

Computational Physics

Gridless quasistatic model for efficient simulation of plasma-based accelerators

Á. Ferran Pousa ^{a,*}, W. M. den Hertog ^{a,b,1}, S. Diederichs ^{a,c}, A. Martinez de la Ossa ^a,
J. L. Ordóñez Carrasco ^{a,d}, A. Sinn ^a, M. Thévenet ^{a,*}

^a Deutsches Elektronen-Synchrotron DESY, Notkestr. 85, Hamburg, 22607, Germany

^b Universidade de Santiago de Compostela, Santiago de Compostela, Spain

^c CERN, Espl. des Particules 1, Geneva, 1211, Switzerland

^d Universidad Carlos III de Madrid, avda. Universidad, 30, Leganés, Madrid, 28911, Spain

ARTICLE INFO

The review of this paper was arranged by Prof. Andrew Hazel

Keywords:

Plasma acceleration
Laser-plasma interaction
Particle-in-cell

ABSTRACT

The accurate modeling of plasma-based accelerators relies on costly numerical simulations due to the complexity of laser-plasma and beam-plasma interactions. Several strategies can highly reduce the computational cost compared to 3D first-principles particle-in-cell simulations, such as exploiting the near axial symmetry and quasistatic nature of plasma wakefields in many practical cases. Here, we propose a quasistatic algorithm that enables the modeling of axially symmetric plasma wakes without the need of a numerical grid. The gridless approach allows extremely fine features to be resolved without a dramatic increase in computational cost. This is critical, e.g., for the design of future plasma-based colliders with nanometer emittance beams. The proposed model has been implemented in the Wake-T code, where it is coupled to a laser envelope solver and a particle beam pusher to enable the efficient simulation of laser- and beam-driven plasma accelerators.

1. Introduction

Plasma-based accelerator (PBA) technology has the potential of shrinking the size of future particle accelerators thanks to its capability of sustaining up to $\sim 100 \text{ GV m}^{-1}$ gradients. The promise of extreme gradients and compactness has driven an increasing interest in the research and development of this technology, with breakthroughs in the achievement of GeV energies [1], high beam quality [2–4], reproducibility [5], and even free-electron lasing [6–8].

Critical to this progress and to its continuity is the use of numerical simulation tools that enable testing new concepts, optimizing the accelerator performance, and improving the understanding of experimental results. However, 3D full electromagnetic particle-in-cell (PIC) simulations are computationally demanding. Due to the Courant-Friedrichs-Lewy (CFL) condition, resolving the sub-micron features of the plasma wake over centimeter to meter distances would result in $\mathcal{O}(10^5)$ to $\mathcal{O}(10^7)$ time steps, respectively, which can require in excess of thousands of node-hours on a high-performance computing cluster [9]. This poses a practical limit to such simulations, and therefore to the development of new solutions to the current challenges of PBA technology. A particularly challenging example is the start-to-end modeling of plasma-

based collider concepts [10], which would require the simulation of tens to hundreds of accelerating stages with extreme resolution due to the nanometer emittance of the accelerated beams.

Developing more efficient algorithms is therefore essential to enable large-scale studies and to agilitize the progress of PBAs. Several approaches have been developed that exploit the symmetries and characteristic properties of plasma wakefield generation to achieve a lower computational cost. These include the quasistatic approximation (QSA) [11–14], the transformation into a Lorentz-boosted frame [15], the use of a laser envelope model [11], and the use of algorithms with cylindrical or quasi-cylindrical geometry [16]. A particularly efficient example is the approach proposed by P. Baxevanis and G. Stupakov [17], which makes use of the QSA and axial symmetry to develop a model for the plasma response that does not require a numerical grid nor a predictor-corrector loop [18]. However, this method was developed to describe the blowout regime in uniform electron plasmas that are driven by particle beams with analytic density profiles, and did not allow for the modeling of more complex and realistic cases.

Here, we propose a generalization of the gridless algorithm that enables the modeling of arbitrary density profiles, multiple plasma species (including the motion of background plasma ions), as well as a

* Corresponding authors.

E-mail addresses: angel.ferran.pousa@desy.de (Á. Ferran Pousa), maxence.thevenet@desy.de (M. Thévenet).

¹ Now at UHasselt, Faculty of Sciences, Agoralaan, 3590 Diepenbeek, Belgium.

non-uniform radial discretization of the plasma. The model has been implemented in the Wake-T code [19], where it has been coupled with a laser envelope solver [20] and a particle beam pusher to enable the full simulation of laser- and beam-driven plasma accelerators [21]. While the laser and particle beam solvers do need a numerical grid, the size and resolution of each grid is independent of the others and of the gridless background plasma. This allows for great flexibility, with grids that can dynamically adapt to the size of the particle beams and that are therefore able to resolve arbitrarily small particle distributions without increasing the number of grid points. This is critical, among others, for the cost-effective simulation of collider concepts. Benchmarks of the proposed algorithm against the quasi-3D full-PIC code FBPIC [22] and the 3D QSA code HiPACE++ [23] demonstrate the high fidelity and drastic cost reduction of the proposed algorithm.

2. Gridless quasistatic model

2.1. Theoretical framework

The underlying model for the wakefield generation is based on the gridless approach originally developed in Ref. [17], which describes the non-linear, axially symmetric response of a uniform electron plasma with a fixed ion background to the space-charge field of analytical electron beam profiles. Our model generalizes this initial derivation to non-uniform plasmas with an arbitrary number of mobile species (both electrons and ions) which respond to the space-charge fields of beam particle distributions as well as the ponderomotive force of a laser pulse. The result is a numerical model of broad applicability that enables the efficient simulation of laser- and beam-driven plasma accelerators when the wakefield geometry can be assumed to be axisymmetric.

We introduce the model starting from Maxwell's equations in axial symmetry under the QSA [11,14], using $\zeta = k_p(z - ct)$ and $\rho = k_p r$ as the normalized longitudinal and radial coordinates, where $k_p^{-1} = \sqrt{m_e \epsilon_0 c^2 / n_0 e^2}$ is the plasma skin depth, n_0 is a reference plasma density, t is time, and z and r are the longitudinal and radial dimensions. This results in the following set of equations:

$$\frac{1}{\rho} \frac{\partial}{\partial \rho} (\rho E_r) + \frac{\partial E_z}{\partial \zeta} = \sum_{s \in S^{(p)}} q_s n_s + \sum_{s \in S^{(b)}} q_s n_s, \quad (1a)$$

$$\frac{\partial E_z}{\partial \rho} = \sum_{s \in S^{(p)}} q_s n_s v_{r,s}, \quad (1b)$$

$$\frac{\partial}{\partial \zeta} (B_\theta - E_r) = - \sum_{s \in S^{(p)}} q_s n_s v_{r,s}, \quad (1c)$$

$$\frac{1}{\rho} \frac{\partial}{\partial \rho} (\rho B_\theta) + \frac{\partial E_z}{\partial \zeta} = \sum_{s \in S^{(p)}} q_s n_s v_{z,s} + \sum_{s \in S^{(b)}} q_s n_s, \quad (1d)$$

where the right-hand side contains the contribution of all N_p plasma particle species $S^{(p)} = \{s_i^{(p)}\}_{i \in [1, N_p]}$ and N_b particle beam species $S^{(b)} = \{s_i^{(b)}\}_{i \in [1, N_b]}$ to the charge density and current. Particle beam species are assumed to be ultra-relativistic and to propagate with no radial motion, i.e., $v_{z,s} = 1$ and $v_{r,s} = 0 \forall s \in S^{(b)}$. All quantities in Eqs. (1a)–(1d) and in the following derivation are dimensionless: velocity v is normalized to the speed of light c , momentum u to $m_e c$, mass m to the electron mass m_e , charge q to the elementary charge e , electric field E to $m_e c \omega_p / e$, magnetic field B to $m_e c \omega_p / e$, scalar potential ϕ to $m_e c^2 / e$, vector potential a to $m_e c / e$, and particle density n to n_0 . For simplicity, we drop the ζ and ρ dependencies from the notation of any field quantity F , such that $F(\zeta, \rho)$ is written F whenever unambiguous.

Introducing the potential $\psi = \phi - a_z$ allows the electric and magnetic fields to be written as

$$E_z = - \frac{\partial \psi}{\partial \zeta}, \quad (2a)$$

$$B_\theta - E_r = \frac{\partial \psi}{\partial \rho}, \quad (2b)$$

which, if introduced into Eqs. (1a)–(1d), yields the following equations for ψ and B_θ :

$$\frac{1}{\rho} \frac{\partial}{\partial \rho} \left(\rho \frac{\partial \psi}{\partial \rho} \right) = - \sum_{s \in S^{(p)}} (1 - v_{z,s}) q_s n_s, \quad (3a)$$

$$\begin{aligned} \frac{\partial}{\partial \rho} \left[\frac{1}{\rho} \frac{\partial}{\partial \rho} (\rho B_\theta) \right] &= - \frac{\partial}{\partial \zeta} \sum_{s \in S^{(p)}} q_s n_s v_{r,s} \\ &+ \frac{\partial}{\partial \rho} \left[\sum_{s \in S^{(p)}} q_s n_s v_{z,s} + \sum_{s \in S^{(b)}} q_s n_s \right]. \end{aligned} \quad (3b)$$

In principle, an additional equation for ψ can be obtained from Eq. (1b), i.e., $\frac{\partial}{\partial \rho} \left(\frac{\partial \psi}{\partial \zeta} \right) = \sum_{s \in S^{(p)}} q_s n_s v_{r,s}$, but we will not be making use of it because integrating Eq. (3a) is already sufficient to find a solution for ψ .

