Kinetic implicit fluid methods

E. Alec Johnson, Stefano Markidis, and Giovanni Lapenta

Nov 26, 2013

Abstract:

Fully explicit discretizations must resolve all three processes that ideal MHD assumes are instantaneous: oscillations, collisions, and light waves. Asymptotic-preserving discretization requires stepping over these processes. Fully implicit discretization allows stepping not only over these three processes but also over MHD waves, but is expensive because one must repeatedly re-push particles with successive iterations of the implicit field solver. We therefore present semi-implicit discretizations of Maxwell's equations that step over subsets of these processes. Discretizing implicitly in the source term steps over plasma oscillations and collisions and allows asymptotic-preserving agreement with (quasi-)relativistic MHD. Discretizing the flux terms of Maxwell's equations implicitly steps over light waves and allows asymptotic-preserving agreement with classical two-fluid MHD.

To facilitate asymptotic-preserving agreement with fluid models and conservation of physical invariants, we consider kinetic fluid closure.

particle evolution:

$$d_t\mathbf{x}_{\mathrm{p}} = \mathbf{v}_{\mathrm{p}},$$
$$d_t\mathbf{u}_{\mathrm{p}} = \frac{q_{\mathrm{p}}}{m_{\mathrm{p}}}\left(\mathbf{v}_{\mathrm{p}} \times \mathbf{B}(\mathbf{x}_{\mathrm{p}}) + \mathbf{E}(\mathbf{x}_{\mathrm{p}})\right),$$
$$\gamma_{\mathrm{p}}^2 := 1 + (\mathbf{u}_{\mathrm{p}}/c)^2,$$
$$\mathbf{v}_{\mathrm{p}} := \mathbf{u}_{\mathrm{p}}/\gamma_{\mathrm{p}}.$$

electromagnetic field:

$$\partial_t\mathbf{B} + \nabla \times \mathbf{E} = 0,$$
$$-c^{-2}\partial_t\mathbf{E} + \nabla \times \mathbf{B} = \mu_0\mathbf{J},$$
$$\nabla \cdot \mathbf{B} = 0, \quad c^{-2}\nabla \cdot \mathbf{E} = \mu_0\sigma,$$

where $\sigma(\mathbf{x}) := \sum_p S_p(\mathbf{x})q_p$ is charge density and $\mathbf{J}(\mathbf{x}) := \sum_p S_p(\mathbf{x})q_p\mathbf{v}_p$ is current density; here $S_{\mathrm{p}}(\mathbf{x}) = S(\mathbf{x} - \mathbf{x}_{\mathrm{p}})$ is the shape function of particle p. To abbreviate, we drop the particle summation index p and the independent variable $\mathbf{x}$ and write

$$\sigma := \sum qS \qquad \text{(charge)},$$
$$\rho := \sum mS \qquad \text{(mass)},$$
$$\mathbf{J} := \sum \mathbf{v}qS \qquad \text{(current)},$$
$$\mathbf{M} := \sum \mathbf{v}mS \qquad \text{(momentum)},$$
$$\mathcal{E} := \sum \tfrac{1}{2}|v|^2 mS \qquad \text{(energy)}.$$

Moment evolution

Fields evolve in response to charge moments, and mass, momentum, and energy are conserved, motivating the use of fluid models. Kinetic closure allows kinetic algorithms to transition efficiently and smoothly to fluid models.

Moment definitions are of the form $\sum \chi S$. To derive fluid equations, we differentiate the moment definition with respect to time and use the basic derivatives to the right.

Charge density evolution.

$$\boxed{\partial_t \sigma + \nabla \cdot \mathbf{J} = 0}\,,$$

using $\partial_t \sigma = \sum \dot{q} S + \sum q \partial_t S$,
$\partial_t S = -\mathbf{v} \cdot \nabla S$, and $\nabla \mathbf{v} = 0$.

General moment evolution

$$\boxed{\begin{array}{l} \partial_t \sum \chi S + \nabla \cdot \sum \mathbf{v} \chi S = \sum \dot{\chi} S, \\ \dot{\chi} = \dfrac{\partial \chi}{\partial \mathbf{u}} \cdot \dot{\mathbf{u}} = \dfrac{\partial \chi}{\partial \mathbf{v}} \cdot \dot{\mathbf{v}}. \end{array}}$$

Shape motion:

$$\boxed{\partial_t S = -\mathbf{v} \cdot \nabla S}\,, \tag{1}$$

where we have used the chain rule.

Lorentz force.

$$\boxed{\dot{\mathbf{u}} = \frac{q}{m} \left(\mathbf{E} + \mathbf{v} \times \mathbf{B}\right);}$$

Energy change.

$$\boxed{\dot{\gamma} = \mathbf{v} \cdot \dot{\mathbf{u}} = \frac{q}{m} \mathbf{v} \cdot \mathbf{E}}\,, \tag{2}$$

because $\gamma^2 = 1 + (\mathbf{u}/c)^2$, so $\gamma \dot{\gamma} = \mathbf{u} \cdot \dot{\mathbf{u}}/c^2$.

Velocity change.

$$\boxed{\dot{\mathbf{v}} = \frac{q}{\gamma m} \left(\mathbf{E} - \frac{\mathbf{v}\mathbf{v}}{c^2} \cdot \mathbf{E} + \mathbf{v} \times \mathbf{B}\right),}$$

which follows from differentiating $\mathbf{u} = \gamma \mathbf{v}$, to get $\dot{\mathbf{u}} = \dot{\gamma}\mathbf{v} + \gamma\dot{\mathbf{v}}$, i.e. $\gamma\dot{\mathbf{v}} = \dot{\mathbf{u}} - \mathbf{v}\mathbf{v} \cdot \dot{\mathbf{u}}$.

Current evolution (Ohm's law, $\chi = q\mathbf{v}$).

$$\partial_t \mathbf{J} + \nabla \cdot \mathcal{P} = \sum S \frac{q^2}{\gamma m} \left(\mathbb{I} - \frac{\mathbf{v}\mathbf{v}}{c^2} \right) \cdot \mathbf{E} + \sum S \frac{q^2}{\gamma m} \times \mathbf{B},$$

where $\mathcal{P} := \sum q\mathbf{v}\mathbf{v}$.

Classical current evolution for species s.

$$\partial_t \mathbf{J}_s + \nabla \cdot \mathcal{P}_s = \frac{q_s}{m_s} (\sigma_s \mathbf{E} + \mathbf{J}_s \times \mathbf{B}),$$

where $\mathcal{P}_s$ restricts to species s.

Relativistic case.

Mass density ($\chi = m$).

$$\partial_t \rho + \nabla \cdot \sum m\mathbf{v} S = 0.$$

Momentum density ($\chi = m\mathbf{u}$).

$$\partial_t \mathbf{M} + \nabla \cdot \sum m\mathbf{v}\mathbf{u} S = \sigma \mathbf{E} + \mathbf{J} \times \mathbf{B}$$

Energy density ($\chi = mc^2\gamma$).

$$\partial_t \mathcal{E} + \nabla \cdot \mathbf{M} = \mathbf{J} \cdot \mathbf{E}$$

Classical case.

