

FLEKS: A Flexible Particle-in-Cell code for Multi-Scale Plasma Simulations

Yuxi Chen^{1a,b}, Gábor Tóth^a, Hongyang Zhou^a, Xiantong Wang^a

^a *Center for Space Environment Modeling, University of Michigan, Ann Arbor, Michigan 48109, USA*

^b *Department of Astrophysical Sciences, Princeton University, Princeton, New Jersey 08540, USA*

Abstract

The magnetohydrodynamics with embedded particle-in-cell (MHD-EPIC) model has been successfully applied to global magnetospheric simulations in recent years. However, the PIC region was restricted to be one or more static boxes, which is not always sufficient to cover the whole physical structure of interest efficiently. The FLEKS (Flexible Exascale Kinetic Simulator), which is a new PIC code and allows a dynamic PIC region of any shape, is designed to break this restriction. FLEKS is usually used as the PIC component of the MHD with adaptively embedded particle-in-cell (MHD-AEPIC) model. FLEKS supports dynamically activating or deactivating cells to fit the regions of interest during a simulation. An adaptive time-stepping scheme is also introduced to improve the accuracy and efficiency of a long simulation. The particle number per cell may increase or decrease significantly and lead to load imbalance and large statistical noise in the cells with fewer particles. A particle splitting scheme and a particle merging algorithm are designed to limit the change of the particle number and hence improve the accuracy of the simulation as well as load balancing. Both particle splitting and particle merging conserve the total mass, momentum, and energy. FLEKS also contains a test-particle module to enable tracking particle trajectories due to the time-dependent electromagnetic field that is obtained from a global simulation.

Keywords: particle-in-cell; particle merging; test particle; global kinetic simulation

¹Corresponding author. Email address: yuxichen@umich.edu

1. Introduction

Multi-scale plasma simulations are challenging due to the limitation of computational resources. Fluid models are efficient for global simulations, but kinetic-scale physics is missing. Fully kinetic codes, such as particle-in-cell (PIC) codes and Vlasov solvers, contain electron and ion scale physics. However, it is extremely computationally expensive to resolve the global scale and the electron scale at the same time for three-dimensional (3D) global simulations. Traditional hybrid models, which usually treat electrons as a massless fluid and simulate ions with a PIC method or a Vlasov solver, incorporate ion-scale physics into global simulations by sacrificing electron-scale kinetic physics. Another class of hybrid methods embeds a kinetic code into a global fluid model so that the kinetic code can resolve the regions where the kinetic physics is important, and the fluid model handles the rest of the domain efficiently. In recent years, independent groups have developed models that couple either a PIC code [1] or a Vlasov solver [2] with a fluid model.

Sugiyama and Kusano [3] demonstrated the concept of coupling a PIC code with a fluid code. The magnetohydrodynamics (MHD) with embedded particle-in-cell (MHD-EPIC) model developed by Daldorff et al. [1] is the first mature coupled model that is capable of running 3D large-scale simulations. The MHD-EPIC model usually covers the dayside or/and the tail magnetic reconnection sites with the PIC code when it is applied to simulate the dynamics of magnetospheres [4, 5, 6, 7]. Multiple isolated PIC domains are supported so that a few regions of interest can be covered by the PIC code in one simulation [4]. However, in an MHD-EPIC simulation, each PIC region is restricted to be a static box, which is not always efficient or suitable to cover the whole physical structure of interest due to either the limitation of computational resources or geometric complexity of the physical region. Recently, Shou et al. [8] developed the magnetohydrodynamics with adaptively embedded particle-in-cell (MHD-AEPIC) model, which allows changing the location of an active PIC region dynamically.

In this paper, we introduce a new code, the FLExible Exascale Kinetic Simulator (FLEKS), which is designed and implemented to be the PIC component of the MHD-AEPIC model. FLEKS shares some similarities with the Adaptive Mesh Particle Simulator (AMPS) used in the work by Shou et

36 al. [8], but FLEKS provides a more flexible grid design. FLEKS uses the
37 parallel data structures provided by the AMReX library [9, 10]. The grid of
38 FLEKS has to be uniform and Cartesian, but the active PIC region is not
39 limited to be a box anymore since the PIC cells can be turned off to fit the
40 region of interest. Furthermore, FLEKS also supports switching on and off
41 grid cells dynamically for MHD-AEPIC simulations.

42 FLEKS employs the Gauss’s law satisfying energy-conserving semi-implicit
43 method (GL-ECSIM) [11] as the base PIC solver. The time step of the
44 semi-implicit PIC methods is limited by the Courant–Friedrichs–Lewy (CFL)
45 condition based on the macro-particle velocities in order to be accurate [12].
46 Since the the speed of macro-particles may change significantly during a long
47 MHD-AEPIC simulation, the simulation will be either too slow or inaccur-
48 rate with a fixed time step. To keep the simulation efficient and accurate
49 at the same time, FLEKS uses an adaptive time-stepping algorithm, which
50 still satisfies the requirement of the energy-conserving semi-implicit method
51 (ECSIM) [13] to keep energy conservation. Section 2 describes the adaptive
52 grid and temporal discretization of FLEKS.

53 The statistical noise of macro-particles is the primary source of numeri-
54 cal error in typical PIC simulations. Dozens to hundreds of particles per cell
55 are usually used to achieve a balance between accuracy and computational
56 cost. Since there are much more macro-particles than grid cells in a ki-
57 netic PIC simulation, particle-related calculations, such as updating particle
58 positions and velocities, dominate the total computational time. In addi-
59 tion, a massively parallel simulation can be significantly slowed down due to
60 the imbalance of macro-particle numbers among the parallel processes. On
61 the other hand, the decrease of the number of macro-particles in some cells
62 increases the statistical noise and reduces the accuracy. A particle resam-
63 pling algorithm that is able to control the macro-particle number per cell
64 is crucial for improving both the simulation efficiency and accuracy. More
65 macro-particles need to be added into the cells that contain fewer macro-
66 particles than required to represent the plasma velocity-space distribution
67 accurately. This goal is usually achieved by splitting particles. In the cells
68 with more macro-particles than some threshold, a particle merging algorithm
69 needs to be applied to reduce the number of macro-particles and speed up the
70 simulation. A particle resampling algorithm is even more crucial for a PIC
71 code with adaptive mesh refinement, where the motion of macro-particles
72 between the coarse and fine cells alters the macro-particle number per cell
73 dramatically [14, 15].

74 Both the particle splitting and particle merging processes replace the
75 original particles with a set of new particles. Lapenta [16] suggested that the
76 replacement should maintain the following properties:

- 77 1. The plasma moments evaluated on the simulation grid, which are used
78 to update electric and magnetic fields, should not be changed by the
79 replacement.
- 80 2. The replacement should keep the original particle phase space distri-
81 butions.

82 It is more challenging to achieve these two goals for a particle merging al-
83 gorithm than for a particle splitting algorithm, because it is inevitable to
84 lose information when replacing original particles with fewer particles. A few
85 algorithms have been designed to merge particles. Lapenta [16] introduced
86 two algorithms to merge particles that are close to each other in the phase
87 space. The algorithm C1 merges two particles into one, and the algorithm
88 C2 merges three particles into two. The algorithm C2 conserves the mass,
89 momentum, and energy of the particles, and also the charge densities on
90 the grid, but it is not straightforward to extend to 2D and 3D. Vranic et
91 al. [17] also proposed an algorithm to merge particles into two new particles
92 while conserving the overall mass, momentum, and energy, and the original
93 particles are chosen by binning particles in the momentum space. Instead
94 of merging a few particles into one or two, the algorithms designed by As-
95 sous et al. [18], Welch et al. [19], Pfeiffer et al. [20], and Faghihi et al.
96 [21] use a set of particles to replace the old ones. Assous et al. [18] and
97 Welch et al. [19] focused on the conservation of the grid quantities, but the
98 fine structures in the velocity space may not be well preserved. Pfeiffer et
99 al. [20] generated the new particle velocities from a distribution function
100 and adjusted the velocities to conserve energy afterward. Faghihi et al. [21]
101 created new particles with a uniform distribution inside a phase space bin,
102 and adjusted the weights to conserve the moments. As a general rule, the
103 particles selected for merging should be close to each other in the phase space
104 to minimize the error that is introduced by merging. Besides the method of
105 binning the velocity space [17, 21], Teunissen and Ebert [22] applied a k-d
106 tree to find the particles that are closest to each other, and Luu et al. [23]
107 showed how to partition particles with the Voronoi diagram. Our new parti-
108 cle merging algorithm implemented into FLEKS searches for 6 particles that
109 are close in phase space and merges them into 5 particles while preserving
110 mass, momentum and energy and also minimizes the change in the phase

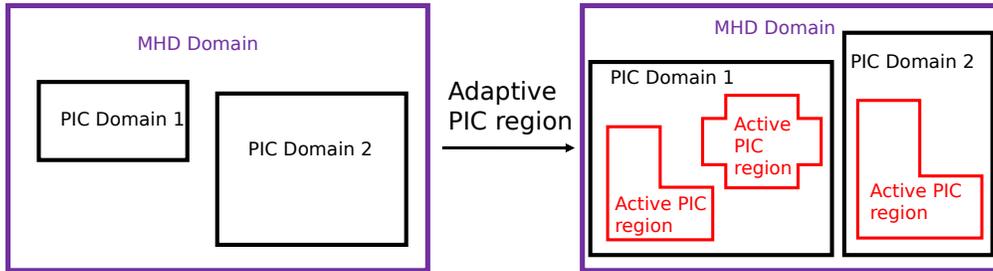


Figure 1: A schematic shows the improvement of the MHD-AEPIC (right) model from the MHD-EPIC (left) model.