The theoretical framework of the wakefield model is completed with the continuity equation for the fluid quantities for each plasma species [17]

$$\frac{1}{\rho} \frac{\partial}{\partial \rho} (\rho n_s v_{r,s}) - \frac{\partial}{\partial \zeta} [n_s (1 - v_{z,s})] = 0 \quad (4)$$

and the constant of motion for an initially cold and unperturbed plasma [11,12]

$$\gamma_s - u_{z,s} + \frac{q_s}{m_s} \psi = 1. \quad (5)$$

This constant of motion represents the conservation of energy in the co-moving frame ζ . It dictates how the potential energy from the wakefield generated by the incoming driver shall be translated into kinetic energy in the background plasma and thereby limits the possible trajectories that a plasma particle might follow, as we will see in the next section.

2.2. Discretization of plasma species into macroparticles

Solving Eqs. (3a) and (3b) requires finding an appropriate expression for the plasma source terms on the right hand side. The beam density distributions $n_s^{(b)}$ are instead assumed to be known and are provided externally. They could be, e.g., an analytical expression or a particle density deposited on a grid. Following the original derivation in Ref. [17], the plasma is discretized into a set of $\mathcal{N}_s$ macroparticles $P_s = \{p_1^s, p_2^s, \dots, p_{\mathcal{N}_s}^s\} \forall s \in S^{(p)}$. Each macroparticle p_j^s is described as a delta function $\delta_j^s(\rho, \zeta) = \delta[\rho - \rho_j^s(\zeta)]$ with a radial and longitudinal position $\rho_j^s(\zeta)$ and ζ_j^s , momenta $u_{r,j}^s(\zeta)$ and $u_{z,j}^s(\zeta)$, and a weight w_j^s that contains the radially integrated particle density represented by the macroparticle. Similar to the fluid quantities, we will avoid writing the explicit dependency of the macroparticle quantities except were needed.

The macroparticles are initialized at the front of the simulation domain ($\zeta = 0$, ahead of all laser and particle beams, where the plasma is unperturbed) with a radial position $\rho_{j0}^s = \rho_j^s(\zeta = 0) = \sum_{k < j} \Delta \rho_k^s + \Delta \rho_j^s / 2$, where $\Delta \rho_j^s$ is the radial extent of plasma represented by the p_j^s macroparticle. In the case of uniform macroparticle spacing (i.e., $\Delta \rho_j^s = \Delta \rho^s$), as considered in Ref. [17], the initial radial distribution simplifies to $\rho_{j0}^s = (1/2 + j) \Delta \rho$. Here, however, we make use of the general expression. This enables a strategy for refining the plasma discretization transversely only where needed (e.g., close to the axis) and therefore minimizing the computational cost of the algorithm. This is discussed in detail in Section 4. The macroparticle weights depend on the initial radial distribution, and are given by:

$$w_j^s = \int_{\rho_{j0}^s - \Delta \rho_j^s / 2}^{\rho_{j0}^s + \Delta \rho_j^s / 2} n_s(0, \rho) \rho d\rho \simeq n_{j0}^s \rho_{j0}^s \Delta \rho_j^s, \quad (6)$$

where we have assumed that the initial particle density profiles of all plasma species are identical (i.e., $n_s(0, \rho) = n(0, \rho) \forall s \in S^{(p)}$) and that the density is uniform (i.e., $n(0, \rho) = n(0, \rho_{j0}^s) = n_{j0}^s$) within the bounds of the integral, i.e., within the immediate surroundings of the macroparticle.

Eqs. (4)–(6) allow the fluid quantities in the right-hand side of Eqs. (3a)–(3b) to be obtained for each species $s \in S^{(p)}$ by summing over

all macroparticles $p_j^s \in P_s$:

$$n_s = \sum_{p_j^s \in P_s} \frac{w_j^s}{\rho_j^s [1 - v_{z,j}^s]} \delta_j^s, \quad (7a)$$

$$n_s(1 - v_{z,s}) = \sum_{p_j^s \in P_s} \frac{w_j^s}{\rho_j^s} \delta_j^s, \quad (7b)$$

$$n_s v_{r,s} = \sum_{p_j^s \in P_s} \frac{w_j^s v_{r,j}^s}{\rho_j^s [1 - v_{z,j}^s]} \delta_j^s, \quad (7c)$$

$$n_s v_{z,s} = \sum_{p_j^s \in P_s} \frac{w_j^s v_{z,j}^s}{\rho_j^s [1 - v_{z,j}^s]} \delta_j^s. \quad (7d)$$

2.3. Equations of motion for plasma macroparticles

After initialization, the distribution of plasma macroparticles can be perturbed by the space-charge field of particle beams, the ponderomotive force of a laser pulse, as well as the resulting fields generated by the plasma response. To simplify the notation and avoid the superindex s in the macroparticle quantities, we will use in the following a single index i to identify each plasma macroparticle independently of its corresponding species. This index is defined such that macroparticles are sorted by increasing radius, i.e., $\rho_{i+1} > \rho_i \forall i$ always. Assuming that Eq. (5) holds, the evolution of each plasma macroparticle is given by the equations of motion

$$\frac{du_{r,i}}{d\zeta} = \frac{q_i}{m_i} \left(\frac{\gamma_i}{\bar{\psi}_i} \partial_\rho \psi_i - B_{\theta,i} - \frac{q_i}{4m_i} \frac{\nabla_\perp |\hat{a}|_i^2}{\bar{\psi}_i} \right), \quad (8a)$$

$$\frac{d\rho_i}{d\zeta} = -\frac{u_{r,i}}{\bar{\psi}_i}, \quad (8b)$$

$$u_{z,i} = \frac{1 + u_{r,i}^2 - \bar{\psi}_i^2 + \frac{q_i^2}{2m_i^2} |\hat{a}|_i^2}{2\bar{\psi}_i}, \quad (8c)$$

$$\gamma_i = \frac{1 + u_{r,i}^2 + \bar{\psi}_i^2 + \frac{q_i^2}{2m_i^2} |\hat{a}|_i^2}{2\bar{\psi}_i}, \quad (8d)$$

where we define $\bar{\psi}_i = 1 - \frac{q_i}{m_i} \psi_i$ and include the ponderomotive force of a linearly polarized laser pulse with a transverse normalized vector potential $a_\perp(\zeta, \rho, t) = \Re[\hat{a}(\zeta, \rho, t) \exp(-ik_0 \zeta)]$. In Eqs. (5)–(8d) and in the following, we adopt the notation $\mathcal{F}_i = \mathcal{F}(\zeta_i, \rho_i)$ and $\partial_v \mathcal{F}_i = \frac{\partial \mathcal{F}}{\partial v}(\zeta_i, \rho_i) \forall v \in \{\zeta, \rho\}$ for the field quantities and their derivatives evaluated at the position of the macroparticles.

2.4. Numerical integration of wakefields

Eqs. (8a)–(8d) can be integrated numerically by discretizing the longitudinal direction with a finite step $\Delta\zeta$. Then, by recovering the plasma fluid quantities with Eqs. (7a)–(7d), a numerical solution of Eqs. (3a) and (3b) can be obtained. Starting with Eq. (3b), the magnetic field can be first separated into two components $B_\theta = B_\theta^{(b)} + B_\theta^{(p)}$ that represent the contribution from the current of the particle beams $B_\theta^{(b)}$ and the plasma species $B_\theta^{(p)}$. In this way, the magnetic field from the beams can be obtained from a given charge distribution (either an analytical profile or a discretized density deposited on a grid) as

$$\frac{\partial}{\partial \rho} \left[\frac{1}{\rho} \frac{\partial}{\partial \rho} (\rho B_\theta^{(b)}) \right] = \frac{\partial}{\partial \rho} \sum_{s \in S^{(b)}} q_s n_s, \quad (9)$$

while $B_\theta^{(p)}$ is given by

$$\frac{\partial}{\partial \rho} \left[\frac{1}{\rho} \frac{\partial}{\partial \rho} (\rho B_\theta^{(p)}) \right] = \sum_{s \in S^{(p)}} q_s \left[\frac{\partial n_s v_{z,s}}{\partial \rho} - \frac{\partial n_s v_{r,s}}{\partial \zeta} \right]. \quad (10)$$

The radial derivative of the longitudinal plasma current is found using Eqs. (5) and (7d)

$$\sum_{s \in S^{(p)}} q_s \frac{\partial n_s v_{z,s}}{\partial \rho} = \sum_i \frac{q_i w_i}{\rho_i} \left(\frac{\gamma_i}{\bar{\psi}_i} - 1 \right) \partial_\rho \delta_i, \quad (11)$$

while the longitudinal derivative of the radial current can be obtained from Eqs. (5), (7c), (8a) and (8b) as [17]

$$\begin{aligned} \sum_{s \in S^{(p)}} q_s \frac{\partial n_s v_{r,s}}{\partial \zeta} = \sum_i \frac{q_i w_i}{\rho_i} \left[\right. \\ \frac{q_i}{m_i} \frac{\delta_i}{\bar{\psi}_i} \left(\frac{\gamma_i \partial_\rho \psi_i}{\bar{\psi}_i} - B_{\theta,i}^{(p)} - \frac{q_i}{m_i} \frac{\partial_\rho |a|_i^2}{\bar{\psi}_i} \right) \\ + \frac{q_i}{m_i} \frac{u_{r,i} \delta_i}{\bar{\psi}_i^2} \left(-\frac{u_{r,i} \partial_\rho \psi_i}{\bar{\psi}_i} + \partial_\zeta \psi_i \right) \\ \left. + \frac{u_{r,i}^2 \partial_\rho \delta_i}{\bar{\psi}_i^2} + \frac{u_{r,i}^2 \delta_i}{\rho_i \bar{\psi}_i^2} \right]. \end{aligned} \quad (12)$$

This expression is equivalent to the explicit form described in [24], and therefore allows for a numerical solution of the plasma evolution without the need of a predictor-corrector loop [25,26].

