2 + \left(\frac{q}{m_0} \right)^2 \frac{|a|^2}{2} \right]^{1/2}, \quad (7)$$

where the $\bar{\mathbf{u}}$ is the time-averaged proper velocity.

The equation of motion for particles couples the evolution of the laser envelope and the plasma response together. The time-averaged equation of motion for plasma particles can be expressed as

$$\frac{d\bar{\mathbf{u}}}{dt} = \frac{q}{m_0} \left(\mathbf{E} + \frac{\bar{\mathbf{u}}}{\bar{\gamma}} \times \mathbf{B} - \frac{1}{4} \frac{q}{\bar{\gamma} m_0} \nabla |a|^2 \right) \quad (8)$$

where \mathbf{E} and \mathbf{B} are the time-averaged electric and magnetic fields associated with the plasma response. The term having $\nabla |a|^2$ is the ponderomotive force from the laser. As long as the particle experiences enough laser oscillation cycles and the laser envelope changes little within a single cycle, the equation above is sufficiently accurate. Under the QSA, we have $d_t \rightarrow (1 - \bar{v}_z) d_\xi$ and the time averaged equation of motion for a plasma particle becomes

$$\frac{d\bar{\mathbf{u}}}{d\xi} = \frac{q/m_0}{\bar{\gamma} - \bar{u}_z} \left(\bar{\gamma} \mathbf{E} + \bar{\mathbf{u}} \times \mathbf{B} - \frac{1}{4} \frac{q}{m_0} \nabla |a|^2 \right). \quad (9)$$

The laser envelope equation [Eq. (5)], the equation of motion for plasma particles [Eq. (9)] and the Maxwell's equations under QSA (see ref. [27]) constitute a complete set of equations, and thus can be adopted as the working equations for a QSA simulation.

3. Laser field solver

3.1. Finite difference solver

Assuming that in a QSA simulation, the physical quantities associated with the plasma are defined on the integer time step $s = n \Delta_s$ while the laser field is defined on the half-integer time step $s = (n + \frac{1}{2}) \Delta_s$. The plasma and laser are then advanced forward in s by the leapfrog method. Approximating $\partial_s a$ with a central difference operator, i.e., $\partial_s a \rightarrow (a^{n+\frac{1}{2}} - a^{n-\frac{1}{2}})/\Delta_s$, Eq. (5) can be written as

$$\left[i k_0 - \frac{1}{4} \Delta_s (\nabla_\perp^2 + \chi^n) \right] a^{n+\frac{1}{2}} + \frac{\partial a^{n+\frac{1}{2}}}{\partial \xi} = \left[i k_0 + \frac{1}{4} \Delta_s (\nabla_\perp^2 + \chi^n) \right] a^{n-\frac{1}{2}} + \frac{\partial a^{n-\frac{1}{2}}}{\partial \xi}. \quad (10)$$

In a QSA PIC code, the asynchronous parallel algorithm (pipelining) [32] is usually employed to improve the parallelism of the simulations. In this approach the quantities at smaller ξ (the front of the driver) are known at larger values of s . Since using the regular central difference for $\partial_\xi a$ requires synchronous data communication among the MPI nodes distributed in the ξ -direction, then it is not possible to use a centered definition for $\partial_\xi a$. Instead, we choose a second-order backward difference operator to approximate $\partial_\xi a$, i.e.,

$$\frac{\partial a_j}{\partial \xi} \rightarrow D_\xi a_j = \frac{1}{2\Delta_\xi} (3a_j - 4a_{j-1} + a_{j-2}), \quad (11)$$

where the index j represents the field defined at $\xi = j\Delta_\xi$. In this way, each MPI node in the ξ -direction only needs to receive data from the upstream node which keeps the value of a_{j-1} and a_{j-2} at $s + \Delta_s$ and s at the MPI boundary (at index j).

Another complication arises from the $\chi^n a^{n+\frac{1}{2}}$ term on the LHS of Eq. (10) for the quasi-3D algorithm. The value of χ^n depends on position, thus the differential equation for a is linear but with “spatially” dependent (non constant) coefficients. In the quasi-3D algorithm, the fields are expanded into azimuthal harmonics, i.e., Fourier modes. The presence of the spatially varying $\chi^n a^{n+\frac{1}{2}}$ term thus leads to cross-product terms between different azimuthal modes of χ^n and $a^{n+\frac{1}{2}}$. This makes the equations for the different azimuthal modes of $a^{n+\frac{1}{2}}$ coupled together, adding extra difficulties to the implementation compared with a fully finite difference method. However, as we describe shortly if the couplings are carefully accounted for then the s advance for each mode requires the inversion of a single constant coefficient matrix.

Therefore, to address both issues we have constructed the following iterative formula for Eq. (10)

$$\begin{aligned} & \left(ik_0 - \frac{1}{4} \Delta_s \nabla_\perp^2 + \frac{3}{2\Delta_\xi} \right) a_j^{n+\frac{1}{2},l} \\ &= \frac{1}{4} \Delta_s \chi_j^n a_j^{n+\frac{1}{2},l-1} + \frac{1}{2\Delta_\xi} (4a_{j-1}^{n+\frac{1}{2}} - a_{j-2}^{n+\frac{1}{2}}) + \left[ik_0 + \frac{1}{4} \Delta_s (\nabla_\perp^2 + \chi_j^n) \right] a_j^{n-\frac{1}{2}} + D_\xi a_j^{n-\frac{1}{2}} \end{aligned} \quad (12)$$

Here l is the iteration step and $a^{n+\frac{1}{2},l} \rightarrow a^{n+\frac{1}{2}}$ when $l \rightarrow \infty$ if the iteration is converged. For this iterative finite difference equation, the LHS (which includes the unknown values of the laser field for the next iteration) only contains finite difference operators that have constant coefficients. We point out that Eq. (12) is applicable to the full 3D Cartesian, full 3D cylindrical, and the quasi-3D descriptions. The primary differences between them lie in the specific form of the transverse Laplacian ∇_\perp^2 and the calculation of the $\chi^n a^{n+\frac{1}{2}}$ term. For the full 3D Cartesian coordinate system, ∇_\perp^2 can be approximated using the classical 5-point formula and thus the linear system solvers like fast FFT and multigrid method are the usual choice. For the quasi-3D algorithm, as we will describe in section 3.3 the ∇_\perp^2 is split into multiple 1D operators in r , and we need to calculate the cross-product terms of $\chi^n a^{n+\frac{1}{2}}$ as stated before.

3.2. Convergence and stability

It is important to understand the stability and the rate of convergence of the proposed iteration loop. In the case where the plasma is absent, i.e., $\chi = 0$, there is no need to iterate over Eq. (12) and the finite difference wave equation reduces to the regular Crank-Nicolson (CN) method which is unconditionally stable [35].

When plasma is present ($\chi \neq 0$), the convergence condition of the iteration over l can be determined when χ^n is not spatially dependent. It can be proven (see Appendix A) that the convergence condition in this situation is

$$k_0^2 \Delta_\xi^2 + \frac{9}{16} > \frac{1}{16} |\chi^n|^2 \Delta_s^2 \Delta_\xi^2. \quad (13)$$

For typical simulation scenarios, we have $\chi \sim O(1)$, $\Delta_\xi \sim O(0.1)$ and the validity of the PGC model requires $k_0 \gg 1$. This allows for a large Δ_s [e.g. $\sim O(1)$ to $\sim O(10)$] according to the convergence condition [Eq. (13)]. In reality χ is not spatially uniform making finding an analytical convergence condition difficult. We can however still qualitatively determine Δ_s through Eq. (13) by replacing χ^n with the maximum χ on the grid.

For an infinite number of iterations ($l \rightarrow \infty$) and for a spatially independent χ , it can be shown that the proposed scheme is stable without any restrictions for Δ_s . However, for the more relevant scenario where l is finite and χ is not spatially uniform, mathematically analyzing the numerical stability becomes very difficult. Nevertheless, numerous tests indicate that the proposed scheme is numerically stable even for a large Δ_s [$\sim O(10)$] for a wide variety of LWFA-relevant problems.