111 space distribution. The details of the splitting and merging algorithms are
 112 described in section 4.

113 Tracking the motion of macro-particles is useful for investigating the par-
 114 ticle trajectories and the energization of particles. FLEKS provides a paral-
 115 lel test particle module to follow the motion of macro-particles and save the
 116 particle trajectory data to disk. The test particle module can be used either
 117 inside the PIC code, or as an independent component directly coupled to
 118 the MHD model. Section 5 describes the implementation details of the test
 119 particle module.

120 The paper is organized as follows. Section 2 describes the grid design of
 121 FLEKS. Section 3 introduces the adaptive time-stepping scheme. Section 4
 122 focuses on the particle splitting and particle merging algorithms. Section 5
 123 discusses the implementation of the test particle module. Section 6 presents
 124 numerical tests to demonstrate the capability of the adaptive active PIC
 125 regions, the role of the particle resampling algorithms, the parallel efficiency
 126 of FLEKS, and examples of global simulations with FLEKS. Finally, section
 127 7 presents the conclusions.

128 2. Adaptive grid

129 Since the MHD-EPIC model was developed by Daldorff et al. [1], we
 130 have developed new features to make it more flexible to use. It now supports
 131 multiple independent PIC domains to cover several regions of interest [4], and
 132 it also allows rotating a PIC box domain to align with the features of interest
 133 [24]. However, a box is not always suitable or efficient to cover the region
 134 of interest. For example, if a PIC box is used to cover the whole dayside

135 magnetopause, which is close to a paraboloid, the box will cut through the
136 planet and introduce extra difficulties, and the PIC box will also contain
137 a large portion of cells, where the kinetic effects are not important, which
138 slows down the simulation. A flexible grid that allows creating an active PIC
139 domain that approximates the shape of a paraboloid to fit the magnetopause
140 can solve this problems. A dynamically adaptive grid is also useful to improve
141 the efficiency of some simulations. For instance, the near-Earth X-line may
142 move from the inner magnetotail to the middle or even far magnetotail [25],
143 and an adaptive grid that only covers the environment around the X-line is
144 much more efficient than a large PIC box that covers the whole magnetotail.
145 The MHD-AEPIC algorithm is designed to solve these problems and FLEKS
146 is the key component. Figure 1 shows the conceptual difference between the
147 MHD-EPIC and MHD-AEPIC models.

148 FLEKS still requires the shape of the full PIC grid to be a box, and the
149 Cartesian grid has to be uniform (this is a requirement of the GL-ECSIM
150 algorithm). But FLEKS allows switching off part of the cells to approx-
151 imately fit a region of any shape. The most straightforward approach is
152 switching on/off each cell independently. However, this approach has several
153 drawbacks, as discussed below, so we make the algorithm a bit more sophis-
154 ticated. We divide the whole PIC domain into patches (Figure 2(a)). Each
155 patch contains N cells in each direction, and one can turn on or turn off
156 each patch. The patch size N is required to be larger or equal to 2. We do
157 not allow $N = 1$ (switching on/off each cell independently) for the follow-
158 ing reasons. FLEKS requires two ghost cell layers for coupling with MHD
159 at the PIC region boundary. If $N = 1$, the boundary ghost cells of an ac-
160 tive region may overlap with the physical cells of another active region, and
161 hence introduces more difficulties to handle the boundary ghost cells. A large
162 patch size also benefits the coupling efficiency. In MHD-AEPIC simulations,
163 the fluid model controls the status of the patches based on geometric and
164 physics-based criteria [25]. The fluid model passes the bit-wise patch status
165 array to FLEKS through the Message Passing Interface (MPI), and the size
166 of this array is reduced significantly with a larger patch size (proportional to
167 N^{-3} in 3D). In this paper, we use the word ‘active’ to describe the patches
168 or cells that are switched on. The active cells do not have to be connected,
169 and the boundary ghost cells of the active regions are filled in with the in-
170 formation obtained from the fluid model [1]. Figure 2(a) shows an example
171 that contains two separated active regions.

172 FLEKS uses the data structures provided by the AMReX library to store

173 the fields and also the particles. After the patch status array is obtained from
174 the fluid model, FLEKS uses the functions provided by the AMReX library
175 to divide the active regions into blocks. AMReX does not require all the
176 blocks to have the same size. We note that the patch and the block are two
177 independent concepts. The patches are only used to activate or deactivate
178 cells. For example, the 'L' shape active region in Figure 2(a) consists of 3
179 patches and it can be divided into 2 blocks (Figure 2(b)).

180 FLEKS allows activating or deactivating patches during a simulation. If
181 the active regions change, FLEKS will produce a new set of blocks to cover the
182 new active regions. With the function provided by AMReX, FLEKS copies
183 the fields and particles from the old blocks to the new ones for the cells that
184 are already active and deletes the information of the newly deactivated cells.
185 The newly activated cells are filled in with the information obtained from
186 the fluid model as what is done for FLEKS initialization.

187 FLEKS has two ghost cell layers, but the outer layer is only used to receive
188 and store the magnetic fields, which are necessary for calculating currents on
189 the nodes of the inner ghost cell layer from $\vec{J} = \nabla \times \vec{B}$ in normalized units.
190 The currents are used to generate particles with correct velocities in the inner
191 layer ghost cells. To simplify the description, we ignore the outer layer in
192 Figure 2(c) and also in the rest of the paper unless otherwise specified. The
193 principle of setting boundary conditions of the electromagnetic fields and the
194 particles is still the same as the MHD-EPIC coupling algorithm [1]. However,
195 the non-box shape of an active region introduces some extra implementation
196 difficulties. There are three types of ghost cells for a block: the internal ghost
197 cells (blue cells in Figure 2(c)), the exclusive boundary ghost cells (gray cells
198 in Figure 2(c)) and the shared boundary ghost cells (cyan cells in Figure 2(c)).
199 The internal ghost cells are not boundary cells, and there is no need to apply
200 boundary conditions. The exclusive boundary ghost cells are not overlapped
201 with any cells of the neighboring blocks, and they should be filled in with new
202 macro-particles as the particle boundary condition. The shared boundary
203 ghost cells are overlapped with the boundary ghost cells of the neighboring
204 blocks, and only one of these blocks should generate boundary particles.
205 Here is the algorithm to choose the block for populating new particles. The
206 first step is to distinguish the boundary ghost cells from the internal ghost
207 cells. Then, for each boundary ghost cell, either the exclusive type or the
208 shared type, we loop through its at most 26 neighboring cells (3D) in a fixed
209 order (we choose to loop through all the face neighbors first, then the edge
210 neighbors, and finally the corner neighbors), skip the nonexistent cells and

211 find out the first neighboring cell that is either a physical cell or an internal
 212 ghost cell. If this neighboring cell is inside the physical domain of this block,
 213 this block should generate particles inside this boundary ghost cell. For
 214 example, in Figure 2(c), C1 and C3 are overlapped with each other. We
 215 loop through the neighboring cells of C1 and find C2 is its first neighboring
 216 cell that is either a physical cell or an internal ghost cell (C2 is a physical
 217 cell), so block-1 should generate particles in C1 since C2 is inside block-1. We
 218 repeat the same procedure for the cell C3, and find C4 is its first neighboring
 219 cell that is either a physical cell or an internal ghost cell (C3 is an internal
 220 ghost cell), but block-2 should not generate particles in C3 since C4 is outside
 221 the physical domain of block-2.

222 The electric fields are node-based in FLEKS. For a node that is shared
 223 by multiple blocks, such as the one indicated by a red-cross in Figure 2(c),
 224 only one block should take care of the shared node when solving the linear
 225 equations of the electric fields. The aforementioned algorithm is also applied
 226 to choose the proper block for a shared node.

227 3. Adaptive time-stepping

228 The time step of the energy-conserving semi-implicit method (ECSIM)
 229 is subject to the accuracy condition $v_{rms}\Delta t/\Delta x < 1$ just as other semi-
 230 implicit PIC methods [12], where v_{rms} is the maximum root mean square of
 231 macro-particle velocities. For a long MHD-AEPIC simulation, v_{rms} may vary
 232 significantly, so an adaptive time-stepping algorithm that adjusts time-step
 233 accordingly will improve the simulation efficiency and accuracy. However,
 234 the energy conservation property of ECSIM is sensitive to the temporal dis-
 235 cretization scheme, and the adaptive time-stepping algorithm should not
 236 break the conservation.