Mass density ($\chi = m$)

$$\partial_t \rho + \nabla \cdot \mathbf{M} = 0,$$

Momentum density ($\chi = m\mathbf{v}$)

$$\partial_t \mathbf{M} + \nabla \cdot \sum m\mathbf{v}\mathbf{v} S = \sigma \mathbf{E} + \mathbf{J} \times \mathbf{B},$$

Energy density ($\chi = \frac{1}{2} m|\mathbf{v}|^2$)

$$\partial_t \mathcal{E} + \nabla \cdot \sum \frac{1}{2} m\mathbf{v} v^2 S = \mathbf{J} \cdot \mathbf{E},$$

Field discretizations

The Implicit Moment Method is an example of a semi-implicit method.

Semi-implicit methods are used to step over fast processes such as plasma oscillations (implicit source) and light waves (IMM) without having to step over all processes.

For plasma simulations, one of four types of discretization is generally used:

Discretization	must resolve. . .	must iterate. . .
1. Explicit	everything (e.g. ω_p)	[no iteration]
2. Implicit Source	light waves	classical: no iteration
		relativistic: source
3. Implicit Moment (IMM)	electron sound waves	fields
4. Fully implicit	[no restriction]	particles [everything]

The Implicit Source discretization naturally suits an asymptotic-preserving transition to relativistic MHD, since ideal relativistic MHD takes the electron charge (gyrofrequency) to infinity. The Implicit Moment Method naturally suits an asympto preserving transition to two-fluid MHD, which takes light speed to infinity.

Explicit discretization

Start with the basic equations:

$$\partial_t \mathbf{B} + \nabla \times \mathbf{E} = 0,$$
$$\partial_t \mathbf{E} - c^2 \nabla \times \mathbf{B} = -\mathbf{J}/\epsilon_0,$$
$$\partial_t \mathbf{x}_p = \mathbf{v}_p,$$
$$\partial_t \mathbf{u}_p = \frac{q_p}{m_p} \left(\mathbf{E}(\mathbf{x}_p) + \mathbf{v}_p \times \mathbf{B}(\mathbf{x}_p) \right).$$

For a second-order discretization, use a leapfrog discretization:

- **B** does not need particle velocities to advance, but **E** does, so particle velocities $\mathbf{u}_p$ and current **J** should be staggered relative to **E**.
- $\partial_t X = (X^2 - X^0)/\Delta t$
- $\partial_t Y = (Y^3 - Y^1)/\Delta t$

In the relativistic case, time-split the velocity update for a symplectic method. In full detail:

$$\mathbf{B}^2 = \mathbf{B}^0 - \Delta t \nabla \times \mathbf{E}^1,$$
$$\mathbf{E}^1 = \mathbf{E}^{-1} + \Delta t \left(c^2 \nabla \times \mathbf{B}^0 - \mathbf{J}^0/\epsilon_0 \right),$$
$$\mathbf{u}_p^* = \mathbf{u}_p^0 + \frac{q_p \Delta t}{2 m_p \gamma_p^0}(\mathbf{u}_p^* + \mathbf{u}_p^0) \times \bar{\mathbf{B}}^1(\mathbf{x}_p^1),$$
$$\mathbf{u}_p^2 = \mathbf{u}_p^* + \Delta t \frac{q_p}{m_p} \mathbf{E}^1(\mathbf{x}_p^1),$$

where $\bar{\mathbf{B}}^1 := \frac{1}{2}\mathbf{B}^0 + \frac{1}{2}\mathbf{B}^2$.

Use initial values for flux terms.
and implicit values in stiff source.

Fields:

$$\partial_t \mathbf{B} + \nabla \times \mathbf{E} = 0,$$
$$\partial_t \mathbf{E} - c^2 \nabla \times \mathbf{B} = -\bar{\mathbf{J}}/\epsilon_0.$$

Particles:

$$\partial_t \mathbf{u}_p = \frac{q_p}{m_p} \left(\bar{\mathbf{E}}(\mathbf{x}_p) + \bar{\mathbf{v}}_p \times \bar{\mathbf{B}}(\mathbf{x}_p) \right),$$
$$\partial_t \mathbf{x}_p = \bar{\mathbf{v}}_p,$$
$$\partial_t \sigma_s + \nabla \cdot \mathbf{J}_s = 0.$$

Cassical current:

$$\partial_t \mathbf{J}_s + \nabla \cdot \mathcal{P}_s = \frac{q_s}{m_s} \left(\bar{\sigma}_s \bar{\mathbf{E}} + \bar{\mathbf{J}}_s \times \bar{\mathbf{B}} \right).$$

We designate $n = 0$ as initial time,
$n + 1 = 1$ as final time, and time
discretization as

$$\partial_t Q \to (Q^1 - Q_0)/\Delta t,$$
$$\nabla F \to F^0,$$
$$\bar{\mathbf{X}} = X^1.$$

Classical case: No source term iteration
happens to be needed, because $\bar{\mathbf{v}}$ is linear
in $\bar{\mathbf{E}}$. Can sum the response over all
particles to eliminate $\bar{\mathbf{J}} = \hat{\mathbf{J}} + \mathbb{A} \cdot \bar{\mathbf{E}}$ in favor
of $\bar{\mathbf{E}}$.

Relativistic case: Must iterate particle
velocity advance, but positions need not be
advanced, so iterative solve involves no
communication between mesh cells.

High-order accuracy: Use an IMEX
Runge-Kutta solver.

Modify the IMM discretization by making all terms implicit:

$$\partial_t \mathbf{B} + \nabla \times \bar{\mathbf{E}} = 0$$

$$\partial_t \mathbf{E} - c^2 \nabla \times \bar{\mathbf{B}} = -\bar{\mathbf{J}}/\epsilon_0$$

$$\partial_t \mathbf{u}_p = \frac{q_p}{m_p} \left(\bar{\sigma}_p \bar{\mathbf{E}}(\bar{\mathbf{x}}_p) + \bar{\mathbf{v}}_p \times \bar{\mathbf{B}}(\bar{\mathbf{x}}_p) \right),$$

$$\partial_t \mathbf{x}_p = \bar{\mathbf{v}}_p,$$

$$\partial_t \sigma_s + \nabla \cdot \bar{\mathbf{J}}_s = 0,$$

$$\partial_t \mathbf{J}_s + \nabla \cdot \bar{\overline{\mathcal{P}}}_s = \frac{q_s}{m_s} \left(\bar{\sigma}_s \bar{\mathbf{E}} + \bar{\mathbf{J}}_s \times \bar{\mathbf{B}} \right).$$

Remarks.

