Contents lists available at ScienceDirect

Journal of Computational Physics

journal homepage: www.elsevier.com/locate/jcp

A Lagrangian velocity-spectral particle method for low-noise plasma and neutral-gas kinetics with conservative distribution exchange

Chang Lu[a,*], Pei Qiu[b]

[a]State Key Laboratory of Industrial Equipment Structural Analysis and Optimization and CAE Software, Dalian University of Technology, Dalian, 116024, China
[b]College of Zhan Tianyou, Dalian Jiaotong University, Dalian, 116028, China

ARTICLE INFO

Article history:

Kinetic plasma simulation, Velocity-spectral particles, Particle-in–Cell, Material point method, Relative transport, BGK relaxation, Conservative particle–mesh exchange

ABSTRACT

A Lagrangian velocity-spectral particle (LVSP) method is developed for plasma and neutral-gas kinetics, combining low sampling noise with the conservative exchange of distribution information. In this method, moving spatial point particles carry continuous velocity distributions rather than random samples. Each point particle moves at its own mean velocity while retaining a simple spatial-point representation. This representation avoids tracking deforming spatial shape functions, as our central innovation describes advection relative to particle motion by combining characteristic transport with distribution exchange. To compute this relative advection, particles first project their relative fluxes onto a background grid. These grid-combined fluxes are then conservatively allocated back to the particles to update each particle's velocity distribution. During the distribution update, transport and collisions are integrated at the same time stages, with exponential weights used for collision relaxation. Together, these mechanisms constitute a material-point (MPM)-style framework that provides a unified treatment for both charged and neutral particles. Furthermore, in the single-velocity and local-equilibrium limits, the particle and weak-moment formulations reduce to a Particle-in-Cell (PIC) method and a fluid MPM description, respectively. Numerical tests validate the features and advantages of the proposed framework.

© 2026 Elsevier Inc. All rights reserved.

1. Introduction

Weak electric fields and small changes in velocity distributions can control plasma dynamics. Resolving these processes requires numerical errors to remain small relative to the physical signal. In conventional full-distribution

*Corresponding author.
e-mail: changlu@dlut.edu.cn (Chang Lu)

Particle-in-Cell (PIC) calculations, finite random velocity samples introduce fluctuations into computed moments and fields [1]. Increasing the number of simulation particles (referred to simply as "particles" throughout this paper) reduces sampling noise but increases computational cost. During long simulations, fluctuating fields can alter particle trajectories. Insufficient spatial resolution can also cause finite-grid instabilities [2]. Numerical drag and diffusion may artificially relax non-Maxwellian distributions [3]. The importance of these errors depends on particle loading, resolution, and field coupling, but they can make sustained weak-signal calculations difficult.

Many strategies have been developed to tackle these challenges. Broadly speaking, they fall into three categories: pure particle methods, deterministic kinetic solvers, and hybrid schemes that combine deterministic evolution with stochastic particle sampling. Representative particle-based techniques include δf variance reduction [4], quiet-start loading [5], wavelet denoising [6], and Particle-in-Fourier methods [7], to name a few. Other particle-based approaches enrich the information carried by particles or reconstruct the particle population [8, 9, 10]. Deterministic kinetic solvers include the unified gas-kinetic scheme (UGKS) and discrete unified gas-kinetic scheme (DUGKS) [11, 12, 13]. Representative hybrid schemes are the unified gas-kinetic particle (UGKP) and wave–particle (UGKWP) methods [15, 16], which integrate deterministic evolution and stochastic particle transport within the UGKS formalism.

Within particle methods, two major strategies stand out for improving distribution representation: enriching particle shapes and periodically remapping the distribution. Assigning shapes in both position and velocity allows each particle to represent a local distribution in phase space. Linearly transformed particles track the deformation of these shapes and have demonstrated low-noise Vlasov–Poisson calculations [8]. Strong deformation can make these shapes harder to represent and transfer accurately. The decorated particle method of Quashie, Burby, Christlieb, and Tang represents another highly inspiring work [9]. It equips particles with phase-space moments and derives their evolution from a Hamiltonian reduction. These moments differ from finite-width particle shapes, but both approaches carry information beyond a single phase-space sample. Remapping instead reconstructs the distribution and assigns it to a new particle population. Adaptive high-order remapping can preserve nonnegativity and improve distribution resolution [10], although repeated transfers can introduce interpolation error and numerical diffusion. These developments motivate us to retain rich velocity-space information to mitigate noise, while seeking to avoid deformation of the particles' spatial shapes.

On the deterministic side, kinetic methods based on characteristic transport provide a complementary idea. UGKS and DUGKS couple transport and collisions, allowing calculations under different collision conditions [11, 12, 13]. Liu et al. [14] applied DUGKS to electrostatic plasmas, avoiding random velocity-sampling fluctuations. Resolving distributions in both position and velocity can nevertheless be expensive in multiple dimensions, motivating adaptive velocity discretizations [18]. UGKP instead couples deterministic updates of macroscopic quantities with stochastic particle transport [15, 16]. UGKWP extends this approach by evaluating part of the distribution's contribution analytically, reducing particle sampling, tracking costs, and statistical noise [16]. This reduction is especially effective near the continuum regime. UGKWP has also been extended to gas mixtures and plasma transport [17]. Its remaining stochastic particle contribution can still produce sampling fluctuations, as reported in the plasma calculations [17].

These developments motivate retaining particle flexibility while describing velocity information deterministically.

Inspired by the above works, namely velocity-resolved particle representations and characteristic transport, we develop Lagrangian velocity-spectral particle (LVSP) method to combine particle flexibility with deterministic velocity distributions. The key insight is to employ characteristic transport and represent the relative flux with respect to particle motion via conservative exchange of distribution information, thereby avoiding the treatment of deforming spatial particle shapes. Each particle carries a continuous velocity distribution and moves with its own mean velocity. Exchange accounts for the remaining transport, allowing particles to retain their spatial point representation. A background grid organizes this exchange in a material point method (MPM) framework [19]. Transport and BGK collisions [20] are included within the same particle time stages. This formulation removes random velocity-sampling fluctuations while preserving the flexibility to choose local velocity functions. Localized flow features can be represented with few coefficients, and compact multivariate bases reduce storage at a fixed polynomial order. Conservative exchange and a common treatment of charged and neutral particles provide a framework for kinetic calculations under different collision conditions. The collisionless, single-velocity specialization recovers the PIC method, while local-equilibrium moments reduce the kinetic description to a fluid MPM formulation. LVSP retains the velocity structure needed between these limits.

The remainder of the paper follows the algorithm's computational steps. Section 2 states the governing problem, derives its material form, and outlines the solution strategy. Section 3 introduces the particle representation and particle-to-grid collection. Section 4 explains how the background basis evaluates transport and collision contributions. Section 5 determines individual particle rates and advances particle positions and velocity distributions. Section 6 discusses low noise, conservation, the common charged–neutral framework, and the prospects for electromagnetic extension. Section 7 presents the numerical verification, followed by conclusions. Velocity bases and detailed operators are given in the appendices.

2. Problem formulation and solution strategy

2.1. The Vlasov–Poisson–BGK problem

Let $\boldsymbol{x} \in \mathcal{D} \subset \mathbb{R}^d$, $\boldsymbol{v} \in \mathbb{R}^d$, and $t \geq 0$, with $d = 1, 2, 3$. The unknown $f(\boldsymbol{x}, \boldsymbol{v}, t)$ denotes the mass distribution, with $\rho = \int_{\mathbb{R}^d} f \, d\boldsymbol{v}$ the corresponding density. Let q and m denote the charge and mass parameters of the modeled species, so that q/m is its charge-to-mass ratio. Let $Q(\boldsymbol{x}, t)$ be the total charge density, ε_0 the permittivity, Ψ the electrostatic potential, $\boldsymbol{E}$ the electric field, and Ω the collision term. The modeled species contributes $(q/m)\rho$ to Q. The coupled problem is

$$\begin{aligned} \partial_t f + \boldsymbol{v} \cdot \nabla_x f + \frac{q}{m} \boldsymbol{E} \cdot \nabla_v f &= \Omega(f), \\ \boldsymbol{E} = -\nabla_x \Psi, \qquad -\varepsilon_0 \Delta \Psi &= Q. \end{aligned} \tag{1}$$

Define bulk velocity $\boldsymbol{u}$, kinetic energy density $\mathcal{E}$, and temperature T by

$$\rho \boldsymbol{u} = \int \boldsymbol{v} f \, d\boldsymbol{v}, \qquad \mathcal{E} = \frac{1}{2} \int |\boldsymbol{v}|^2 f \, d\boldsymbol{v}, \qquad T = \frac{2\mathcal{E}/\rho - |\boldsymbol{u}|^2}{d}. \tag{2}$$

All unspecified velocity integrals are over $\mathbb{R}^d$, and the gas constant is absorbed into T. For $\rho, T > 0$, the normalized distribution, normalized Maxwellian, full equilibrium, and collision operator are

$$\phi = f/\rho, \quad g(\boldsymbol{v}; \boldsymbol{u}, T) = (2\pi T)^{-d/2} e^{-|\boldsymbol{v}-\boldsymbol{u}|^2/(2T)}, \quad G = \rho g, \quad \Omega(f) = (G - f)/\tau. \tag{3}$$

The function g has unit velocity integral, whereas $G = \rho g$ includes the density. Here $\tau > 0$ is constant and $\int(1, \boldsymbol{v}, |\boldsymbol{v}|^2)\Omega \, d\boldsymbol{v} = 0$. The collisionless limit sets $1/\tau = 0$. The neutral-gas specialization sets $q = 0$, decoupling kinetic evolution from the electrostatic problem. Initial data are $f(\boldsymbol{x}, \boldsymbol{v}, 0) = f_0 \geq 0$, with finite mass and energy. For periodic kinetic and potential boundary conditions, the Poisson problem requires the compatibility condition $\int_{\mathcal{D}} Q \, d\boldsymbol{x} = 0$. The potential is fixed by the gauge $\int_{\mathcal{D}} \Psi \, d\boldsymbol{x} = 0$. At an absorbing kinetic boundary, $f = 0$ for incoming velocities $\boldsymbol{v} \cdot \boldsymbol{n} < 0$, where $\boldsymbol{n}$ is the outward normal. Electrostatic problems on nonperiodic domains additionally require boundary conditions for the potential or electric flux.

2.2. Derivation of the material-frame equation

Integrating Eq. (1) in velocity gives mass continuity: the electric-force term integrates to zero under sufficiently rapid decay of f at velocity infinity, and BGK collisions conserve mass.

$$\partial_t \rho + \nabla_x \cdot (\rho \boldsymbol{u}) = 0. \tag{4}$$

With $f = \rho\phi$ and $D/Dt = \partial_t + \boldsymbol{u} \cdot \nabla_x$, using the continuity equation in the product rule gives

$$\partial_t f + \nabla_x \cdot (\boldsymbol{u} f) = \rho \frac{D\phi}{Dt} + \phi[\partial_t \rho + \nabla_x \cdot (\rho \boldsymbol{u})] = \rho \frac{D\phi}{Dt}. \tag{5}$$

Adding and subtracting the bulk flux in Eq. (1) gives the conservative form

$$\partial_t f + \nabla_x \cdot (\boldsymbol{u} f) = -\nabla_x \cdot [(\boldsymbol{v} - \boldsymbol{u}) f] - \nabla_v \cdot \left(\frac{q}{m} \boldsymbol{E} f\right) + \Omega(f). \tag{6}$$

Equation (5) then yields the normalized-distribution form

$$\rho \frac{D\phi}{Dt} = -\nabla_x \cdot [(\boldsymbol{v} - \boldsymbol{u}) f] - \nabla_v \cdot \left(\frac{q}{m} \boldsymbol{E} f\right) + \Omega(f). \tag{7}$$

The conservative form is used for particle substitution below. The normalized form identifies relative spatial transport, acceleration, and collisions along the collective bulk flow.

2.3. Solution strategy and organization of the algorithm

LVSP assigns bulk motion to moving spatial point particles. Here a particle is a numerical particle represented by a moving point in physical space and carrying a continuous velocity distribution. Different microscopic velocity components would transport distribution information to distinct spatial locations. To retain the delta-function spatial representation of particles, this relative transport is approximated via information exchange among particles. The bulk velocity $\boldsymbol{u}$ in Eq. (7) is the mean velocity of the collective distribution. However, in LVSP, each numerical particle instead moves with the mean velocity $\boldsymbol{U}_p$ of the distribution it carries, defined below. This choice assigns the particle's

mean motion to its position update and the remaining transport to the relative flux $(\boldsymbol{v} - \boldsymbol{U}_p)\psi_p$. The corresponding particle balance is derived from the total particle representation in Section 3.

The algorithm has three stages, developed in the same order in the next three sections. Particle-to-grid (PtoG) collection (Section 3) constructs collective moments, equilibria, charge, and relative-flux information from the current particles. Background evaluation (Section 4) tests the relative flux against the spatial basis and forms the constraints on the collective change; the electric field and collision gain are also determined from this stage's particle information. Grid-to-particle (GtoP) allocation and advancement (Section 5) determine individual particle rates and advance positions and velocity distributions. These operations are repeated at each particle time stage, and only particle positions and velocity-distribution coefficients persist as the kinetic state:

$$\text{particles} \xrightarrow{\text{PtoG}} \text{fields and fluxes} \xrightarrow{\text{background evaluation}} \text{rate constraints} \xrightarrow{\text{GtoP}} \text{particle rates and update.}$$

3. Particle representation and PtoG projection

3.1. Particle representation of the distribution function

Particle p carries a position $\boldsymbol{x}_p(t)$ and a velocity mass distribution $\psi_p(\boldsymbol{v}, t)$. Its contribution and the complete population are

$$f_p = \psi_p(\boldsymbol{v}, t)\delta^{(d)}(\boldsymbol{x} - \boldsymbol{x}_p(t)), \qquad f^P = \sum_p f_p. \tag{8}$$

To relate this notation to a normalized velocity shape, write

$$\omega_p = \int \psi_p \, d\boldsymbol{v}, \qquad \phi_p = \psi_p/\omega_p, \qquad \boldsymbol{U}_p = \frac{\int \boldsymbol{v}\psi_p \, d\boldsymbol{v}}{\omega_p}, \qquad \dot{\boldsymbol{x}}_p = \boldsymbol{U}_p. \tag{9}$$

For $\omega_p > 0$, $\phi_p = \psi_p/\omega_p$ is the normalized shape, whereas ψ_p is the velocity mass distribution. The desired distribution is nonnegative. Some velocity-basis coefficients may initially be zero, including those of an additional component used to represent collision gain. The particle as a whole still carries a nonzero distribution with positive mass and moves with its mean velocity $\boldsymbol{U}_p$. Each particle represents only part of the full distribution, and its mean velocity $\boldsymbol{U}_p$ is computed from its own internal velocity distribution. In f^P, the superscript P denotes the total particle representation, while the subscript p labels an individual particle. The derivative $\partial_t\psi_p$ describes the time evolution of the velocity distribution carried by particle p, with $\boldsymbol{v}$ held fixed; the particle's position evolves separately.