3.3. Azimuthal Fourier decomposition

In order to make the proposed finite difference scheme applicable to QPAD, we need to further expand Eq. (12) into azimuthal modes and derive the finite difference equations for each azimuthal harmonic. The complex-valued iterative equation [Eq. (12)] is thus split into two coupled real-valued equations corresponding to the real and imaginary parts respectively. The real and imaginary parts of a , i.e., a_R and a_I , are expanded using the following azimuthal Fourier series,

$$a_{\langle \cdot \rangle} = a_{\langle \cdot \rangle}^0 + 2 \sum_{m=1}^{+\infty} \Re \{ a_{\langle \cdot \rangle}^m \} \cos(m\phi) - 2 \sum_{m=1}^{+\infty} \Im \{ a_{\langle \cdot \rangle}^m \} \sin(m\phi) \quad (14)$$

where $\langle \cdot \rangle$ denotes “R” or “I” and m is the harmonics number. The complex Fourier coefficients a_R^m and a_I^m obey the following equations

$$\left(-\frac{1}{4} \Delta_s \Delta_m + \frac{3}{2\Delta_\xi} \right) a_{R,j}^{m,l} - k_0 a_{I,j}^{m,l} = \frac{1}{4} \Delta_s (\chi_j a_{R,j})^{m,l-1} + \frac{1}{2\Delta_\xi} (4a_{R,j-1}^m - a_{R,j-2}^m) + S_{R,j}^m, \quad (15)$$

$$\left(-\frac{1}{4} \Delta_s \Delta_m + \frac{3}{2\Delta_\xi} \right) a_{I,j}^{m,l} + k_0 a_{R,j}^{m,l} = \frac{1}{4} \Delta_s (\chi_j a_{I,j})^{m,l-1} + \frac{1}{2\Delta_\xi} (4a_{I,j-1}^m - a_{I,j-2}^m) + S_{I,j}^m, \quad (16)$$

where the Laplacian for m -th harmonic $\Delta_m \equiv \frac{1}{r} \frac{\partial}{\partial r} (r \frac{\partial}{\partial r}) - \frac{m^2}{r^2}$. Without introducing ambiguity, we have left out the time indices to simplify the denotations. Here, the symbols a_R , a_I and χ explicitly appearing in Eq. (15) and (16) are defined at $s = (n + \frac{1}{2})\Delta_s$ while those included in S_R^m and S_I^m , given by

$$S_{R,j}^m = \frac{1}{4} \Delta_s (\Delta_m a_{R,j}^m + (\chi_j a_{R,j})^m) - k_0 a_{I,j}^m + D_\xi a_{R,j}^m \quad (17)$$

$$S_{I,j}^m = \frac{1}{4} \Delta_s (\Delta_m a_{I,j}^m + (\chi_j a_{I,j})^m) + k_0 a_{R,j}^m + D_\xi a_{I,j}^m, \quad (18)$$

are defined at $s = (n - \frac{1}{2})\Delta_s$. The nonlinear terms in Eqs. (15) to (18) are calculated via the following truncated series

$$(\chi a_{\langle \cdot \rangle})^m = \sum_{k=m-M}^M \chi^k a_{\langle \cdot \rangle}^{m-k} \quad (19)$$

where M is the maximum harmonic number involved in the simulation.

Using a 3-point discretization, the Δ_m operator with second-order precision can be written as

$$\Delta_m a_{\langle \cdot \rangle}^m \rightarrow \beta_i^- a_{\langle \cdot \rangle, i-1}^m - \alpha_i a_{\langle \cdot \rangle, i}^m + \beta_i^+ a_{\langle \cdot \rangle, i+1}^m \quad (20)$$

where

$$\beta_i^\pm = \frac{1}{\Delta_r^2} \left(1 \pm \frac{1}{2i} \right), \quad \alpha_i = \frac{1}{\Delta_r^2} \left(2 + \frac{m^2}{i^2} \right). \quad (21)$$

The subscript i of $a_{\langle \cdot \rangle}^m$ means it is defined at the radial position $r = i\Delta_r$. The coupled equations (15) and (16) can be recast into a penta-diagonal linear system $PX = \text{RHS}$, where the unknown vector X is constructed by alternatively placing a_R and a_I

$$X = (\dots, a_{R,i}^m, a_{I,i}^m, a_{R,i+1}^m, a_{I,i+1}^m, \dots). \quad (22)$$

The penta-diagonal coefficient matrix P is given by

$$\left(\begin{array}{ccccccc} \ddots & \ddots & \ddots & \ddots & \ddots & & \\ & B_i^- & 0 & A_i & -k_0 & B_i^+ & & \\ & B_i^- & k_0 & A_i & 0 & B_i^+ & & \\ & B_{i+1}^- & 0 & A_{i+1} & -k_0 & B_{i+1}^+ & & \\ & B_{i+1}^- & k_0 & A_{i+1} & 0 & B_{i+1}^+ & & \\ & & \ddots & \ddots & \ddots & \ddots & \ddots & \end{array} \right) \quad (23)$$

where $B_i^\pm = \frac{1}{4} \Delta_s \beta_i^\pm$ and $A_i = -\frac{1}{4} \Delta_s \alpha_i + \frac{3}{2\Delta_\xi}$. This linear system can be efficiently solved using the cyclic reduction method [36].

3.4. Boundary conditions

The boundary conditions at $r = 0$ are different for the $m = 0$ and $m > 0$ modes. For $m = 0$ mode, $\partial_r a_{\langle \cdot \rangle}^0|_{r=0} = 0$ and the Laplacian $\Delta_0 = \frac{1}{r} \frac{\partial}{\partial r} (r \frac{\partial}{\partial r}) \rightarrow 2 \frac{\partial^2}{\partial r^2}$ when $r \rightarrow 0$. The former indicates that $a_{\langle \cdot \rangle, i=-1} = a_{\langle \cdot \rangle, i=1}$. In this situation, the stencil coefficients for $i = 0$ become $\beta_0^- = 0$, $\beta_0^+ = \alpha_0 = 2/\Delta_r^2$. For $m > 0$ modes, the symmetry requires $a_{\langle \cdot \rangle}^m = 0$ at $r = 0$, and thus $a_{\langle \cdot \rangle, i>0}^m$ can be solved without changing the stencil coefficients.

A Dirichlet boundary condition is used for the laser envelope of all the harmonics at $r = r_{\max}$. As long as the laser envelope does not extend to the boundaries, a Dirichlet boundary condition can work properly without introducing unphysical results.

4. Plasma susceptibility deposition

In order to solve the laser envelope equations for each harmonic [Eqs. (15) to (18)], the plasma susceptibility χ^m needs to be deposited onto the grid based on the particle information, i.e., charge, position, momentum and mass. The deposition scheme for χ^m basically follows that used for the source terms (charge density, current density, etc.) as described in Section 3.4 of the original QPAD paper [29]. A more convenient form for χ [Eq. (6)] when dealing with particle data is

$$\chi = -\frac{q}{m_0} \frac{\rho - J_z}{\bar{\gamma} - \bar{u}_z}, \quad (24)$$

which is obtained by multiplying both the numerator and denominator of Eq. (6) by $(1 - \bar{v}_z)$.

We start from the general formula for depositing χ as a sum over particles,

$$\chi = \frac{1}{\text{vol. } m_0} \frac{q}{\sum_i} \sum_i \frac{q_i}{\bar{\gamma}_i - \bar{u}_{z,i}} \frac{1}{r} S_r(r - r_i) S_\phi(\phi - \phi_i), \quad (25)$$

where q_i , r_i , ϕ_i , $\bar{\gamma}_i$ and $\bar{u}_{z,i}$ are the charge, radial position, azimuthal angle, time-averaged Lorentz factor and longitudinal proper velocity of the i -th particle. The particle shape function $\frac{1}{r} S_r(r - r_i) S_\phi(\phi - \phi_i)$ is used to interpolate the particle information at position (r_i, ϕ_i) onto the grid position (r, ϕ) . Both S_r and S_ϕ should satisfy the normalization conditions $\int dr S_r = 1$ and $\int d\phi S_\phi = 1$. Since QPAD uses the gridless description in ϕ , we need to expand the azimuthal shape function S_ϕ into azimuthal harmonics,

$$S_\phi(\phi - \phi_i) = \sum_m S_\phi^m(\phi_i) e^{im\phi} \quad (26)$$

where the Fourier coefficient is expressed as

$$S_\phi^m(\phi_i) = \frac{1}{2\pi} \int_0^{2\pi} d\phi' S_\phi(\phi' - \phi_i) e^{-im\phi'}. \quad (27)$$

In QPAD we adopt the Dirac delta function as the particle shape function in ϕ , i.e., $S_\phi(\phi - \phi_i) = \delta(\phi - \phi_i)$ and thus $S_\phi^m(\phi_i) = e^{-im\phi_i}/(2\pi)$. Therefore, the deposition formula for the m -th harmonic of χ can be written as

$$\chi^m = \frac{1}{\text{vol. } m_0} \sum_i \frac{q_i}{\bar{\gamma}_i - \bar{u}_{z,i}} \frac{1}{2\pi r} S_r(r - r_i) e^{-im\phi_i}. \quad (28)$$

In practice, it is not necessary to calculate each harmonic but only the $m = 0$ mode from each particle. The contribution to any $m > 0$ mode from an individual particle can be calculated from the $m = 0$ contribution by simply multiplying a constant phase factor, i.e., through the relation $\chi^m = \chi^0 e^{-im\phi_i}$ or recursively through $\chi^m = \chi^{m-1} e^{-i\phi_i}$.