- Particle advance must be redone with successive iterations of the field solver.
- This discretization is fully symmetric in time, so by Noether's theorem conserves energy if iterated to convergence.

To linearize the Maxwell source term, modify the fully implicit discretization by replacing the updated charge densities and magnetic field with explicitly evolved values.

$$\partial_t \mathbf{B} + \nabla \times \bar{\mathbf{E}} = 0$$

$$\partial_t \mathbf{E} - c^2 \nabla \times \bar{\mathbf{B}} = -\bar{\mathbf{J}}/\epsilon_0,$$

$$\partial_t \mathbf{v}_p = \frac{q_p}{m_p}\left(\bar{\mathbf{E}}(\bar{\mathbf{x}}_p') + \bar{\mathbf{v}}_p \times \bar{\mathbf{B}}'(\bar{\mathbf{x}}_p')\right),$$

$$\partial_t \mathbf{x}_p = \bar{\mathbf{v}}_p,$$

$$\partial_t \sigma_s + \nabla \cdot \bar{\mathbf{J}}_s = 0,$$

$$\partial_t \mathbf{J}_s + \nabla \cdot \mathcal{P}_s = \frac{q_s}{m_s}\left(\sigma_s' \bar{\mathbf{E}} + \bar{\mathbf{J}}_s \times \mathbf{B}'\right)$$

For second-order accuracy in relevant terms, use time averages for the implicit terms:

$$\partial_t Q \to (Q^1 - Q^0)/\Delta t,$$

$$\nabla F \to F^0, \quad \bar{\mathbf{X}} = \tfrac{1}{2}X^0 + \tfrac{1}{2}X^1,$$

$$\bar{\mathbf{v}} = \tfrac{1}{2}\mathbf{v}^0 + \tfrac{1}{2}\mathbf{v}^1,$$

$$\bar{\mathbf{J}} = \tfrac{1}{2}\mathbf{J}^0 + \tfrac{1}{2}\mathbf{J}^1,$$

$$\bar{\mathbf{E}} = \tfrac{1}{2}\mathbf{E}^0 + \tfrac{1}{2}\mathbf{E}^1.$$

Observe that in this discretization,

$$\partial_t X = 2(\bar{\mathbf{X}} - X^0)/\Delta t.$$

Divergence constraints. Eliminating $\mathbf{B}$ from this field discretization and assuming mimetic operators, the field solve is exactly equivalent to

$$c^{-2}\partial_t \mathbf{E} = (\tfrac{1}{2}\Delta t)\left(\nabla^2 \bar{\mathbf{E}} - \nabla(\nabla \cdot \bar{\mathbf{E}})\right) + \left(\nabla \times \mathbf{B}^0 - \mu_0 \bar{\mathbf{J}}\right);$$

following [RicciLapentaBrackbill02], to ensure the divergence constraint error is damped, substitute $(\nabla \cdot \bar{\mathbf{E}}) \to (\mu_0 c^2 \bar{\sigma})$, where $\bar{\sigma} := \sigma^0 - \tfrac{1}{2}\Delta t \bar{\mathbf{J}}$.

High-order accuracy. Use implicit Euler discretization and plug scheme into Runge-Kutta method.

Classical case. Can eliminate $\bar{\mathbf{J}}$ in favor of $\bar{\mathbf{E}}$ by putting current evolution in the form $\bar{\mathbf{J}} = \hat{\mathbf{J}} + \mathbb{A} \cdot \bar{\mathbf{E}}$. Defining $\mathbf{B}' = \mathbf{B}^0 - \frac{\Delta t}{2}\nabla \times \mathbf{E}^0$, $\bar{\mathbf{B}}' = \bar{\mathbf{B}}$, $\sigma_s' = \sigma_s^0 - \frac{\Delta t}{2}\nabla \cdot \mathbf{J}_s^0$ would yield full second-order accuracy in time. IMM in literature uses $\mathbf{B}' = \bar{\mathbf{B}}' = \mathbf{B}^0$ and $\sigma_s' = \sigma_s^0$, Particle advance is solved iteratively, initialized with $\bar{\mathbf{x}}_p' \leftarrow \mathbf{x}_p^0$; two iterations is enough for second-order accuracy.

Relativistic case. Implicit source seems preferable to IMM for the fully relativistic case for multiple reasons:

- Why step over light waves but not over relativistic sound waves or fluid speed? An implicit source time step is much cheaper and involves no long-distance communication.

- The source term system responds nonlinearly to $\bar{\mathbf{E}}$ and is not closed, so an implicit particle velocity advance must be repeated with successive iterations of the field solve.

An IMEX-RK (Implicit-Explicit Runge-Kutta) scheme is designed to solve an ODE with two forcing terms:

$$\partial_t Q(t) = \mathcal{F}_{\mathrm{E}}(Q) + \mathcal{F}_{\mathrm{I}}(Q),$$

where $\mathcal{F}_{\mathrm{E}}$ is discretized explicitly and $\mathcal{F}_{\mathrm{I}}$ is discretized implicitly. A first-order IMEX update would be

$$Q^1 = Q^0 + \Delta t \, \mathcal{F}_{\mathrm{E}}(Q^0) + \Delta t \, \mathcal{F}_{\mathrm{I}}(Q^1).$$

Usually this is written as an explicit Euler update combined with an implicit Euler update:

$$Q^* = Q^0 + \Delta t \, \mathcal{F}_{\mathrm{E}}(Q^0),$$

$$Q^1 = Q^* + \Delta t \, \mathcal{F}_{\mathrm{I}}(Q^1).$$

DIRK schemes are built from compositions and linear combinations of these two updates.

Plasma equations fit the general form

$$\partial_t q + \mathcal{F}(q) = S(q),$$

where $\mathcal{F}$ represent the flux term. The source-implicit discretization is of the form

$$Q^1 = Q^0 - \Delta t \, \mathcal{F}(Q^0) + \Delta t \, S(Q^1),$$

which clearly fits the DIRK framework.

Can we make the Implicit Moment Method fit the DIRK framework?

High-order-in-time IMM via DIRK

The plasma equations are of the form

$$\partial_t Q_c + \mathcal{F}_c(Q_s) = 0,$$
$$\partial_t Q_s + \mathcal{F}_\mathbf{J}(Q_s) + \mathcal{F}_s(Q_c) = S_s(Q_c, Q),$$

where the "conserved" quantities are $Q_c = (\mathbf{B}, \sigma, \mathbf{x}_p)$ and the quantities forced by the oscillatory source term are $Q_s = (\mathbf{E}, \mathbf{J}, \mathbf{v}_p)$. The implicit moment method discretizes this as

$$Q_c^1 = Q_c^0 - \Delta t \mathcal{F}_c(Q_s^1),$$
$$Q_s^1 = Q_s^0 - \Delta t \mathcal{F}_\mathbf{J}(Q_s^0)$$
$$\qquad - \Delta t \mathcal{F}_s(Q_c^1) + \Delta t S_s(Q_c', Q^1),$$

where choosing $Q_c' = Q_c^0$ avoids a nonlinear system; unfortunately, this destroys the DIRK framework, which assumes that every term is fully implicit or fully explicit.