For numerical implementation, $\psi_p = \sum_k c_{pk}\chi_{pk}(\boldsymbol{v})$, where c_{pk} are stored expansion coefficients and $\chi_{pk}(\boldsymbol{v})$ are velocity basis functions. The basis centres, scaling factors, and truncated polynomial orders may vary across particle groups. Particles within the same group share the same velocity basis so that their distribution information can be combined directly, while each particle keeps its own coefficients and spatial position. The specific basis choices adopted in this work are given in §Appendix A.

Particle exchange and collision processes can modify ω_p. The scheme evolves ψ_p directly and conserves the relevant population totals. Each particle retains a spatial Dirac-delta representation at $\boldsymbol{x}_p(t)$: its position moves in time, and no finite spatial width or shape-deformation matrix is tracked.

3.2. Substitution of the particle representation into the kinetic equation

Let $\delta_p = \delta^{(d)}(\boldsymbol{x} - \boldsymbol{x}_p(t))$ within this derivation. Differentiating Eq. (8) gives

$$\partial_t f^P = \sum_p (\partial_t \psi_p)\delta_p - \sum_p \psi_p \dot{\boldsymbol{x}}_p \cdot \nabla_x \delta_p. \tag{10}$$

To substitute this expression and f^P into Eq. (6), first expand its relative spatial flux:

$$-\nabla_x \cdot [(\boldsymbol{v} - \boldsymbol{u}) f^P] = -\sum_p \psi_p \boldsymbol{v} \cdot \nabla_x \delta_p + \nabla_x \cdot (\boldsymbol{u} f^P). \tag{11}$$

The substituted equation, before cancelling any terms, is therefore

$$\begin{aligned} &\sum_p (\partial_t \psi_p)\delta_p - \sum_p \psi_p \dot{\boldsymbol{x}}_p \cdot \nabla_x \delta_p + \nabla_x \cdot (\boldsymbol{u} f^P) \\ &\quad = -\sum_p \psi_p \boldsymbol{v} \cdot \nabla_x \delta_p + \nabla_x \cdot (\boldsymbol{u} f^P) - \sum_p \nabla_v \cdot \left[\frac{q}{m}\boldsymbol{E}(\boldsymbol{x}_p)\psi_p\right]\delta_p + \frac{G - f^P}{\tau}. \end{aligned} \tag{12}$$

The identical term $\nabla_x \cdot (\boldsymbol{u} f^P)$ now appears on both sides and cancels in full, including the contribution from spatial variations of $\boldsymbol{u}$. Moving the term containing $\dot{\boldsymbol{x}}_p$ to the right and setting $\dot{\boldsymbol{x}}_p = \boldsymbol{U}_p$ combines the two remaining spatial terms into $-\sum_p (\boldsymbol{v} - \boldsymbol{U}_p)\psi_p \cdot \nabla_x \delta_p$. Thus the population relation is

$$\sum_p (\partial_t \psi_p)\delta_p = -\sum_p (\boldsymbol{v} - \boldsymbol{U}_p)\psi_p \cdot \nabla_x \delta_p - \sum_p \nabla_v \cdot \left[\frac{q}{m}\boldsymbol{E}(\boldsymbol{x}_p)\psi_p\right]\delta_p + \frac{G - f^P}{\tau}. \tag{13}$$

Thus particle motion accounts for $\boldsymbol{U}_p\psi_p$, and exchange must account for the remaining spatial flux $(\boldsymbol{v} - \boldsymbol{U}_p)\psi_p$.

Equation (13) is a formal identity for the full particle ensemble and is not amenable to direct pointwise evaluation. In general, adjusting the weights of a finite set of Dirac masses cannot exactly reproduce the spatial derivatives of Dirac masses appearing on the right-hand side. For this reason, we adopt a weak-form discretization of the aggregate particle relation. The resulting allocation constraints specify the evolution of individual particles, and equilibrium Maxwellian states are built from the reconstructed collective moments rather than raw singular point densities.

3.3. Grid basis functions and collective information

Let the nonnegative spatial basis functions $N_i(\boldsymbol{x})$ have local support and satisfy $\sum_i N_i = 1$. Their gradients are evaluated within each element. Define $\mu_i = \int_{\mathcal{D}} N_i \, d\boldsymbol{x}$. For the basic PtoG reconstruction, the distribution coefficients at a node are

$$F_i(\boldsymbol{v}) = \frac{\sum_p N_i(\boldsymbol{x}_p)\psi_p(\boldsymbol{v})}{\mu_i}. \tag{14}$$

Applying the moment definitions in Section 2 yields

$$\rho_i = \int F_i \, d\boldsymbol{v}, \quad \boldsymbol{u}_i = \frac{\int \boldsymbol{v} F_i \, d\boldsymbol{v}}{\rho_i}, \quad T_i = \frac{\int |\boldsymbol{v} - \boldsymbol{u}_i|^2 F_i \, d\boldsymbol{v}}{d\rho_i}, \quad G_i = \rho_i g(\boldsymbol{v}; \boldsymbol{u}_i, T_i). \tag{15}$$

These collective moments supply the collision equilibrium. For charged systems, the collected charge also supplies the Poisson problem; the field gathered to a particle is $\boldsymbol{E}_h(\boldsymbol{x}_p) = \sum_i N_i(\boldsymbol{x}_p)\boldsymbol{E}_i$.

Let g_p label a particle group. For this group, the relative-flux field used in the basic reconstruction is

$$\boldsymbol{J}_{g_p,h}(\boldsymbol{x},\boldsymbol{v}) = \sum_i \frac{N_i(\boldsymbol{x})}{\mu_i} \sum_{p\in g_p} N_i(\boldsymbol{x}_p)(\boldsymbol{v}-\boldsymbol{U}_p)\psi_p(\boldsymbol{v}). \tag{16}$$

This reconstruction operates on these individual relative-flux contributions from each particle, rather than evaluating $(\boldsymbol{v}-\boldsymbol{u}_h)\sum_i N_i F_i$.

4. Background evaluation of transport and collision contributions

4.1. Weak balance of relative transport

Multiplying Eq. (13) by the test function N_i and integrating over space yields the exact weakly tested interior transport sum $\sum_p \nabla_x N_i(\boldsymbol{x}_p)\cdot(\boldsymbol{v}-\boldsymbol{U}_p)\psi_p$. Instead of evaluating this singular flux directly, LVSP uses the smooth flux field reconstructed during the PtoG stage. For particles in group g_p, the relative-transport constraint at node i is

$$R_{i,g_p} = \int_{\mathcal{D}} \nabla_x N_i \cdot \boldsymbol{J}_{g_p,h}\, d\boldsymbol{x} - \int_{\partial\mathcal{D}} N_i \boldsymbol{J}_{g_p,h}\cdot \boldsymbol{n}\, dS. \tag{17}$$

This expression gives a spatial approximation to the weakly tested particle relation. The constraint restricts the weighted sum of particle evolution rates, while the individual particle rates remain undetermined at this stage. Boundary integrals cancel out under periodic conditions. At absorbing boundaries, incoming kinetic flux is set to zero; the outgoing relative-flux treatment is detailed in §Appendix B.3. Shock simulations further adopt a conservative slope-limited edge term, defined in §Appendix B.

4.2. Electric acceleration and collective collision gain

The field obtained from the current particle charge acts on each particle velocity function through

$$\mathcal{A}_p = -\nabla_v \cdot \left[\frac{q}{m}\boldsymbol{E}_h(\boldsymbol{x}_p)\psi_p\right]. \tag{18}$$

The same spatial weights deposit charge and gather the field. The electric acceleration is applied directly to velocity coefficients.

For collisions, each particle retains the BGK loss $-\psi_p/\tau$. The gain is formed from the collective equilibrium in Eq. (15). Write G_{i,g_p} for the portion of G_i assigned to group g_p, with $\sum_{g_p} G_{i,g_p} = G_i$. In the collisional mixture tests, each particle group includes an equilibrium component whose coefficients are initially zero. It receives the group's local nodal mass fraction of G_i, as specified in §Appendix A. With the basic PtoG rule, the group weak balance is

$$\sum_{p\in g_p} N_i(\boldsymbol{x}_p)\partial_t\psi_p = R_{i,g_p} + \sum_{p\in g_p} N_i(\boldsymbol{x}_p)\mathcal{A}_p + \frac{\mu_i G_{i,g_p} - \sum_{p\in g_p} N_i(\boldsymbol{x}_p)\psi_p}{\tau}. \tag{19}$$

The gain is coupled to neighbouring particles through their collective moments.

4.3. Consistency of reconstruction and constraints

Because $\boldsymbol{U}_p$ is the mean velocity of each particle's own distribution, $\int(\boldsymbol{v}-\boldsymbol{U}_p)\psi_p\,d\boldsymbol{v}=0$. The weak reconstruction preserves this moment property. Even so, collision exchange and the stabilization scheme described below can redistribute mass among particles. Conservation constraints are enforced on the full particle population rather than on individual particles. Summing Eq. (17) over all nodes cancels the interior contributions via the partition of unity property. For the collision invariants $1, \boldsymbol{v}, |\boldsymbol{v}|^2$, total gains and losses balance one another when their velocity moments and spatial loads are consistent. We verify these conservation properties after projecting the distribution onto the selected velocity basis. All quantities in this section are computed from the current particle state. They are re-evaluated whenever the particle state evolves during time integration, and provide constraints on the particle evolution rates.

5. GtoP projection and particle advancement

5.1. Determining individual particle rates

The collective constraints alone do not uniquely specify how each particle evolves. To resolve this ambiguity, we select the relative-transport rates $\mathcal{T}_p$ at each velocity by solving a positive-volume-weighted minimization problem:

$$\min_{\mathcal{T}_p}\frac{1}{2}\sum_{p\in g_p}\frac{\mathcal{T}_p^2}{V_p}\quad\text{subject to}\quad\sum_{p\in g_p}N_i(\boldsymbol{x}_p)\mathcal{T}_p=R_{i,g_p}. \tag{20}$$

Here $V_p>0$ is a spatial allocation weight and is distinct from the particle mass. Under the basic PtoG rule, V_p partitions the nodal volumes supported by the particle group:

$$V_p=\sum_{i:,\sum_{q\in g_p}N_i(\boldsymbol{x}_q)>0}\frac{N_i(\boldsymbol{x}_p)\mu_i}{\sum_{q\in g_p}N_i(\boldsymbol{x}_q)},\qquad M_{ij,g_p}=\sum_{p\in g_p}V_pN_i(\boldsymbol{x}_p)N_j(\boldsymbol{x}_p). \tag{21}$$

Stationarity of the minimization problem yields the closed-form expression for the transport rates:

$$\mathcal{T}_p=V_p\sum_{i,j}N_i(\boldsymbol{x}_p)(M_{g_p}^{-1})_{ij}R_{j,g_p}. \tag{22}$$

The matrix inverse represents a linear solve. On nodes covered by the particle group, the sampled basis must have full row rank; unsupported nodes carry zero target residual. Since the objective function is strictly convex, any feasible constraint set yields a unique set of particle rates. Substituting this solution back and summing with weighting $N_i(\boldsymbol{x}_p)$ exactly recovers the residual R_{i,g_p}. This shows that our single-particle evolution rule comes from enforcing the group-level balance together with an explicit allocation principle.

We apply this same allocation rule for equilibrium collision gain, by replacing R_{j,g_p} with the corresponding equilibrium load μ_jG_{j,g_p}:

$$G_p^L(\boldsymbol{v})=V_p\sum_{i,j}N_i(\boldsymbol{x}_p)(M_{g_p}^{-1})_{ij}\mu_jG_{j,g_p}(\boldsymbol{v}). \tag{23}$$

Here G_p^L denotes the resulting velocity-space mass gain lifted onto individual particles; the superscript L indicates this lifting operation. The particle evolution equations read

$$\dot{\boldsymbol{x}}_p=\boldsymbol{U}_p,\qquad\partial_t\psi_p=\mathcal{T}_p+\mathcal{A}_p+C_p,\qquad C_p=\frac{G_p^L-\psi_p}{\tau}. \tag{24}$$

When summed over nodes, this recovers Eq. (19). For $\tau = \infty$, the collision term C_p vanishes. We note that neither G_p^L nor this allocation scheme guarantees nonnegativity by default.

5.2. Advancing particle positions and velocity functions

We use second-order time integration for the particle evolution in Eq. (24). We first consider collisionless problems, where the collision term C_p vanishes, and we advance particle positions and velocity-distribution coefficients together. Denote these particle variables by Y. Their rates of change, denoted by $L(Y)$, are determined by $\dot{\boldsymbol{x}}_p = \boldsymbol{U}_p$ and $\partial_t \psi_p = \mathcal{T}_p + \mathcal{A}_p$. Each evaluation of $L(Y)$ uses the current particle positions and distributions to perform PtoG reconstruction, evaluate transport and electric acceleration, and allocate the transport rates back to particles.

The second-order explicit midpoint method (RK2) first predicts the particle state at half a time step. The rates evaluated at this predicted state then advance the particles over the full step:

$$Y^* = Y^n + \tfrac{1}{2}\Delta t L(Y^n), \qquad Y^{n+1} = Y^n + \Delta t L(Y^*). \tag{25}$$

The RK2 method evaluates the particle rates twice per step. In this method, each intermediate state consists of particle positions and particle-carried velocity distributions, and the grid is used to evaluate their rates of change.