5. Particle pusher

In QPAD, all the physical quantities associated with the plasma are defined on the integer time step $s = n\Delta_s$. Within a time step, the shape of the laser or particle beam is assumed unchanged, and the ponderomotive force or the Coulomb like force of a drive or trailing beam is used to update the plasma response. The quantities associated with plasma, including the information of plasma particles, the charge and current density and the plasma-induced electromagnetic fields, are updated from the front to the end of the moving window for each transverse slice. The particle information and the fields are solved using Eq. (9) and the Maxwell's equations with QSA described in ref. [27]. Before numerically integrating Eq. (9), one needs to interpolate \mathbf{E} , \mathbf{B} and $\nabla|a|^2$ from the grid points onto the position of each particle. Instead of calculating $\nabla|a|^2$ on the grid points first and then interpolating onto the particles' positions, we choose to interpolate a and ∇a separately. The $|a|^2$ term in $\bar{\gamma}$ for an individual particle is calculated from the interpolated a , and the $\nabla|a|^2$ at the location for individual particles is calculated via $\nabla|a|^2 = 2\Re\{a^* \nabla a\}$. This approach will reduce the number of floating point operations especially when there are many Fourier harmonics. The formula for interpolating a from the grid points onto the i -th particle position (r_i, ϕ_i) can be derived as follows

$$\begin{aligned}
a(r_i, \phi_i) &= \int dr \int d\phi a(r, \phi) S_r(r - r_i) S_\phi(\phi - \phi_i) \\
&= \frac{1}{2\pi} \int dr \int d\phi \sum_m a^m(r) e^{im\phi} \sum_{m'} e^{im'\phi_i} e^{-im'\phi} S_r(r - r_i) \\
&= \sum_m \int dr a^m(r) e^{im\phi_i} S_r(r - r_i).
\end{aligned} \tag{29}$$

The interpolation of \mathbf{E} and \mathbf{B} can be done in a similar way. Since $\nabla_\perp a^m = \frac{\partial a^m}{\partial r} \mathbf{e}_r + \frac{im}{r} a^m \mathbf{e}_\phi$, the formula for interpolating ∇a is given by

$$\nabla a(r_i, \phi_i) = \sum_m \int dr \left(\frac{\partial a^m}{\partial r} \mathbf{e}_r + \frac{im}{r} a^m \mathbf{e}_\phi + \frac{\partial a^m}{\partial \xi} \mathbf{e}_z \right) e^{im\phi_i} S_r(r - r_i). \tag{30}$$

In QPAD, the plasma particle positions are defined on the integer grid points in ξ , i.e., $\xi = j\Delta_\xi$, so that the fields and ponderomotive force felt by particles are also located at $\xi = j\Delta_\xi$. The particle momenta are defined on the half-integer grid points, i.e., $\xi = (j + \frac{1}{2})\Delta_\xi$. After calculating the fields and ponderomotive force at the particle positions, we can numerically integrate Eq. (9) using the Boris method [37]. It should be noted that when advancing $\bar{\mathbf{u}}^{j-\frac{1}{2}}$ to $\bar{\mathbf{u}}^{j+\frac{1}{2}}$, one needs to know the values of \bar{y} and $\bar{y} - \bar{u}_z$ at $\xi = j\Delta_\xi$ to make the Boris method applicable. These two values can be obtained during the so-called predictor-corrector iteration embedded in the loop for updating the plasma.

6. Benchmark and example simulations

In this section, we present comparisons of LWFA related simulations based on QPAD with the PGC and the ones based on full 3D OSIRIS. These benchmark comparisons include the propagation of a Gaussian laser pulse in vacuum and a standard LWFA operating in the nonlinear blowout regime with a matched spot size [38]. For these two examples we assume azimuthal symmetry for QPAD and thus only the $m = 0$ mode is used. The OSIRIS simulation is full 3D and any physical instabilities or physics that involves higher order modes would be included. We also include a simulation where the LWFA is driven by a Laguerre-Gaussian beam. In the QPAD simulation up to the $m = 2$ mode is included. Such a laser driver has been proposed to accelerate positron beams using LWFA [39].

6.1. Laser pulse propagation in vacuum

We start by simulating the propagation of a laser pulse with a transverse Gaussian profile (a lowest order Laguerre-Gaussian mode). The laser field evolution has an analytical expression which is the solution to the paraxial Helmholtz equation. The complex amplitude of a certain ξ slice is given by

$$a(r, z) = a_0 \frac{w_0}{w(z)} \exp\left(\frac{-r^2}{w(z)^2}\right) \exp\left[-i\left(\frac{k_0 r^2}{2R(z)} - \Psi(z)\right)\right] \tag{31}$$

where z (equivalently s) is the distance from the focal waist, w_0 is the waist radius, $w(z) = w_0 \sqrt{1 + z^2/z_R^2}$ is the beam radius at z , $R(z) = z(1 + z_R^2/z^2)$ is the radius of curvature, and $\Psi = \arctan(z/z_R)$ is the Gouy phase. In the benchmark tests, a laser pulse with $k_0 w_0 = 20$ and $a_0 = 1$ starts from the beam focal waist ($z = 0$), and travels a distance of $600k_0^{-1}$, i.e., three Rayleigh lengths, $3 z_R$. The complex amplitude initially has a longitudinal profile of the form $\cos^2[\pi(\xi - \xi_0)/(2\tau_{FWHM})]$ with central position $\xi_0 = 0$ and duration $\omega_0 \tau_{FWHM} = 60$. In the OSIRIS simulation, a moving window with spatial resolution of $\Delta_x = \Delta_y = 2k_0^{-1}$ and $\Delta_z = 0.2k_0^{-1}$ is used. The time step is $\Delta_t = 0.125\omega_0^{-1}$. To accurately simulate the group velocity, we used the numerical-dispersion-free field solver [14] for all the OSIRIS simulations described in this paper. In the QPAD simulation, $\Delta_r = 1.56k_0^{-1}$ and $\Delta_\xi = 0.78k_0^{-1}$. We varied the time steps from $\Delta_s = 20\omega_0^{-1} = 0.1z_R/c$ to $\Delta_s = 200\omega_0^{-1} = z_R/c$.

Fig. 1 summarizes the simulation results. In Fig. 1(a) to (d), the snapshots of the laser field at four propagation distances (s) are presented. The top and bottom halves correspond to the OSIRIS and QPAD simulation results, respectively. Unless otherwise specified, the laser pulse in the simulations here and in the remaining sections moves from the right to left. We carried out a series of simulations to examine the impact of Δ_s . The lineouts at $s = 3z_R/c$ for various Δ_s are shown in Fig. 1(e). The results for $\Delta_s = 0.1, 0.2, 0.5z_R/c$ fully converge to that of OSIRIS while for $\Delta_s = z_R/c$ there is a large deviation. This comparison indicates that the time step needs to properly resolve the evolution of the laser which in vacuum is the Rayleigh length.