A trick that restores the DIRK framework is to duplicate the conserved variables in the source term that cannot be fully implicit in the source. So we solve the system

$$\partial_t \tilde{Q}_c + \mathcal{F}_c(Q_s) = 0,$$
$$\partial_t Q_c + \mathcal{F}_c(Q_s) = 0,$$
$$\partial_t Q_s + \mathcal{F}_\mathbf{J}(Q_s) + \mathcal{F}_s(Q_c) = S_s(\tilde{Q}_c, Q),$$

using the discretization

$$\tilde{Q}_c^1 = \tilde{Q}_c^0 - \Delta t \mathcal{F}_c(Q_s^0),$$
$$Q_c^1 = Q_c^0 - \Delta t \mathcal{F}_c(Q_s^1),$$
$$Q_s^1 = Q_s^0 - \Delta t \mathcal{F}_s(Q_c^1)$$
$$\qquad - \Delta t \mathcal{F}_\mathbf{J}(Q_s^0) + \Delta t S_s(\tilde{Q}_c^1, Q^1);$$

this is a DIRK Euler update, in which every term is fully implicit or fully explicit.

Outline

Calculating particle and current advance in response to the electromagnetic field requires solving equations of the form

$$\mathbf{U} = \mathbf{V} + \mathbf{U} \times \boldsymbol{\Omega} \qquad (3)$$

To solve for $\mathbf{U}$, dot and cross both sides with $\boldsymbol{\Omega}$ to get the equations:

$$\mathbf{U} \cdot \boldsymbol{\Omega} = \mathbf{V} \cdot \boldsymbol{\Omega},$$

$$\mathbf{U} \times \boldsymbol{\Omega} = \mathbf{V} \times \boldsymbol{\Omega} + \boldsymbol{\Omega}\boldsymbol{\Omega} \cdot \mathbf{U} - |\boldsymbol{\Omega}|^2 \mathbf{U}$$

$$= \mathbf{V} \times \boldsymbol{\Omega} + \boldsymbol{\Omega}\boldsymbol{\Omega} \cdot \mathbf{V} - |\boldsymbol{\Omega}|^2 \mathbf{U}.$$

So eliminating $\mathbf{U} \times \boldsymbol{\Omega}$ in (3),

$$\mathbf{U}(1 + |\boldsymbol{\Omega}|^2) = (\mathbb{I} - \boldsymbol{\Omega} \times \mathbb{I} + \boldsymbol{\Omega}\boldsymbol{\Omega}) \cdot \mathbf{V}.$$

That is,

$$\boxed{\mathbf{U} = \frac{\mathbb{I} - \boldsymbol{\Omega} \times \mathbb{I} + \boldsymbol{\Omega}\boldsymbol{\Omega}}{1 + |\boldsymbol{\Omega}|^2} \cdot \mathbf{V}}.$$

Equation (3) says that

$$(\mathbb{I} + \boldsymbol{\Omega} \times \mathbb{I}) \cdot \mathbf{U} = \mathbf{V}$$

We thus infer that

$$\boxed{(\mathbb{I} + \boldsymbol{\Omega} \times \mathbb{I})^{-1} = \frac{\mathbb{I} - \boldsymbol{\Omega} \times \mathbb{I} + \boldsymbol{\Omega}\boldsymbol{\Omega}}{1 + |\boldsymbol{\Omega}|^2}}.$$

Recall classical current evolution:

$$\partial_t \mathbf{J}_s + \nabla \cdot \mathcal{P}_s = \frac{q_s}{m_s}\left(\sigma_s'\bar{\mathbf{E}} + \bar{\mathbf{J}}_s \times \mathbf{B}'\right).$$

Discretize $\partial_t \mathbf{J}_s$ as $\frac{\mathbf{J}_s^1 - \mathbf{J}_s^0}{\Delta t} = \frac{\bar{\mathbf{J}}_s - \mathbf{J}_s^0}{\Delta t/2}$. So

$$\bar{\mathbf{J}}_s = \mathbf{J}_s^0 - \frac{\Delta t}{2}\nabla \cdot \mathcal{P}_s + \beta_s(\sigma_s'\bar{\mathbf{E}} + \bar{\mathbf{J}}_s \times \mathbf{B}'),$$

where $\beta_s := \frac{q_s \Delta t}{2 m_s}$. This is of the form (3),

$$\mathbf{U} = \mathbf{V} + \mathbf{U} \times \boldsymbol{\Omega},$$

where

$$\mathbf{U} = \bar{\mathbf{J}}_s,$$
$$\mathbf{V} = \mathbf{J}_s^0 - \frac{\Delta t}{2}\nabla \cdot \mathcal{P}_s + \beta_s\sigma_s'\bar{\mathbf{E}},$$
$$\boldsymbol{\Omega} = \beta_s \mathbf{B}.$$

Thus, the linear response of average current to average electric field is given by:

$$
\begin{aligned}
\bar{\mathbf{J}} &= \widehat{\mathbf{J}} + \mathbb{A} \cdot \bar{\mathbf{E}}, \\
\mathbb{A} &:= \sum_s \beta_s \sigma_s' \Pi_s, \\
\widehat{\mathbf{J}} &:= \sum_s \widehat{\mathbf{J}}_s, \\
\widehat{\mathbf{J}}_s &:= \Pi_s \cdot \left(\mathbf{J}_s^0 - \frac{\Delta t}{2}\nabla \cdot \mathcal{P}_s\right), \\
\Pi_s &:= \frac{\mathbb{I} - \boldsymbol{\Omega}_s \times \mathbb{I} + \boldsymbol{\Omega}_s \boldsymbol{\Omega}_s}{1 + |\boldsymbol{\Omega}_s|^2}, \\
\boldsymbol{\Omega}_s &:= \beta_s \mathbf{B}', \\
\beta_s &:= \frac{q_s \Delta t}{2 m_s}, \\
\sigma_s^0 &:= \sum_{p \in s} S(\mathbf{x} - \mathbf{x}_p^0) q_p, \\
\mathbf{J}_s^0 &:= \sum_{p \in s} S(\mathbf{x} - \mathbf{x}_p^0) q_p \mathbf{v}_p^0, \\
\mathcal{P}_s &:= \sum_{p \in s} S(\mathbf{x} - \mathbf{x}_p^0) q_p \mathbf{v}_p^0 \mathbf{v}_p^0.
\end{aligned}
$$