For collisional problems with constant $0 < \tau < \infty$, we retain the exponential relaxation associated with the loss term $-\psi_p/\tau$. Integrating Eq. (24) over an interval of length h gives

$$\psi_p(t+h) = e^{-h/\tau}\psi_p(t) + \int_0^h e^{-(h-s)/\tau}[\mathcal{T}_p + \mathcal{A}_p + G_p^L/\tau](t+s)\,ds. \tag{26}$$

This identity is exact: all terms in the integral are evaluated along the evolving particle state. In particular, G_p^L depends on the collective particle distribution and may change during the step. The numerical update below approximates this time-dependent integral using a half-step particle state and a recomputed equilibrium gain.

For an interval-to-relaxation-time ratio z, define

$$c(z) = \frac{z}{e^z - 1}, \qquad a(z) = 1 - c(z), \qquad b(z) = \frac{z}{1 - e^{-z}} - 1, \tag{27}$$

with $c(0) = 1$ and $a(0) = b(0) = 0$. In the collisionless limit $\tau \to \infty$, the update below therefore reduces to the explicit midpoint method in Eq. (25). These weights satisfy $c/(1+b) = e^{-z}$ and $(a+b)/(1+b) = 1 - e^{-z}$. Thus, when transport and acceleration vanish and the equilibrium gain is constant, the update reproduces exponential relaxation exactly. For a changing gain, its accuracy also depends on when the gain and transport rates are evaluated; the weights alone do not make the full integral exact.

In the following steps, a superscript $*$ denotes the half-step stage, and a tilde denotes a provisional distribution before the collision correction. A particle state Y contains all particle positions and velocity-distribution coefficients. The notation $(G_p^L)[Y]$ means: reconstruct collective moments from Y, form the equilibrium, and apply Eq. (23) to obtain its particle gain.

Half-step prediction. Starting from Y^n, evaluate $\boldsymbol{U}_p^n$, $\mathcal{T}_p^n$, $\mathcal{A}_p^n$, and $(G_p^L)^n = (G_p^L)[Y^n]$. Set $h = \Delta t/2$ and calculate

$$\begin{aligned} \boldsymbol{x}_p^* &= \boldsymbol{x}_p^n + h\boldsymbol{U}_p^n, \\ \widetilde{\psi}_p^* &= c(h/\tau)\psi_p^n + h(\mathcal{T}_p^n + \mathcal{A}_p^n) + a(h/\tau)(G_p^L)^n. \end{aligned} \tag{28}$$

The predicted positions $\boldsymbol{x}_p^*$ and coefficients of $\widetilde{\psi}_p^*$ together define the provisional state $\widetilde{Y}^*$. This prediction already includes relative transport, electric acceleration, and collision contributions.

Half-step collision correction. Recompute the equilibrium gain from $\widetilde{Y}^*$ and complete the half-step distribution update:

$$\psi_p^* = \frac{\widetilde{\psi}_p^* + b(h/\tau)(G_p^L)[\widetilde{Y}^*]}{1 + b(h/\tau)}. \tag{29}$$

The positions remain $\boldsymbol{x}_p^*$. Together with the corrected distributions ψ_p^*, they define Y^*. Here "collision correction" refers to Eq. (29), not to the nonnegativity correction discussed in the next subsection.

Full-step prediction using the half-step rates. From Y^*, calculate the particle mean velocities $\boldsymbol{U}_p^*$ and repeat the reconstruction and allocation to obtain $\mathcal{T}_p^*$. For charged particles, reconstruct the electric field from the half-step charge and calculate $\mathcal{A}_p^*$; for neutral particles, this term is zero. These quantities all use the corrected half-step distributions. Advance from the original state at t^n over the full interval Δt:

$$\begin{aligned} \boldsymbol{x}_p^{n+1} &= \boldsymbol{x}_p^n + \Delta t \boldsymbol{U}_p^*, \\ \widetilde{\psi}_p^{n+1} &= c(\Delta t/\tau)\psi_p^n + \Delta t(\mathcal{T}_p^* + \mathcal{A}_p^*) + a(\Delta t/\tau)(G_p^L)^n. \end{aligned} \tag{30}$$

This update originates from Y^n and advances over a full time step using the rate computed at Y^*. The new positions and provisional distributions define $\widetilde{Y}^{n+1}$. The initial gain $(G_p^L)^n$ is retained here; the endpoint gain is incorporated in the next equation.

Full-step collision correction. Reconstruct the collective equilibrium from $\widetilde{Y}^{n+1}$, allocate its gain, and set

$$\psi_p^{n+1} = \frac{\widetilde{\psi}_p^{n+1} + b(\Delta t/\tau)(G_p^L)[\widetilde{Y}^{n+1}]}{1 + b(\Delta t/\tau)}. \tag{31}$$

The final state Y^{n+1} contains $\boldsymbol{x}_p^{n+1}$ and the coefficients of ψ_p^{n+1}. At open boundaries, remove particles that have left the domain and record their carried mass, momentum, and energy before reconstructing this final equilibrium gain.

We note that, for smooth states at fixed τ, the weights have the small-step expansions $c(z) = 1 - z/2 + O(z^2)$ and $a(z) = b(z) = z/2 + O(z^2)$, providing the weights of a second-order collision treatment. Consistent gain allocation preserves the collective collision-invariant moments during each collision correction.

5.3. Unresolved particle information, admissibility, and boundaries

Different adjustments to particle coefficients can cancel one another when projected onto all background nodes. To suppress such unresolved particle modes, we add a stabilization rate to $\mathcal{T}_p$:

$$\lambda \left[V_p \sum_{i,j} N_i(\boldsymbol{x}_p)(M_{g_p}^{-1})_{ij} \sum_{q \in g_p} N_j(\boldsymbol{x}_q)\psi_q - \psi_p \right]. \tag{32}$$

Here $\lambda \geq 0$ is a numerical stabilization rate that controls the redistribution of velocity information among particles without changing the nodal distributions. Its value is specified in §Appendix B.2. This term sums to zero over all nodes, and therefore does not alter the constrained nodal distribution rates enforced by our balance conditions. It only modifies how the total rate is distributed among particles and can alter the resulting position–velocity correlations.

The specific stabilization form is provided in §Appendix B, and it is assembled together with Eq. (22) within the same spatial linear solve.

For shock simulations, we further introduce a constrained correction that minimizes changes to the particle coefficients while preserving mass, momentum and energy whenever feasible. When a single particle cannot satisfy the constraints on its own, moments are redistributed among nearby particles such that the total moments of the particle group remain unchanged. The full set of constraints, Bernstein bounds and the neighborhood fallback strategy are documented in §Appendix C. We omit this correction for smooth test cases, where lower-bound constraints are only examined diagnostically.

At absorbing boundaries, particles that exit the computational domain are removed from the simulation. We accumulate the moments carried by these removed particles together with the integrated outward relative flux, leading to the global balance for $q = 0$:

$$\sum_{p\ \text{active}} \int (1, \boldsymbol{v}, |\boldsymbol{v}|^2)\psi_p^{n+1}\, d\boldsymbol{v} + \boldsymbol{J}_{\text{out}}^{n+1} = \sum_{p\ \text{initial}} \int (1, \boldsymbol{v}, |\boldsymbol{v}|^2)\psi_p^{0}\, d\boldsymbol{v}. \tag{33}$$

Here $\boldsymbol{J}_{\text{out}}$ denotes the cumulative budget of escaped raw moments, which serves as a diagnostic balance relation and does not modify particle evolution. Periodic simulations have no outward flux contribution.

6. Theoretical properties of LVSP

6.1. Low-noise simulation and compact velocity-space representation

In LVSP, velocity moments can be computed from particle distribution functions either analytically or via deterministic quadrature. This deterministic loading procedure eliminates Monte Carlo variance from velocity sampling. Nevertheless, spatial quadrature, finite velocity-space expansions, field interpolation, time integration, and correction steps still introduce deterministic approximation errors.

The choice of velocity-space representation balances accuracy and computational cost and can vary across physical problems. For example, a local basis centered and scaled to a narrow velocity peak can represent that peak using only a small number of coefficients. Bulk motion at the particle's own mean velocity accounts for its primary translational motion, while the exchange term captures the remaining relative transport and shape deformation. Distributions with well-separated peaks may be handled by independent particle groups. A broad but smooth distribution can also admit a compact representation after appropriate scaling. The number of degrees of freedom required is governed by the complexity of the distribution shape and the target accuracy.

The total-degree basis introduced in §Appendix A.1 retains cross-velocity coupling terms while excluding high-degree monomial combinations along individual velocity coordinates. For polynomial order r in d_v velocity dimensions, the number of basis functions is

$$K = \binom{r + d_v}{d_v}. \tag{34}$$

For $r = 6$, $K = 7, 28, 84$ in one, two, and three velocity dimensions, respectively. By contrast, a tensor-product basis ($K_{\text{tensor}} = (r+1)^{d_v}$) with degree six on each coordinate yields $7, 49, 343$ modes. Multiplication and differentiation

link the retained polynomial terms to one another. We can build this basis and its operators directly, without first generating the full tensor-product basis. For a fixed polynomial order r, the number of basis functions grows more slowly as the velocity dimension d_v increases.

As dimensions increase, memory and computational costs tend to grow inherently, and fine structures and other factors can further raise the overhead. Nevertheless, GtoP operator reuse and compact velocity-space representations mitigate the additional cost brought by dimensional expansion. Section 7.3 compares the accuracy, storage, and computational cost of total-degree and tensor-product velocity representations within LVSP for one-, two-, and three-dimensional transport problems.

6.2. Conservation through particle–mesh exchange

Proposition 1. *Assume a partition of unity, feasible full-rank allocation constraints, exact evaluation of retained invariant moments, consistent equilibrium loads, and conservative boundary and correction operations. For a neutral periodic system, the semidiscrete particle scheme conserves total mass, momentum, and kinetic energy. The time-stepping schemes introduced above conserve these linear global invariants. When absorbing boundaries are used, the update obeys the global balance in Eq.* (33).

Proof. Summing the allocation constraints over all nodes is equivalent to summing the particle rate terms, owing to $\sum_i N_i(\boldsymbol{x}_p) = 1$. Interior contributions from relative transport vanish upon summation because $\sum_i \nabla N_i = 0$. The stabilization term in Eq. (32) sums to zero on every node. The collective equilibrium gain shares the same invariant moments as the corresponding loss, and corrections are designed to preserve the total moments of individual particles or particle groups. Explicit multi-stage Runge–Kutta updates preserve linear invariants. In the collisional update, $c + a = 1$, so changes to the provisional invariants come solely from transport; the gain-recovery step preserves these provisional invariants, since the equilibrium gain is reconstructed from the provisional state. Outward relative flux and removed particles together supply the moment budget in Eq. (33). □

Particle masses are allowed to evolve. Where the mass is positive, $\phi_p = \psi_p / \int \psi_p$ is normalized by definition, and particle masses evolve through exchange, while the total mass remains conserved. Importantly, satisfaction of the global invariant balances alone does not guarantee positivity across the full velocity domain, convergence, or entropy stability; the conditional Bernstein bounds used to enforce admissibility are valid only within their prescribed intervals.

For electrostatic acceleration, integration by parts applied to Eq. (18) yields the particle force and power moments $(q/m)\omega_p \boldsymbol{E}_h(\boldsymbol{x}_p)$ and $(q/m)\omega_p \boldsymbol{U}_p \cdot \boldsymbol{E}_h(\boldsymbol{x}_p)$. The diagnostic total energy reads

$$\mathcal{H} = \frac{1}{2} \sum_p \int |\boldsymbol{v}|^2 \psi_p \, d\boldsymbol{v} + \frac{\varepsilon_0}{2} \int_{\mathcal{D}} |\boldsymbol{E}_h|^2 \, d\boldsymbol{x}. \tag{35}$$

The field and particle time discretizations used in this work do not enforce exact cancellation between the work done on particles and changes in the field energy.

6.3. A unified framework for charged and neutral particles across collision regimes

LVSP decomposes the microscopic flux into advection by particle mean velocity and relative exchange: $\boldsymbol{U}_p\psi_p + (\boldsymbol{v}-\boldsymbol{U}_p)\psi_p = \boldsymbol{v}\psi_p$. The flexibility of this representation arises from the amount of velocity-space information retained: a particle can behave like a single velocity sample in conventional PIC, or carry internal energy and stress consistent with a fluid material point. Setting $q = 0$ disables electric acceleration while preserving the entire transport architecture. As a result, the LVSP framework supports both charged and neutral particles within the same formulation. In what follows, we derive the asymptotic reduction of the scheme to standard PIC in the collisionless limit and to MPM in the continuum regime.

Single-velocity, collisionless limit. Take $\psi_p = \omega_p\delta^{(d)}(\boldsymbol{v} - \boldsymbol{U}_p)$ and $\tau^{-1} = 0$. Then $(\boldsymbol{v} - \boldsymbol{U}_p)\psi_p = 0$, so relative transport vanishes. With exchange stabilization and distribution corrections disabled, the force moments give

$$\dot{\omega}_p = 0, \qquad \dot{\boldsymbol{x}}_p = \boldsymbol{U}_p, \qquad \dot{\boldsymbol{U}}_p = \frac{q}{m}\sum_i N_i(\boldsymbol{x}_p)\boldsymbol{E}_i. \tag{36}$$

Together with the same charge deposition, field interpolation, and RK2 update, these are the PIC equations and their matched discrete update. The velocity centres are advanced directly. This specialization yields the single-velocity limit, which is a velocity Dirac delta distribution.