6.2. LWFA driven by a Gaussian laser pulse

In this subsection, we provide results of simulations for a standard LWFA (unloaded) operating in the self guided blowout regime [38]. In the blowout regime the wake is created by expelling the plasma electrons forward and radially outward

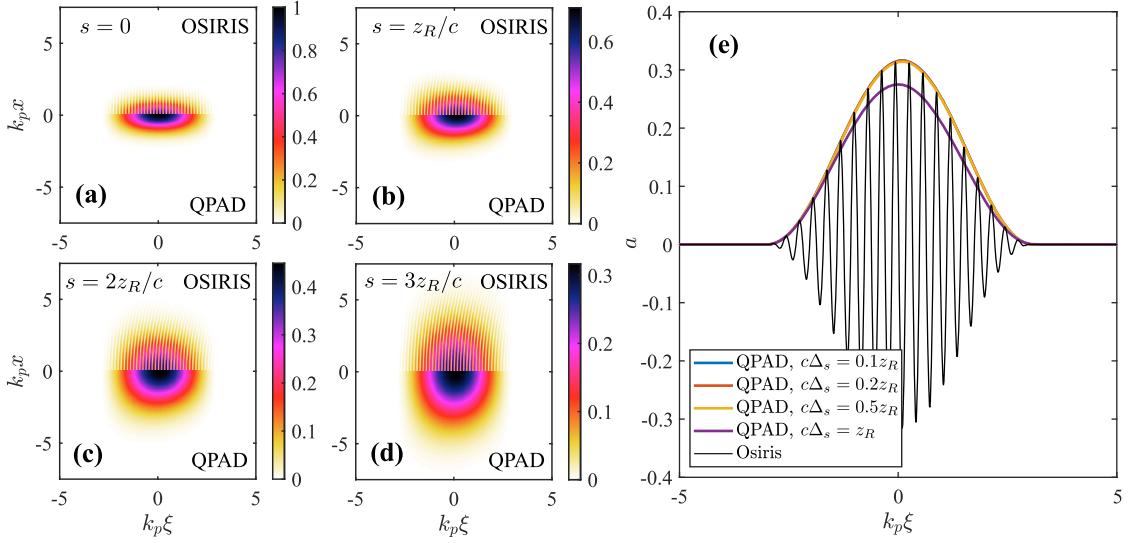


Fig. 1. Comparison simulations of a freely drifting laser pulse between OSIRIS and QPAD PGC. Snapshots of the laser fields at (a) $s = 0$, (b) $s = z_R$, (c) $s = 2z_R$ and (d) $s = 3z_R$ are presented. The time step of QPAD simulations from (a) to (d) is $\Delta_s = 0.2z_R/c$. (e) The impact induced by the time step sizes in QPAD.

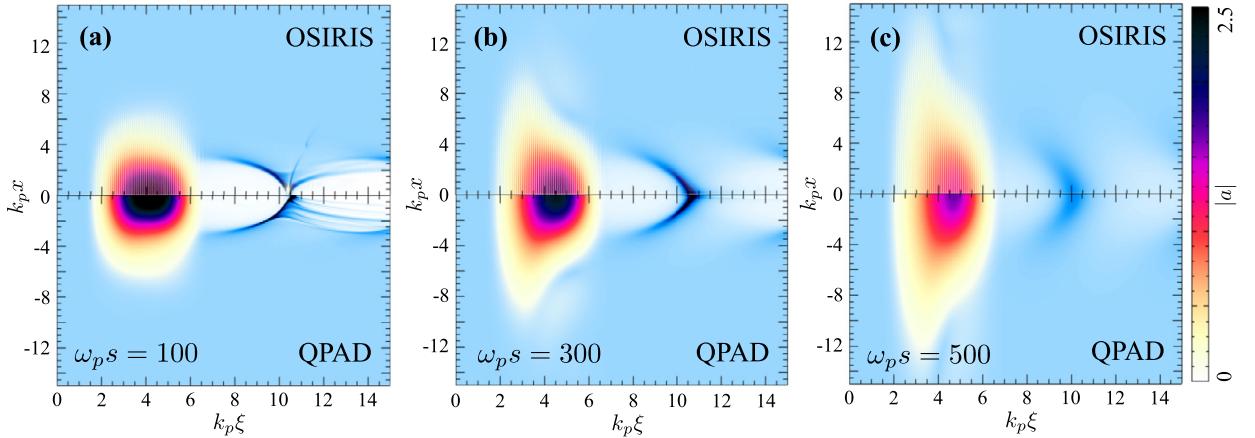


Fig. 2. Snapshots of the laser fields and background electron density of QPAD and OSIRIS simulations at (a) $s = 100\omega_p^{-1} \simeq 0.6z_R/c$, (b) $s = 300\omega_p^{-1} \simeq 1.8z_R/c$, and (c) $s = 500\omega_p^{-1} \simeq 2.9z_R/c$.

where they form a narrow sheath that surrounds a plasma column [40,41]. In the nonlinear, self-guide LWFA regime a matched spot size

$$k_p w_0 \simeq 2\sqrt{a_0} \quad (32)$$

for the laser is assumed and the laser intensity is relativistic $a_0 > 1$. Under this condition, the ponderomotive force felt by the plasma background electrons that have been blown out is roughly balanced by the Coulomb force from the ions, and hence the wakefield structure as well as the laser propagation can remain stable without large oscillation in the wake and laser amplitudes. In the simulation, we initialized a laser pulse with $k_p w_0 = 3.2$, $k_0/k_p = 34$ (the Rayleigh length is thus $k_p z_R \simeq 170$), $\omega_p \tau_{\text{FWHM}} = 2.26$ and $a_0 = 2.5$ in vacuum and it is then focused to the beginning of a plasma density upramp of length $20k_p^{-1}$. We used 8 particles per cell in r and 16 particles distributed in ϕ . The plasma density was constant at the end of the ramp. In the OSIRIS simulation, the spatial resolution of simulation box is $\Delta_x = \Delta_y = 0.1k_p^{-1}$ and $\Delta_z = 0.2k_0^{-1}$. The time step is $\Delta_t = 0.136\omega_0^{-1}$. In the QPAD simulations $\Delta_r = 0.1k_p^{-1}$ and $\Delta_\xi = 0.015k_p^{-1}$. Again, only $m = 0$ mode is considered due to the assumed azimuthal symmetry.

Fig. 2(a) to (c) shows snapshots of the laser field and the background electron density at different propagation distances from both the QPAD and OSIRIS simulations. The time step of the QPAD simulation is $\Delta_s = 10\omega_p^{-1} \simeq 0.058z_R/c$ and 3 iterations for the PGC laser solver are used. The full set of numerical parameters of QPAD and OSIRIS simulations can be found in Table 1. It can be clearly seen that the laser diffraction, the erosion of the laser front and the wake shape in both QPAD and OSIRIS agree well with each other.

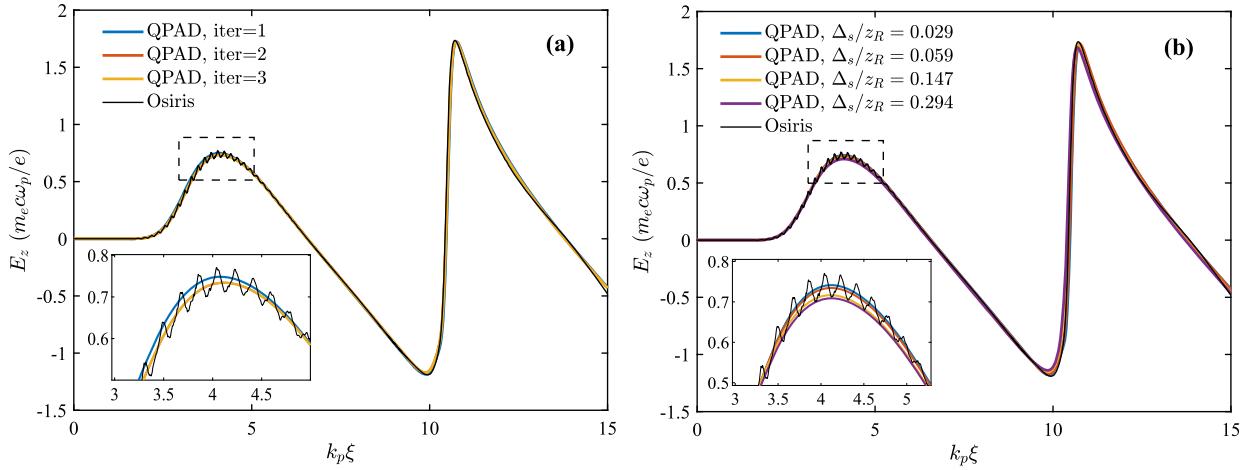


Fig. 3. Benchmark test results between OSIRIS and QPAD PGC simulations. The impacts induced by (a) the number of iterations for PGC laser solver and (b) the time step sizes in QPAD are also presented. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