IMM implicit field solver

The implicit moment method differences Maxwell's evolution equations implicitly as:

$$\nabla \times \bar{\mathbf{E}} + \frac{\bar{\mathbf{B}} - \mathbf{B}^0}{\theta \Delta t} = 0,$$

$$\nabla \times \bar{\mathbf{B}} - \frac{\bar{\mathbf{E}} - \mathbf{E}^0}{c^2 \theta \Delta t} = \mu_0 \bar{\mathbf{J}};$$

where $\bar{\mathbf{X}} := \theta \mathbf{X}^1 + (1 - \theta) \mathbf{X}^0$, so $\theta(\mathbf{X}^1 - \mathbf{X}^0) = \bar{\mathbf{X}} - \mathbf{X}^0$. To eliminate $\bar{\mathbf{B}}$ and get an equation implicit in $\bar{\mathbf{E}}$, take the curl of the first equation:

$$(c\theta \Delta t)^2 \nabla \times \nabla \times \bar{\mathbf{E}} + \bar{\mathbf{E}}$$
$$= \mathbf{E}^0 + c^2 \theta \Delta t (\nabla \times \mathbf{B}^0 - \mu_0 \bar{\mathbf{J}}).$$

The implicit moment method assumes that average current responds linearly to average electric field:

$$\bar{\mathbf{J}} = \hat{\mathbf{J}} + \frac{\boldsymbol{\chi}}{\mu_0 c^2 \theta \Delta t} \cdot \bar{\mathbf{E}},$$

where the "**implicit susceptibility**" tensor $\boldsymbol{\chi}$ is defined so as to be unitless.

Substituting for $\bar{\mathbf{J}}$ yields the field equation used in [KamimuraMBLT92]:

$$(c\theta \Delta t)^2 \nabla \times \nabla \times \bar{\mathbf{E}} + (\bar{\mathbf{E}} + \boldsymbol{\chi} \cdot \bar{\mathbf{E}}) \tag{4}$$
$$= \mathbf{E}^0 + c^2 \theta \Delta t (\nabla \times \mathbf{B}^0 - \mu_0 \hat{\mathbf{J}}).$$

Including the approximate identities

$$\nabla \times \nabla \times \bar{\mathbf{E}} = -\nabla^2 \bar{\mathbf{E}} + \nabla \nabla \cdot \bar{\mathbf{E}},$$

$$\nabla \cdot \bar{\mathbf{E}} = \mu_0 c^2 \bar{\sigma}, \tag{5}$$

where $\bar{\sigma} := \sigma^0 - \theta \Delta t \nabla \cdot \bar{\mathbf{J}}$

gives a numerically overdetermined system; we modify the method to damp diverge error by invoking the discrete divergence constraint (5). With these assumptions,

$$\nabla \cdot \bar{\mathbf{E}} = \mu_0 c^2 \hat{\sigma} - \nabla \cdot (\boldsymbol{\chi} \cdot \bar{\mathbf{E}}),$$

where $\boxed{\hat{\sigma} := \sigma^0 - \theta \Delta t \nabla \cdot \hat{\mathbf{J}}}$.

Substituting for $\nabla \times \nabla \times \bar{\mathbf{E}}$ in equation (4),

$$\boxed{\begin{array}{l}(\bar{\mathbf{E}} + \boldsymbol{\chi} \cdot \bar{\mathbf{E}}) - (c\theta \Delta t)^2 (\nabla^2 \bar{\mathbf{E}} + \nabla \nabla \cdot (\boldsymbol{\chi} \cdot \bar{\mathbf{E}})) \\ = \mathbf{E}^0 + c^2 \theta \Delta t (\nabla \times \mathbf{B}^0 - \mu_0 (\hat{\mathbf{J}} + c^2 \nabla \hat{\sigma})).\end{array}}$$

Divergence error is damped by this equation, whereas it neither grows nor decays for (4).

Having found $\bar{\mathbf{E}}$,

$$\mathbf{B}^1 = \mathbf{B}^0 + \Delta t \nabla \times \bar{\mathbf{E}},$$
$$\mathbf{E}^1 = \theta^{-1} \bar{\mathbf{E}} + (1 - \theta^{-1}) \mathbf{E}^0.$$

Observe that this magnetic field update maintains the condition that the magnetic field is a discrete curl up to machine precision independent of whether mimetic operators are used.

Accurate closure.

This scheme is second-order accurate for $\theta = 1/2$ and first-order accurate for $\frac{1}{2} < \theta \leq 1$; it is unstable for $\theta < \frac{1}{2}$.

Needed information:

- $\boldsymbol{\chi} := (\mu_0 c^2 \Delta t) \, \mathbb{A}$
- $\hat{\mathbf{J}}, \sigma^0$
- $\mathbf{B}^0, \mathbf{E}^0$

Particle position and velocity are differenced as

$$\mathbf{x}_{\mathrm{p}}^1 = \mathbf{x}_{\mathrm{p}}^0 + \bar{\mathbf{v}}_{\mathrm{p}} \Delta t,$$
$$\bar{\mathbf{v}}_{\mathrm{p}} := \tfrac{1}{2}\mathbf{v}_{\mathrm{p}}^1 + \tfrac{1}{2}\mathbf{v}_{\mathrm{p}}^0,$$
$$\mathbf{v}_{\mathrm{p}}^1 = \mathbf{v}_{\mathrm{p}}^0 + 2\beta_{\mathrm{s}}\left(\mathbf{E}_{\mathrm{p}}^\theta + \bar{\mathbf{v}}_{\mathrm{p}} \times \mathbf{B}_{\mathrm{p}}^\vartheta \right),$$

where s is the species of particle p, ϑ might equal θ or 0, and $\beta_{\mathrm{s}} := \frac{q_{\mathrm{s}}\Delta t}{2m_{\mathrm{s}}}$.

Choosing $\mathbf{B}_{\mathrm{p}}^\vartheta := \mathbf{B}^\vartheta(\mathbf{x}_{\mathrm{p}}^0)$ yields an explicit particle advance. Choosing $\mathbf{B}_{\mathrm{p}}^\theta := \mathbf{B}^\vartheta(\bar{\mathbf{x}}_{\mathrm{p}})$ and $\mathbf{E}_{\mathrm{p}}^\theta := \mathbf{E}^\theta(\bar{\mathbf{x}}_{\mathrm{p}})$, where $\bar{\mathbf{x}}_{\mathrm{p}} := \tfrac{1}{2}\mathbf{x}_{\mathrm{p}}^1 + \tfrac{1}{2}\mathbf{x}_{\mathrm{p}}^0$, defines an implicit particle advance. Use two iterations beginning with the explicit advance for second-order accuracy.