Local-equilibrium fluid limit. Taking velocity moments gives material momentum balance,

$$\rho\frac{D\boldsymbol{u}}{Dt} = -\nabla_x\cdot\boldsymbol{\Pi} + \frac{q}{m}\rho\boldsymbol{E}. \tag{37}$$

where $\boldsymbol{\Pi} = \int(\boldsymbol{v}-\boldsymbol{u})\otimes(\boldsymbol{v}-\boldsymbol{u})f\,d\boldsymbol{v}$. Under collective local Maxwellian equilibrium, $\boldsymbol{\Pi} = \rho T\mathsf{I}$ and the heat flux vanishes. For a neutral gas, the remaining mass and internal-energy balances are

$$\frac{D\rho}{Dt} = -\rho\nabla_x\cdot\boldsymbol{u}, \qquad \rho\frac{d}{2}\frac{DT}{Dt} = -\rho T\nabla_x\cdot\boldsymbol{u}. \tag{38}$$

The connection to the particle weak balance follows directly from the Maxwellian relative-flux moments. For $\psi_p = \omega_p g(\boldsymbol{v}; \boldsymbol{U}_p, T_p)$, with particle temperature T_p,

$$\int \boldsymbol{v}\otimes(\boldsymbol{v}-\boldsymbol{U}_p)\psi_p\,d\boldsymbol{v} = \omega_p T_p\mathsf{I}, \qquad \int \tfrac{1}{2}|\boldsymbol{v}|^2(\boldsymbol{v}-\boldsymbol{U}_p)\psi_p\,d\boldsymbol{v} = \omega_p T_p\boldsymbol{U}_p. \tag{39}$$

Thus the relative-transport exchange becomes pressure-force and pressure-work exchange. Moving points carry the fluid information, and the background grid evaluates these interactions, giving a fluid material-point description [19]. This is the sense in which the local-equilibrium moment equations reduce to fluid MPM.

This asymptotic reduction recovers the expected limiting equations, though different MPM transfer rules may produce distinct trajectories under finite resolution. The present pressure reconstruction, volume-weighted allocation, and redistribution of individual particle masses differ from conventional fixed-mass, lumped-grid MPM. Appendix B.4 states these differences; Section 7.5 compares the resulting fluid response. Small-τ acoustic tests support this response.

Between the two limits, particle velocity functions retain nonequilibrium stress and heat transport. The same architecture covers the separate charged and neutral cases tested here. Interacting charged–neutral species additionally require collision operators conserving their combined momentum and energy.

6.4. Extensibility to electromagnetic problems

Mathematically, the framework of LVSP can be extended to electromagnetic problems by replacing the electrostatic force term with the Lorentz force. The intrinsic properties of the Lorentz force are compatible with our velocity polynomial operators. For three-dimensional velocity space, the formal extension replaces Eq. (18) with

$$\mathcal{A}_p = -\nabla_v \cdot \left\{ \frac{q}{m} [\boldsymbol{E}_h(\boldsymbol{x}_p) + \boldsymbol{v} \times \boldsymbol{B}_{\mathrm{em},h}(\boldsymbol{x}_p)] \psi_p \right\}. \tag{40}$$

Lorentz acceleration has zero divergence in velocity space, and magnetic forces do not change particle kinetic energy. Multiplication and differentiation can be applied to the velocity expansion coefficients. Nevertheless, full electromagnetic simulations require additional supporting modules such as a Maxwell solver, current deposition scheme, and time integration suitable for magnetized plasmas, which are not validated in the present work.

7. Numerical experiments

The tests follow the properties discussed in Section 6: low noise and compact velocity representation first, conservation second, and the kinetic-to-fluid connection and multidimensional applications last. The charged tests are electrostatic.

7.1. Common settings and diagnostics

All kinetic tests use basic PtoG, RK2 particle stages with exponential collision weights when needed, particle-mean motion, and volume-weighted GtoP. Spatial bases are P1 in one dimension and multilinear in two and three dimensions. Initial particle positions and velocity coefficients are deterministic. There is no spatial reseeding. The single-velocity PIC specialization is stated separately below.

Both solutions are averaged over the same comparison cells. Let C_j be comparison cell j and $|C_j|$ its volume. The reconstructed LVSP field is $F_h(\boldsymbol{x}, \boldsymbol{v}) = \sum_i N_i(\boldsymbol{x}) F_i(\boldsymbol{v})$, giving $\widehat{f_j} = |C_j|^{-1} \int_{C_j} F_h \, d\boldsymbol{x}$ and $\widehat{\rho}_j = \int \widehat{f_j} \, d\boldsymbol{v}$. The reference cell averages are f_j^{ref} and $\rho_j^{\mathrm{ref}} = \int f_j^{\mathrm{ref}} \, d\boldsymbol{v}$. At the comparison time, their total mass in the comparison region is $M_{\mathrm{ref}} = \sum_j |C_j| \rho_j^{\mathrm{ref}}$. The relative L^1 errors are

$$e_f = \frac{\sum_j |C_j| \int |\widehat{f_j} - f_j^{\mathrm{ref}}| \, d\boldsymbol{v}}{M_{\mathrm{ref}}}, \qquad e_\rho = \frac{\sum_j |C_j| |\widehat{\rho}_j - \rho_j^{\mathrm{ref}}|}{M_{\mathrm{ref}}}. \tag{41}$$

Neutral conservation checks include mass, momentum, and energy leaving absorbing boundaries. Electrostatic energy includes the field contribution in Eq. (35). Smooth distributions are checked at selected velocities against the normalized-shape tolerance -10^{-8}. The shock also uses the conservative correction in §Appendix C.

7.2. Low-noise electrostatic signals

These tests examine the weak-signal resolution discussed in Section 6.1. On $[0, 4\pi)$, take $q = -1$, $m = \varepsilon_0 = 1$, $\tau = \infty$, and background charge one, so $Q = 1 - \rho$. Initial data are

$$\begin{aligned} &\text{Landau:} && [1 + 10^{-4} \cos(x/2)] g(v; 0, 1), \quad t_{\mathrm{end}} = 12, \\ &\text{Two stream:} && [1 + 10^{-3} \cos(x/2)] [g(v; -1, 0.25) + g(v; 1, 0.25)]/2, \quad t_{\mathrm{end}} = 8. \end{aligned} \tag{42}$$

Landau damping uses 128 cells, 512 particles, Hermite order up to 64, and $\Delta t = 0.0025$. Two stream uses 128 cells, 256 particles carrying two order-32 components, and $\Delta t = 0.005$. Both allow the checked order reduction in §Appendix A. The spectral reference uses 256 spatial points, 512 velocities on $[-8, 8)$, and Strang semi-Lagrangian steps of 0.01; refinement checks make its electric-mode changes small compared with the reported LVSP error.

Both PIC loadings use 64 cells, 16384 particles, linear deposition and interpolation, Fourier Poisson inversion, and kick–drift–kick steps of 0.01. Quiet loading combines 64 positions with 256 velocity quantiles; four random loadings independently sample the velocity mixture. For N_x uniformly spaced points x_i on $[0, 4\pi)$, let $E_i(t)$ be the sampled electric field. The first Fourier mode has wavenumber $1/2$; its complex amplitude uses the factor-two convention below. With i the imaginary unit, the mode and its relative discrete L^2 history error are

$$\widehat{E}_1(t) = \frac{2}{N_x} \sum_i E_i(t) e^{-\mathrm{i}x_i/2}, \qquad e_E = \frac{\left(\sum_n |\widehat{E}_1(t_n) - \widehat{E}_1^{\mathrm{ref}}(t_n)|^2\right)^{1/2}}{\left(\sum_n |\widehat{E}_1^{\mathrm{ref}}(t_n)|^2\right)^{1/2}}. \tag{43}$$

Here t_n are common diagnostic times spaced by 0.1, and superscript ref denotes the spectral reference evaluated at those times. LVSP follows the reference damping and growth while random-start PIC fluctuations can obscure the weak signal (Fig. 1). Quiet-start PIC reduces these fluctuations, providing a useful additional comparison.

Table 1: Particle-stage electrostatic results. Phase errors compare $f - \langle f \rangle_x$, where $\langle f \rangle_x$ is the spatial average, with the reference on 128 common cells. Energy drift is the maximum absolute change in Eq. (35).

Problem	Particles	$100e_E$	Phase error (%)	Energy drift
Landau	512	0.11214	0.65035	4.90×10^{-11}
Two stream	256	0.15326	0.25441	4.48×10^{-9}

Figure 2 applies the same 64 periodic spatial hats and 192 velocity bins on $[-6, 6]$ to all solutions, without extra PIC smoothing. Figure 3 shows $f - \langle f \rangle_x$ to expose the small kinetic perturbations. Together, these figures show that the smooth LVSP result retains the physical perturbation rather than simply suppressing all variation.

7.3. Compact velocity representation and multidimensional cost

This subsection examines the representation choices in Section 6.1. For collisionless transport, use $f_0 = [1 + 0.2\cos(2\pi x)][g(v; -U, T) + g(v; U, T)]/2$ and compare with $f_0(x - vt, v)$ at $t = 0.5$ on 256 cells. The broad case $(U, T) = (1, 0.3)$ uses 128 cells, 4096 particles, order up to 16, and $\Delta t = 0.0003125$. The narrow case $(U, T) = (3, 10^{-6})$ uses 32 cells, 128 particles, order two, and $\Delta t = 0.05$. Both resolve the transported distribution, while the narrow peaks require only 384 coefficients. Particle motion carries their main translation.

To assess dimensional cost separately, compare total-degree and tensor-product bases within LVSP. The initial mixture has centres $\pm(1, 0.7, -0.4)$ restricted to the chosen dimensions, covariance diagonal 0.3, and off-diagonal entries $0.03/|j - k|$. Spatial waves follow each coordinate axis and, in multiple dimensions, $(1, \ldots, 1)$, each with amplitude $0.2/\sqrt{\text{number of waves}}$. All cases use 13^d cells and four particles per cell, with each velocity group at relative positions $(1/4, \ldots, 1/4)$ and $(3/4, \ldots, 3/4)$. Both bases have nominal degree eight and use $\Delta t = 0.005$ to $t = 0.1$, with identical reconstruction, motion, and stabilization settings.

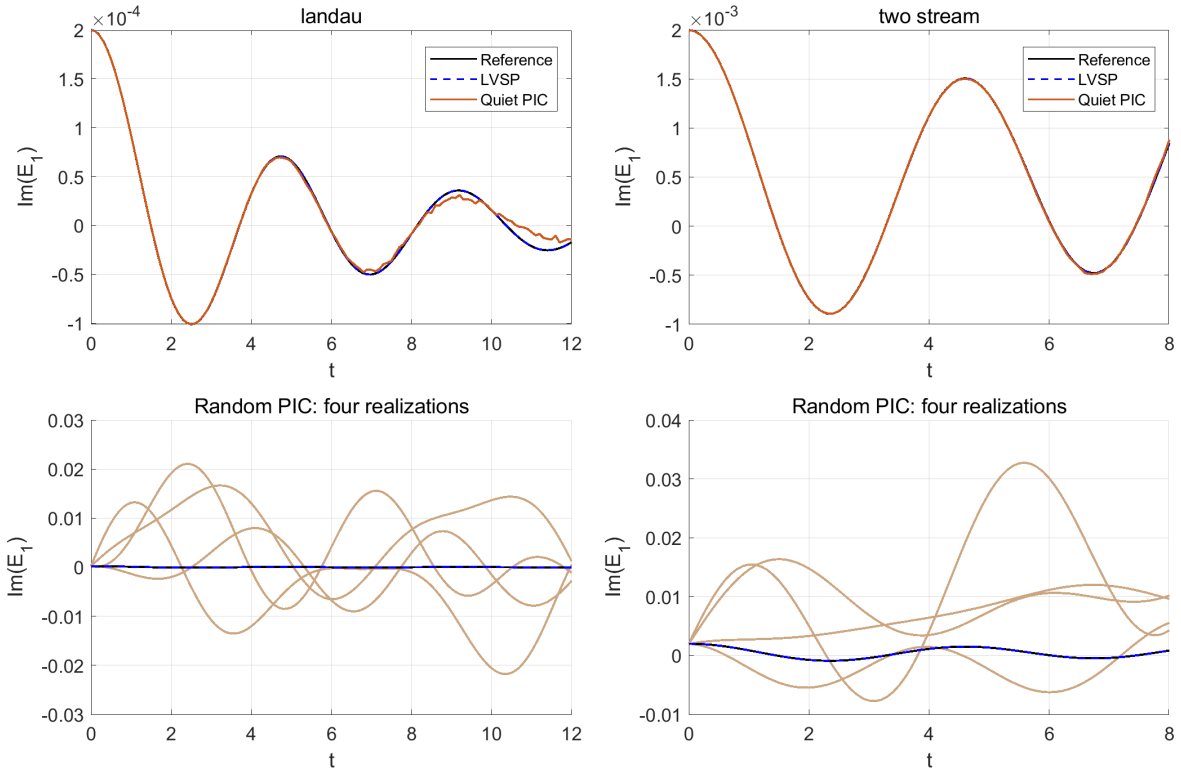


Fig. 1: First electric-mode histories. Top: spectral reference, particle-stage LVSP, and quiet-start PIC. Bottom: four random-start PIC realizations with reference and LVSP overlaid. Vertical scales differ between rows.

Table 2: Internal LVSP comparison on the same spatial discretization. Coefficients count all evolved particle distribution coefficients. Time includes initialization and evolution, excluding error integration; medians use five serial single-thread runs, or three in 3D.

Dimension	Velocity basis	Particles	Coefficients	$100e_f$	Time (s)
1D1V	Total degree	52	468	0.506	0.0352
1D1V	Tensor product	52	468	0.506	0.0268
2D2V	Total degree	676	30420	0.720	0.118
2D2V	Tensor product	676	54756	0.720	0.173
3D3V	Total degree	8788	1450020	0.887	28.7
3D3V	Tensor product	8788	6406452	0.887	124

The results are shown in Table 2. The two bases coincide in one dimension. In higher dimensions, total-degree functions retain fewer mixed high-degree terms, reducing storage and run time while giving similar distribution errors. Spatial particle counts still grow with dimension. The comparison demonstrates that, at equal polynomial degree and spatial discretization, total-degree bases reduce the number of coefficients and computational cost.

7.4. Conservation in collision relaxation and open-boundary transport

This subsection checks the conservation properties discussed in Section 6.2, first for collision relaxation without spatial transport and then for discontinuous transport with boundary outflow.