237 Our adaptive time-stepping algorithm is summarized in Figure 3. At the
 238 end of one cycle, both the electromagnetic fields and the particle velocities are
 239 at time stage t^n , and the particle locations are at the staggered stage $t^{n+1/2}$.
 240 The difference between $t^{n+1/2}$ and t^n is $t^{n+1/2} - t^n = \Delta t^n/2$. The maximum
 241 speed v_{rms} can be obtained with the particle velocities at t^n , and a new
 242 time step Δt^{n+1} can be calculated from $\Delta t^{n+1} = CFL \cdot \Delta x/v_{rms}$. However,
 243 during the next cycle of updating the electromagnetic fields and particle
 244 velocities from t^n to t^{n+1} , the time step should be Δt^n instead of Δt^{n+1} , so
 245 that the particle location $X^{n+1/2}$ is still at the middle of t^n and t^{n+1} , and
 246 the energy conservation property of ECSIM is preserved. In order to adjust

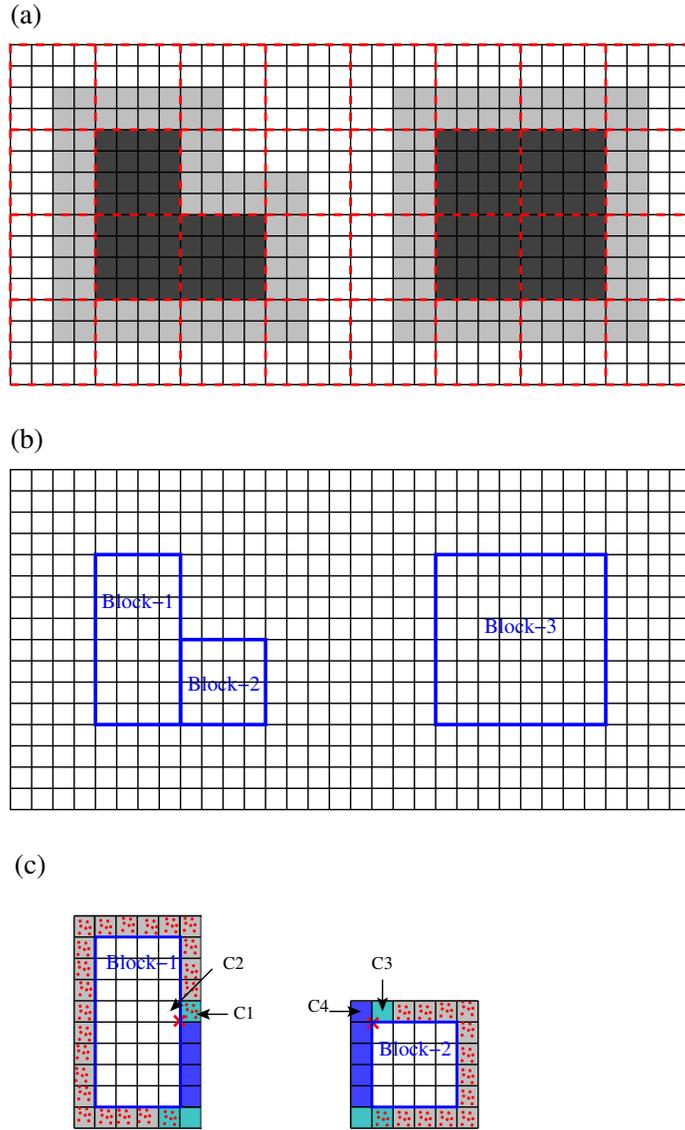


Figure 2: The black lines represent the cells of a PIC domain. The red dashed lines in (a) show the patches, and one patch contains 4×4 cells in this example. In (a), the active patches/cells are colored by dark gray, and light gray area represents the ghost cells of the active PIC regions. (b) shows the blocks of the active regions. (c) shows the inner layer of the ghost cells of two blocks, and the red dots represent the macro-particles that are generated in the ghost cells as the particle boundary condition. Blue ghost cells are internal ghost cells, which are overlapped with the physical cells of the neighboring blocks. The gray cells are exclusive boundary ghost cells, and they should be filled in with macro-particles as the boundary condition. The cyan cells are also boundary ghost cells, but they are overlapped with the boundary ghost cells of the neighboring blocks, and only one of the blocks should generate boundary particles. The C1...C4 labels and the two red crosses are used in the main text describing the algorithm.

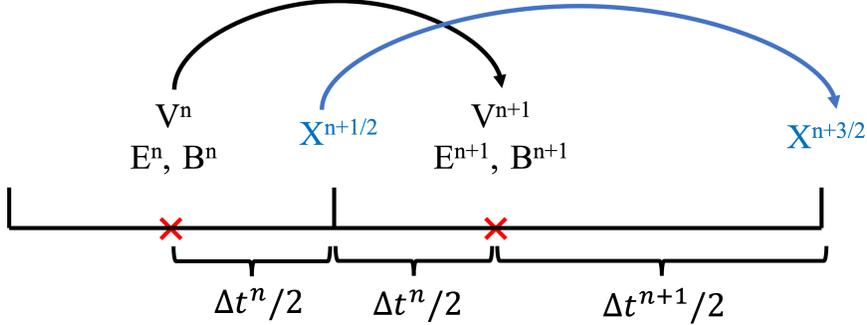


Figure 3: The adaptive temporal discretization.

247 the time step for the next cycle, we use the time step $(\Delta t^n + \Delta t^{n+1})/2$ for
 248 updating the particle location from $X^{n+1/2}$ to $X^{n+3/2}$. The velocity V^{n+1}
 249 at t^{n+1} is not centered exactly between $t^{n+1/2}$ and $t^{n+3/2}$, but the deviation
 250 $(\Delta t^{n+1} - \Delta t^n) \propto (\partial \ln v_{rms} / \partial t)(\Delta t^n)^2$ is second order since the v_{rms} used
 251 to calculate the time steps changes continuously with time. Therefore the
 252 second-order accuracy of updating particle locations is still satisfied.

253 4. Particle resampling

254 Particle resampling algorithms are used to control the macro-particle
 255 number of each cell. At the end of every computational cycle, a particle
 256 splitting (merging) algorithm is applied to generate (remove) macro-particles
 257 for the cells that contain fewer (more) particles than the splitting (merging)
 258 threshold. The goal of splitting and merging is to stop the number of parti-
 259 cles per cell (ppc) from dropping or increasing continually. Essentially, the
 260 particle resampling algorithms use a new set of particles to replace the old
 261 ones. Our guiding principle of designing the algorithms is that the replace-
 262 ment should preserve the original particle phase space distribution as much as
 263 possible. In order to conveniently apply the resampling algorithms, FLEKS
 264 stores the particle data cell by cell.

265 4.1. Particle splitting

266 Our particle splitting algorithm is essentially the same as the one in-
 267 troduced by Lapenta [16], in which one particle is split into two children

268 particles. The children particles have the same velocity as their parent par-
269 ticle, but their locations are oppositely displaced slightly along the velocity
270 direction. By displacing the new particles along the velocity direction, the
271 orbits of the new particles are still close to the orbit of the old particle.

272 The particle splitting will be triggered for the cells with ppc less than the
273 splitting threshold, which is 80% of the initial ppc by default, and we will
274 use this number for all the simulations presented in section 6. Initially, the
275 particles that are close to each other have similar weights, but the weights
276 may become quite different later due to the transport of particles and also the
277 particle splitting and merging. For each cell, we choose to split the heaviest
278 N particles to minimize the particle weight variance, where N is the difference
279 between the current ppc and the splitting threshold.

280 *4.2. Particle merging*

281 The essence of particle merging is replacing a set of particles with a new
282 set, which contains fewer particles than the old set. Particle merging reduces
283 the particle number in some cells and improves simulation speed. In gen-
284 eral, particle merging has a negative impact on the accuracy of a simulation
285 because (1) the replacement introduces errors, and (2) fewer particles lead
286 to larger statistical noise in the subsequent simulation. The statistical noise
287 increasing is inevitable, but the errors caused by the replacement can be
288 minimized with a proper merging algorithm.

289 Our particle merging algorithm consists of two steps: (1) selecting 6 par-
290 ticles that are close to each other in the phase space, and (2) merging these
291 6 particles into 5. In the following subsections, we will describe the merging
292 step first, and what follows is the selecting step.

293 *4.2.1. Merging particles*

294 Once the 6 old particles for merging have been obtained from the select-
295 ing step, we use 5 new particles to replace them. The replacement should
296 not alter the original phase space distribution significantly. However, it is
297 not trivial to quantitatively measure the change of the phase space. The
298 conservation of total mass, momentum, and energy can be used as a guid-
299 ance and indicator of preserving phase space structure. However, satisfying
300 the conservation property is not good enough, it is still possible that the
301 new particle set occupies a very different velocity space volume than the old
302 set. Previous methods [26, 17] usually generate new particles with velocities
303 that are different from the velocities of the old particles, and extra actions

304 are usually required to ensure the new particles are not too far away from
 305 the old ones in the velocity space. To avoid this difficulty, we choose to
 306 delete one of the 6 old particles and distribute its mass to the remaining 5
 307 particles under the constraint of conserving total mass, momentum vector,
 308 and energy. The weights of these 5 particles change, but their velocities are
 309 inherited from the old ones, so the new particle set occupies almost the same
 310 phase space volume as the old set.