We compare the impact of the choices for number of iterations and time step on the wakefield in Fig. 3 where central lineouts for the axial electric field E_z are shown. For $\Delta_s = 10\omega_p^{-1} \simeq 0.058z_R/c$, we can see in Fig. 3(a), where E_z at $s = 50\omega_p^{-1}$ is shown that the results for using 1, 2 and 3 PGC solver iterations are nearly identical and they also agree well with results from OSIRIS simulations. A closer examination of the E_z field in the dashed box (the inset) indicates that using only one iteration overestimates the E_z amplitude slightly (blue line). In Fig. 3(b), results from various time steps but with three iterations for the PGC laser solver are presented. The results for all time steps, $\Delta_s = 0.029, 0.059, 0.147$ and $0.294z_R/c$ are roughly converged and agree well with the OSIRIS result. From the inset in Fig. 3(b), it can be seen that the results for the two cases with larger Δ_s underestimate the wakefield amplitude slightly when compared to the two cases with smaller Δ_s . We also find that increasing the iteration number of the PGC laser solver will not significantly impact the computational performance of simulation. Therefore, in order to guarantee the simulation accuracy, one should avoid using a single iteration for PGC solver and choose a time step that can sufficiently resolve the Rayleigh length.

Particularly noteworthy is the excellent speedup achieved in these QPAD PGC simulations. Both the QPAD and OSIRIS simulations were carried out on the Cori cluster “Haswell” nodes at NERSC (Intel Xeon E5-2698 v3 @ 2.30 GHz). The numerical configurations and the computation time of both simulations are summarized in Table 1. Due to the QSA the number of time steps in the QPAD simulation is dramatically decreased compared to the OSIRIS simulation. The azimuthal Fourier decomposition reduces the computational complexity from 3D to 2D, leading to orders of magnitude reduction in the number of cells and particles needed. Therefore, the final core hours consumed by the QPAD simulation are $\sim 10^4$ less than that of the OSIRIS simulation. In this comparison we used 16 cores for the QPAD simulation and 2048 for the OSIRIS simulation so in this specific example the wall clock time speedup was ~ 200 . The performance of QPAD-PGC is also evaluated and the detailed scaling test results can be found in the Appendix C. For a typical LWFA case, QPAD-PGC can be well scaled to over 10^3 cores.

6.3. LWFA driven by a Laguerre-Gaussian beam

In this subsection, we demonstrate the capability of the azimuthal mode expansion for modeling physics that requires a finite but small number azimuthal modes via an example where a LWFA is driven by a higher order Laguerre-Gaussian laser pulse. The complex amplitude of the general Laguerre-Gaussian modes is given by [42]

$$a(r, z) = a_0 C_{p,l} \frac{w_0}{w(z)} \left(\frac{r}{w(z)} \right)^{|l|} \exp \left(\frac{-r^2}{w(z)^2} \right) L_p^{|l|} \left(\frac{2r^2}{w^2(z)} \right) \exp \left[-i \left(\frac{k_0 r^2}{2R(z)} + l\phi - \Psi(z) \right) \right] \quad (33)$$

where p and l are the radial and azimuthal indices, $L_p^{|l|}$ is the generalized Laguerre polynomial, $C_{p,l}$ is the normalized amplitude of the mode, and the definitions for $w(z)$ and $R(z)$ are identical to those of the fundamental Gaussian mode. The Gouy phase shift of a Laguerre-Gaussian beam is exaggerated by the factor $2p + |l|$, i.e., $\Psi(z) = (2p + |l| + 1) \arctan(z/z_R)$. For the modes $p = 0$ and $|l| > 0$, the laser intensity vanishes at $r = 0$, and increases and then decreases radially, presenting a donut-like intensity distribution. The associated ponderomotive force will push the background electrons inside the donut (the region where the $\frac{\partial|a|^2}{\partial r}$ is positive) inward, which then forms a high-density electron column on the axis. This particular field structure can provide a focusing force for positive charged particles and hence has been considered for positron beam acceleration [39].

Table 1

Configurations and computation time of OSIRIS and QPAD simulations for the LWFA.

OSIRIS	cell sizes ($\Delta_x, \Delta_y, \Delta_z$)	$0.1 k_p^{-1}, 0.1 k_p^{-1}, 5.88 \times 10^{-3} k_p^{-1}$ ($0.2 k_0^{-1}$)
	number of cells	$300 \times 300 \times 2550$
	particles per cell (N_x, N_y, N_z)	2, 2, 1
	number of particles	9.18×10^8
	time step (Δ_t)	$0.004 k_p^{-1}$ ($0.136 k_0^{-1}$)
	number of time steps	1.25×10^4
	number of processors	2048
QPAD	core hours	$\sim 2 \times 10^4$ hrs.
	cell sizes (Δ_r, Δ_ξ)	$0.1 k_p^{-1}, 0.015 k_p^{-1}$
	number of cells	256×1000
	particles per cell (N_r, N_ϕ)	8, 16
	number of particles (slice)	4096
	number of particles (total)	4.096×10^6
	time step (Δ_t)	$10 k_p^{-1}$
	number of time steps	50
	number of azimuthal mode	1
	number of processors	16
	core hours	0.8

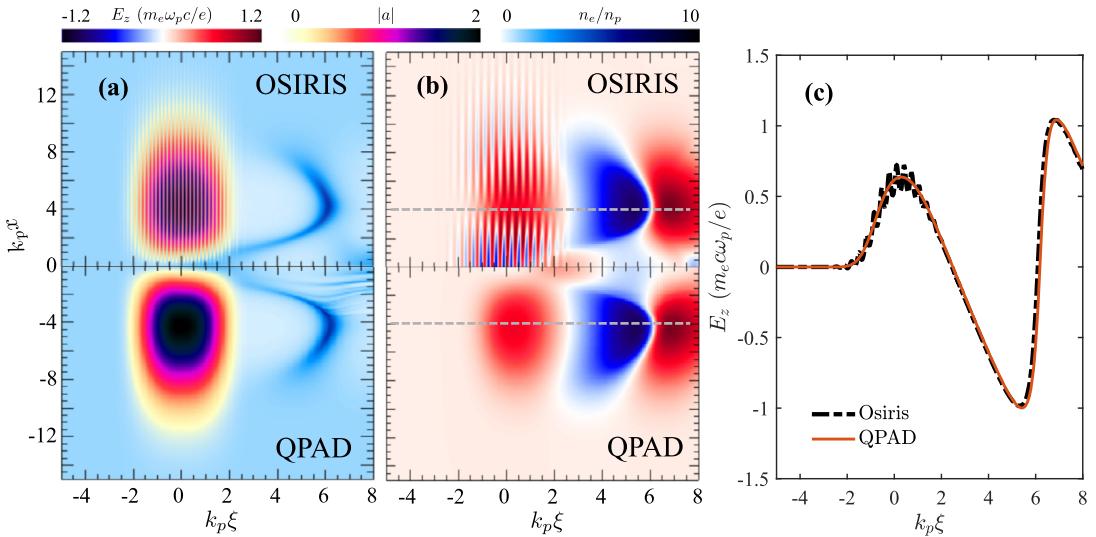


Fig. 4. QPAD-PCC simulation of a LWFA driven by a Laguerre-Gaussian laser pulse with $l = 1$ and $p = 0$. Snapshots of (a) the laser fields and plasma electron density, (b) axial field E_z and (c) lineouts of E_z at $k_p r = 4$ are taken at $\omega_p s = 40$.