Following our assumption that $\rho_{i+1} > \rho_i$, the magnetic field between macroparticles i and $i+1$ can be expressed as

$$B_\theta^{(p)} = a_i \rho + \frac{b_i}{\rho}, \quad (13)$$

which follows from the fact that there are no plasma currents in the space between the two plasma macroparticles (delta functions), i.e., $\partial_r [\partial_r (r B_\theta^{(p)}) / r] = 0$. The a_i and b_i coefficients are obtained by solving the recursive system

$$\begin{pmatrix} a_i \\ b_i \end{pmatrix} = \begin{pmatrix} 1 + \frac{A_i \rho_i}{2} & \frac{A_i}{2\rho_i} \\ -\frac{A_i \rho_i^3}{2} & 1 - \frac{A_i \rho_i}{2} \end{pmatrix} \begin{pmatrix} a_{i-1} \\ b_{i-1} \end{pmatrix} + \begin{pmatrix} \frac{1}{4} (2B_i + A_i C_i) \\ \frac{1}{4} \rho_i (4C_i - 2B_i \rho_i - A_i C_i \rho_i) \end{pmatrix} \quad (14)$$

with boundary conditions $b_0 = 0$ and $a_N = 0$ [17]. The coefficients A_i , B_i and C_i result from rewriting Eq. (10) as $\frac{\partial}{\partial \rho} \left[\frac{1}{\rho} \frac{\partial}{\partial \rho} (\rho B_\theta^{(p)}) \right] = \sum_i \frac{w_i}{\rho_i} (A_i \delta_i B_\theta^{(p)} + B_i \delta_i + C_i \partial_\rho \delta_i)$ after substituting Eqs. (11) and (12). Thus, they are given by

$$A_i = \frac{w_i}{\rho_i} \frac{q_i^2}{m_i} \frac{1}{\bar{\psi}_i}, \quad (15a)$$

$$\begin{aligned} B_i = & -\frac{w_i}{\rho_i} \frac{q_i^2}{m_i} \frac{\gamma_i \partial_\rho \psi_i}{\bar{\psi}_i^2} + \frac{w_i}{\rho_i} \frac{q_i^2}{m_i} \frac{B_{\theta,i}^{(b)}}{\bar{\psi}_i} \\ & + \frac{w_i}{\rho_i} \frac{q_i^3}{m_i^2} \frac{\partial_\rho |\hat{a}|_i^2}{\bar{\psi}_i^2} - \frac{q_i w_i}{\rho_i^2} \frac{u_{r,i}^2}{\bar{\psi}_i^2} \\ & + \frac{w_i}{\rho_i} \frac{q_i^2}{m_i} \frac{u_{r,i}^2 \partial_\rho \psi_i}{\bar{\psi}_i^3} - \frac{w_i}{\rho_i} \frac{q_i^2}{m_i} \frac{u_{r,i} \partial_\zeta \psi_i}{\bar{\psi}_i^2}, \end{aligned} \quad (15b)$$

$$C_i = \frac{q_i w_i}{\rho_i} \left(\frac{\gamma_i}{\bar{\psi}_i} - 1 \right) - \frac{q_i w_i}{\rho_i} \frac{u_{r,i}^2}{\bar{\psi}_i^2}. \quad (15c)$$

Details about how to solve this recursive system of equations can be found in Appendix A.

A numerical solution for ψ and $\partial\psi/\partial\rho$ can be found by substituting Eq. (7b) into Eq. (3a), which, after integration from the axis to ρ , yields

$$\frac{\partial\psi}{\partial\rho} = -\frac{1}{\rho} \sum_{\rho_i < \rho} q_i w_i. \quad (16)$$

By integrating again Eq. (16) with boundary condition $\psi(\rho \rightarrow \infty) = 0$ and assuming that the plasma is initially neutral ($\sum_i q_i w_i = 0$) it is found




Fig. 1. Plasma response to (a) an electron and (b) laser beam as calculated with the gridless quasistatic model. The trajectory of each plasma electron macroparticle is shown in blue. These trajectories can be used to calculate the resulting electromagnetic fields at any location using Eqs. (13)–(18). In this case, the E_x and E_z fields are shown for both the beam- and laser-driven cases in panels (c)–(f). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

that

$$\psi = - \sum_i q_i w_i \ln(\rho_i) - \sum_{\rho_i < \rho} q_i w_i \ln\left(\frac{\rho}{\rho_i}\right). \quad (17)$$

Finally, taking the longitudinal derivative of Eq. (17) results in

$$\frac{\partial \psi}{\partial \zeta} = - \sum_i \frac{q_i w_i}{\rho_i} \frac{u_{r,i}}{\psi_i} + \sum_{\rho_i < \rho} \frac{q_i w_i}{\rho_i} \frac{u_{r,i}}{\psi_i}, \quad (18)$$

where Eq. (8b) has been used.

Eqs. (13)–(18) allow the calculation of $B_\theta^{(p)}$, ψ , $\partial\psi/\partial\rho$, $\partial\psi/\partial\zeta$ (and therefore of E_r and E_z) at any location ρ within the radial plasma column if the position and momentum of the macroparticles, as well the source terms $B_\theta^{(b)}$, $|\hat{a}|_i^2$ and $\partial_\rho |\hat{a}|_i^2$ are known. Thus, they can be used to determine the fields experienced by the plasma macroparticles and to calculate their motion according to Eqs. (8a) and (8b), which can be solved numerically to obtain the response of the plasma to particle and laser beams, as illustrated in Fig. 1. Numerical algorithms for this task include the Runge-Kutta (used in Ref. [17]) or Adams-Bashforth methods [27], where the latter can be more computationally efficient by not requiring the calculation of intermediate steps.

Note that Eqs. (15a)–(18) reduce to those in Ref. [17] in the case of an electron plasma ($q_i = -1$, $m_i = 1$) with no laser source terms ($|\hat{a}|^2 = 0$) and changing to the longitudinal variable $\xi = -\zeta$.

A remarkable property of this numerical model is that Eqs. (13)–(18) allow the plasma fields to be sampled with arbitrary radial resolution. Compared to a typical PIC algorithm, the plasma macroparticles do not interact with each other through a numerical grid. Instead, the fields are calculated directly at the location of each macroparticle, so that the resolution is only limited by the inter-macroparticle distance and the longitudinal step $\Delta\zeta$ of the numerical integrator. Thus, the presented method can resolve extremely fine features in regions of high macroparticle concentration without the need of a high-resolution grid, and therefore without increasing the computational cost of the algorithm. This is useful, for example, for modeling the blowout sheath, which is a narrow region where the plasma electrons expelled by the driver accumulate and that features a sharp spike in density and current whose shape is

important for determining the blowout fields [28]. In these cases, the gridless method is therefore expected to reach convergence at a fraction of the computational cost, compared to an equivalent gridded algorithm.

While the gridless method is fully kinetic, it inherits the limitations of the quasistatic approximation: wavebreaking is not properly captured, and the results of a strong blowout regime should be taken with care behind the back of the first bubble.

2.5. Key algorithmic adjustments

Special care should be taken when evaluating the fields at the exact location of the plasma macroparticles. This is due to their shape being a delta function, which results in two main issues that should be handled.

The first one is that the radial sums in Eqs. (13)–(18) present a discontinuity at $\rho = \rho_i$, i.e., where they need to be evaluated to obtain the field values at the location of each plasma macroparticle. Reference [17] addresses this issue by taking the average of the two neighboring values, but we have found this approach to still introduce significant error, particularly for the macroparticles closest to the axis. Instead, we propose to compute the fields as described by the equations above, but using the modified weight $\hat{w}_i = \int_{\rho_{i,0}-\Delta\rho_i/2}^{\rho_{i,0}} n(0, \rho_i) \rho d\rho = w_i/2 - \Delta\rho_i^2/8$ instead of Eq. (6) for the macroparticle at which the field is being calculated. The modified weight $\hat{w}_i$, compared to w_i , takes into account only the charge of the macroparticle below ρ_i , which is the charge that would actually be contributing to the radial integral being computed.

The second issue comes from having to consider multiple plasma species of opposite charge. This is in contrast to Ref. [17], where only the plasma electrons were discretized and a uniform ion background was assumed. The challenge here is to achieve a field that is radially smooth even though the underlying plasma is a non-smooth distribution of delta functions. For example, assuming that the same number of electron and ion macroparticles are initialized at exactly the same positions, Eq. (16) would correctly result in $\partial_\rho \psi(\rho) = 0$. However, the tiniest displacement between an electron and its corresponding ion would lead to large, nonphysical gradients between them. The solution we have adopted to this issue is to solve Eqs. (13)–(18) separately for each plasma species, and afterwards to gather the fields of a species s_A into another species s_B using linear interpolation. More specifically, this means that the field F_A (any of ψ , $\partial_\rho \psi$, $\partial_\zeta \psi$ or $B_\theta^{(p)}$) of species s_A at the position $\rho_{i,B}$ of a macroparticle from species s_B is given by $F_A(\rho_{i,B}) = F_{A,j} + \frac{F_{A,j+1} - F_{A,j}}{\rho_{j+1,A} - \rho_{j,A}} (\rho_{i,B} - \rho_{j,A})$, where j is the index of the last macroparticle from A with $\rho_{j,A} < \rho_{i,B}$. This method allows for smoother interpolation without resorting to a grid. Another assumption we make to simplify the algorithm is to compute $B_\theta^{(p)}$ in Eq. (13) by taking only the electrons into account, since the ion currents are negligible in comparison to those of the electrons due to their much bigger mass ($v_{\text{ion}}/v_e \propto m_e/m_{\text{ion}} < 10^{-3}$).