311 The total mass, momentum and energy of the old particle set are:

$$w_t = \sum_{i=1}^{N_{old}} w_i, \quad \mathbf{p}_t = \sum_{i=1}^{N_{old}} w_i \mathbf{v}_i, \quad e_t = \sum_{i=1}^{N_{old}} \frac{1}{2} w_i v_i^2, \quad (1)$$

312 where $N_{old} = 6$. From these 6 particles, we find the pair that is closest to each
 313 other in the 6D phase space (Figure 4(d)), then remove the lighter one of this
 314 pair and adjust the weights of the rest 5 particles to satisfy the conservation
 315 requirement:

$$w_t = \sum_{i=1}^{N_{new}} w_{i,new}, \quad \mathbf{p}_t = \sum_{i=1}^{N_{new}} w_{i,new} \mathbf{v}_i, \quad e_t = \sum_{i=1}^{N_{new}} \frac{1}{2} w_{i,new} v_i^2 \quad (2)$$

$$w_{i,new} > 0 \quad (3)$$

316 where we choose $N_{new} = 5$ since there are 5 quantities to conserve. The
 317 velocities \mathbf{v}_i are known, and the new weights $w_{i,new}$ are the 5 unknowns of
 318 the linear equations (2) under the constraint of positivity (3). If the solution
 319 does not satisfy the constraint, we skip this merging.

320 To minimize the impact of the merging on the phase space distribution,
 321 we need to quantify a distance in the 6D phase space. The actual definition
 322 will be described in the next subsection, here we simply assume that the
 323 appropriate distance function exists.

324 By deleting the lighter particle from a pair that is closest to each other
 325 in the phase space, it is likely that its heavier neighbor will gain most of the
 326 weight and the other 4 particles adjust their weights relatively slightly. By
 327 inheriting the velocities and locations from the old particles, the new particles
 328 occupy almost the same phase space volume as the old particles (Figure 4(d)
 329 and (e)), so there is no room for the phase space structure to change dramati-
 330 cally. Compared to the schemes that allow choosing new particle velocity
 331 with fewer restrictions [17], our merging algorithm is less efficient to reduce

332 the particle number because (1) the new particle set still contains 5 parti-
 333 cles, and (2) the merging may fail when the constraint $w_{i,new} > 0$ can not
 334 be satisfied. If it is required, our algorithm can be easily modified to use a
 335 N_{old} that is larger than 6 by deleting $N_{old} - 5$ particles. However, as it can
 336 be seen from the following numerical test section, merging 6 particles into 5
 337 is already efficient enough for our typical applications.

338 4.2.2. *Selecting particles*

339 To minimize the phase space change, the particles selected for merging
 340 should be close to each other in the 6D phase space. Several strategies have
 341 been proposed for selecting particles, including binning particles in the phase
 342 space [21], partitioning phase space with Voronoi diagram [23], and using k-d
 343 tree data structure [22]. For the sake of simplicity, we choose the binning
 344 strategy. The dimension of the 6D phase space is so high that even only
 345 splitting each direction into 3 pieces leads to $3^6 = 729$ bins in total. Our
 346 typical simulations use about 100 particles per cell initially, and it is likely
 347 few phase space bins contain enough particles for merging with 3^6 bins. To
 348 avoid this problem, we only bin particles in the 3D velocity space and skip the
 349 merging if the 6D volume occupied by the selected particles is too large. This
 350 approach takes into account the spatial distribution of the selected particles,
 351 but also implies reducing the variance in the velocity space is more crucial
 352 than controlling the spatial location variance. Because all the particles are
 353 already in the same spatial cell, i.e., they can not be too far away from each
 354 other in the spatial dimensions. Inside each velocity space bin, we choose 6
 355 particles that are closest to each other for merging.

356 The particle merging algorithm needs to calculate the distance between
 357 two macro-particles in the 6D phase space. The distance is defined as:

$$d = \frac{\Delta s}{\Delta x} + c_1 \frac{\Delta v}{v_{th}} \quad (4)$$

358 where Δs is the spatial distance and Δv is the distance in the velocity space
 359 between the two particles. The normalization in space is the simulation cell
 360 size Δx . The velocity is normalized to the thermal velocity in the cell

$$v_{th} = \frac{1}{N_p} \sum_{i=1}^{N_p} |\mathbf{v}_i - \mathbf{v}|^2 \quad (5)$$

361 where \mathbf{v} it the bulk velocity of the cell and N_p is the number of particles in
 362 the cell. We note that the thermal velocity defined above is used to measure

363 the separation of particles in the velocity space, so the particle weight is not
 364 involved in the calculation. The constant c_1 in (4) determines the relative
 365 importance of the spatial distance Δl_{3D} and the velocity distance Δv . We
 366 choose $c_1 = 2$ based on our experience with many numerical tests.

367 At the end of each computational cycle, the following algorithm is per-
 368 formed to select particles for merging if the ppc of a cell is larger than the
 369 merging threshold, which is 1.5 times of the initial ppc by default:

- 370 1. Bin the particles in the velocity space. For each spatial cell (Fig-
 371 ure 4(a)), we create a grid in the velocity space ranging from $(v_x -$
 372 $v_{th}, v_y - v_{th}, v_z - v_{th})$ to $(v_x + v_{th}, v_y + v_{th}, v_z + v_{th})$, and assign par-
 373 ticles to velocity bins (Figure 4(b)). The velocity grid is divided into
 374 $n_{bin} = \lceil 0.8N_p^{1/3} \rceil$ bins in each direction, where N_p is the current ppc of
 375 the cell and the constant 0.8 is chosen based on numerical experiments.
 376 We note that each bin contains a buffer region (Figure 4(c)), and the
 377 particles in the buffer region may also belong to other bins. We use
 378 $1/8$ of the velocity space bin size as the width of the buffer region (Fig-
 379 ure 4(c)). Due to the existence of the buffer region, one particle may
 380 belong to multiple bins, but it can only be selected for merging at most
 381 once during one cycle.
- 382 2. Select particles from a bin. If there are more than 6 particles inside
 383 a bin, including the buffer region, we choose a cluster of 6 particles
 384 from them. For each velocity bin, we calculate the velocity center of
 385 the associated particles (black cross in Figure 4(c)), and find the 6
 386 particles closest to the center in the 3D velocity space. If a particle in
 387 the buffer region has been selected for merging by a neighboring bin,
 388 this particle should not be selected again.
- 389 3. Limit the 6D distance. The previous step selects particles only based
 390 on the distance in the velocity space. This step ensures the selected
 391 particles are also close to each other in the 6D phase space. We find
 392 the 6D center (blue cross in Figure 4(d)) of these 6 particles, and the
 393 6D distance d of all the 6 particles to the center should be less than
 394 0.6. Again, the constant 0.6 is chosen based on numerical experiments.
- 395 4. Merge 6 particles into 5 with the algorithm described in section 4.2.1.

396 The particle selection method used in step 2 prefers selecting particles in the
 397 center of a bin. Without applying the buffer region in step 1, the particles
 398 near the edge of a bin are less likely to be chosen for merging. On the other

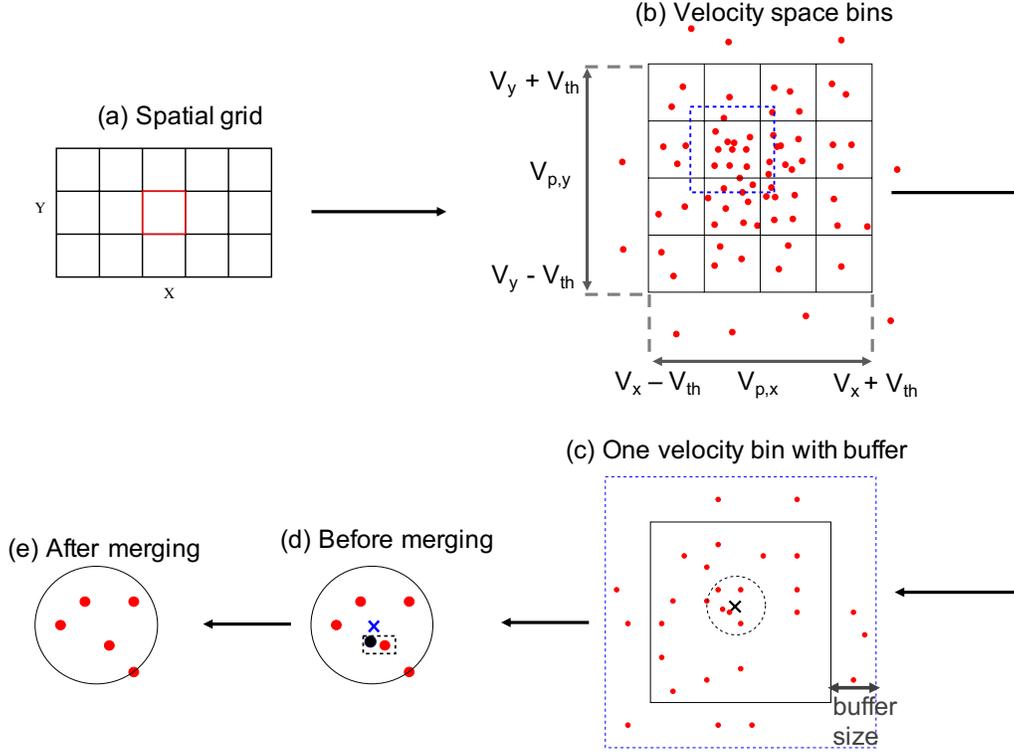


Figure 4: Schematics of the algorithm of merging macro-particles. See text for details.