Fig. 4 shows the plasma wake excited by a laser pulse with $(p, l) = (0, 1)$, $a_0 = 2.1$, $k_p w_0 = 6$, $\omega_0/\omega_p = 20$ and $\omega_p \tau_{FWHM} = 2.3$. The simulation box has a dimension of $25.6 k_p^{-1} \times 13 k_p^{-1}$ in the radial and axial directions and the spatial resolution is $\Delta_r = 0.1 k_p^{-1}$ and $\Delta_\xi = 0.04 k_p^{-1}$. The time step is $\Delta_s = 5 \omega_p^{-1}$. To well resolve the asymmetry due to the $e^{-il\phi}$ term, the azimuthal mode expansion is truncated at $m = 2$. Figs. 4(a) to (c) present the background electron density, axial field E_z at $\omega_p s = 40$ and the lineouts of E_z at $k_p r = 4$, respectively. In Fig. 4(a) and (b), the top and bottom subfigures correspond to the results from 3D OSIRIS and QPAD-PGC simulations, respectively. As expected, the laser with a donut-shaped intensity profile simultaneously repels the ambient electrons inward and outward, forming an electron-rarefied toroidal wake behind. From the Fig. 4(c), it is clear that the lineout of E_z at $k_p r = 4$ [denoted by the dashed lines in Fig. 4(b)] from the QPAD simulation is in excellent agreement with that from a full 3D OSIRIS simulation.

7. Conclusion

In summary, we have described a new quasi-static based particle-in-cell algorithm that utilizes the ponderomotive guiding center approach together with an azimuthal mode expansion. The proposed algorithm is not restricted by the regular CFL condition for the time step as is the case for a standard fully electromagnetic PIC code. The algorithm can also be asyn-

chronously parallelized (pipelined). The algorithm has been implemented into the quasi-static PIC code QPAD. In QPAD, the laser envelope equation, the Maxwell's equations and all the physical quantities are decomposed into azimuthal harmonics, i.e., Fourier modes. With this hybrid approach where a PIC description is used in $r - z$ and a gridless description is used in ϕ , QPAD with PGC model can more efficiently simulate short pulse laser-plasma interactions than full 3D PIC, quasi-3D explicit PIC, or full 3D quasi-static codes. The implementation of the plasma particle pusher and the plasma susceptibility deposition are also described. This new algorithm was benchmarked and compared against the results from the full 3D PIC code OSIRIS for a few sample cases. Excellent agreement was achieved for the simulations of a laser propagating freely in vacuum and a standard LWFA problem. Enormous reduction of in core-hours (~ 4 orders of magnitude) compared to a full 3D PIC code were obtained while maintaining high fidelity. We also simulated a LWFA driven by a Laguerre-Gaussian laser pulse as an example to show the capability of multi-mode simulations. With the PGC algorithm, QPAD can also very effectively model PWFA and LWFA with trailing particle beams and photon acceleration. Future directions for algorithm enhancement will include mesh refinement and improving the envelope solver to handle pump depletion distances.

CRediT authorship contribution statement

Fei Li: Conceptualization, Formal analysis, Methodology, Software, Visualization, Writing – original draft. **Weiming An:** Methodology. **Frank S. Tsung:** Resources. **Viktor K. Decyk:** Writing – review & editing. **Warren B. Mori:** Formal analysis, Funding acquisition, Resources, Supervision, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgements

This work was supported in parts by the US Department of Energy through SciDAC FNAL subcontract 644405 and contract number DE-SC0010064 and US National Science Foundation grant number 2108970. The work of W.A. was supported by the National Natural Science Foundation of China (NSFC) grant number 12075030. Simulations were carried out on the Cori Cluster of the National Energy Research Scientific Computing Center (NERSC).

Appendix A. Derivation of iterative convergence conditions for a uniform χ^n

Given a simulation time step $s = n\Delta_s$, the quantities $a_{j-1}^{n+\frac{1}{2}}$, $a_{j-2}^{n+\frac{1}{2}}$ and $a_j^{n-\frac{1}{2}}$ in Eq. (12) can be viewed as constants since here we only explore how the solution converges as increasing l . Subtracting Eq. (12) for l and $l-1$ gives,

$$\left(ik_0 - \frac{1}{4}\Delta_s\nabla_{\perp}^2 + \frac{3}{2\Delta_{\xi}}\right)r^l = \frac{1}{4}\Delta_s\chi^n r^{l-1} \quad (\text{A.1})$$

where $r^l \equiv a_j^{n+\frac{1}{2},l} - a_j^{n+\frac{1}{2},l-1}$ and χ^n is assumed to have no spatial dependence. We define $i[\mathbf{k}_{\perp}]$ as the counterpart to ∇_{\perp} in the \mathbf{k} -space such that in the Cartesian coordinates, $[\mathbf{k}_{\perp}]$ is given by

$$[\mathbf{k}_{\perp}] = \frac{\sin(k_x\Delta_x/2)}{\Delta_x/2}\mathbf{e}_x + \frac{\sin(k_y\Delta_y/2)}{\Delta_y/2}\mathbf{e}_y, \quad (\text{A.2})$$

where Δ_x , Δ_y are the cell sizes of the mesh-grid in the x - y plane. Using the ansatz $r^l = \tilde{r}^l e^{i\mathbf{k}_{\perp}\cdot\mathbf{x}_{\perp}}$, the ratio of the residuals of two adjacent iterations is given by

$$\left| \frac{\tilde{r}^l}{\tilde{r}^{l-1}} \right|^2 = \frac{\frac{1}{16}\Delta_s^2|\chi^n|^2}{\left(\frac{1}{4}\Delta_s[\mathbf{k}_{\perp}]^2 + \frac{3}{2\Delta_{\xi}}\right)^2 + k_0^2} \leq \frac{\frac{1}{16}\Delta_s^2|\chi^n|^2}{\frac{9}{4\Delta_{\xi}^2} + k_0^2}, \quad (\text{A.3})$$

A necessary condition for convergence is $|\tilde{r}^l/\tilde{r}^{l-1}| < 1$ which leads to

$$k_0^2\Delta_{\xi}^2 + \frac{9}{4} > \frac{1}{16}|\chi^n|^2\Delta_s^2\Delta_{\xi}^2. \quad (\text{A.4})$$

Appendix B. Proof of numerical stability for a uniform χ^n and $l \rightarrow \infty$

We carry out a von Neumann analysis to show that the proposed laser solver is unconditionally stable when χ has no spatial dependence for each time step and the number of iterations approaches infinity. In this situation, the $a^{n+\frac{1}{2},l}$ and $a^{n+\frac{1}{2},l-1}$ converge to $a^{n+\frac{1}{2}}$ and Eq. (12) actually becomes Eq. (10). The round-off error $\epsilon^n = a_N^n - a^n$ (a_N^n is the numerical solution obtained in finite precision arithmetic and a^n is precise solution of the difference equation) still satisfies Eq. (10) in the sense of machine precision. Using the ansatz $\epsilon^n = g^n e^{ik_\xi \xi} e^{ik_\perp \cdot \mathbf{x}_\perp}$ where g is the error growth factor, we can obtain

$$\left(ik_0 + \frac{1}{4}\Delta_s([k_\perp]^2 - \chi^n)\right)\epsilon^{n+\frac{1}{2}} + i[k_\xi]\epsilon^{n+\frac{1}{2}} = \left(ik_0 - \frac{1}{4}\Delta_s([k_\perp]^2 - \chi^n)\right)\epsilon^{n-\frac{1}{2}} + i[k_\xi]\epsilon^{n-\frac{1}{2}} \quad (\text{B.1})$$

where $i[k_\xi]$ is the counterpart of $\frac{\partial}{\partial \xi}$ in the \mathbf{k} -space. For the backward difference operator in ξ used in this paper, $[k_\xi]$ is given by

$$[k_\xi] = \left(\frac{3}{2}e^{-i\frac{1}{2}k_\xi \Delta_\xi} - \frac{1}{2}e^{-i\frac{3}{2}k_\xi \Delta_\xi}\right) \frac{\sin(k_\xi \Delta_\xi/2)}{\Delta_\xi/2}. \quad (\text{B.2})$$

The growth rate of the numerical error is

$$g = \frac{\epsilon^{n+\frac{1}{2}}}{\epsilon^{n-\frac{1}{2}}} = \frac{i(\Re\{[k_\xi]\} + k_0) - \Im\{[k_\xi]\} - \frac{1}{4}\Delta_s([k_\perp]^2 - \chi^n)}{i(\Re\{[k_\xi]\} + k_0) - \Im\{[k_\xi]\} + \frac{1}{4}\Delta_s([k_\perp]^2 - \chi^n)}. \quad (\text{B.3})$$