Eliminating $\mathbf{v}_{\mathrm{p}}^1$ in favor of $\bar{\mathbf{v}}_{\mathrm{p}}$,

$$\bar{\mathbf{v}}_{\mathrm{p}} = \mathbf{v}_{\mathrm{p}}^0 + \beta_{\mathrm{s}}\left(\mathbf{E}_{\mathrm{p}}^\theta + \bar{\mathbf{v}}_{\mathrm{p}} \times \mathbf{B}_{\mathrm{p}}^\vartheta \right).$$

This is of the form

$$\mathbf{U} = \mathbf{V} + \mathbf{U} \times \mathbf{\Omega},$$

where

$$\mathbf{U} = \bar{\mathbf{v}}_{\mathrm{p}},$$
$$\mathbf{V} = \mathbf{v}_{\mathrm{p}}^0 + \beta_{\mathrm{s}}\mathbf{E}_{\mathrm{p}}^\theta,$$
$$\mathbf{\Omega} = \beta_{\mathrm{s}}\mathbf{B}_{\mathrm{p}}^\vartheta.$$

Thus,

$$\boxed{\begin{aligned}
\bar{\mathbf{v}}_{\mathrm{p}} &= \Pi_{\mathrm{p}}^\vartheta \cdot (\mathbf{v}_{\mathrm{p}}^0 + \beta_{\mathrm{s}}\mathbf{E}_{\mathrm{p}}^\theta), \\
\Pi_{\mathrm{p}}^\vartheta &:= \frac{\mathbb{I} - \mathbf{\Omega}_{\mathrm{p}} \times \mathbb{I} + \mathbf{\Omega}_{\mathrm{p}}\mathbf{\Omega}_{\mathrm{p}}}{1 + |\mathbf{\Omega}_{\mathrm{p}}|^2}, \\
\mathbf{\Omega}_{\mathrm{p}} &:= \beta_{\mathrm{s}}\mathbf{B}_{\mathrm{p}}^\vartheta, \\
\beta_{\mathrm{s}} &:= \frac{q_{\mathrm{p}}\Delta t}{2m_{\mathrm{p}}}.
\end{aligned}}$$

Observe that $|\mathbf{\Omega}_{\mathrm{p}}|$ is half the gyrofrequency for particle p. Note that $\vartheta \in \{0, \theta\}$ is chosen based on whether the updated magnetic field is already known. A first-order-accurate field predictor allows for a fully second-order-accurate solve with $\vartheta = \theta = \frac{1}{2}$.

Advance the position and velocity of each particle p via

$$\mathbf{u}^1 = \mathbf{u}^0 + 2\beta_s\left(\mathbf{E}_p^\theta + \bar{\mathbf{v}} \times \mathbf{B}_p^\vartheta\right),$$
(6)

$$\bar{\mathbf{v}} = \bar{\mathbf{u}}/\bar{\gamma},$$

$$\mathbf{x}_p^1 = \mathbf{x}_p^0 + \bar{\mathbf{v}}\Delta t,$$

where $\bar{\mathbf{u}} := \frac{1}{2}\mathbf{u}^1 + \frac{1}{2}\mathbf{u}^0$, $\mathbf{u} = \mathbf{u}_p$ (etc.), $\bar{\gamma} := \frac{1}{2}\gamma^1 + \frac{1}{2}\gamma^0$, and $\beta_s := \frac{q_s\Delta t}{2m_s}$; choosing $\mathbf{B}_p^\vartheta := \mathbf{B}^\vartheta(\mathbf{x}_p^0)$ yields an explicit advance, and choosing $\mathbf{B}_p^\vartheta := \mathbf{B}^\vartheta(\bar{\mathbf{x}}_p)$ and $\mathbf{E}_p^\theta := \mathbf{E}^\theta(\bar{\mathbf{x}}_p)$ defines an implicit advance, where

$$\bar{\mathbf{x}}_p := \frac{1}{2}\mathbf{x}_p^1 + \frac{1}{2}\mathbf{x}_p^0.$$

Use 2 iterations beginning with the explicit advance for second-order accuracy. Since $\mathbf{u}^1 = 2\bar{\mathbf{u}} - \mathbf{u}^0$,

$$\bar{\mathbf{u}} = \mathbf{u}^0 + \beta_s(\mathbf{E}_p^\theta + \bar{\mathbf{v}} \times \bar{\gamma}^{-1}\mathbf{B}_p^\vartheta).$$

This is of the form

$$\mathbf{U} = \mathbf{V} + \mathbf{U} \times \mathbf{\Omega}, \qquad (7)$$

where $\mathbf{U} = \bar{\mathbf{u}}$, $\mathbf{V} = \mathbf{u}^0 + \beta_s\mathbf{E}_p^\theta$, and, $\mathbf{\Omega} = \bar{\gamma}^{-1}\beta_s\mathbf{B}_p^\vartheta$ is half the gyrofrequency vector. Thus (restoring index p),

$$\boxed{\begin{aligned}\bar{\mathbf{u}}_p &= \Pi_p^\vartheta \cdot (\mathbf{u}_p^0 + \beta_s\mathbf{E}_p^\theta), \\ \Pi_p^\vartheta &:= \frac{\mathbb{I} - \mathbf{\Omega}_p \times \mathbb{I} + \mathbf{\Omega}_p\mathbf{\Omega}_p}{1 + |\mathbf{\Omega}_p|^2}, \\ \mathbf{\Omega}_p &:= \frac{\beta_p}{\bar{\gamma}_p}\mathbf{B}_p^\vartheta, \\ \beta_s &:= \frac{q_s\Delta t}{2m_s},\end{aligned}}$$

This system would be an explicit solution if $\bar{\mathbf{x}}_p$ and $\bar{\gamma}_p$ were known. In the classical case, $\bar{\gamma}_p = 1$. In the relativistic case, $\gamma^2 = 1 + (\mathbf{u}/c)^2$, so $\bar{\gamma}d\gamma = \bar{\mathbf{u}} \cdot d\mathbf{u}/c^2$, i.e., $d\gamma = \bar{\mathbf{v}} \cdot d\mathbf{u}/c^2$.

Substituting into (6),

$$\boxed{\bar{\gamma}_p = \gamma_p^0 + \beta_p\mathbf{E}_p^\theta \cdot \bar{\mathbf{v}}'/c^2,}$$

where this equality holds if $\bar{\mathbf{v}}' = \bar{\mathbf{v}}_p$. The initialization $\bar{\mathbf{v}}' \leftrightarrow \mathbf{v}_p^0$ yields an iterative solver. The relativistic implicit moment method is based on this first iteration.

Note that $\vartheta \in \{0, \theta\}$ is chosen based on whether the updated magnetic field is already known. A first-order-accurate field predictor allows for a fully second-order-accurate solve with $\vartheta = \theta = \frac{1}{2}$.

Divergence constraints: consistency of fields with fluid moments

A predictor-corrector strategy can be used to maintain the divergence constraints to machine precision while maintaining order of accuracy for arbitrarily high order:

1. Evolve the fields with sufficient accuracy.
2. Evolve face-normal field fluxes using the evolved fields.
3. Enforce consistency of the evolved fields with the evolved face-normal field fluxes.