3. Integration into wake-T

Wake-T (Wakefield particle Tracker) is an open-source simulation code for charged particle beams that is specialized in the modeling of PBAs [19]. It aims to provide a fast and inexpensive alternative to general-purpose PIC codes by making use of reduced physical models for the plasma wake. The approach followed by Wake-T is similar to that of other tools used for the modeling of conventional accelerators, such as ASTRA [29], whose focus is on the numerical integration of the motion of beams of charged particles in the presence of electromagnetic fields. The fields experienced by the beam particles can be given as analytical expressions or as multidimensional grids whose value is interpolated to the location of the beam particles. The motion of the beam particles in these fields is integrated using either the Boris method [30,31] or a classical fourth order Runge-Kutta method. Both of them are parallelized using multi-threading to take advantage of multi-core CPUs with shared memory (i.e., in a single node). The code is written in pure




Fig. 2. Simplified view of the main Wake-T particle tracking loop (left) and the implementation of the gridless model (right). The main loop evolves each of the elements in the simulation from their initial state up to a maximum final time $t_{\max}$. Since each element can have a different time step, every iteration begins by determining the next element to evolve by comparing the current simulation time t_{sim} with the current time t_e and time step Δt_e of each element. When updating the plasma fields with the gridless model, the loop shown on the right is executed. The plasma macroparticles are initialized at the right boundary $\zeta = 0$ and are iteratively evolved until the left boundary $\zeta_{\min}$ with a step $\Delta\zeta$. Every iteration over longitudinal slices involves sorting the macroparticles radially, computing the required quantities at the location of the macroparticles, and evolving their radial motion. A grid is used to communicate between the plasma and the laser and macroparticle beams, as explained in the main text. The steps that involve communication with a grid are highlighted with a blue gridded background. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

PYTHON and uses NUMBA [32] to compile the most compute-intensive functions into fast machine code using LLVM [33]. Being part of the BLAST toolkit [34], Wake-T is fully compliant with the openPMD data standard [35] for both input and output of the simulations.

The particle tracking loop in Wake-T evolves the simulation by advancing the beams, the fields, and the output diagnostics in time. Each of these elements can advance with a different time step, and this time step can be fixed or vary throughout the simulation. This is done in an iterative process where, at each step, the next element to advance is determined based on the current simulation time, the time step of the elements, and the last time they were advanced. In this context, element is one of beam, field and diagnostics. This is illustrated in the left diagram of Fig. 2 (Main Wake-T loop).

Several reduced wakefield models were previously implemented in Wake-T, including the original gridless [17] model generalized to support parabolic plasma profiles and a laser envelope solver [20], which allowed for fast optimization of laser-plasma accelerators (LPAs) [36] or multistage accelerators [21]. Key features enabled by the present work include the use of multiple plasma species (and thus the modeling of ion motion) and arbitrary density profiles. Furthermore, adaptive grids (see Section 4) and non-uniform plasma macroparticles spacing enable the modeling of arbitrarily thin beams at constant computational cost. These features are key, for example, to enable the simulation of future plasma-based collider concepts, such as HAHLF [37].

A high-level overview of how the new gridless algorithm was implemented in Wake-T is shown in Fig. 2. At a given time step, the plasma

column is iteratively evolved from the front to the back of the domain with a longitudinal step $\Delta\zeta$. Each iteration starts by sorting the plasma macroparticles radially and gathering the beam and laser source terms at each macroparticle. This information is then used to calculate the wakefield potential and its derivatives with Eqs. (16)–(18). Knowing ψ_i and the source terms then allows $u_{z,i}$ and γ_i to be updated using Eqs. (8c) and (8d), and to calculate $B_{\theta,i}^{(p)}$ using Eqs. (13)–(15c). Finally, the plasma macroparticles are evolved to the next longitudinal slice (i.e., by $-\Delta\zeta$) using an Adams-Bashforth algorithm of second order. Once this iterative process is completed ($\zeta = \zeta_{\min}$), the plasma response is known in the whole domain at this time step. Its effect on the laser pulses (refraction) and particle beams (electromagnetic force) can then be accounted for when advancing these by one time step Δt .

As indicated in Fig. 2, it should be noted that while the algorithm to evolve the plasma macroparticles is gridless, it still needs to be able to communicate with an external grid to handle the laser-plasma and beam-plasma interaction. The use of a laser grid is required by the envelope solver, which implies that the laser source terms need to be interpolated from the grid onto the plasma macroparticles, and that the macroparticles need to deposit the plasma susceptibility $\chi = n/\gamma$ into the laser grid. The use of a regular grid for the beam-plasma communication is not strictly required, but it allows for a simpler and potentially more efficient algorithm by using well-established interpolation methods and avoiding the need to sort the beam macroparticles radially and longitudinally at every step. With this approach, Eqs. (13)–(18) are used to calculate the plasma fields at the grid nodes, and their value at the beam

macroparticles is later found through interpolation. An implementation with direct beam-plasma interaction (i.e., without an intermediate grid) could be explored in the future. The laser and beam grids are independent from each other and can have different extents and resolutions. The time step of the laser envelope solver can also be set to be an integer fraction of the time step for updating the fields $\Delta t_{\text{laser}} = \Delta t_{\text{fields}}/N_{\text{sub}}$ by performing N_{sub} subcycles using the same χ . All of this allows for great flexibility when choosing which parameters to refine in a given simulation. The longitudinal spacing of the beam grid is always set to match that of the plasma macroparticle solver.

4. Adaptive grids for particle beams

A potential downside of using an intermediate grid for the beam-plasma communication is that the resolution of the fields seen by the beams is limited by that of the grid. This would be especially problematic for beams with narrow transverse size σ_r , as it would require a grid with radial resolution $\Delta r \ll \sigma_r$ and therefore a large number of radial grid points N_r that would increase the cost of the simulation. This is, for example, one of the main challenges for electromagnetic PIC codes to simulate beams with collider-relevant parameters. Due to the nm-range normalized emittance and TeV-range energy, σ_x would be orders of magnitude smaller than the transverse size of the wake, resulting in an extremely high resolution and, therefore, computational cost. Fortunately, the gridless method can overcome this issue because the radial extent and resolution of the beam grid is independent of the background plasma. This allows us to introduce *adaptive grids* (AGs), i.e., grids with a fixed N_r but whose radial and longitudinal extent dynamically adapts to the size of the particle beams. The number of points per beam width can therefore remain constant, without increasing the simulation cost. This could in principle be implemented per slice. In order for the fields in the AGs to be accurate, there needs to be sufficient plasma macroparticles in the domain of the AG. As a general rule, the macroparticle spacing should be smaller than that of the AG, so that the field resolution is not limited by the AG. Here, we take advantage of a non-uniform initialization of plasma macroparticles enabled by the gridless method presented in Section 2: the number of plasma macroparticles can be increased only where needed to keep the computational cost low. For example, in a simulation with collider-relevant parameters ($\sigma_x \sim \mathcal{O}(1 \text{ nm})$) only the region close to the axis would require a higher macroparticle density.

An example of using AGs is shown in Fig. 3, which corresponds to a case with two electron beams (one driver, one witness) in a uniform hydrogen plasma with a density of 10^{16} cm^{-3} . Both beams have an initial energy of 1 GeV and a relative energy spread of 1 %. The driver has a charge of 1 nC, an RMS duration of 100 fs, a normalized emittance of $10 \mu\text{m}$ and a transverse size of $10 \mu\text{m}$, resulting in a peak particle density of $\sim 1.3 \times 10^{17} \text{ cm}^{-3}$. The witness beam is centered $200 \mu\text{m}$ behind the center of the driver and has a charge of 0.5 nC, an RMS duration of 50 fs, a normalized emittance of $1 \mu\text{m}$ and a transverse size of $3 \mu\text{m}$. The peak particle density is $\sim 1.5 \times 10^{18} \text{ cm}^{-3}$. The transverse phase space is mismatched to the focusing fields in the plasma, causing its transverse size – and therefore the extent of its AG – to oscillate throughout the simulation, as observed in Fig. 3(a) and (b). The plasma macroparticles have a non-uniform initial distribution, with 8 macroparticles per micron for $r < 10 \mu\text{m}$ and 2 macroparticles per micron in the rest of the domain. This makes sure that the ion motion arising close to the axis when the witness beam contracts is accurately modeled.

By default, the extent of the AGs is such that all macroparticles in the beam are within the grid. Optionally, a maximum or a fixed radial extent can be set to prevent them from getting too large if, e.g., a few beam macroparticles escape from the plasma wake. Macroparticles that escape the AGs will gather the fields from a lower-resolution base grid covering the full domain of the simulation, just as they would if AGs were not used.

The space charge field of the beams $B_{\theta}^{(b)}$ is also contained within their respective AGs, and can be extended beyond them. If a plasma




Fig. 3. Use of adaptive grids (AGs) for gathering the plasma fields into the beam macroparticles. The grids dynamically adapt to the size of the macroparticle beams at (a) $z = 0$ and (b) $z = 1.5 \text{ cm}$. For improved visibility, the grids in (a) and (b) show only a fraction of the grid resolution used in the simulation. The AGs allow for high-resolution field sampling only where needed (i.e., in the region the macroparticle beams), as seen in (c) and (d). The trajectories of the plasma electrons and ions are shown in blue and red, respectively. There, plasma is initialized with more macroparticles close to the axis to capture more accurately the ion motion due to the witness beam. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

macroparticle with a radius larger than the extent of the AG needs to gather $B_{\theta}^{(b)}$, the beam field at the macroparticle will be given by $B_{\theta,i}^{(b)} = B_{\theta,N_r}^{(b)} r_{N_r}/r_i$, where r_{N_r} and $B_{\theta,N_r}^{(b)}$ are, respectively, the radial position and value of the $B_{\theta}^{(b)}$ at the last grid node.