According to Eq. (B.2) the imaginary part of $[k_\xi]$

$$\Im\{[k_\xi]\} = \left[\frac{1}{2}\sin\left(\frac{3}{2}k_\xi \Delta_\xi\right) - \frac{3}{2}\sin\left(\frac{1}{2}k_\xi \Delta_\xi\right)\right] \frac{\sin(k_\xi \Delta_\xi/2)}{\Delta_\xi/2} < 0 \quad (\text{B.4})$$

for any real k_ξ . Since $\chi^n < 0$ according to Eq. (6), $[k_\perp]^2 - \chi^n > 0$. Therefore,

$$|g|^2 = \frac{(\Re\{[k_\xi]\} + k_0)^2 + (\frac{1}{4}\Delta_s([k_\perp]^2 - \chi^n) + \Im\{[k_\xi]\})^2}{(\Re\{[k_\xi]\} + k_0)^2 + (\frac{1}{4}\Delta_s([k_\perp]^2 - \chi^n) - \Im\{[k_\xi]\})^2} \leq 1 \quad (\text{B.5})$$

and the equality holds only when $k_\xi = 0$. $|g| \leq 1$ means the round-off error will not grow as n increases and the algorithm is thus numerically stable. This conclusion does not depend on the specific forms of $[\mathbf{k}_\perp]$ which implies that the numerical stability is independent of the coordinate system and specific form of the difference operator in the transverse plane.

Appendix C. Parallel scaling

Like the original QPAD, QPAD-PGC is also parallelized using MPI and runs on distributed memory systems. In this appendix, we present the scalings of QPAD-PGC on the Cori cluster at NERSC. For a QSA code, the longitudinal and transverse parallel scalability differ from each other due their use of fundamentally different parallelization schemes. In the longitudinal direction the code is parallelized using a pipelining algorithm which has shown excellent scalability in previous work [32]. In the transverse direction the code is parallelized using domain decomposition and the parallelization does not scale linearly since the field and laser solvers use the cyclic reduction method. To quantify the parallel scalability in the different directions, we carried out transverse, longitudinal and bi-directional scaling tests. In the transverse and longitudinal tests, the simulation window is partitioned only in the transverse and longitudinal directions respectively, while in the bi-directional test the window is partitioned in both directions and the number of cores in both directions are set to be equal.

An LWFA simulation with the same physical parameters as Section 6.2 is used for the scaling tests. In the strong scaling test the total problem size is fixed and the relation between the runtime and the number of cores is examined. The simulation window has $2^{14} \times 2^{13}$ cells in the transverse and longitudinal directions, and there are 8 macro-particles within a radial cell size Δ_r and 8 macro-particles distributed in 2π angle. In the weak scaling test, the same physical problem is simulated but the resolution is varied from $2^9 \times 2^8$ cells to $2^{14} \times 2^{13}$ cells where the numbers of MPI partitions are varied from 1 to 1,024 respectively. Thus, each MPI partition is fixed with 2^{17} cells (the partition shape may vary for different total number of cores) in each case and the same number of particles per cell as for the strong scaling test is used. The results of the scaling benchmarks are presented in Fig. C.5. Due to the excellent scalability of pipelining algorithm, the pure longitudinal partitioning (blue lines) in both strong and weak scaling tests are very close to the ideal performance (dashed line). For the strong scaling the pure transverse partitioning (red lines) start to deviate significantly from the ideal performance when the number of cores is roughly larger than 100 (or the number of cells per core becomes less than $\sim 10^6$). The performance of the bi-directional partitioning (yellow lines) lies between the pure transverse and longitudinal partitioning when the number of cores $\gtrsim 20$. The one-step single-core run of the strong scaling test ($2^{14} \times 2^{13}$ cells) costs 3300 core seconds. All the above tests excluded the file I/O.

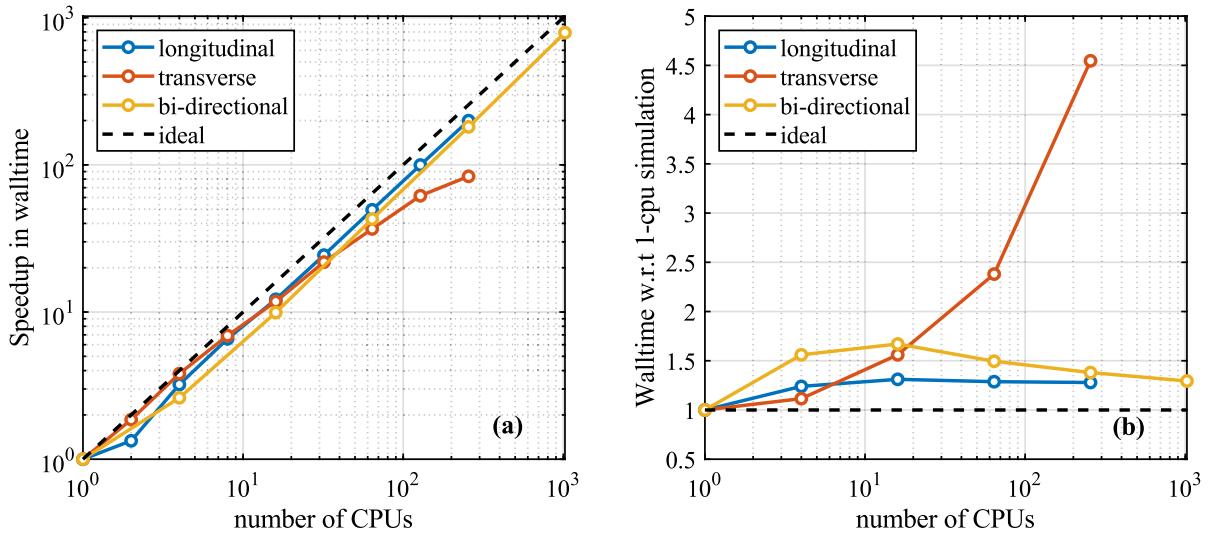


Fig. C.5. (a) Strong scaling and (b) weak scaling of QPAD-PGC.