The homogeneous relaxation test starts from $[g(v;-1,0.3)+g(v;1,0.3)]/2$, with $\tau = 0.02$ and equilibrium $g(v;0,1.3)$. It uses 16 cells, 64 particles, degree 12, and an additional collision-gain component. Single-step tests

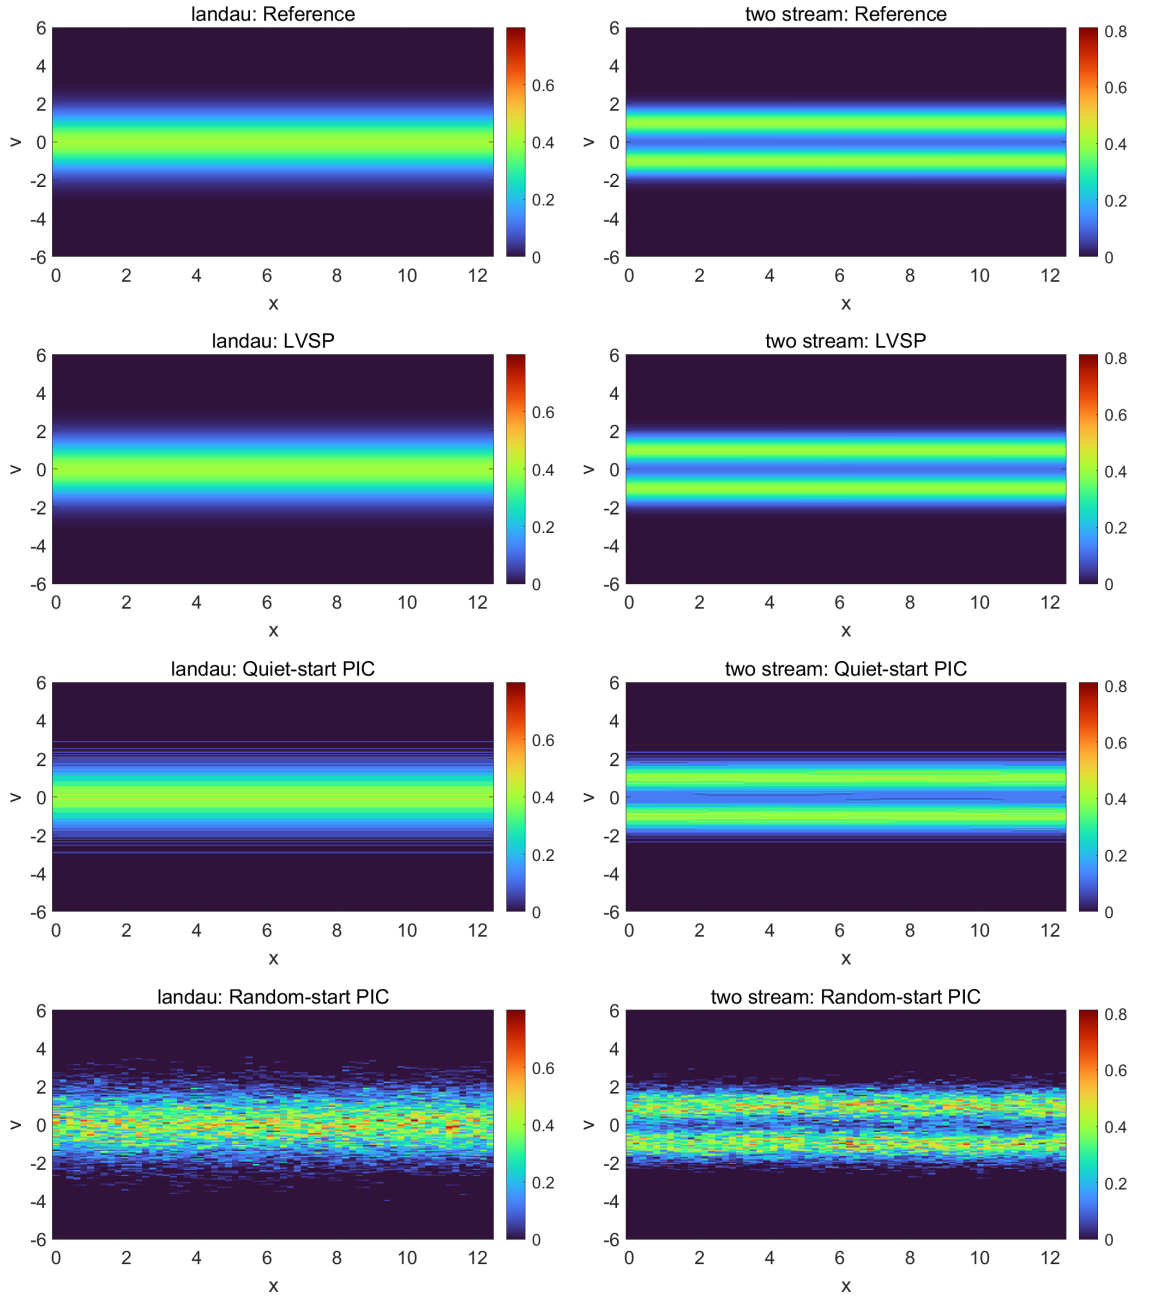


Fig. 2: Terminal phase-space distributions: Landau at $t = 12$ (left) and two stream at $t = 8$ (right). Rows: reference, LVSP, quiet-start PIC, and random-start PIC. Each column shares one colour range and observation operator.

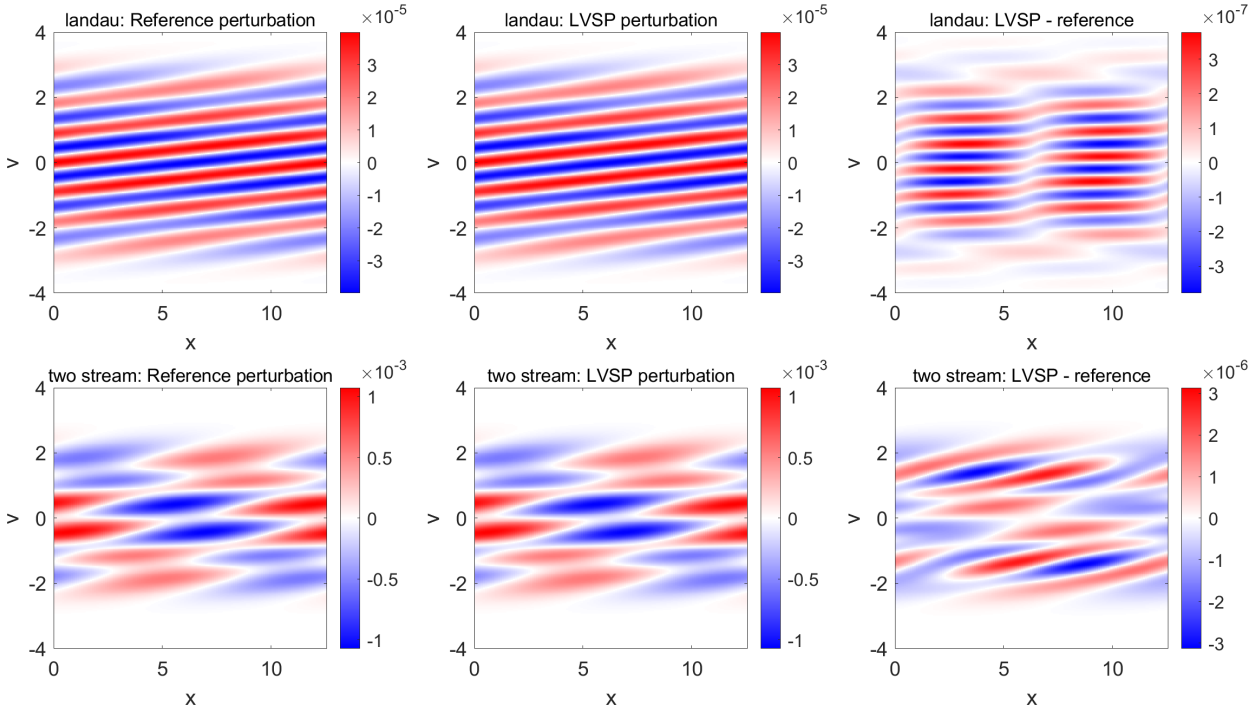


Fig. 3: Terminal perturbations $f - \langle f \rangle_x$ on 128 common observation cells. Columns: reference, LVSP, and their signed difference. Reference and LVSP share a scale; error scales are separate.

with fixed τ and $\Delta t/\tau = 0.1, 1, 10, 100$ recover exact exponential relaxation and conserve mass, momentum, and energy to roundoff (Fig. 4). This isolates collisions; Section 7.5 tests their coupling with transport.

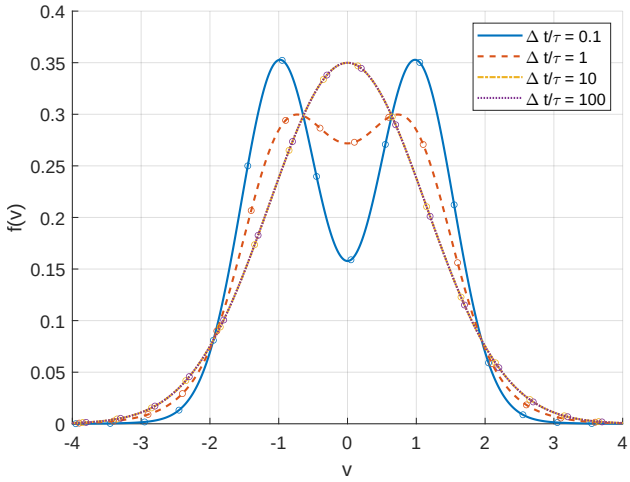


Fig. 4: Homogeneous collision relaxation. Four single-step solutions are compared with the exact relaxation, shown by lines, and LVSP results, shown by circles. The solutions for $\Delta t/\tau = 10$ and 100 nearly coincide.

We next consider a shock tube with absorbing boundaries to test transport across discontinuities and conservation in the presence of particle outflow. The initial distributions are Maxwellians on $[0, 1]$ with

$$(\rho, u, T) = \begin{cases} (1, 0, 1), & x < 0.5, \\ (0.125, 0, 0.8), & x \geq 0.5, \end{cases} \qquad \tau = 0.02, \quad t_{\text{end}} = 0.15. \tag{44}$$

Both methods impose zero incoming distribution. LVSP uses 256 cells, 1024 initial particles, two degree-12 half-range components, and $\Delta t = 0.000125$, with limited edge transport and conservative nonnegativity correction. The DVM reference uses 2048 cells, 321 velocities on $[-8, 8]$, and upwind transport with exact BGK relaxation in a Strang composition; its step is $0.4\Delta x_{\text{ref}}/8$.

Macroscopic profiles and local velocity distributions consistent with the reference solution are recovered throughout expansion and compression regions (Fig. 5). By accounting for both outgoing particle moments and boundary relative flux, the global mass, momentum, and energy budgets are closed to roundoff error.

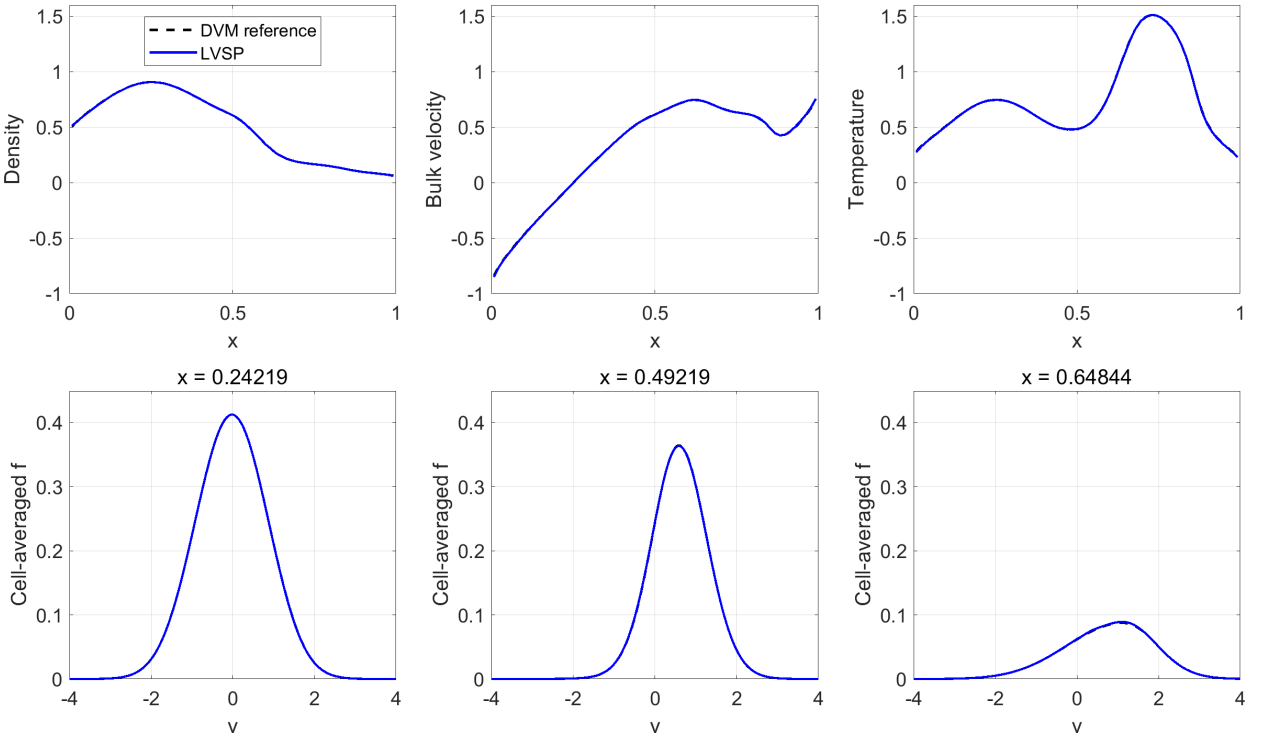


Fig. 5: Absorbing shock at $t = 0.15$. Top: density, bulk velocity, and temperature. Bottom: velocity distributions at common-cell centres $x = 0.24219, 0.49219, 0.64844$, using identical distribution and velocity ranges. Both solutions are observed on 64 common cells.