399 hand, it is more likely that a bin extended with a buffer region contains
 400 more than 6 particles, which improves the merging efficiency. Based on
 401 our numerical experiments, applying the buffer region does not improve the
 402 simulation results significantly, but it is still kept by default to avoid the
 403 aforementioned potential issues.

404 5. Test particle module

405 An independent test particle (TP) module is designed to track the mo-
 406 tion of the macro-particles for FLEKS. It can be used either as an auxil-
 407 iary component of the PIC algorithm or as an independent component. The TP
 408 module uses the same algorithm to move particles as the GL-ECSIM algo-
 409 rithm. When the TP module is used with the PIC component together, the

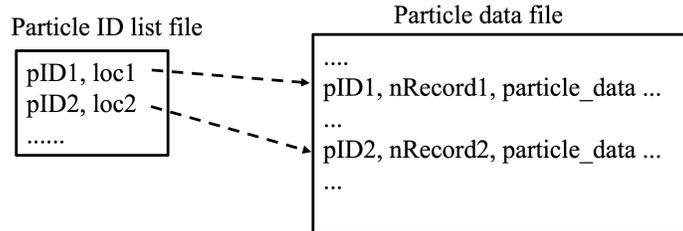


Figure 5: The file structure for storing test particles.

410 TP module shares the same grid layout as the PIC component and uses the
 411 electromagnetic fields calculated by PIC to update test particles. When the
 412 PIC component is turned off, FLEKS becomes a pure test particle code, and
 413 the TP module can directly obtain the grid structure and electromagnetic
 414 fields from the MHD model. Compared to the embedded PIC simulations,
 415 the pure test particle simulations are only one-way coupled, i.e., the MHD
 416 model provides the electromagnetic fields for FLEKS, but there is not any
 417 feedback from FLEKS to the MHD model.

418 In a 3D simulation, it is common to track the motion of millions of test
 419 particles, and a few thousand steps of the update will easily produce a few
 420 hundred Gigabytes of particle trajectory data. The test particle module
 421 should organize the data properly to improve both the IO performance of
 422 writing data to disk and also the efficiency of reading the trajectory of a
 423 particle for data analysis. To reduce the IO frequency, the TP module of
 424 FLEKS saves the particle trajectory data every 100 cycles, and all the pro-
 425 cessors write to the same file with MPI-IO APIs. We note that if a test
 426 particle moves from one processor to another in the middle of two IO op-
 427 erations, its trajectory data should also be transferred to the destination
 428 processor. Besides the particle trajectory data file, a particle ID list file,
 429 which maps a particle ID to its data location in the particle data file, is also
 430 created. An example of these two files is shown in Figure 5. With this file
 431 structure, it is efficient to find all the trajectory data of a particle for data
 432 analysis.

433 **6. Numerical tests**

434 *6.1. Two-dimensional fast magnetosonic wave propagation with adaptive PIC*
 435 *region*

436 We use a two-dimensional (2D) fast magnetosonic wave propagation test
 437 to demonstrate the capability of FLEKS's adaptive grids. The same initial
 438 condition as what is described in [1] is applied here to produce a propagat-
 439 ing fast magnetosonic wave. The simulation domain of the MHD code is
 440 $-160/3 < x < 160/3$ and $-40 < y < 40$. Two independent PIC domains
 441 are used. The left domain in Figure 6 covers the region of $-40 < x < 0$ and
 442 $-20 < y < 20$ with a grid resolution of $\Delta x = \Delta y = 1/16$. The right domain
 443 covers the region of $20 < x < 40$ and $-10 < y < 10$ with a grid resolution
 444 of $\Delta x = \Delta y = 1/8$. All cells of the right PIC domain are always switched
 445 on during the simulation. For the left domain, only the cells that satisfy the
 446 following conditions are switched on:

$$\begin{aligned}
 & r < \frac{L_x}{10} \\
 & \text{or} \\
 & r < \frac{L_x}{4} + \frac{L_x}{4} \frac{t \bmod 200}{200} \quad \text{and} \quad r > \frac{L_x}{8} + \frac{L_x}{10} \frac{t \bmod 200}{200},
 \end{aligned}
 \tag{6}$$

447 where r is the distance to the center of the PIC domain, L_x is the length of
 448 the PIC domain, which is 40 in this case, and t is the simulate time. The
 449 central PIC cells ($r < L_x/10$) are always switched on, and the outer shell of
 450 active PIC cells keep changing during the simulation. A movie that shows
 451 the adaptation of the active PIC region is provided as an online supplement.
 452 We note that the simulation parameters for these two PIC domains can be
 453 specified independently. For example, the cell size is different for these two
 454 PIC domains as it is described above, and the ion-electron mass ratio m_i/m_e
 455 is 25 for the left domain and it is 100 for the right domain. Both PIC domains
 456 use CFL = 0.2, and 900 particles per cell (ppc) per species.

457 Figure 6 shows the plasma velocity U_x and the area of the active PIC
 458 cells at the beginning and at $t = 400$. The interface between the active
 459 PIC region and the MHD region is smooth, and there is not any significant
 460 artificial effect observed.

461 *6.2. One-dimensional non-linear magnetosonic wave evolution*

462 The evolution of the magnetosonic wave is non-linear. The wave may
 463 finally evolve into a shock, where the plasma phase space distributions may

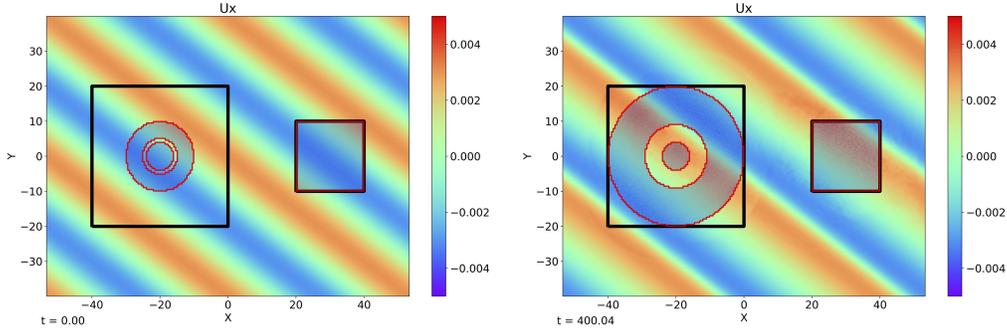


Figure 6: The velocity U_x of the 2D fast magnetosonic wave test at the beginning (left) and at $t = 400$ (right). The black rectangles show the area of the PIC grids. Inside each PIC grid, the semi-transparent area, which is enclosed by red lines, represents the active PIC region. Since all PIC cells are active for the right PIC grid, the black lines and the red lines are overlapped.

464 become non-Maxwellian. So the non-linear evolution of the magnetosonic
 465 wave simulation is suitable for testing the particle resampling algorithms.

466 In section 6.1, the wave vector is perpendicular to the background mag-
 467 netic field direction. To make the particle phase space distribution further
 468 away from Maxwellian and hence more challenging for the particle resampling
 469 algorithms, we use a more general setting that the background magnetic field
 470 is neither perpendicular nor parallel to the wave vector in this 1D test. The

471 initial condition of the 1D magnetosonic wave is:

$$\begin{aligned}
B_x(x) &= B_0 \cos(\theta) \\
B_y(x) &= B_0 [\sin(\theta) + \delta \sin(kx - \omega t)] \\
B_z(x) &= 0 \\
u_x(x) &= \delta \sin(\theta) \frac{v_A^2 v_p}{v_p^2 - v_s^2} \sin(kx - \omega t) \\
u_y(x) &= \delta \cos(\theta) \frac{v_A^2}{v_p} \sin(kx - \omega t) \\
u_z(x) &= 0 \\
\rho(x) &= \rho_0 \left[1 + \delta \sin(\theta) \frac{v_A^2}{v_p^2 - v_s^2} \sin(kx - \omega t) \right] \\
p(x) &= p_0 \left[1 + \gamma \delta \sin(\theta) \frac{v_A^2}{v_p^2 - v_s^2} \sin(kx - \omega t) \right],
\end{aligned} \tag{7}$$

472 where γ is the specific heat ratio, $v_A = \frac{B_0}{\sqrt{\rho_0}}$ is the Alfvén speed, $v_s = \sqrt{\frac{\gamma p_0}{\rho_0}}$
473 is the sound speed, and θ is the angle between the wave vector, which is
474 the x-direction here, and the background magnetic field. The phase speed
475 $v_p = \omega/k$ is the fast magnetosonic speed:

$$v_p^2 = \frac{1}{2} \left\{ v_A^2 + v_s^2 + [(v_A^2 + v_s^2)^2 - 4v_s^2 v_A^2 \cos^2 \theta]^{1/2} \right\}, \tag{8}$$

476 In this paper, we use $\gamma = 5/3$, $B_0 = 0.1$, $\theta = 30^\circ$, $\rho_0 = 1$, $p_0 = 0.0001$,
477 $k = 2\pi/\lambda = 2\pi/64$, and $\delta = 0.5$. We note that the perturbation $\delta = 0.5$
478 is not small so that the solution will evolve to the nonlinear stage soon.
479 Since the goal of this test is to compare the simulation results with and
480 without particle resampling, it is suitable and acceptable to use such a large
481 perturbation.