5. Validation benchmarks

The accuracy, performance and applicability of the gridless model implemented in Wake-T has been assessed with a laser-driven and a beam-driven benchmarks that are also simulated with the quasi-cylindrical electromagnetic PIC code FBPIC and the 3D quasistatic PIC code HiPACE++, respectively.

The laser-driven setup consists of a plasma stage with a longitudinal density profile given by a 10 cm plateau with an on-axis density $n_0 = 2 \times 10^{17} \text{ cm}^{-3}$ and a Gaussian up- and downramp with a decay length $\sigma_{z,R} = 0.3 \text{ mm}$ and a total length of 1 mm. The radial profile is parabolic, with a density given by $n(r) = n_0[1 + 4r^2/(k_p^2 w_0^4)]$. The laser driver has an energy of 10 J, a wavelength $\lambda_0 = 800 \text{ nm}$, a peak normalized vector potential $a_0 = 2.65$, a duration $\tau_{\text{FWHM}} = 25 \text{ fs}$ and a focal spot size $w_0 = 40 \mu\text{m}$ at the entrance to the plateau. The electron beam is externally injected and has a trapezoidal current profile with a total charge of 114.6 pC that has been optimized for beam loading [38]. It has a current of 4.26 kA (head) and 3.50 kA (tail), a head-to-tail duration of 21.2 fs, and a time delay between the head and the laser centroid of 184.0 fs. The beam profile is smoothed with a $1 \mu\text{m}$ Gaussian ramp both at the head and the tail. It has an initial energy of 200 MeV with an RMS relative energy spread of 0.1 %, a normalized emittance of $1 \mu\text{m}$ in both planes and a transverse size of $0.92 \mu\text{m}$ that is matched [39] to the focusing fields in the plateau.

The Wake-T simulations are performed using a simulation box with length $L_{\text{box}} = 120 \mu\text{m}$, where the plasma wakefield and the laser en-

velope are updated with a time step $c \Delta t_{\text{field}} = L_{\text{box}}/2$. The plasma extends radially up to $2.5w_0$ and the base grid used both for the electron beam and the laser envelope extends up to $4w_0$. No adaptive grid is needed for the electron beam in this case. The larger extent of the grid is needed due to the reflective boundary conditions of the envelope solver. The numerical convergence of the algorithm is tested by varying the radial and longitudinal resolutions. In the radial direction, the resolutions $k_p^{-1}/\Delta r = 10, 20, 40$ and 80 are tested while making sure that the underlying gridless plasma features the equivalent of 2 macroparticles per cell. The longitudinal direction is discretized using the resolutions $c\tau_{\text{FWHM}}/\Delta z = 10, 20, 40, 80$ and 160 , which also determine the longitudinal step of the plasma particle pusher. In the FBPIC case, the simulations are performed in the Lorentz boosted frame [15] using $\gamma_{\text{boost}} = 21$, which makes the full plasma stage fit within the boosted simulation window. The dimensions of the simulation box, the reflective boundary conditions, and the radial resolutions tested are the same as in Wake-T. In the longitudinal direction, the resolutions used are instead $\lambda_0/\Delta z = 10, 20, 40, 80, 160$ and 320 due to the need to capture the laser wavelength. In all cases, the time step in the non-boosted frame is given by $\Delta z/c$, and the number of macroparticles per cell is 1, 2 and 4 in the longitudinal, radial and azimuthal direction, respectively. Due to the different algorithms used by FBPIC and Wake-T, there is a slight difference in how the laser and particle beam are initialized in the simulations. In Wake-T, due to the quasistatic approximation, the laser and witness beam are initialized inside the plasma, at the very beginning of the entrance ramp. In FBPIC, they are both initialized before the plasma such that the simulation box in the non-boosted frame has its right edge (in the direction of propagation) at the beginning of the entrance ramp. The laser is injected using a virtual antenna, and the witness beam is propagated ballistically until the entrance of the plasma. The smooth 1 mm-long upramp helps mitigate any potential differences in beam evolution between these two initialization strategies.

An overview of the simulation results is shown in Fig. 4. The simulations converge to a final witness beam energy and energy spread of 2.91 GeV and 0.37 % in Wake-T, and of 2.91 GeV and 0.40 % in FBPIC. The laser evolves considerably in this long plasma stage (the electron beam is accelerated by 2.7 GeV), which is well captured by Wake-T. Wake-T reaches convergence with a longitudinal resolution that is at least a factor 20 lower than in FBPIC. The radial convergence is also faster in Wake-T, as indicated by the final beam parameters and by the shape of the blowout sheath highlighted in Fig. 4 (a). While the sheath at the back of the wake appears identical in all Wake-T simulations, its shape changes in FBPIC unless the resolution is increased above $k_p^{-1}/\Delta r = 80$. The faster radial convergence is likely a feature of the gridless model, which can naturally resolve regions of high radial macroparticle density. As a result of the faster convergence and simpler physical model, the converged Wake-T simulations (using $k_p^{-1}/\Delta r = 40$ and $c\tau_{\text{FWHM}}/\Delta z = 80 \Rightarrow k_p^{-1}/\Delta z \sim 127$) can be performed in only 7 min on a single CPU core, while an FBPIC simulation of comparable convergence level (using $k_p^{-1}/\Delta r = 80$ and $\lambda_0/\Delta z = 160 \Rightarrow k_p^{-1}/\Delta z \sim 2376$) would require 9.8 h making full use of an NVIDIA A100 GPU.

Besides the quasistatic gridless model, the faster longitudinal convergence of Wake-T can mostly be attributed to the use of laser envelope solver, as expected from Ref. [40]. Using 80 points per laser period is typical to resolve the laser dynamics with high accuracy. This corresponds to $k_p^{-1}/\Delta z = 1200$, which is where FBPIC reaches few-percent-level convergence. Isolating the specific speedup due to laser envelope solver is out of the scope of this work as it also impacts the boosted frame method [41].

Finally, FBPIC assumes quasi-cylindrical symmetry: comparison with a 3D PIC code would yield a larger speedup and is out of scope of the present work. This considerable reduction in computational cost and runtime enables, e.g., fast prototyping of LPA design, simulation of long multi-stage setups [21], large-scale optimization, and stability and jitter studies.




Fig. 4. Simulation results of the laser-driven study. (a) Side-by-side comparison of the highest-resolution results obtained from FBPIC (top) and Wake-T (bottom) at the same propagation distance in the simulation. The colored lines at the back to the wake indicate the shape of the blowout sheath for different radial resolutions. The inset shows the final longitudinal phase space of the witness as obtained from both simulation codes. The bottom axes show the convergence of the final beam energy spread, measured as the median absolute deviation (MAD), in terms of (b) radial and (c) longitudinal resolution.

The beam-driven benchmark consists of a 4-meter-long lithium stage with a uniform density of 10^{16} cm^{-3} preionized to the first level. The choice of lithium is motivated by the large energy difference between the first and second ionization levels, which prevents a narrow, high-intensity witness beam with collider-like parameters from further ionizing the plasma. The driver beam is initially Gaussian in all directions and has an energy of 20 GeV, an energy spread of 1 %, a charge of 2 nC, an RMS duration $\sigma_t = \omega_p^{-1} \simeq 177.3 \text{ fs}$, a normalized emittance of $10 \mu\text{m}$ and a matched transverse size of $1.95 \mu\text{m}$. The peak particle density is $\sim 3.9 \times 10^{18} \text{ cm}^{-3}$. The witness has a longitudinal trapezoidal profile that, as in the laser-driven case, has been optimized for beam loading. It has a total charge of 833 pC, a head-to-tail duration of 217.3 fs, a current ratio of 0.32 between the tail and the head, and a time delay between the head and the driver centroid of 694.9 fs. The initial energy is 175 GeV with an energy spread of 0.35 %. The normalized emittance is 135 nm in both planes, with a matched transverse size of 132 nm, resulting in a peak particle density of $\sim 1.1 \times 10^{21} \text{ cm}^{-3}$. These parameters are in the range of what could be expected from an acceleration stage identified in the European Strategy for Particle Physics roadmap [42], or a plasma-based Higgs factory, such as the HALHF proposal [37] if round beams were used [43].