7.5. PIC and fluid MPM limits

This subsection tests the reductions discussed in Section 6.3. For the single-velocity PIC check, use 16 periodic cells on $[0, 2\pi)$ and 256 particles with $x_p = (p - 1/2)2\pi/256$, $\omega_p = (2\pi/256)(1 + 0.01\cos x_p)$, and $U_p = 0.4\sin[2\pi(p - 1)/7]$. Take $q/m = -1$, background charge one, and $t = 0.1$. Both implementations use the same Fourier Poisson solve, P1 coupling, and RK2 rule, with independent charge deposition and force gathering. Exchange stabilization is disabled and velocity centres advance directly. Stage and endpoint comparisons agree to roundoff (Table 3), verifying the single-velocity specialization.

An acoustic wave tests the fluid limit. Initialize $(\rho, u, T) = (1 + \eta, \sqrt{3}\eta, 1 + 2\eta)$ with $\eta = 10^{-3}\sin(2\pi x)$ on $[0, 1)$. LVSP uses 128 cells, 512 particles, degree four, $\Delta t = 0.002$, and $t = 0.04$, for $\Delta t/\tau = 0.1, 1, 10, 100$. The collisional update and stabilization are retained; no Maxwellian is imposed after a stage. Spatial checks use 64 and 256 cells

Table 3: Single-velocity LVSP versus matched PIC: maximum absolute differences over time.

Δt	Charge or field	Position or velocity
0.002	6.67×10^{-16}	1.12×10^{-16}
0.001	8.89×10^{-16}	5.56×10^{-17}

with four particles per cell.

The fluid-MPM comparison uses the same initial particle moments, RK2 step, and particle-mean motion. It retains fixed material masses and uses ideal-gas pressure, particle pressure quadrature, lumped nodal masses, and pressure work from the nodal velocity gradient. Fields are compared on 256 common cells. The linear Euler wave with speed $\sqrt{3}$ is an independent small-amplitude reference, not an exact finite-τ BGK solution.

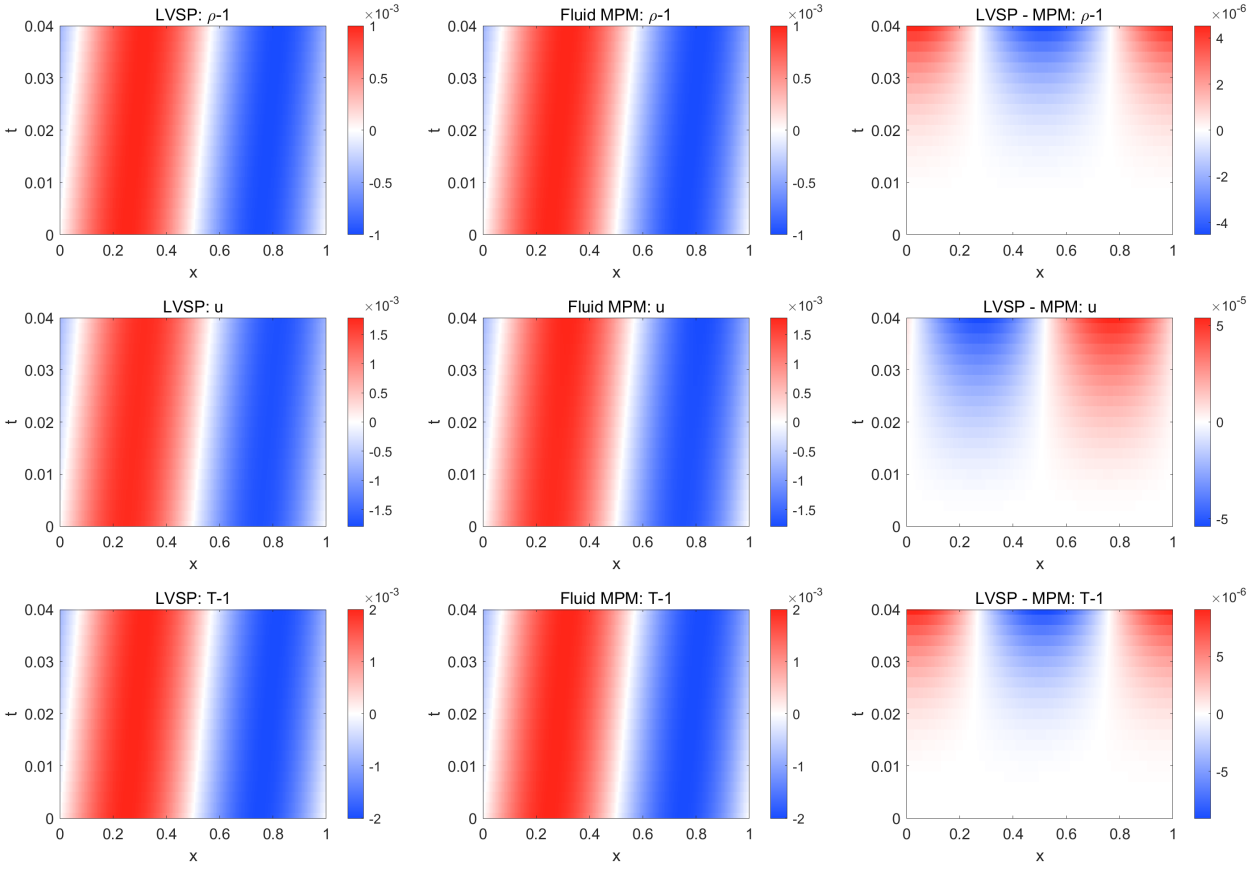


Fig. 6: Acoustic space–time fields at $\Delta t/\tau = 100$: LVSP, fluid MPM, and signed difference. Rows show density perturbation, velocity, and temperature perturbation. The first two columns share each row's colour range; difference scales are separate. Uniform backgrounds are removed to expose the weak response.

Figure 6 shows closely matching propagation. As collisions strengthen, temperature approaches the fluid response, while density and velocity differences plateau. Both methods follow the linear Euler wave; their remaining difference includes both discretization errors. This supports a common fluid response, rather than identical transfer rules.

7.6. Multidimensional kinetic transport

The final tests examine collisional 2D and collisionless 3D transport within the framework of Section 6.3. On a periodic unit square or cube, initialize

$$f_0(\boldsymbol{x},\boldsymbol{v}) = \left[1 + A\prod_{j=1}^{d}\cos(2\pi x_j)\right]\frac{1}{d}\sum_{r=1}^{d} g(\boldsymbol{v};\boldsymbol{e}_r,0.2), \qquad d = 2,3. \tag{45}$$

Here $\boldsymbol{e}_r$ is a coordinate unit vector. The mixture couples velocity coordinates rather than evolving independent one-dimensional problems. Both cases use total degree two, $\Delta t = 0.005$, and $t = 0.05$. Complementary subcell positions give four particles per cell in 2D and eight in 3D. The collisional 2D case adds an equilibrium component to each particle, giving 12 coefficients; the collisionless 3D case uses ten.

The 2D case has $A = 0.2$, $\tau = 0.02$, 40^2 cells, and 6400 particles. The reference combines Fourier free transport and exact local relaxation with Strang splitting, using 64^2 spatial points, 64^2 velocities on $[-6,6)^2$, and step 0.0005. Both results are averaged over 32^2 cells, with 64 velocity integration points per coordinate. Figure 7 shows agreement in both density and the two-dimensional velocity distribution, including the redistribution produced by collisions. The

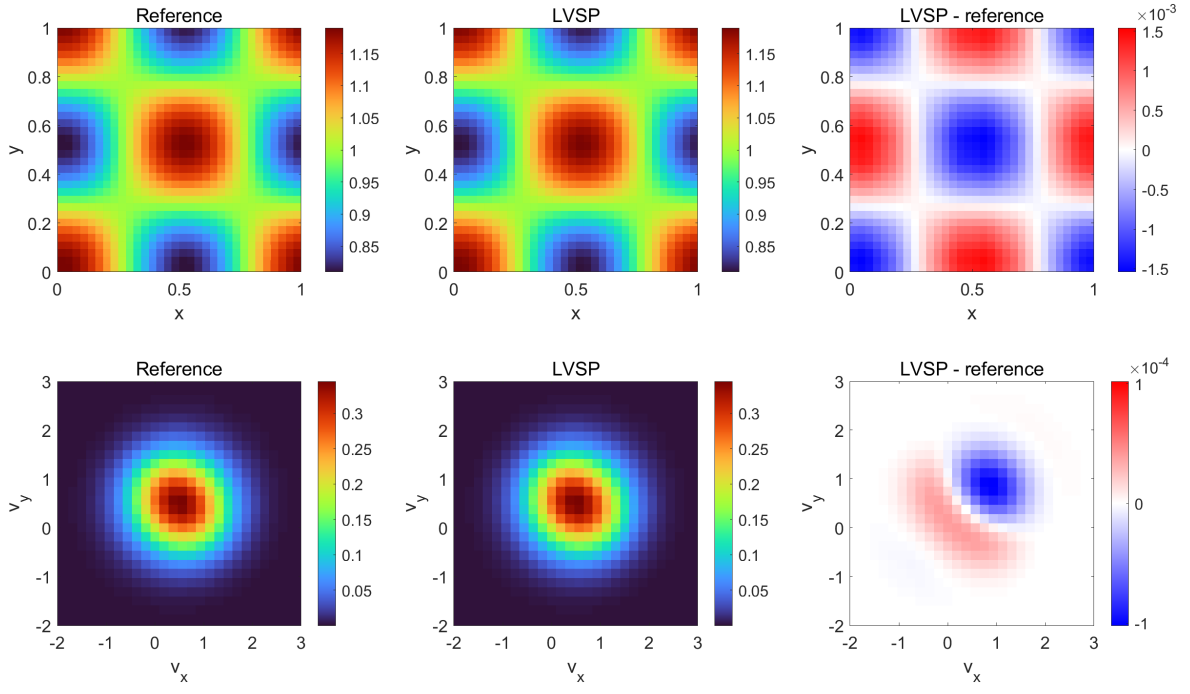


Fig. 7: Collisional crossing beams at $t = 0.05$. Columns: reference, LVSP, and signed difference. Top: density on 32^2 common cells; bottom: joint velocity distribution in the cell centred at $(0.23438, 0.23438)$. LVSP evolves on 40^2 cells. Reference–LVSP pairs have identical scales.

3D case has $A = 0.1$, $\tau = \infty$, 16^3 cells, and 32768 particles. The reference is the exact free-transport solution $f_0(\boldsymbol{x} - \boldsymbol{v}t, \boldsymbol{v})$. Comparisons use 8^3 cells and mixture-weighted Gaussian quadrature with 24 points per velocity coordinate. Figure 8 compares density and the full three-dimensional velocity distribution. Spatial and velocity structures agree with the references, and mass, momentum, and energy remain conserved to roundoff.

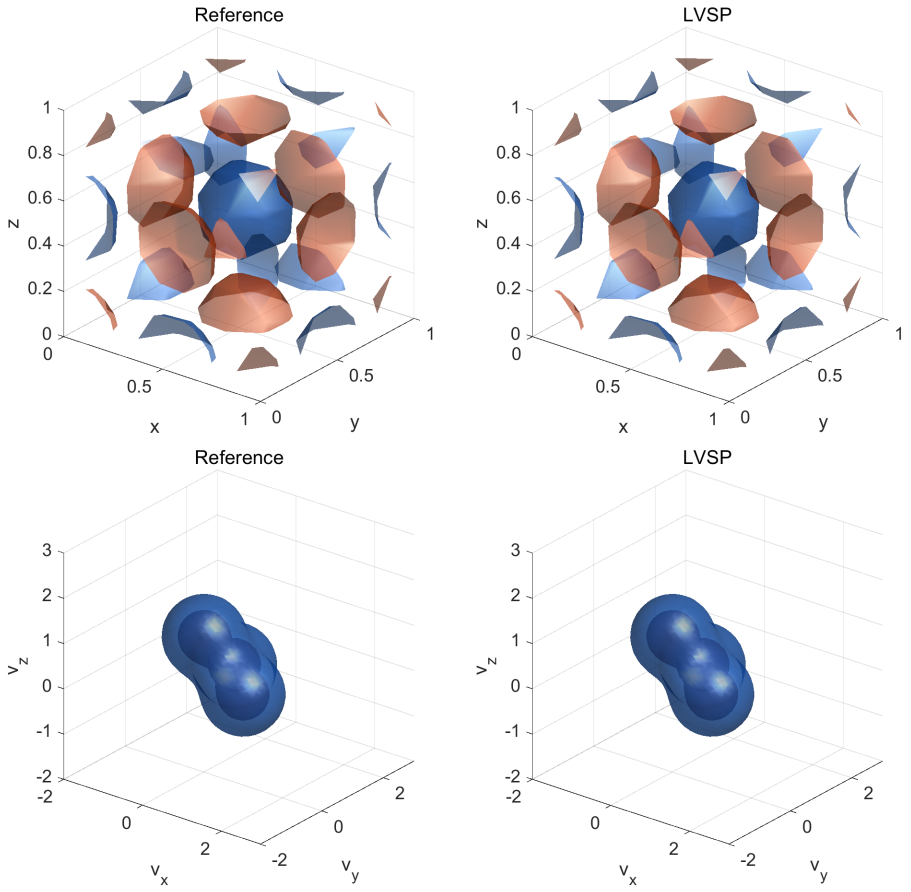


Fig. 8: Collisionless 3D transport at $t = 0.05$: reference (left) and LVSP (right). Top: density isosurfaces $\rho = 0.97, 1.03$ on the common 8^3 observation cells. Bottom: joint-distribution isosurfaces $f = 0.04, 0.12$ at $(0.1875, 0.1875, 0.1875)$. Levels, rendering, opacity, and view are identical within each pair.

8. Conclusions

LVSP combines moving spatial point particles with deterministic velocity distributions. Particle motion carries the mean transport, while conservative exchange accounts for motion relative to each particle's mean velocity. A temporary background grid organizes this exchange and supplies collective moments and fields. Positions and velocity functions evolve through particle time stages that include transport and collisions, without advancing a separate grid distribution.