482 The 1D simulation domain is $-32 < x < 32$ with a cell size $\Delta x = 0.05$.
483 The initial number of particles per cell per species is 900, and $\text{CFL} = 0.2$.
484 The simulation results at $t = 200$ are presented in Figure 7. To distinguish
485 between the physical density and macro-particle number per cell, we use
486 'mass density' to represent the physical density, and 'number density' is the
487 number of macro-particles per simulation cell or phase space bin. At $t = 200$,
488 the wave already evolves into a non-linear state, and the velocity shows a
489 sharp gradient near $x = 20$. The minimum and maximum number of ppc are

490 about 500 and 3140, respectively, for the simulation without applying particle
 491 resampling. For the simulation with particle resampling, the minimum ppc
 492 is about 750 and the maximum ppc is about 1360, and these numbers are
 493 close to the splitting limit $0.8 * 900 = 720$ and the merging limit $1.5 * 900 =$
 494 1350 . It suggests that the particle resampling algorithms are effective in
 495 controlling particle numbers. Except for the particle number, the physical
 496 quantities of these two simulations are very similar to each other. The only
 497 noticeable difference is that the electric field E_y of the simulation with particle
 498 resampling is noisier near $x = 20$ due to the reduction of particle number.
 499 Figure 7(b) shows the ion phase space distribution for particles between
 500 $x = 21$ and $x = 21.2$. The two mass density distributions are comparable
 501 even though the particle number densities are quite different.

502 Figure 8 shows the simulation speed, which represents the number of PIC
 503 cells that are updated per second per CPU core. For the first 700 cycles, both
 504 simulations become slower and slower due to the imbalance of the particle
 505 number per CPU core. Later, the minimum and maximum ppc reach the
 506 splitting and merging thresholds and the particle splitting algorithms start
 507 controlling the further change of the minimum and maximum ppc, so the
 508 simulation speed stops dropping for the simulation with particle resampling.
 509 At the end of the simulation, the simulation with particle resampling is almost
 510 twice faster than the one without particle resampling.

511 *6.3. Two-dimensional double-current-sheet magnetic reconnection*

512 Magnetic reconnection is regarded as one of the most important physical
 513 processes for energy transfer between magnetic field and plasma in the space
 514 plasma environment, so it is also widely used to benchmark the performance
 515 of a kinetic plasma modeling code. Here, we use a two-dimensional (2D)
 516 asymmetric magnetic reconnection problem to test the particle resampling
 517 algorithms, because the particles distributions near the reconnection site can
 518 be non-Maxwellian. It is crucial to demonstrate that the particle resampling
 519 algorithms preserve the non-Maxwellian distributions.

520 A double current sheet is used to initialize the simulation so that the
 521 whole system is symmetric, and periodic boundary conditions can be applied
 522 in all directions. The simulation domain is $-64 < x < 64$ and $-16 < y < 16$.
 523 The background magnetic field is initialized as:

$$B_x(y) = \left(\frac{B_1 + B_2}{2} \right) \left[\tanh \left(\frac{y + L_y/4}{\delta} \right) - \tanh \left(\frac{y - L_y/4}{\delta} \right) \right] - B_2, \quad (9)$$

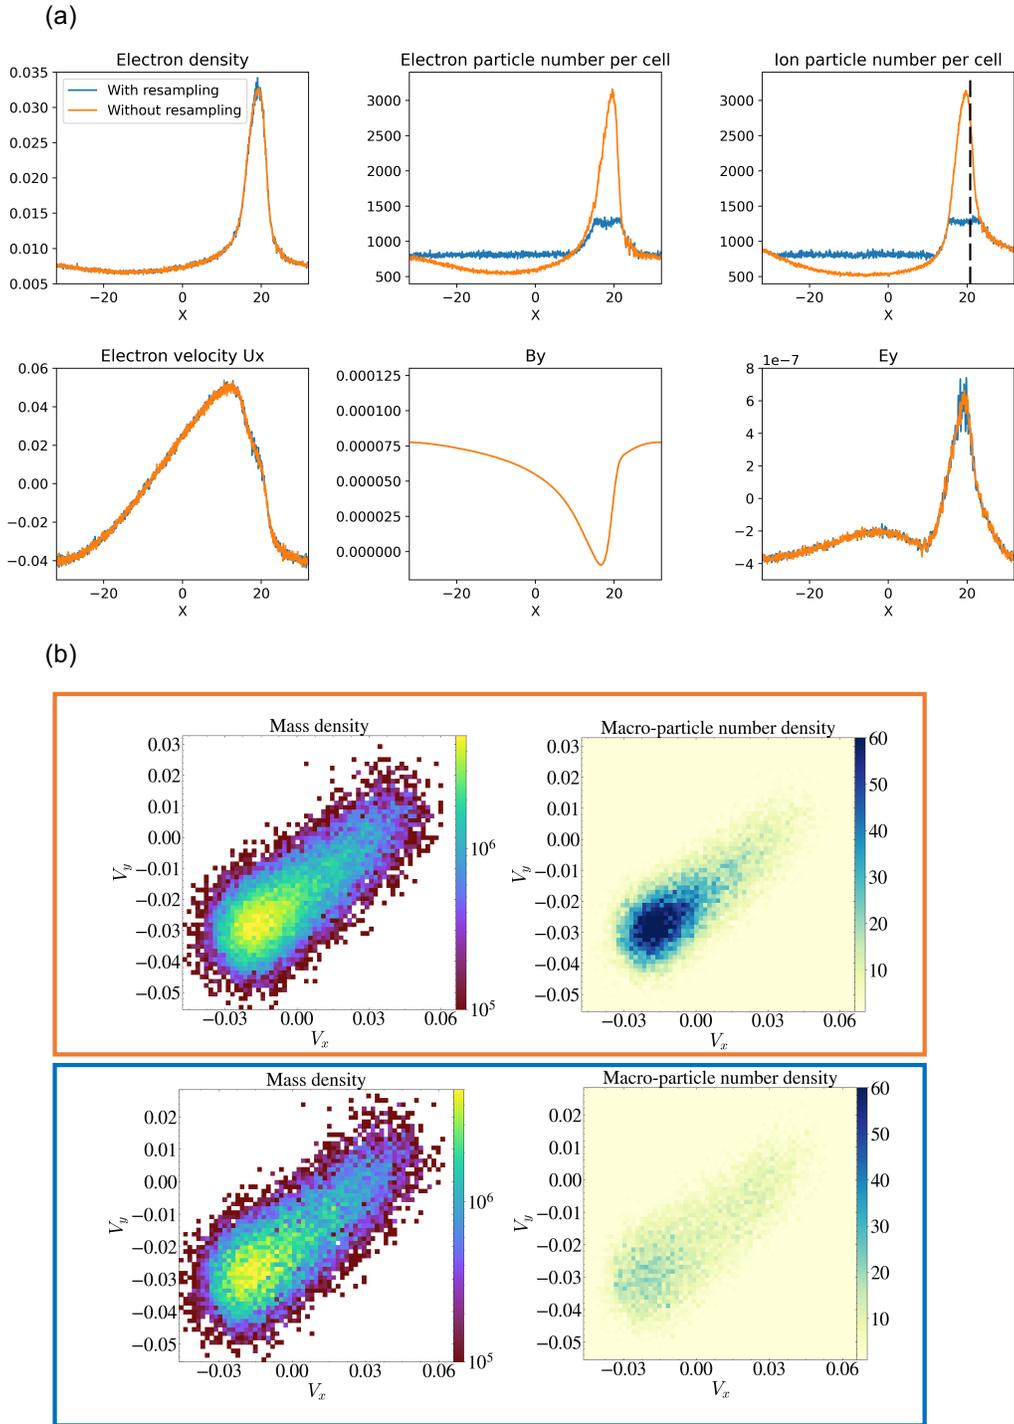


Figure 7: The 1D magnetosonic wave simulation results at $t = 200$. (b) shows phase space distributions at $x = 21$ that is marked with a dashed black line in (a). The upper panel of (b) shows the results without particle resampling, and the lower panel shows the results with particle resampling.