Due to the small transverse size of the witness, extreme scale discrepancies make this scenario very challenging even with the most advanced quasi-cylindrical electromagnetic PIC codes like FBPIC, and the use of AGs in Wake-T and of mesh refinement (MR) in HiPACE++ [44,45] is required to improve the efficiency and reduce the run time. In or-




Fig. 5. (a) Simulation of a beam-driven plasma stage with collider-relevant parameters performed with Wake-T. The use of adaptive grids (AGs) is advantageous here due to the small transverse size of the witness beam. The insets in (a) show a detailed view of the witness beam and its AG. The fields in the AG feature a much higher resolution than in the rest of the domain, and are able to resolve the ion motion cone. A transverse lineout of $W_x = E_x - B_y$ in the base grid and the AG is shown in the top right inset. (b) Final emittance growth for different effective radial resolution as obtained from Wake-T (with and without AGs) and HiPACE++ (with and without MR), including a case where ion motion is disabled. (c) Comparison of the total runtime of the Wake-T simulations with and without AGs.

der to assess the advantages of AGs, the Wake-T simulations are performed with and without AGs. Similarly, ion motion is expected to play a significant role due to the high intensity of the witness. Thus, simulations with and without mobile ions are included in the study. The simulation box in Wake-T has a length $L_{\text{box}} = 650 \mu\text{m}$ and a plasma column that extends up to $r_{\text{max}} = 300 \mu\text{m}$. The longitudinal resolution of all grids and of the plasma pusher is $L_{\text{box}}/\Delta z = 600$, while the radial resolution is varied to perform a convergence study. The time step for updating the plasma wakefields is $c\Delta t_{\text{field}} = 2 \text{ mm}$. When not using AGs, both driver and witness share the same base grid, which has the same radial extent as the plasma. The radial resolutions tested for this case are $r_{\text{max}}/\Delta r = 256, 512, 1024, 2048, 4096, 8192, 16384$ and 32768 . When making use of AGs, each beam gets an associated grid with a radial extent $r_{\text{max},w} = 1 \mu\text{m}$ for the witness and $r_{\text{max},d} = 20 \mu\text{m}$ for the driver. The radial resolutions tested for the witness AG are $r_{\text{max},w}/\Delta r_w = 4, 8, 16, 32, 64, 128$ and 256 ; and $r_{\text{max},d}/\Delta r_d = 16r_{\text{max},w}/\Delta r_w$ for the driver. The radial size of both AGs is fixed throughout the computation to keep Δr_w and Δr_d constant and therefore simplify the analysis of the convergence study. The macropar-

ticle distributions of both beams are also identical in all simulations, having 8×10^5 and 4×10^5 macroparticles for the driver and witness, respectively. Although this number highly exceeds the requirements of the lower resolution simulations (which could therefore be proportionally less expensive), using the exact same distribution allows the convergence scan to exhibit much smoother behavior. Otherwise, the behavior of the beam parameters would be hindered by differently random beam distributions.

The 3D simulations carried out with HiPACE++ use the same box length and longitudinal resolution as those from Wake-T, while the radial extent is $300 \mu\text{m}$. The transverse resolution when not using MR is varied between $r_{\text{max}}/\Delta r = 256, 512, 1024, 2048$ and 4096 . The cases with MR use a base grid with $r_{\text{max}}/\Delta r_0 = 4096$ and an on-axis refined patch extending 7 radial cells of the base grid with a resolution of $r_{\text{max}}/\Delta r_1 = 4096, 8192, 16384, 32768, 65536$ and 131072 . Note that HiPACE++ resolutions are discussed in terms of radius for direct comparison with Wake-T: for the same resolution, HiPACE++ uses twice as many points in x and y as Wake-T does in r . One macroparticle per cell is used in all dimensions, both in the base grid and in the refined patch. As in Wake-T, the same driver and witness beams are used across all simulations, both with 10^8 macroparticles.

In terms of computational resources, the HiPACE++ simulations used 16 nodes on the JUWELS Booster supercomputer, each equipped with AMD EPYC 7402 CPU sockets and 4 NVIDIA A100 GPUs. The Wake-T simulations ran on a single core of an AMD EPYC 7552 CPU.

The results of the simulation study are summarized in Fig. 5. The witness beam experiences an energy gain of 17.1 GeV and an emittance growth of 25% , where the latter occurs due to the strong ion motion caused by the witness beam itself. Both Wake-T and HiPACE++ converge at the same rate and to the same result. This confirms that the ion motion is correctly modeled in Wake-T, and therefore the validity of the multi-species implementation in the presented gridless model. Another highlight is that the use of AGs in Wake-T and of MR in HiPACE++ allows for convergence to the same results while allowing for strong reduction of the computational cost. The use of AGs allows the electron beams to experience the fields with high resolution (see Fig. 5(a)), reaching convergence with simulations that are an order of magnitude faster (25 min vs 2.8 h) for the same equivalent resolution. It should be noted that the cost of the Wake-T simulations is artificially dominated by the beam particle pusher, due to the large number of beam macroparticles. Since convergence is already reached with $r_{\text{max}}/\Delta r \sim 20000$, which is 4 times smaller than the highest resolution tested, the number of beam macroparticles could also be safely reduced by the same factor. Thus, when adjusting the number of beam macroparticles accordingly, a converged simulation can be performed in 10 min on a single CPU core. The high efficiency demonstrated by the gridless model therefore enables the study of future plasma-based collider concepts with accurate, axisymmetric simulations of the plasma wake.

Despite the differences in computing architecture (CPU vs GPU), we can obtain a crude comparison between the performance of both codes by calculating the time required to compute a longitudinal slice per time step on 1 compute unit (a single AMD EPYC 7552 CPU core for Wake-T; one NVIDIA A100 GPU for HiPACE++). For the challenging simulation with 8192 radial cells and no AG/MR, we estimate that Wake-T required $\sim 2 \text{ ms}$ per longitudinal slice per time step, while HiPACE++ required $\sim 100 \text{ ms}$ per slice per timestep. This difference is expected from the added dimension in HiPACE++, which requires much more operations using a more potent processing unit. Incidentally, the number of cells computed per second, normalized to the peak performance of the hardware used ($0.05/19.5 \text{ TFLOPS}$ for the EPYC/A100) is approximately 0.1 cell per nanosecond per peak TFLOPS for both codes.

Finally, Wake-T can also be compared with QPAD [46], another quasistatic code that makes use of reduced geometry. While Wake-T assumes full cylindrical symmetry, QPAD – like FBPIC – makes use of an azimuthal mode decomposition to capture slight deviations from cylindrical symmetry at the cost of solving additional 2D grids (one per az-

imuthal mode). While care should be taken when comparing different codes on different platforms and use cases, data from Li et al. [47] (cf. Table 1) suggests that QPAD demands more computational resources than Wake-T, albeit using fewer radial cells (256 vs. 8192). This can be attributed to the azimuthal decomposition, which provides added accuracy at the cost of more computation.

6. Conclusion

The presented gridless quasistatic model makes it possible to simulate axially symmetric plasma wakes with high fidelity at a reduced cost. Our implementation of the model in the Wake-T code is coupled to a laser envelope solver and a particle beam pusher, and proves to be an accurate and cost-effective alternative to typical PIC codes. In a laser-driven case, it converges to equivalent results with a resolution that is at least a factor ~ 40 lower (factor ~ 20 in longitudinal direction and factor ~ 2 in radial direction) than with the cutting-edge quasi-cylindrical electromagnetic PIC code (FBPIC), which uses the most advanced numerical schemes. The faster convergence of the gridless model coupled with the laser envelope solver reduces the computation runtime from 9.8 h on an NVIDIA A100 GPU (FBPIC) to 7 min on a single CPU core (Wake-T). Another advantage of the gridless algorithm is that it allows for a straightforward implementation of local refinement of the field resolution by initializing more plasma macroparticles in the region of interest (e.g., close to the axis), and is devoid of numerical artifacts on the refined grid boundaries. This keeps the computational cost down while allowing for a high effective resolution.

Although the plasma solver is in itself gridless, the current implementation makes use of an intermediate regular grid for the communication between the plasma and the laser and particle beams. With this approach, the plasma fields are first precomputed in the regular grid and then interpolated onto the beam macroparticles. In order to keep the advantages of local field refinement with this approach, the concept of adaptive grids has been introduced. These grids dynamically adapt to the size of the beams so that the fields are only sampled where needed while keeping a constant resolution relative to the beam size.

Simulations with and without adaptive grids demonstrate convergence to the same solution, and are in full agreement with the results obtained with the 3D quasistatic PIC code HiPACE++. Furthermore, the use of adaptive grids together with on-axis refinement yields significant savings in computational time (factor ~ 10 in the presented study) when simulating extremely narrow beams, such as those needed for collider applications (nm-level emittance). As an example, a ~ 20 GeV, 4 m-long plasma stage with collider-relevant parameters can be simulated in 10 min on a single CPU core. This enables simulation studies that would otherwise be prohibitively expensive with current tools, and which are critically needed to solve the challenges of plasma acceleration towards becoming a viable technology for future high-energy colliders.

CRediT authorship contribution statement

Á. Ferran Pousa: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Investigation, Conceptualization; **W. M. den Hertog:** Software, Investigation; **S. Diederichs:** Validation, Investigation; **A. Martínez de la Ossa:** Writing – review & editing, Investigation; **J. L. Ordóñez Carrasco:** Software, Investigation; **A. Sinn:** Software, Investigation; **M. Thévenet:** Writing – review & editing, Supervision, Investigation, Funding acquisition, Conceptualization.

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.

Acknowledgments

We thank P. Baxevanis and G. Stupakov for useful discussions and for providing access to the PLEBS code; R. Lehe, A. Huebl and J.-L. Vay for fruitful discussions regarding the gridless algorithm; and C. Benedetti for advice in the implementation of the laser envelope model in Wake-T. This research was supported in part through the Maxwell computational resources operated at DESY, Hamburg, Germany. The authors gratefully acknowledge the Gauss Centre for Supercomputing e.V. (www.gauss-centre.eu) for funding this project by providing computing time through the John von Neumann Institute for Computing (NIC) on the GCS Supercomputer JUWELS [48] at Jülich Supercomputing Centre (JSC). This work was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)—491245950. No large language models (LLMs) or other generative-AI tools were used in the writing of this manuscript.

Data availability

Data will be made available on request.