Enforcing $\nabla \cdot \mathbf{B} = 0$:

Integrating magnetic field evolution over a time step and over a mesh cell face A gives

$$\int_A \widehat{\mathbf{n}} \cdot \mathbf{B}^1 = \int_A \widehat{\mathbf{n}} \cdot \mathbf{B}^0 - \Delta t \left(\oint_{\partial A} \widehat{\tau} \cdot \bar{\mathbf{E}} \right),$$

where $\bar{\mathbf{E}}$ is the time-averaged value of the the electric field on the edges of the face. A predictor step can be used to supply high-order-accurate values of the average value of the parallel component of $\mathbf{E}$ along each edge, giving a high-order-accurate update of the total magnetic flux perpendicular to each face. For any order of accuracy, there is at least enough freedom to enforce that the magnetic field representation satisfies these constraints, and when extra freedom is available one can do so so as to minimize a norm.

Enforcing $\nabla \cdot \mathbf{E} = \sigma/\epsilon_0$:

Integrating electric field evolution over a time step and over a mesh cell face A gives

$$\int_A \widehat{\mathbf{n}} \cdot \mathbf{E}^1 = \int_A \widehat{\mathbf{n}} \cdot \mathbf{E}^0 + c^2 \Delta t \left(\oint_{\partial A} \widehat{\tau} \cdot \bar{\mathbf{B}} \right) - \Delta t \int_A \widehat{\mathbf{n}} \cdot \bar{\mathbf{J}}/\epsilon_0;$$

to maintain consistency, define $\bar{\mathbf{J}}$ using the flux of charge across face A, update the face-normal electric field flux accordingly, and modify the electric field representation to enforce consistency.

Conservation framework

To enforce exact conservation of momentum and energy, use Maxwell's equations to put evolution in conservation form.

Conservation of momentum.

To put momentum equation in conservation form, rewrite the source term as the sum of a time derivative and a spatial derivative:

$$-(\sigma\mathbf{E} + \mathbf{J}\times\mathbf{B})/\epsilon_0 = \partial_t(\mathbf{E}\times\mathbf{B}) + \tfrac{1}{2}\nabla(E^2 + c^2 B^2) - \nabla\cdot(\mathbf{EE} + c^2\mathbf{BB}),$$

where we have used all four Maxwell's equations and a vector product rule.

Conservation of energy.

To put energy evolution in conservation form, rewrite the source term as

$$-\mathbf{J}\cdot\mathbf{E}/\epsilon_0 = \tfrac{1}{2}\partial_t(E^2 + c^2 B^2) + \nabla\cdot(\mathbf{E}\times\mathbf{B}),$$

where we have used Maxwell's evolution equations and the identity
$\mathbf{E}\cdot\nabla\times\mathbf{B} = \mathbf{B}\cdot\nabla\times\mathbf{E} - \nabla\cdot(\mathbf{E}\times\mathbf{B})$.

Positivity.

Conservation form allows an explicit fluid code to enforce positivity using the framework of *Outflow Positivity Limiting*, as described in [JoRo12]. Here we assume that the thermal energy, defined as the gas-dynamic energy in the reference frame in which the gas-dynamic momentum is zero, is a concave function of the conserved state variables ($\rho, \widetilde{\mathbf{M}}, \widetilde{\mathcal{E}}, \mathbf{B}, \mathbf{E}$), where $\widetilde{\mathcal{E}} := \mathcal{E} + \tfrac{1}{2}(E^2 + c^2 B^2)$ is defined to be the total energy and $\widetilde{\mathbf{M}} := \mathbf{M} + \mathbf{E}\times\mathbf{B}$ is defined to be the total momentum.

Particle representation

To use kinetic closure, we need (1) a means of representing the particle distribution and (2) a means to approximate closing moments at cell boundaries.

Blobs. Traditional PIC approximates the distribution of particles with computational particles that are rigid blobs of charge. This restricts agreement of PIC with the Vlasov equation to low-order accuracy in phase space. In a limit involving an increasing number of computational particles per mesh cell, it would still be possible to get higher-order accuracy (of fluid moments) in physical space; this would need integration of flux values along faces. Traditionally, one gives up on higher-order accuracy, simply computing the average value of closing moments in each cell and then interpolating.

Samples. In higher-order PIC (HOPIC), each particle carries a value of phase-space density that is convected with the particle. High-order-accurate evolution can be achieved by interpolating these values at quadrature points.

Samples at random points. The approach to HOPIC taken by [EdwardsBridson11] is to create particles at randomly chosen points and sample the representation with a least-squares polynomial fit. Such a fit requires solving a linear system (10 by 10 in the classical case) at each sampled point.

Samples at fixed points. An alternative approach to HOPIC advocated here is to resample the particles with every time step at predetermined points so that no resampling is needed to sum closing moments. If particles live at predetermined points, then the matrix of the linear system that must be solved to interpolate distribution values can be inverted once for all at the beginning of the simulation and stored. Similarly, the contribution of each particle to the closing moments is its charge or mass times a quantity that can be precomputed at startup. Thus, computation of each closing moment is simply a weighted sum over all particles, and interpolation simply multiplies particle data by a precomputed weighting matrix.

To ensure that particles live at predetermined points when resampling particles, one can resample with each time step by predetermining the locations of the resampled particles at the updated time step and tracking them backwards in time to determine sampling locations. This needs an inital estimate of the updated electromagnetic fields. Alternating iteratively between field and particle updates converges to a self-consistent fully implicit update.

In the relativistic case, the value of γ^{-1} needs to be computed once for each particle. Fortunately, accelerators such as the MIC and GPU implement the square root and multiplicative inverse operations very efficiently.

Positivity of distribution.

To avoid instability, it is desirable for moments to be calculated with a positive distribution. We can guarantee this in HOPIC if sampled values are always positive. An easy way to ensure this is for each particle to carry the *logarithm* of the distribution. Calculating moments requires exponentiating each sampled distribution value. Fortunately, exponentiation is cheap on the MIC and the GPU.

Consistency with mass moments.

Evolved field moments must not drift without limit from the moments of the distribution with successive updates. To enforce consistency, one can calculate values of the evolved moments from the particles and then modify the particle data to enforce consistency.

Mass consistency.

Assume that the evolved mass ρ is positive. To make the mass density of the particles consistent with the evolved mass density, simply rescale the phase space density of the particles appropriately at each point. If the logarithm of particle density is being represented, then at each point in physical space one simply shifts the logarithmic particle values by a constant.

Momentum consistency.

To enforce momentum consistency, one can apply an appropriate constant shift to all particle velocities at each point in physical space. One simply reinterprets particle velocities when sampling particle values by applying this shift.

In the relativistic case, determining the appropriate shift to enforce exact consistency would require a nonlinear (iterative) solve, but an approximate solve using a single iteration could be used to damp inconsistency and prevent its growth. An alternative is to simply *define* particle velocities relative to the evolved fluid velocity, i.e., to represent thermal velocities.

Energy consistency.