References

- [1] C. Joshi, T. Katsouleas, Plasma accelerators at the energy frontier and on tabletops, *Phys. Today* 56 (6) (2003) 47–53.
- [2] E. Esarey, C.B. Schroeder, W.P. Leemans, Physics of laser-driven plasma-based electron accelerators, *Rev. Mod. Phys.* 81 (3) (2009) 1229–1285.
- [3] C. Joshi, S. Corde, W.B. Mori, Perspectives on the generation of electron beams from plasma-based accelerators and their near and long term applications, *Phys. Plasmas* 27 (7) (2020) 70602.
- [4] J.M. Dawson, Particle simulation of plasmas, *Rev. Mod. Phys.* 55 (2) (1983) 403–447.
- [5] R.W. Hockney, J.W. Eastwood, *Computer Simulation Using Particles*, CRC Press, 1988.
- [6] C.K. Birdsall, A.B. Langdon, *Plasma Phys. Comput. Simul.* (2005).
- [7] S. Wilks, T. Katsouleas, J.M. Dawson, J.J. Su, Beam loading efficiency in plasma accelerators, Part. Accel. 22 (1987) 100–102.
- [8] M. Tzoufras, W. Lu, F.S. Tsung, C. Huang, W.B. Mori, T. Katsouleas, J. Vieira, R.A. Fonseca, L.O. Silva, Beam loading in the nonlinear regime of plasma-based acceleration, *Phys. Rev. Lett.* 101 (14) (2008) 145002.
- [9] B.B. Godfrey, J.-L. Vay, I. Haber, Numerical stability analysis of the pseudo-spectral analytical time-domain PIC algorithm, *J. Comput. Phys.* 258 (2014) 689–704.
- [10] X. Xu, P. Yu, S.F. Martins, F.S. Tsung, V.K. Decyk, J. Vieira, R.A. Fonseca, W. Lu, L.O. Silva, W.B. Mori, Numerical instability due to relativistic plasma drift in EM-PIC simulations, *Comput. Phys. Commun.* 184 (11) (2013) 2503–2514.
- [11] R. Lehe, A. Lifschitz, C. Thaury, V. Malka, X. Davoine, Numerical growth of emittance in simulations of laser-wakefield acceleration, *Phys. Rev. Spec. Top., Accel. Beams* 16 (2) (2013) 21301.
- [12] F. Li, P. Yu, X. Xu, F. Fiuzza, V.K. Decyk, T. Dalichaouch, A. Davidson, A. Tableman, W. An, F.S. Tsung, R.A. Fonseca, W. Lu, W.B. Mori, Controlling the numerical Cerenkov instability in PIC simulations using a customized finite difference Maxwell solver and a local FFT based current correction, *Comput. Phys. Commun.* 214 (2017) 6–17.
- [13] X. Xu, F. Li, F.S. Tsung, T.N. Dalichaouch, W. An, H. Wen, V.K. Decyk, R.A. Fonseca, M.J. Hogan, W.B. Mori, On numerical errors to the fields surrounding a relativistically moving particle in PIC codes, *J. Comput. Phys.* 413 (2020) 109451.
- [14] F. Li, K.G. Miller, X. Xu, F.S. Tsung, V.K. Decyk, W. An, R.A. Fonseca, W.B. Mori, A new field solver for modeling of relativistic particle-laser interactions using the particle-in-cell algorithm, *Comput. Phys. Commun.* 258 (2021) 107580.
- [15] Top 500 supercomputer sites, <https://www.top500.org/lists/top500/2021/11/>.
- [16] C.D. Decker, W.B. Mori, Group velocity of large amplitude electromagnetic waves in a plasma, *Phys. Rev. Lett.* 72 (4) (1994) 490.
- [17] J.L. Vay, Noninvariance of space- and time-scale ranges under a Lorentz transformation and the implications for the study of relativistic interactions, *Phys. Rev. Lett.* 98 (13) (2007) 130405.
- [18] P. Yu, X. Xu, A. Davidson, A. Tableman, T. Dalichaouch, F. Li, M.D. Meyers, W. An, F.S. Tsung, V.K. Decyk, F. Fiuzza, J. Vieira, R.A. Fonseca, W. Lu, L.O. Silva, W.B. Mori, Enabling Lorentz boosted frame particle-in-cell simulations of laser wakefield acceleration in quasi-3D geometry, *J. Comput. Phys.* 316 (2016) 747–759.
- [19] P. Mora, T.M. Antonsen, Kinetic modeling of intense, short laser pulses propagating in tenuous plasmas, *Phys. Plasmas* 4 (1) (1997) 217–229.
- [20] D.F. Gordon, W.B. Mori, T.M. Antonsen, A ponderomotive guiding center particle-in-cell code for efficient modeling of laser-plasma interactions, *IEEE Trans. Plasma Sci.* 28 (4) (2000) 1135–1143.
- [21] C. Benedetti, C.B. Schroeder, C.G. Geddes, E. Esarey, W.P. Leemans, An accurate and efficient laser-envelope solver for the modeling of laser-plasma accelerators, *Plasma Phys. Control. Fusion* 60 (1) (2017) 014002.
- [22] A.F. Lifschitz, X. Davoine, E. Lefebvre, J. Faure, C. Rechatin, V. Malka, Particle-in-cell modelling of laser-plasma interaction using Fourier decomposition, *J. Comput. Phys.* 228 (5) (2009) 1803–1814.
- [23] A. Davidson, A. Tableman, W. An, F.S. Tsung, W. Lu, J. Vieira, R.A. Fonseca, L.O. Silva, W.B. Mori, Implementation of a hybrid particle code with a PIC description in r-z and a gridless description in ϕ into OSIRIS, *J. Comput. Phys.* 281 (2015) 1063–1077.
- [24] R. Lehe, M. Kirchen, I.A. Andriyash, B.B. Godfrey, J.-L. Vay, A spectral, quasi-cylindrical and dispersion-free particle-in-cell algorithm, *Comput. Phys. Commun.* 203 (2016) 66–82.
- [25] K.V. Lotov, Fine wakefield structure in the blowout regime of plasma wakefield accelerators, *Phys. Rev. Spec. Top., Accel. Beams* 6 (6) (2003) 9–14.
- [26] C. Huang, V.K. Decyk, C. Ren, M. Zhou, W. Lu, W.B. Mori, J.H. Cooley, T.M. Antonsen, T. Katsouleas, QUICKPIC: a highly efficient particle-in-cell code for modeling wakefield acceleration in plasmas, *J. Comput. Phys.* 217 (2) (2006) 658–679.
- [27] W. An, V.K. Decyk, W.B. Mori, T.M. Antonsen, An improved iteration loop for the three dimensional quasi-static particle-in-cell algorithm: QuickPIC, *J. Comput. Phys.* 250 (2013) 165–177.

- [28] T. Mehrling, C. Benedetti, C.B. Schroeder, J. Osterhoff, HiPACE: a quasi-static particle-in-cell code, *Plasma Phys. Control. Fusion* 56 (8) (2014) 84012.
- [29] F. Li, W. An, V.K. Decyk, X. Xu, M.J. Hogan, W.B. Mori, A quasi-static particle-in-cell algorithm based on an azimuthal Fourier decomposition for highly efficient simulations of plasma-based acceleration: QPAD, *Comput. Phys. Commun.* 261 (2021) 107784.
- [30] T. Wang, V. Khudik, G. Shvets, WAND-PIC: a three-dimensional quasi-static particle-in-cell code with parallel multigrid solver and without predictor-corrector, <https://arxiv.org/abs/2012.00881v1>, dec 2020.
- [31] S. Diederichs, C. Benedetti, A. Huebl, R. Lehe, A. Myers, A. Sinn, J.L. Vay, W. Zhang, M. Thévenet, HiPACE++: a portable, 3D quasi-static particle-in-cell code, *Comput. Phys. Commun.* 278 (2022) 108421.
- [32] B. Feng, C. Huang, V. Decyk, W.B. Mori, P. Muggli, T. Katsouleas, Enhancing parallel quasi-static particle-in-cell simulations with a pipelining algorithm, *J. Comput. Phys.* 228 (15) (2009) 5340–5348.
- [33] R.A. Fonseca, L.O. Silva, F.S. Tsung, V.K. Decyk, W. Lu, C. Ren, W.B. Mori, S. Deng, S. Lee, T. Katsouleas, J.C. Adam, OSIRIS: a three-dimensional, fully relativistic particle in cell code for modeling plasma based accelerators, in: P. Sloot, A. Hoekstra, C. Tan, J. Dongarra (Eds.), *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, in: LNCS, vol. 2331, 2002, pp. 342–351.
- [34] R.G. Hemker, Particle-in-cell modeling of plasma-based accelerators in two and three dimensions, arXiv:1503.00276, 2015.
- [35] J.W. Thomas, *Numerical Partial Differential Equations: Finite Difference Methods*, vol. 22, Springer Science & Business Media, 2013.
- [36] C. Levit, Parallel solution of pentadiagonal systems using generalized odd-even elimination, 1989, pp. 333–336.
- [37] J. Boris, Relativistic plasma simulation-optimization, in: 4th Conference on Numerical Simulation of Plasma, 1970, p. 3.
- [38] W. Lu, M. Tzoufras, C. Joshi, F.S. Tsung, W.B. Mori, J. Vieira, R.A. Fonseca, L.O. Silva, Generating multi-GeV electron bunches using single stage laser wakefield acceleration in a 3D nonlinear regime, *Phys. Rev. Spec. Top., Accel. Beams* 10 (6) (2007) 61301.
- [39] J. Vieira, J.T. Mendonça, Nonlinear laser driven donut wakefields for positron and electron acceleration, *Phys. Rev. Lett.* 112 (21) (2014) 215001.
- [40] J.B. Rosenzweig, B. Breizman, T. Katsouleas, J.J. Su, Acceleration and focusing of electrons in two-dimensional nonlinear plasma wake fields, *Phys. Rev. A* 44 (10) (1991) R6189–R6192.
- [41] W. Lu, C. Huang, M. Zhou, W.B. Mori, T. Katsouleas, Nonlinear theory for relativistic plasma wakefields in the blowout regime, *Phys. Rev. Lett.* 96 (16) (2006) 165002.
- [42] L. Allen, M.W. Beijersbergen, R.J.C. Spreeuw, J.P. Woerdman, Orbital angular momentum of light and the transformation of Laguerre-Gaussian laser modes, *Phys. Rev. A* 45 (1992) 8185–8189.