Appendix A. Solve recursive system

The recursive system in Eq. (14) cannot be solved directly, as we would need to know the values of both a_0 and a_N (or, alternatively, of b_0 and b_N). Instead, we have the conditions $a_N = 0$ and $b_0 = 0$, which are at opposite ends of the iteration. To solve this, Eq. (14) can be rewritten so that a_i and b_i are linear functions of a_0

$$a_i = K_i a_0 + T_i, \quad (\text{A.1a})$$

$$b_i = U_i a_0 + P_i, \quad (\text{A.1b})$$

where the coefficients K_i , U_i , T_i , P_i are determined by solving the recursive systems

$$\begin{pmatrix} K_i \\ U_i \end{pmatrix} = \begin{pmatrix} 1 + \frac{A_i \rho_i}{2} & \frac{A_i}{2\rho_i} \\ -\frac{A_i \rho_i^3}{2} & 1 - \frac{A_i \rho_i}{2} \end{pmatrix} \begin{pmatrix} K_{i-1} \\ U_{i-1} \end{pmatrix} \quad (\text{A.2})$$

and

$$\begin{pmatrix} T_i \\ P_i \end{pmatrix} = \begin{pmatrix} 1 + \frac{A_i \rho_i}{2} & \frac{A_i}{2\rho_i} \\ -\frac{A_i \rho_i^3}{2} & 1 - \frac{A_i \rho_i}{2} \end{pmatrix} \begin{pmatrix} T_{i-1} \\ P_{i-1} \end{pmatrix} + \begin{pmatrix} \frac{1}{4}(2B_i + A_i C_i) \\ \frac{1}{4}\rho_i(4C_i - 2B_i \rho_i - A_i C_i \rho_i) \end{pmatrix} \quad (\text{A.3})$$

which, in this case, do have a set of initial conditions $K_0 = 1$, $U_0 = 0$, $T_0 = 0$, $P_0 = 0$ all at $i = 0$.

Once the system of equations has been solved iteratively to obtain the value of K_N and T_N , we have that $a_N = K_N a_0 + T_N$. Since $a_N = 0$, this implies that

$$a_0 = -\frac{T_N}{K_N}, \quad (\text{A.4})$$

which allows a_i and b_i to be obtained from Eqs. (A.1a) and (A.1b), and therefore to calculate $B_\theta^{(p)}$ according to Eq. (13).

This approach is equivalent to using a linear shooting method [49]: if Eq. (14) is the recursive system arising from the boundary value problem $\frac{\partial}{\partial \rho} \left[\frac{1}{\rho} \frac{\partial}{\partial \rho} (\rho B_\theta^{(p)}) \right] = \sum_i \frac{w_i}{\rho_i} (A_i \delta_i B_\theta^{(p)} + B_i \delta_i + C_i \partial_\rho \delta_i)$ with $B_\theta^{(p)}(0) = 0$ and $B_\theta^{(p)}(\infty) = 0$, then Eq. (A.3) is the recursive system stemming from the initial value problem $\frac{\partial}{\partial \rho} \left[\frac{1}{\rho} \frac{\partial}{\partial \rho} (\rho y_1) \right] = \sum_i \frac{w_i}{\rho_i} (A_i \delta_i y_1 + B_i \delta_i + C_i \partial_\rho \delta_i)$ with $y_1(0) = 0$ and $\frac{\partial}{\partial \rho} y_1(0) = 0$ and Eq. (A.2) is that from the initial value problem $\frac{\partial}{\partial \rho} \left[\frac{1}{\rho} \frac{\partial}{\partial \rho} (\rho y_2) \right] = \sum_i \frac{w_i}{\rho_i} A_i \delta_i y_2$ with $y_2(0) = 0$ and $\frac{\partial}{\partial \rho} y_2(0) = 1$. The solution can then be expressed as $B_\theta^{(p)}(\rho) = y_1(\rho) - \frac{T_N}{K_N} y_2(\rho)$, which stems from $\lim_{\rho \rightarrow \infty} \frac{y_1(\rho)}{y_2(\rho)} = \frac{T_N}{K_N}$, and which is equivalent to Eq. (13) combined with Eqs. (A.1a), (A.1b) and (A.4).