Assume that the evolved energy density $\mathcal{E}$ is positive. To enforce energy consistency, one can at each point in physical space apply an appropriate constant rescaling to the thermal velocity. In the classical case, this rescaling factor is simply the square root of the ratio of evolved energy to calculated energy. In the relativistic case, determining the appropriate scaling to enforce exact consistency would require a nonlinear (iterative) solve, but an approximate solve using a single iteration could be used to damp inconsistency and prevent its growth.

Consistency of particle distribution with charge (and mass) moments

In addition to enforcing consistency of the evolved electric field with the evolved charge density, one also needs to enforce consistency of the evolved charge density with the evolved particle distribution.

To enforce agreement both with mass density and with charge density, one can rescale the positive charges with a factor α^+ and the negative charges with a different factor α^-. Indeed, the linear system

$$\begin{bmatrix} \rho^+ & \rho^- \\ \sigma^+ & \sigma^- \end{bmatrix} \cdot \begin{pmatrix} \alpha^+ \\ \alpha^- \end{pmatrix} = \begin{pmatrix} \rho \\ \sigma \end{pmatrix}$$

has positive determinant if charges of both species are present in all mesh cells; ensuring this condition involves some type of particle resampling (of which particle splitting is one example).

If particles are (e.g. rigid) blobs (as opposed to distribution samples), then one can alternatively track the flow of charge across cell boundaries to determine exact charge flux.
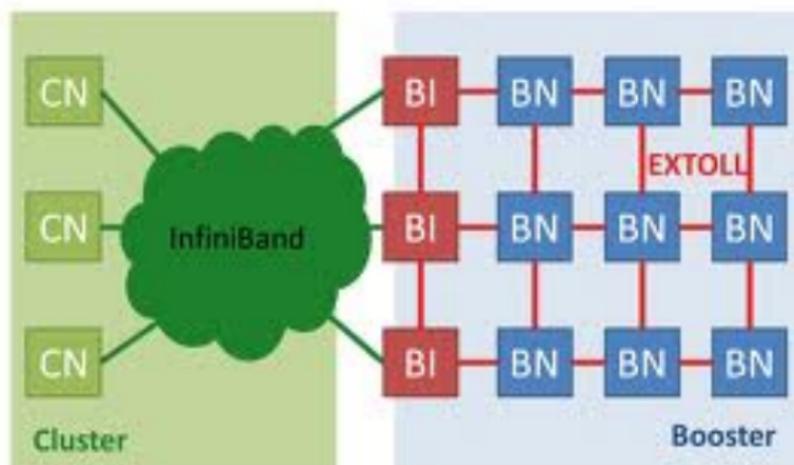
The architecture of DEEP consists of a 128-node, 21.3 teraflop **Cluster** of 16-core 2.7 GHz Intel Xeon processors connected to an infiniband network and will be augmented with a a 512-node, 590 teraflop **Booster** each of whose nodes is a 1.2 GHz 60-core Xeon Phi MIC (Many Integrated Core) accelerator.

The DEEP architecture was developed with the idea of accelerating codes that run on the cluster by making it easy to offload computationally intensive parts of the code to the booster while leaving complex, communication-intensive code on the cluster.

FieldSolver on Cluster, ParticleSolver on Booster

iPic3D implements the implicit moment method and consists of a cycle of three steps:

1. (ParticleSolver): Sum implicit moments $\widehat{\mathbf{J}}^0$ and σ_s^0 of particles.
   Communicate implicit moments to FieldSolver.

2. (FieldSolver): Advance the electromagnetic field, $\mathbf{E}^0 \to \mathbf{E}^1$, using implicit moments.
   Communicate fields to ParticleSolver.

3. (ParticleSolver): Advance the particles using the already advanced fields.

Due to its implicit nature, the FieldSolver makes many all-to-all communications and naturally is best suited to the Cluster.

In contrast, since particles can be pushed in parallel, the ParticleSolver only needs to communicate particles to neighboring processors in the torus.

To increase resolution and demonstrate converged solutions, the number of particles per mesh cell must increase. Therefore, in the high resolution limit, computing particles increasingly dominates PIC codes. The port of iPic3D to DEEP aims for unprecedented resolution of turbulent reconnection on long time scales in the classical regime.

A beneficial side-effect of the DEEP port is the clean separation of the particle solver and the field solver. This allows replacing one or the other as appropriate:

- The FieldSolver could be replaced with an IMEX field solver or could be iterated to yield a fully implicit discretization.
- The ParticleSolver could be replaced with a fluid model or a gyrokinetic or relativistic mover.

References

The following approach to HOPIC represents the phase space distribution function with particles at random positions.

[EdwardsBridson11] Essex Edwards and Robert Bridson, *A high-order accurate particle-in-cell method.* International Journal for Numerical Methods in Engineering 2012; 90:1073-1088
http://www.doi.org/10.1002/nme.3356

The following paper describes a general framework to enforce positivity in a fluid model.

[JoRo12] E.A. Johnson and J.A. Rossmanith, *Outflow positivity limiting: Part 1, framework and recipe.*
http://arxiv.org/abs/1212.4695

[KamimuraMBLT92] T. Kamimura, E. Montalvo, D.C. Barnew, J.N. Leboeuf, and T. Tajima. *Implicit Particle Simulation of Electromagnetic Plasma Phenomena.* Journal of Computational Physics 100, 77-90 (1992). http://dx.doi.org/10.1016/0021-9991(92)90311-L.

[KumarMishra11] Harish Kumar and Siddartha Mishra, *Entropy stable numerical schemes for two-fluid plasma equations,* J. Sci. Comput. (2012), http://dx.doi.org/10.1007/s10915-011-9554-7.

The following describes solving implicitly the momentum evolution equations for each species with kinetic closure in conjunction with Maxwell's equations.

[MarkidisHLRHML13] Stefano Markidis, Pierre Henri, Giovanni Lapenta, Kjell Rönnmark, Maria Hamrin, Zakaria Meliani, Erwin Laure. *The Fluid-Kinetic Particle-in-Cell Solver for Plasma Simulations.* Submitted to Journal of Computational Physics (2013).
http://arxiv.org/abs/1306.1089.

[NoguchiTroZucLap07] Koichi Noguchi, Cesare Tronci, Gianluca Zuccaro, and Giovanni Lapenta. *Formulation of the relativistic moment implicit particle-in-cell method.* Physics of Plasmas 14, 042308 (2007).
http://dx.doi.org/10.1063/1.2721083.

The following article derives the implicit Maxwell scheme.

[RicciLapentaBrackbill02] Paolo Ricci, Giovanni Lapenta, and J.U. Brackbill. *A simplified implicit Maxwell solver.* Journal of Computational Physics **183**, 117–141 (2002).

The following article derives an approximation for **J** using a Taylor expansion.

[VuBrackbill92] H.X. Vu and J.U. Brackbill. *CELEST1D: an implicit, fully kinetic model for low-frequency, electromagnetic plasma simulation.* Computer Physics Communications 69 (1992) 253–276.