The numerical tests demonstrate low-noise resolution of weak electrostatic signals, conservative relaxation and open-boundary transport, and a common treatment of neutral-gas dynamics across collision conditions. Compact local velocity functions resolve narrow peaks with few coefficients. At equal polynomial degree, total-degree bases reduce storage and computational cost in multiple dimensions, while the two- and three-dimensional transport tests verify the corresponding kinetic evolution. Beyond this, this framework connects single-velocity kinetic particles with fluid material points. With collisions and exchange stabilization disabled, the single-velocity specialization reduces to a matched PIC discretization, as confirmed by the stage-by-stage comparison. Under collective local equilibrium, the weak velocity moments reduce to fluid material-point equations with pressure force and pressure

work. The collisional acoustic results recover a response close to fluid MPM. These results support LVSP as a flexible particle framework linking kinetic and fluid descriptions. Further work will address long-time velocity adaptation, electromagnetic coupling, and interactions between charged and neutral species.

Acknowledgments

This work is supported by Doctoral Scientific Research Start-up Project of Liaoning Provincial Natural Science Foundation(2026-BS-0078, 2026-BS-0402).

Appendix A. Velocity-shape representations

This appendix specifies the velocity bases and the operations used for transport, acceleration, and collision gain. Gaussian weights combined with polynomial expansions represent nonequilibrium velocity distributions.

The indices p and i label particles and nodes; coordinate directions and secondary indices are specified where used. The highest polynomial degree is r, and K is the number of coefficients. Group and component labels are omitted for a single group or component. Symbols inherited from the main text retain their meanings; named matrices denote coefficient operations.

Appendix A.1. Local Hermite functions and their operators

Consider one velocity-basis component with centre $\boldsymbol{\xi}$ and variance scale $\theta > 0$. The function $g(\boldsymbol{v};\boldsymbol{\xi},\theta)$ is the normalized Gaussian of Eq. (3) with basis parameters in place of the physical mean and temperature. It serves as the weight in the velocity basis. In d_v velocity dimensions, let $\boldsymbol{\alpha} = (\alpha_1,\dots,\alpha_{d_v})$ list the polynomial degrees in each direction. For one component, the particle expansion in Section 3 becomes

$$\psi_p(\boldsymbol{v},t) = \sum_{|\boldsymbol{\alpha}|\le r} c_{p,\boldsymbol{\alpha}}(t)\chi_{\boldsymbol{\alpha}}(\boldsymbol{v}).$$

The multi-index replaces the basis index k in c_{pk}. Basis parameters remain fixed during each time stage, while coefficients evolve. The basis functions are

$$\chi_{\boldsymbol{\alpha}}(\boldsymbol{v}) = g(\boldsymbol{v};\boldsymbol{\xi},\theta)\prod_{j=1}^{d_v}\frac{\mathrm{He}_{\alpha_j}((v_j-\xi_j)/\sqrt{\theta})}{\sqrt{\alpha_j!}}, \qquad |\boldsymbol{\alpha}| = \sum_j \alpha_j \le r, \tag{A.1}$$

where He_k is the probabilists' Hermite polynomial. All mixed terms satisfying the total-degree bound are retained, including velocity correlations. For several components, sum their expansions; each component has its own centre, width, and coefficients. The following operations are first described for one component.

From basis functions to transport and acceleration operators. The Hermite recurrence and the derivative of the Gaussian weight give

$$v_j\chi_{\boldsymbol{\alpha}} = \xi_j\chi_{\boldsymbol{\alpha}} + \sqrt{\theta(\alpha_j+1)}\chi_{\boldsymbol{\alpha}+\boldsymbol{e}_j} + \sqrt{\theta\alpha_j}\chi_{\boldsymbol{\alpha}-\boldsymbol{e}_j},$$
$$\partial_{v_j}\chi_{\boldsymbol{\alpha}} = -\sqrt{(\alpha_j+1)/\theta}\chi_{\boldsymbol{\alpha}+\boldsymbol{e}_j},$$

where $\boldsymbol{e}_j$ increases the degree in direction j by one; the lowering term is zero if $\alpha_j = 0$. Substitution into the particle expansion gives the following coefficient operators after truncation. For particle coefficient row $\boldsymbol{c}_p = (c_{pk})_k$, the rows $\boldsymbol{c}_p\mathsf{Vel}_j$ and $\boldsymbol{c}_p\mathsf{Vder}_j$ represent $v_j\psi_p$ and $\partial_{v_j}\psi_p$, respectively, after truncation to the retained space. Rows index input basis functions and columns index output basis functions. The nonzero entries are

$$\begin{aligned}(\mathsf{Vel}_j)_{\boldsymbol{\alpha},\boldsymbol{\alpha}} &= \xi_j,\\ (\mathsf{Vel}_j)_{\boldsymbol{\alpha},\boldsymbol{\alpha}+\boldsymbol{e}_j} = (\mathsf{Vel}_j)_{\boldsymbol{\alpha}+\boldsymbol{e}_j,\boldsymbol{\alpha}} &= \sqrt{\theta(\alpha_j+1)},\\ (\mathsf{Vder}_j)_{\boldsymbol{\alpha},\boldsymbol{\alpha}+\boldsymbol{e}_j} &= -\sqrt{(\alpha_j+1)/\theta}.\end{aligned} \tag{A.2}$$

Omitting terms outside the retained space is the finite-degree approximation. Relative transport uses $\boldsymbol{c}_p\mathsf{Vel}_j - U_{p,j}\boldsymbol{c}_p$; electric acceleration uses the derivative operator. Terms through degree two determine mass, momentum, and energy.

From collective equilibrium to velocity coefficients. Collisions require the Maxwellian computed from the collective nodal moments to be expressed in the same basis. In one velocity dimension, define

$$\beta_k = \int_{\mathbb{R}} g(v;u,T)\frac{\mathrm{He}_k((v-\xi)/\sqrt{\theta})}{\sqrt{k!}}\,dv, \qquad \rho g(v;u,T) \approx \rho\sum_{k=0}^{r}\beta_k\chi_k(v).$$

Orthogonality makes $\rho\beta_k$ the expansion coefficient; the approximation is the finite-degree truncation. Gaussian integration by parts, together with the Hermite recurrence, gives the following efficient evaluation, starting from $\beta_0 = 1$ and $\beta_1 = (u-\xi)/\sqrt{\theta}$:

$$\beta_{k+1} = \frac{(u-\xi)\beta_k}{\sqrt{\theta(k+1)}} + \sqrt{\frac{k}{k+1}}\left(\frac{T}{\theta}-1\right)\beta_{k-1}. \tag{A.3}$$

For the isotropic nodal equilibrium, multiply the coordinate moments and then multiply by ρ_i. These coefficients form G_i; the allocation in Eq. (23) gives the particle collision gains. Equation (A.3) therefore supplies the equilibrium coefficients in the basis of Eq. (A.1).

For correlated Gaussian initial data, $\beta_{\boldsymbol{\alpha}}$ is the expectation of $\prod_j \mathrm{He}_{\alpha_j}((v_j-\xi_j)/\sqrt{\theta})/\sqrt{\alpha_j!}$ under that Gaussian. These joint moments, evaluated by multivariate Gaussian integration by parts, include correlations between velocity directions. For isotropic covariance $T\mathsf{I}$, where I is the identity matrix, they reduce to the coordinate products above. Initial coefficients are multiplied by the particle masses from deterministic spatial quadrature.

For the collisional mixture tests, each particle also contains a basis component centred at the initial mixture equilibrium. Its coefficients start at zero; the original component carries the nonzero initial distribution. At node i, group g_p receives the fraction $\sum_{p\in g_p} N_i(\boldsymbol{x}_p)\omega_p / \sum_p N_i(\boldsymbol{x}_p)\omega_p$ of the collective equilibrium gain. These fractions sum to one. The homogeneous relaxation and collisional two-dimensional tests use this additional component. The single-group acoustic and shock tests use their original basis components.

Appendix A.2. Half-range functions for absorbing transport

At an absorbing boundary, incoming and outgoing velocities must be selected separately. For the neutral one-dimensional shock, we therefore replace the full-range Hermite basis by two half-range components:

$$\psi_p(v,t) = \sum_{k=0}^{r} c_{p,+,k}(t)\chi_{+,k}(v) + \sum_{k=0}^{r} c_{p,-,k}(t)\chi_{-,k}(v).$$

The extra sign labels the two coefficient blocks of the same particle. Their basis functions are

$$\chi_{\pm,k}(v) = 2g(v;0,\theta)\,\mathbf{1}_{\{\pm v>0\}}\mathrm{Poly}_{\pm,k}(v), \qquad k = 0,\dots,r, \tag{A.4}$$

where $\mathrm{Poly}_{+,k}$ and $\mathrm{Poly}_{-,k}$ are degree-k polynomials orthonormal on the positive and negative velocity half-axes, respectively, with weight $2g(v;0,\theta)$. Thus

$$\int_{\{\pm v>0\}} 2g(v;0,\theta)\mathrm{Poly}_{\pm,k}(v)\mathrm{Poly}_{\pm,\ell}(v)\,dv = \begin{cases} 1, & k=\ell, \\ 0, & k\neq\ell. \end{cases}$$

Here k, ℓ are polynomial degrees. Each indicator equals one half at zero. Starting from $\mathrm{Poly}_{\pm,0} = 1$, orthogonalization gives a three-term recurrence on each half-axis. The multiplication entries and equilibrium coefficients are

$$(\mathsf{Vel}_1)_{k\ell} = \int_{\{\pm v>0\}} 2g(v;0,\theta)\,v\,\mathrm{Poly}_{\pm,k}(v)\mathrm{Poly}_{\pm,\ell}(v)\,dv,$$
$$\text{coefficient of } \chi_{\pm,k} \text{ in } G_i = \int_{\{\pm v>0\}} G_i(v)\mathrm{Poly}_{\pm,k}(v)\,dv.$$

These integrals replace Eqs. (A.2) and (A.3). The neutral shock needs no electric-acceleration operator. Boundary selection retains the appropriate half-range block (§Appendix B.3).

For numerical construction, the recurrence and integrals above use Gaussian quadrature on $[-12\sqrt{\theta}, 0]$ and $[0, 12\sqrt{\theta}]$, using $\max(200, 8(r+1))$ points on each interval. The omitted tails introduce a representation approximation. The shock uses $\theta = 1$ and $r = 12$, giving 26 coefficients per particle.

Appendix A.3. Removing high-order coefficients

Coefficient truncation reduces storage while retaining all terms through degree two. Before each collisionless step, initially zero degrees are added as needed to accommodate the neighbouring-degree couplings in every time stage. Every five steps, high-order coefficients are checked for removal.

For one component, the squared $1/g$-weighted L^2 norm of the discarded function equals the sum of its squared coefficients. Since $\int g\,d\boldsymbol{v} = 1$, its L^1 norm is also bounded by this coefficient norm. Removal requires every active particle to satisfy its mass-scaled truncation tolerance and sampled lower-bound check. The error allowance is distributed by elapsed time. The coefficient norm estimates the distribution error introduced by each truncation.

Collisional calculations and calculations with nonnegativity corrections or edge viscosity retain fixed degree. Basis centres and widths remain fixed in these tests. The dimensional comparisons use fixed degree and measure the savings from the total-degree basis.

Appendix B. Spatial reconstruction and coefficient operators

Appendix B.1. Particle-to-grid reconstruction

For one velocity-basis group, omit its group index and define $(\mathsf{PtoG})_{ip} = N_i(\boldsymbol{x}_p)$. Rows correspond to nodes and columns to particles. Let row p of Z contain the coefficients of ψ_p. Substitution into Eq. (14) gives $\mathsf{Node} = \mathrm{diag}(\mu_i)^{-1}\mathsf{PtoG}Z$, whose row i contains the coefficients of F_i. Similarly, row i of Load contains the coefficients of R_i. The equilibrium load is $\mu_i G_i$; μ_i and V_p retain their main-text definitions.

Appendix B.2. Relative transport and rate allocation

The relative-transport calculation has three operations. First collect the coefficients of $(v_j - U_{p,j})\psi_p$ at nodes as $\mathsf{PtoG}[Z\mathsf{Vel}_j - \mathrm{diag}(U_{p,j})Z]$. Second divide each nodal row by μ_i to reconstruct the flux. Third insert this reconstructed flux into the weak integral in Eq. (17). Integrating the product of the test-function derivative and a reconstruction basis gives the spatial matrix, whose entries are $(\mathsf{Xder}_j)_{i\ell} = \int_{\mathcal{D}} \partial_{x_j} N_i N_\ell \, d\boldsymbol{x}$. Sum over spatial directions to obtain Load, including the boundary and shock contributions described below. Here Xder_j acts in physical space; Vder_j in Eq. (A.2) acts in velocity space.

Rate allocation uses $M = \mathsf{PtoG}\,\mathrm{diag}(V_p)\mathsf{PtoG}^T$, the matrix of Eq. (21) with its group index suppressed. The minimum-change solution in Eq. (22) applies $\mathrm{diag}(V_p)\mathsf{PtoG}^T M^{-1}$ to Load to give the particle coefficient rates. Collecting these rates back at nodes verifies the constraint:

$$\mathsf{PtoG}\,[\mathrm{diag}(V_p)\mathsf{PtoG}^T M^{-1}\mathsf{Load}] = MM^{-1}\mathsf{Load} = \mathsf{Load}.$$

Under the supported-node rank conditions in the main text, this verifies the constraint. Equilibrium loads use the same allocation. Rates enter the particle stages of Section 5, with geometry updated when particles move.

Stabilization of unresolved particle information. The stabilization in Eq. (32) has coefficient form

$$\lambda[\mathrm{diag}(V_p)\mathsf{PtoG}^T M^{-1}\mathsf{PtoG}Z - Z].$$

Collecting it gives $\lambda[MM^{-1}\mathsf{PtoG}Z - \mathsf{PtoG}Z] = 0$, so it preserves every nodal coefficient.