References

- [1] A.J. Gonsalves, K. Nakamura, J. Daniels, C. Benedetti, C. Pieronek, T.C.H. de Raadt, S. Steinke, J.H. Bin, S.S. Bulanov, J. van Tilborg, C.G.R. Geddes, C.B. Schroeder, C. Tóth, E. Esarey, K. Swanson, L. Fan-Chiang, G. Bagdasarov, N. Bobrova, V. Gasilov, G. Korn, P. Sasorov, W.P. Leemans, Petawatt laser guiding and electron beam acceleration to 8 GeV in a laser-heated capillary discharge waveguide, *Phys. Rev. Lett.* 122 (2019) 084801. <https://doi.org/10.1103/PhysRevLett.122.084801>
- [2] S. Jalas, M. Kirchen, P. Messner, P. Winkler, L. Hübner, J. Dirkwinkel, M. Schnepf, R. Lehe, A.R. Maier, Bayesian optimization of a laser-plasma accelerator, *Phys. Rev. Lett.* 126 (2021) 104801. <https://doi.org/10.1103/PhysRevLett.126.104801>
- [3] M. Kirchen, S. Jalas, P. Messner, P. Winkler, T. Eichner, L. Hübner, T. Hülsenbusch, L. Jeppe, T. Parikh, M. Schnepf, A.R. Maier, Optimal beam loading in a laser-plasma accelerator, *Phys. Rev. Lett.* 126 (2021) 174801. <https://doi.org/10.1103/PhysRevLett.126.174801>
- [4] C.A. Lindström, J.M. Garland, S. Schröder, L. Boulton, G. Boyle, J. Chappell, R. D'Arcy, P. Gonzalez, A. Knetisch, V. Libov, G. Loisch, A. Martinez de la Ossa, P. Niknejadi, K. Pöder, L. Schaper, B. Schmidt, B. Sheeran, S. Wesch, J. Wood, J. Osterhoff, Energy-spread preservation and high efficiency in a plasma-wakefield accelerator, *Phys. Rev. Lett.* 126 (2021) 014801. <https://doi.org/10.1103/PhysRevLett.126.014801>
- [5] A.R. Maier, N.M. Delbos, T. Eichner, L. Hübner, S. Jalas, L. Jeppe, S.W. Jolly, M. Kirchen, V. Leroux, P. Messner, M. Schnepf, M. Trunk, P.A. Walker, C. Werle, P. Winkler, Decoding sources of energy variability in a laser-plasma accelerator, *Phys. Rev. X* 10 (2020) 031039. <https://doi.org/10.1103/PhysRevX.10.031039>
- [6] W. Wang, K. Feng, L. Ke, C. Yu, Y. Xu, R. Qi, Y. Chen, Z. Qin, Z. Zhang, M. Fang, J. Liu, K. Jiang, H. Wang, C. Wang, X. Yang, F. Wu, Y. Leng, J. Liu, R. Li, Z. Xu, Free-electron lasing at 27 nanometres based on a laser wakefield accelerator, *Nature* 595 (7868) (2021) 516–520. <https://doi.org/10.1038/s41586-021-03678-x>
- [7] R. Pompili, D. Alesini, M.P. Anania, S. Arjmand, M. Behtouei, M. Bellaveglia, A. Biagioni, B. Buonomo, F. Cardelli, M. Carpanese, E. Chiadroni, A. Cianchi, G. Costa, A. Del Dotto, M. Del Giorgio, F. Dipace, A. Doria, F. Filippi, M. Galletti, L. Giannessi, A. Giribono, P. Iovine, V. Lollo, A. Mostacci, F. Nguyen, M. Oromolla, E. Di Palma, L. Pellegrino, A. Petralia, V. Petrillo, L. Piersanti, G. Di Piro, S. Romeo, A.R. Rossi, J. Scifo, A. Selce, V. Shpakov, A. Stella, C. Vaccarezza, F. Villa, A. Zigler, M. Ferrario, Free-electron lasing with compact beam-driven plasma wakefield accelerator, *Nature* 605 (7911) (2022) 659–662. <https://doi.org/10.1038/s41586-022-04589-1>
- [8] M. Labat, J.C. Cabadağ, A. Ghaith, A. Irman, A. Berlioux, P. Berteaud, F. Blache, S. Bock, F. Bouvet, F. Briquez, Y.-Y. Chang, S. Corde, A. Debus, C. De Oliveira, J.-P. Duval, Y. Dietrich, M. El Ajjouri, C. Eisenmann, J. Gautier, R. Gebhardt, S. Grams, U. Helbig, C. Herbeaux, N. Hubert, C. Kitegi, O. Kononenko, M. Kuntzsch, M. LaBerge, S. Lê, B. Leluan, A. Loulergue, V. Malka, F. Marteau, M.H.N. Guen, D. Oumbarek-Espinos, R. Pausch, D. Pereira, T. Püschel, J.-P. Ricaud, P. Rommeluere, E. Roussel, P. Rousseau, S. Schöbel, M. Sebdaoui, K. Steiniger, K. Tavakoli, C. Thaur, P. Ufer, M. Valléau, M. Vandenbergh, J. Vétérin, U. Schramm, M.-E. Couprie, Seeded free-electron laser driven by a compact laser plasma accelerator, *Nat. Photonics* 17 (2) (2023) 150–156. <https://doi.org/10.1038/s41566-022-01104-w>
- [9] A. Myers, A. Almgren, L.D. Amorim, J. Bell, L. Fedeli, L. Ge, K. Gott, D.P. Grote, M. Hogan, A. Huebl, R. Jambunathan, R. Lehe, C. Ng, M. Rowan, O. Shapoval, M. Thévenet, J.-L. Vay, H. Vincenti, E. Yang, N. Zam, W. Zhang, Y. Zhao, E. Zoni, Porting WarpX to GPU-accelerated platforms, *Parallel Comput.* 108 (2021) 102833. <https://doi.org/10.1016/j.parco.2021.102833>
- [10] C.A. Lindström, Self-correcting longitudinal phase space in a multistage plasma accelerator, *arXiv:2104.14460* (2021). <https://arxiv.org/ftp/arxiv/papers/2104/2104.14460.pdf>
- [11] P. Mora, T.M. Antonsen, Jr., Kinetic modeling of intense, short laser pulses propagating in tenuous plasmas, *Phys. Plasmas* 4 (1) (1997) 217–229. https://pubs.aip.org/aip/pop/article-pdf/4/1/217/12532941/217_1_online.pdf, <https://doi.org/10.1063/1.872134>
- [12] P. Mora, T.M. Antonsen, Electron cavitation and acceleration in the wake of an ultraintense, self-focused laser pulse, *Phys. Rev. E* 53 (1996) R2068–R2071. <https://doi.org/10.1103/PhysRevE.53.R2068>
- [13] P. Sprangle, E. Esarey, A. Ting, Nonlinear interaction of intense laser pulses in plasmas, *Phys. Rev. A* 41 (1990) 4463–4469. <https://doi.org/10.1103/PhysRevA.41.4463>
- [14] P. Sprangle, E. Esarey, A. Ting, Nonlinear theory of intense laser-plasma interactions, *Phys. Rev. Lett.* 64 (1990) 2011–2014. <https://doi.org/10.1103/PhysRevLett.64.2011>
- [15] 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 (2007) 130405. <https://doi.org/10.1103/PhysRevLett.98.130405>
- [16] 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. <https://doi.org/10.1016/j.jcp.2008.11.017>
- [17] P. Baxevanis, G. Stupakov, Novel fast simulation technique for axisymmetric plasma wakefield acceleration configurations in the blowout regime, *Phys. Rev. Accel. Beams* 21 (2018) 071301. <https://doi.org/10.1103/PhysRevAccelBeams.21.071301>
- [18] 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. <https://doi.org/10.1016/j.jcp.2013.05.020>
- [19] A. Ferran Pousa, R. Assmann, A. Martinez de la Ossa, Wake-T: a fast particle tracking code for plasma-based accelerators, *J. Phys. Conf. Ser.* 1350 (1) (2019) 012056. <https://doi.org/10.1088/1742-6596/1350/1/012056>
- [20] C. Benedetti, C.B. Schroeder, C.G.R. Geddes, E. Esarey, W.P. Leemans, An accurate and efficient laser-envelope solver for the modeling of laser-plasma accelerators, *Plasma Phys. Contr. Fusion* 60 (1) (2017) 014002. <https://doi.org/10.1088/1361-6587/aa8977>
- [21] A. Ferran Pousa, A. Sinn, M. Thévenet, M. Garten, A. Huebl, R. Lehe, R. Sandberg, J.-L. Vay, W.M. den Hertog, J.L. Ordóñez Carrasco, Efficient simulation of multistage plasma accelerators, in: *Proc. 14th International Particle Accelerator Conference (IPAC'23)*, JACoW Publishing, Geneva, Switzerland, 2023, pp. 1533–1536. <https://doi.org/10.18429/JACoW-IPAC2023-TUPA093>
- [22] 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. <https://doi.org/10.1016/j.cpc.2016.02.007>
- [23] 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. <https://doi.org/10.1016/j.cpc.2022.108421>
- [24] T. Wang, V. Khudik, J. Kim, G. Shvets, Accelerated three-dimensional quasistatic particle-in-cell code, *Phys. Rev. Accel. Beams* 25 (2022) 104603. <https://doi.org/10.1103/PhysRevAccelBeams.25.104603>
- [25] 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. <https://doi.org/10.1016/j.jcp.2006.01.039>
- [26] T. Mehring, C. Benedetti, C.B. Schroeder, J. Osterhoff, HiPACE: a quasi-static particle-in-cell code, *Plasma Phys. Contr. Fusion* 56 (8) (2014) 084012. <https://doi.org/10.1088/0741-3335/56/8/084012>
- [27] J.C. Butcher, *Numerical Methods for Ordinary Differential Equations*, John Wiley & Sons, 2016.
- [28] S.A. Yi, V. Khudik, C. Siemon, G. Shvets, Analytic model of electromagnetic fields around a plasma bubble in the blow-out regime, *Phys. Plasmas* 20 (1) (2013) 013108. https://pubs.aip.org/aip/pop/article-pdf/doi/10.1063/1.4775774/13429408/013108_1_online.pdf, <https://doi.org/10.1063/1.4775774>
- [29] K. Flöttmann, ASTRA: a space charge tracking algorithm, 2011. <https://www.desy.de/~mpyflo/>
- [30] J.P. Boris, et al., Relativistic plasma simulation-optimization of a hybrid code, in: *Proc. Fourth Conf. Num. Sim. Plasmas*, 1970, pp. 3–67.
- [31] B. Ripperda, F. Bacchini, J. Teunissen, C. Xia, O. Porth, L. Sironi, G. Lapenta, R. Keppens, A comprehensive comparison of relativistic particle integrators, *Astrophys. J. Suppl. Ser.* 235 (1) (2018) 21. <https://doi.org/10.3847/1538-4365/aab114>
- [32] S.K. Lam, A. Pitrou, S. Seibert, Numba: a LLVM-based python JIT compiler, in: *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC, LLVM '15*, Association for Computing Machinery, New York, NY, USA, 2015. <https://doi.org/10.1145/2833157.2833162>
- [33] C. Latner, V. Advé, LLVM: a compilation framework for lifelong program analysis & transformation, in: *International Symposium on Code Generation and Optimization*, 2004. CGO 2004., 2004, pp. 75–86. <https://doi.org/10.1109/CGO.2004.1281665>
- [34] BLAST: beam, plasma & accelerator simulation toolkit, (<https://blast.lbl.gov/>). Accessed: 2024-05-21.
- [35] A. Huebl, R. Lehe, J.-L. Vay, D.P. Grote, I. Szalzarini, S. Kuschel, M. Bussmann, openPMD 1.0.0: a meta data standard for particle and mesh based data, 2015. <https://doi.org/10.5281/zenodo.33624>
- [36] A. Ferran Pousa, S. Jalas, M. Kirchen, A. Martinez de la Ossa, M. Thévenet, S. Hudson, J. Larson, A. Huebl, J.-L. Vay, R. Lehe, Bayesian optimization of laser-plasma accelerators assisted by reduced physical models, *Phys. Rev. Accel. Beams* 26 (2023) 084601. <https://doi.org/10.1103/PhysRevAccelBeams.26.084601>
- [37] B. Foster, R. D'Arcy, C.A. Lindström, A hybrid, asymmetric, linear Higgs factory based on plasma-wakefield and radio-frequency acceleration, *New J. Phys.* 25 (9) (2023) 093037. <https://doi.org/10.1088/1367-2630/acf395>
- [38] S.W.T. Katsouleas, J. Su, Beam loading efficiency in plasma accelerators, *Part. Accel.* 22 (1987) 81–99.
- [39] T. Mehring, J. Grebenyuk, F.S. Tsung, K. Floetmann, J. Osterhoff, Transverse emittance growth in staged laser-wakefield acceleration, *Phys. Rev. ST Accel. Beams* 15 (2012) 111303. <https://doi.org/10.1103/PhysRevSTAB.15.111303>
- [40] D. Terzani, C. Benedetti, C.B. Schroeder, E. Esarey, Accuracy of the time-averaged ponderomotive approximation for laser-plasma accelerator modeling, *Phys. Plasmas* 28 (6) (2021) 063105. https://pubs.aip.org/aip/pop/article-pdf/doi/10.1063/5.0050580/13316487/063105_1_online.pdf, <https://doi.org/10.1063/5.0050580>
- [41] F. Massimo, A. Beck, J. Derouillat, M. Grech, M. Lobet, F. Pérez, I. Zemzemi, A. Specka, Efficient start-to-end 3D envelope modeling for two-stage laser wakefield acceleration experiments, *Plasma Phys. Contr. Fusion* 61 (12) (2019) 124001. <https://doi.org/10.1088/1361-6587/ab49cf>
- [42] C. Adolphsen, D. Angal-Kalinin, T. Arndt, M. Arnold, R. Assmann, B. Auchmann, K. Aulenbacher, A. Ballarín, B. Baudouy, P. Baudreghien, et al., European strategy for particle physics–accelerator R&D roadmap, *arXiv:2201.07895* (2022).
- [43] S. Diederichs, C. Benedetti, A.F. Pousa, A. Sinn, J. Osterhoff, C.B. Schroeder, M. Thévenet, Resonant emittance mixing of flat beams in plasma accelerators, *arXiv:2403.05871* (2024).
- [44] M. Thévenet, S. Diederichs, A. Huebl, A. Sinn, Small-scale particle accelerators for large-scale science thanks to high-performance computing, *NIC Symposium 2025 Proceedings* 52 (2025) 414–422. <https://doi.org/10.34734/FZJ-2025-01965>
- [45] M. Thévenet, Affordable simulations of collider-relevant plasma-based accelerators (and more) with mesh refinement, 2023, (<https://agenda.infn.it/event/35577/contributions/208606/>). [Online].
- [46] 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. <https://doi.org/10.1016/j.cpc.2020.107784>
- [47] F. Li, W. An, F.S. Tsung, V.K. Decyk, W.B. Mori, Integrating a ponderomotive guiding center algorithm into a quasi-static particle-in-cell code based on azimuthal mode

- decomposition, J. Comput. Phys. 470 (2022) 111599. <https://doi.org/10.1016/j.jcp.2022.111599>
- [48] Jülich Supercomputing Centre, JUWELS cluster and booster: exascale pathfinder with modular supercomputing architecture at Juelich supercomputing centre, J. Large-Scale Res. Facil. 7 (A138) (2021). <https://doi.org/10.17815/jlsrf-7-183>
- [49] H.B. Keller, Numerical Solution of Two Point Boundary Value Problems, Society for Industrial and Applied Mathematics, 1976. <https://doi.org/10.1137/1.9781611970449>