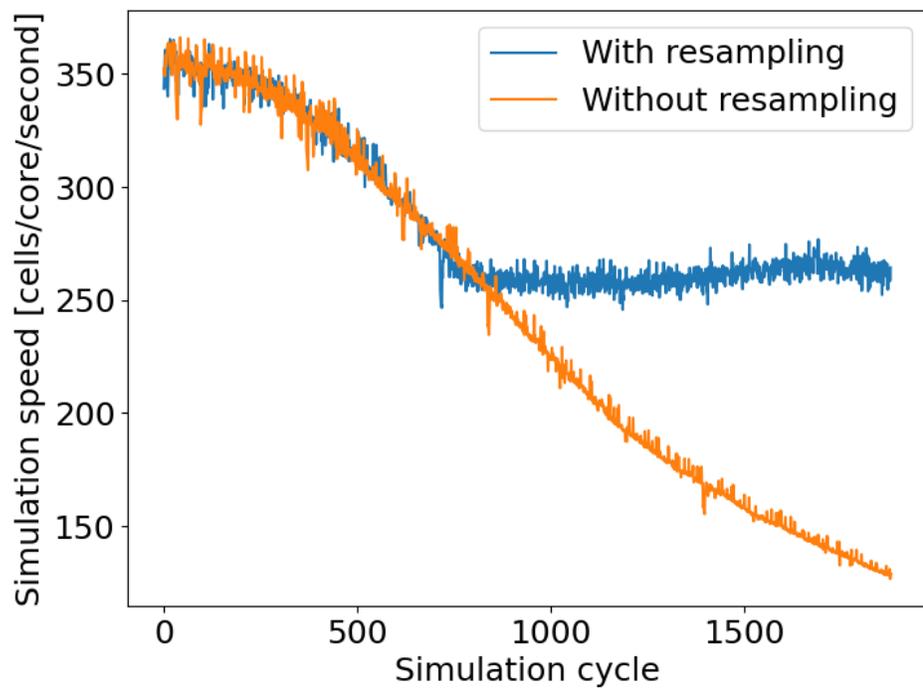


Figure 8: The simulation speed of the 1D magnetosonic wave simulations.

524 where $B_1 = 1$ and $B_2 = 2$ are the asymptotic magnetic field amplitudes.
525 $L_y = 32$ is the width of the simulation domain, and the centers of the two
526 current sheets are at $y = -8$ and $y = 8$, respectively. The plasma pressure
527 is set to balance the magnetic field pressure $p_B = B^2/2$. To mimic the
528 plasma environment of Earth's magnetopause, the asymptotic plasma beta
529 $\beta = (p_i + p_e)/p_B$ are 3.6 and 0.15 on the "1" and "2" sides, respectively. The
530 initial pressure ratio between electrons and ions is $p_i/p_e = 5$ in the whole
531 simulation domain. The ion temperature is:

$$T_i(y) = \left(\frac{T_{i,1} + T_{i,2}}{2} \right) \left[\tanh \left(\frac{y + L_y/4}{\delta} \right) - \tanh \left(\frac{y - L_y/4}{\delta} \right) \right] + T_{i,2}. \quad (10)$$

532 $T_{i,1} = 1.33$ and $T_{i,2} = 3.33$ are used in the simulation. With the pressure
533 and temperature given above, the corresponding densities and ion inertial
534 lengths are $n_1 = 1.127$, $n_2 = 0.0736$, $d_{i,1} = 0.942$ and $d_{i,2} = 3.69$. For all the
535 simulations presented in this subsection, the grid resolution is $\Delta x = 1/16$,
536 and the CFL is 0.4.

537 Figures 9 and 10 show the fields near the reconnection site at $t = 20$ with
538 100 and 400 initial ppc, respectively. The left (right) columns of Figures 9
539 and 10 are the results without (with) applying particle resampling. Due to
540 the magnetic reconnection plasma flow, the electron ppc around the current
541 sheet increases to about 250 (950), and the minimum ppc in the inflow region
542 reduces to less than 50 (200) in Figure 9 (Figure 10) without applying
543 the particle resampling algorithms. After applying the particle resampling
544 algorithms, the electron ppc becomes more uniform in the whole domain.
545 With the threshold parameters described in section 4, the particle splitting
546 (merging) threshold ppc is 80 (150) and 320 (600) for the simulations with
547 the initial ppc of 100 and 400, respectively. The minimum electron ppc in the
548 right column of Figure 9 (Figure 10) is about 83 (325), and the maximum ppc
549 is about 190 (630). The minimum ppc in the simulate is just a few particles
550 more than the splitting threshold since the splitting algorithm is effective in
551 generating new particles. The difference between the maximum ppc and the
552 merging threshold is larger, but the maximum ppc is still much smaller than
553 that in the simulation without applying particle resampling.

554 Figures 9 and 10 also compare the physical quantities of the simula-
555 tions. All simulations show essentially the same structures, including the
556 off-diagonal electron tensor. It demonstrates that the particle resampling
557 algorithms do not introduce any significant artificial effect.

558 Figure 11 shows the electron phase space distributions from three sam-
 559 pling locations near the reconnection site. These three sampling locations
 560 are marked with black rectangles in the first rows of Figures 9 and 10. From
 561 top to bottom, we label these three sampling boxes as box-A, box-B, and
 562 box-C. In Figure 11, rows (a) and (b) show distributions from box-A, rows
 563 (c) and (d) show distributions from box-B, and rows (e) and (f) show dis-
 564 tributions from box-C. Each column shows the distributions from the same
 565 simulation, and the simulation parameters, i.e., the initial ppc and turning
 566 on/off the particle resampling algorithms, are described at the top of Fig-
 567 ure 11. Rows (a), (c) and (e) show the density distributions, and rows (b),
 568 (d) and (f) show the macro-particle number distributions in phase space.
 569 Figure 11 demonstrates that the particle resampling algorithms preserve the
 570 phase space distributions well. The particle resampling does not change the
 571 particle number too much at the sampling location box-A (row (b)), and
 572 the 'U'-shape density distribution is well preserved (row (a)). From rows
 573 (d) and (f), it is clear that the particle resampling significantly reduces the
 574 particle number around the distribution centers, so the centers of the density
 575 distribution (rows (c) and (e)) with particle resampling are noisier. But the
 576 density distribution structure, which consists of a core and a crescent distri-
 577 bution, is still clearly preserved in the row (c) and also in (e1) and (e2). The
 578 distributions of (e3) and (e4) are also very similar to each other.

579 *6.4. Strong and weak parallel scalings*

580 3D asymmetric magnetic reconnection simulations are used to test the
 581 strong and weak scaling of FLEKS. The setup of the 3D test is similar to
 582 the 2D simulation in the previous subsection, and it is uniform in the z-
 583 direction. Since FLEKS uses the parallel field and particle data structures
 584 provided by AMReX, the scaling results largely depend on the performance
 585 of AMReX [9, 10]. Figure 12 shows the weak scalings. With 8^3 cells per core,
 586 the performance is still good with up to about 10k cores. With 16^3 cells per
 587 CPU core, it reaches good performance even with 28,672 cores. Figure 13
 588 shows the strong scalings of two problems. The speedup is not too far away
 589 from the ideal scaling up to about 7k (Figure 13(a)) or 14k (Figure 13(b))
 590 CPUs for these problems.

591 *6.5. Magnetospheric simulations*

592 Magnetospheric simulations represent the most important application of
 593 MHD-AEPIC. Here, we show examples of how FLEKS benefits magneto-

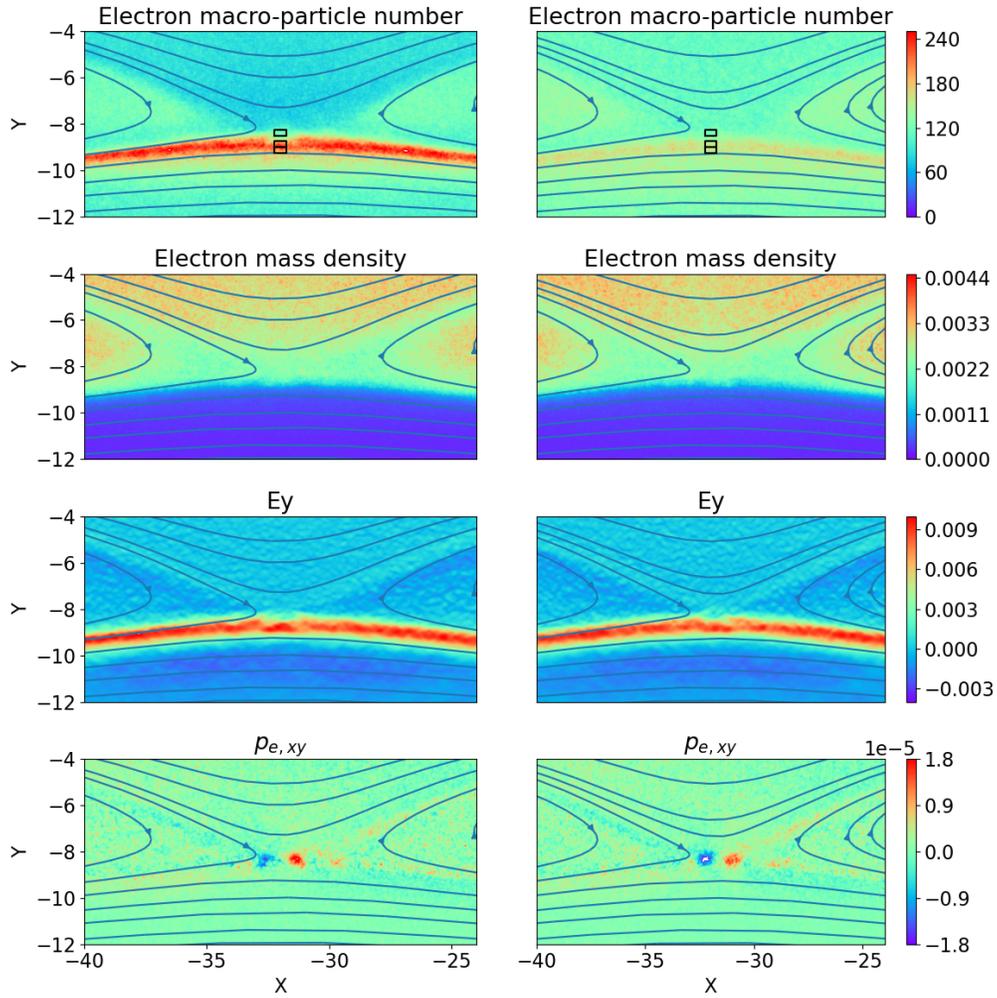


Figure 9: 2D magnetic reconnection results with (right column) or without (left column) particle resampling. The initial particle number per cell is 100. The black boxes in the top row indicate where the distribution functions shown in Figure 11 are taken from.