To specify its value, restore the group index g_p and let ν label a velocity-basis component within that group. Its centre, scale, and fixed reference degree are $\boldsymbol{\xi}_{g_p,\nu}$, $\theta_{g_p,\nu}$, and $r_{g_p,\nu}$. The reference centre $\boldsymbol{\xi}_{g_p}$ is the centre of the group's original component, before adding a collision-gain component. Let d_x be the spatial dimension and Δx the uniform mesh spacing, equal in all directions in these tests. With $\|\cdot\|_1$ denoting the sum of absolute components, the fixed rate is

$$\lambda = \frac{1}{2\Delta x} \max_{g_p,\nu} \left[\|\boldsymbol{\xi}_{g_p,\nu} - \boldsymbol{\xi}_{g_p}\|_1 + 2d_x \sqrt{\theta_{g_p,\nu}(r_{g_p,\nu}+1)} \right]. \tag{B.1}$$

The first term measures the centre offset; the second estimates the represented velocity width. The reference parameters and λ remain fixed during a run. The rate is chosen empirically, with the time steps checked numerically.

Electric acceleration. The one-dimensional electrostatic tests use $q/m = -1$. Fourier inversion of Poisson's equation imposes zero mean potential and field; the real-grid Nyquist derivative is set to zero. Charge deposition follows Eq. (14), and field gathering uses the same P1 functions. Applying Eq. (18) to the expansion in §Appendix A.1 converts its velocity derivative to Vder_1. The acceleration contribution to coefficient row $\boldsymbol{c}_p = (c_{pk})_k$ is $-(q/m)E_h(x_p)\boldsymbol{c}_p\mathsf{Vder}_1$. It is recomputed at every particle time stage.

Appendix B.3. Absorbing boundaries and shock correction

At a one-dimensional absorbing boundary, outgoing velocities satisfy $v < 0$ on the left and $v > 0$ on the right. For one group, select these velocities from $\sum_p (\mathsf{PtoG})_{ip}(v - U_p)\psi_p/\mu_i$ at the boundary node and set the incoming trace

to zero. The signed boundary load enters Eq. (17); half-range blocks perform the selection. Endpoint basis volumes are $\Delta x/2$ and interior volumes are Δx. Intermediate exterior positions use clamped endpoint basis values. The budget in Eq. (33) includes both boundary loads integrated with the time-stage weights and moments carried by removed particles. The remaining moments plus this budget are compared with the initial moments.

For the shock, a slope-limited edge correction damps oscillations near the discontinuity. In one dimension, write the symmetric velocity-multiplication matrix as $\mathsf{Vel}_1 = W\Lambda W^T$, where the columns of W are orthonormal eigenvectors and Λ is the diagonal eigenvalue matrix. In these eigenvector coordinates, velocity multiplication is diagonal, so each coefficient can be limited according to its relative propagation speed. Rotate the nodal coefficient matrix to $\mathsf{Rot} = \mathsf{Node}W$, and denote its row at node i by Rot_i. The limited slope is

$$\mathsf{Slope}_i = \operatorname{minmod}\left(2(\mathsf{Rot}_i - \mathsf{Rot}_{i-1}), \frac{\mathsf{Rot}_{i+1} - \mathsf{Rot}_{i-1}}{2}, 2(\mathsf{Rot}_{i+1} - \mathsf{Rot}_i)\right).$$

The minmod function returns the smallest magnitude when all arguments have the same sign, and zero otherwise; it acts on each coefficient separately. Endpoint slopes are zero. At a fixed edge $(i, i+1)$, let $\mathsf{Speed} = \operatorname{diag}(\Lambda) - (u_i + u_{i+1})/2$ be the vector of relative speeds, where u_i is the mass-weighted particle mean at the node. The limited reconstructions on the two sides of the edge are $\mathsf{Rot}_i + \mathsf{Slope}_i/2$ and $\mathsf{Rot}_{i+1} - \mathsf{Slope}_{i+1}/2$. Their jump gives the first term of the dissipative correction below; the second is the slope-dependent correction used with it. This correction provides numerical stabilization for the shock. Rotating back gives

$$\left[\frac{1}{2}\left(\mathsf{Rot}_{i+1} - \mathsf{Rot}_i - \frac{\mathsf{Slope}_i + \mathsf{Slope}_{i+1}}{2}\right)|\mathsf{Speed}| - \frac{1}{4}(\mathsf{Slope}_i - \mathsf{Slope}_{i+1})\mathsf{Speed}\right] W^T. \tag{B.2}$$

Products with Speed act componentwise. Add the flux to the left nodal load and subtract it from the right one, preserving their sum. This shock-only correction is evaluated from the current particle state.

Appendix B.4. Collision weights and fluid-MPM comparison

The weights a, b, c in Eq. (27) are evaluated by small-argument series near zero and stable exponential differences otherwise. The final equilibrium gain in Eq. (31) is reconstructed once from the provisional particles. Matching its collision-invariant moments preserves those moments in the correction.

Relation to fluid-MPM pressure forces. For one group on a periodic mesh, the LVSP reconstruction and direct MPM particle quadrature give the following nodal momentum loads, labelled by method

$$\begin{aligned} &\text{LVSP:} \quad \sum_\ell \frac{\int_{\mathcal{D}} \nabla_x N_i N_\ell \, d\boldsymbol{x}}{\mu_\ell} \sum_p N_\ell(\boldsymbol{x}_p)\omega_p T_p, \\ &\text{MPM:} \quad \sum_p \omega_p T_p \nabla_x N_i(\boldsymbol{x}_p). \end{aligned} \tag{B.3}$$

LVSP allocates this load by Eq. (22). Lumped-mass MPM divides by nodal mass and interpolates acceleration. Collision allocation can also redistribute mass among LVSP particles while preserving the total.

Appendix C. Conservative nonnegativity correction

A shock-stage candidate may violate the prescribed distribution lower bound. The correction is triggered when its scaled Bernstein constraints fall below the nominal tolerance. We seek a nearby coefficient row satisfying both moment and lower-bound constraints. For smooth tests, the lower bound is checked diagnostically.

Single-particle correction. Fix a particle p with coefficient row $\boldsymbol{c}_p = (c_{pk})_k$. The positive reference weight w_p^0 is the initial density at $\boldsymbol{x}_p^0$ multiplied by its spatial quadrature volume, and remains fixed. It provides a fixed scale for the coefficient correction.

Row k of Mom is $\int (1, v, v^2)\chi_k(v)\,dv$, so multiplication by this matrix returns mass, momentum, and twice the kinetic energy. Here k runs over both half-range blocks. Let Bern map it to the scaled Bernstein coefficients used for the lower-bound check below. For tolerance $\mathrm{tol} > 0$, solve

$$\begin{aligned} \min_{\widetilde{\boldsymbol{c}}_p} \quad & \frac{1}{2}\left\| \frac{\widetilde{\boldsymbol{c}}_p - \boldsymbol{c}_p}{w_p^0} \right\|_2^2, \\ \text{subject to} \quad & (\widetilde{\boldsymbol{c}}_p - \boldsymbol{c}_p)\mathsf{Mom} = 0, \qquad \frac{\widetilde{\boldsymbol{c}}_p}{w_p^0}\mathsf{Bern} \geq -\mathrm{tol}. \end{aligned} \tag{C.1}$$

The equality preserves the three moments; the inequality bounds the scaled distribution from below. The solution $\widetilde{\boldsymbol{c}}_p$ is the corrected particle coefficient row. The equality is enforced in the nullspace of Mom^T, and all constraints are checked before acceptance.

Lower-bound construction and tolerances. On each half-axis, convert the polynomial factor to Bernstein form on intervals of width $0.5\sqrt{\theta}$ up to $12\sqrt{\theta}$. Bernstein basis functions are nonnegative and sum to one, so the polynomial is bounded below by its smallest Bernstein coefficient. Scale each polynomial-to-Bernstein column by the largest Gaussian envelope on that interval to form Bern. This also bounds negative coefficients after multiplication by the envelope, giving an interval-wide distribution bound within the finite checked range.

Use $\mathrm{tol} = 10^{-11}$, allowing scaled coefficients down to $-\max(2\mathrm{tol}, 10^{-10})$. For a single particle, the largest absolute defect in the moments of $\boldsymbol{c}_p/w_p^0$ must be below $10^{-10}\max(1, \|\boldsymbol{c}_p\mathsf{Mom}/w_p^0\|_\infty)$, where $\|\cdot\|_\infty$ denotes the largest absolute component.

Neighbourhood correction. If the single-particle constraints are infeasible or the solve fails, apply a correction that preserves neighbourhood moments. Average $\boldsymbol{c}_p/w_p^0$ with weights w_p^0, first in the current cell and then within one, two, or three cells if needed. Correct this pooled row using Eq. (C.1) with reference weight one. Blend every scaled row towards the corrected pool with the same smallest fraction satisfying the lower-bound checks, then multiply back by w_p^0.

The pool correction preserves its moments. A common blend fraction therefore preserves the neighbourhood totals:

$$\sum_{p \text{ in neighbourhood}} (\widetilde{\boldsymbol{c}}_p - \boldsymbol{c}_p)\mathsf{Mom} = 0.$$

The largest absolute defect in these total moments must be below 10^{-10}. Individual particle moments may change. If no admissible pool is found, the correction reports failure.

References

[1] M. Touati, R. Codur, F. Tsung, V. K. Decyk, W. B. Mori, L. O. Silva, Kinetic theory of particle-in-cell simulation plasma and the ensemble averaging technique, *Plasma Physics and Controlled Fusion* **64** (2022), 115014. https://doi.org/10.1088/1361-6587/ac9016.

[2] A. B. Langdon, Effects of the spatial grid in simulation plasmas, *Journal of Computational Physics* **6** (1970), 247–267. https://doi.org/10.1016/0021-9991(70)90024-0.

[3] S. Jubin, A. T. Powis, W. Villafana, D. Sydorenko, S. Rauf, A. V. Khrabrov, S. Sarwar, I. D. Kaganovich, Numerical thermalization in 2D PIC simulations: Practical estimates for low-temperature plasma simulations, *Physics of Plasmas* **31** (2024), 023902. https://doi.org/10.1063/5.0180421.

[4] S. E. Parker, W. W. Lee, A fully nonlinear characteristic method for gyrokinetic simulation, *Physics of Fluids B* **5** (1993), 77–86. https://doi.org/10.1063/1.860870.

[5] A. Sternlieb, Initial noise and quiet start (Q.S.) in electromagnetic particle simulation codes, *Journal of Computational Physics* **25** (1977), 118–132. https://doi.org/10.1016/0021-9991(77)90016-X.

[6] R. Nguyen van yen, D. del-Castillo-Negrete, K. Schneider, M. Farge, G. Chen, Wavelet-based density estimation for noise reduction in plasma simulations using particles, *Journal of Computational Physics* **229** (2010), 2821–2839. https://doi.org/10.1016/j.jcp.2009.12.010.

[7] M. S. Mitchell, M. T. Miecnikowski, G. Beylkin, S. E. Parker, Efficient Fourier basis particle simulation, *Journal of Computational Physics* **396** (2019), 837–847. https://doi.org/10.1016/j.jcp.2019.07.023.

[8] M. Campos Pinto, E. Sonnendrücker, A. Friedman, D. P. Grote, S. M. Lund, Noiseless Vlasov–Poisson simulations with linearly transformed particles, *Journal of Computational Physics* **275** (2014), 236–256. https://doi.org/10.1016/j.jcp.2014.06.032.

[9] M. B. Quashie, J. W. Burby, A. J. Christlieb, Q. Tang, A Structure-Preserving Decorated Particle Method for the Vlasov–Poisson System, arXiv:2605.21794v1 (2026). https://arxiv.org/abs/2605.21794.

[10] B. Wang, G. Miller, P. Colella, An adaptive, high-order phase-space remapping for the two dimensional Vlasov–Poisson equations, *SIAM Journal on Scientific Computing* **34** (2012), B909–B924. https://doi.org/10.1137/120872954.

[11] K. Xu, J.-C. Huang, A unified gas-kinetic scheme for continuum and rarefied flows, *Journal of Computational Physics* **229** (2010), 7747–7764. https://doi.org/10.1016/j.jcp.2010.06.032.

[12] Z. Guo, K. Xu, R. Wang, Discrete unified gas kinetic scheme for all Knudsen number flows: Low-speed isothermal case, *Physical Review E* **88** (2013), 033305. https://doi.org/10.1103/PhysRevE.88.033305.

[13] Z. Guo, R. Wang, K. Xu, Discrete unified gas kinetic scheme for all Knudsen number flows. II. Thermal compressible case, *Physical Review E* **91** (2015), 033313. https://doi.org/10.1103/PhysRevE.91.033313.

[14] H. Liu, L. Quan, Q. Chen, S. Zhou, Y. Cao, Discrete unified gas kinetic scheme for electrostatic plasma and its comparison with the particle-in-cell method, *Physical Review E* **101** (2020), 043307. https://doi.org/10.1103/PhysRevE.101.043307.

[15] W. Li, C. Liu, Y. Zhu, J. Zhang, K. Xu, A unified gas-kinetic particle method for multiscale photon transport, arXiv:1810.05984 (2018). https://doi.org/10.48550/arXiv.1810.05984.

[16] C. Liu, Y. Zhu, K. Xu, Unified gas-kinetic wave-particle methods I: Continuum and rarefied gas flow, *Journal of Computational Physics* **401** (2020), 108977. https://doi.org/10.1016/j.jcp.2019.108977.

[17] C. Liu, K. Xu, Unified gas-kinetic wave-particle methods IV: Multi-species gas mixture and plasma transport, *Advances in Aerodynamics* **3** (2021), 9. https://doi.org/10.1186/s42774-021-00062-1.

[18] S. Chen, K. Xu, C. Lee, Q. Cai, A unified gas kinetic scheme with moving mesh and velocity space adaptation, *Journal of Computational Physics* **231** (2012), 6643–6664. https://doi.org/10.1016/j.jcp.2012.05.019.

[19] D. Sulsky, Z. Chen, H. L. Schreyer, A particle method for history-dependent materials, *Computer Methods in Applied Mechanics and Engineering* **118** (1994), 179–196.

[20] P. L. Bhatnagar, E. P. Gross, M. Krook, A model for collision processes in gases. I. Small amplitude processes in charged and neutral one-component systems, *Physical Review* **94** (1954), 511–525. https://doi.org/10.1103/PhysRev.94.511.

[21] Z. Guo, J. Li, K. Xu, Unified preserving properties of kinetic schemes, *Physical Review E* **107** (2023), 025301. https://doi.org/10.1103/PhysRevE.107.025301.