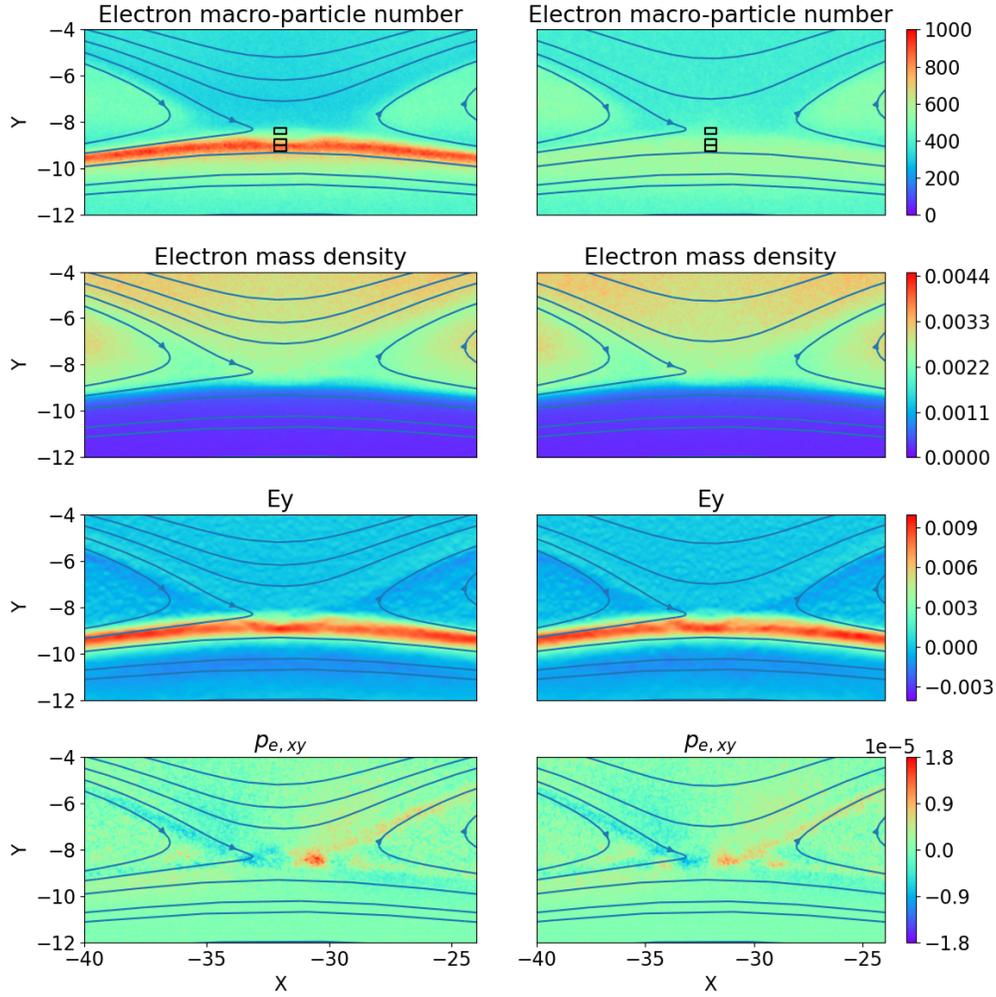


Figure 10: 2D magnetic reconnection results with (right column) or without (left column) particle resampling. The initial particle number per cell is 400 that is 4 times more than in Figure 9. The black boxes in the top row indicate where the distribution functions shown in Figure 11 are taken from.

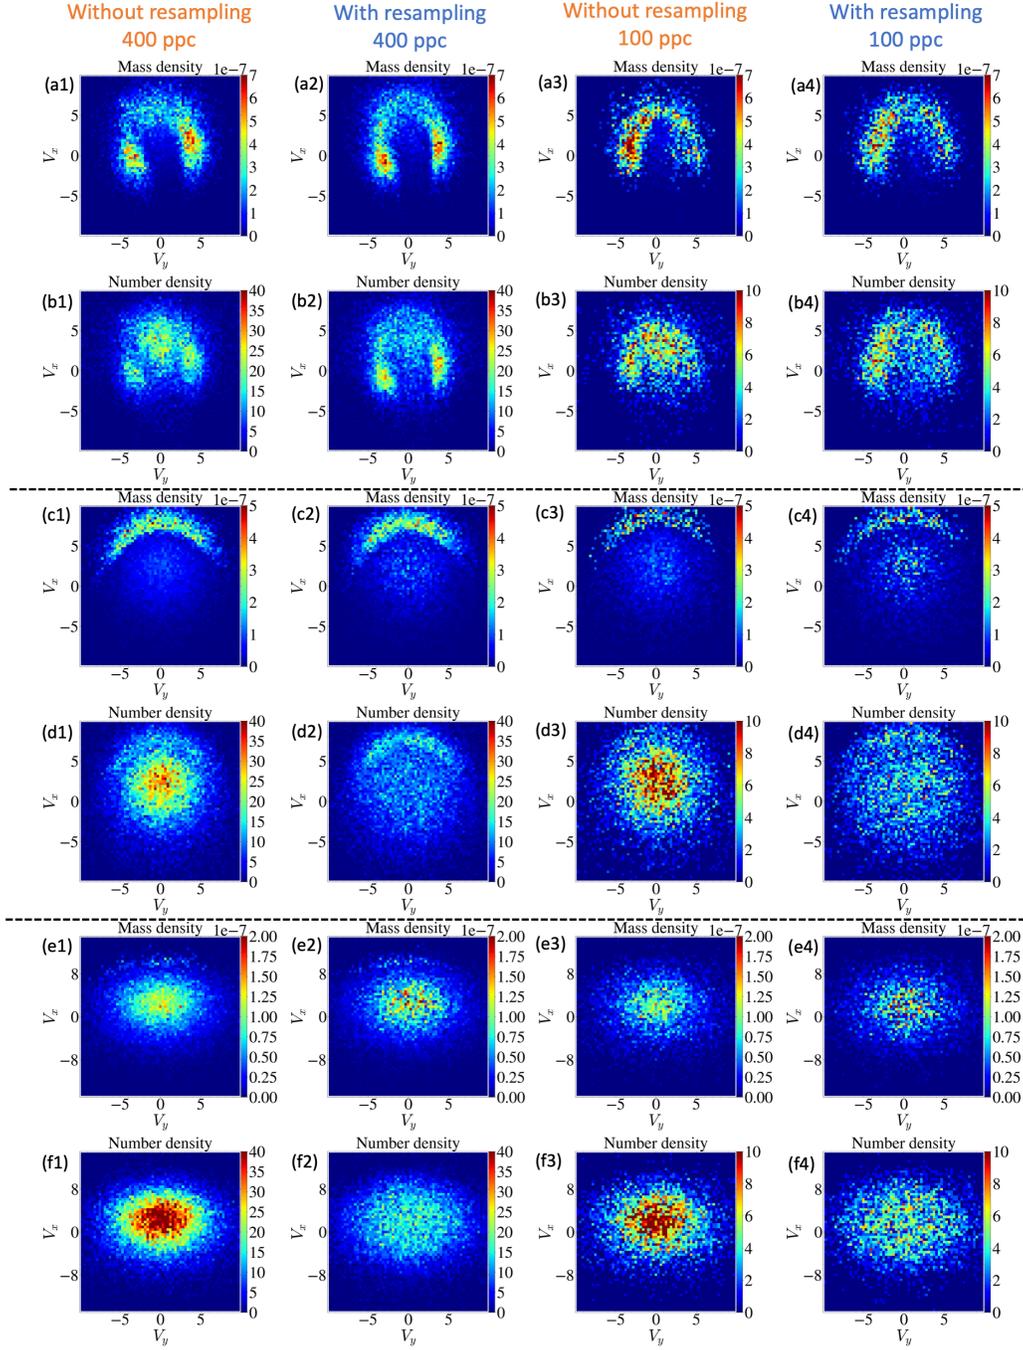


Figure 11: Each column shows the phase space distributions from the same simulation. The first two rows, middle two rows and the last two rows represent the distributions of box-A, box-B and box-C, respectively. From top to bottom, the black rectangles in Figures 9 and 10 show the locations of box-A, box-B and box-C. Rows (a), (c) and (e) are physical density distributions. Rows (b), (d) and (f) show particle number per phase space bin.

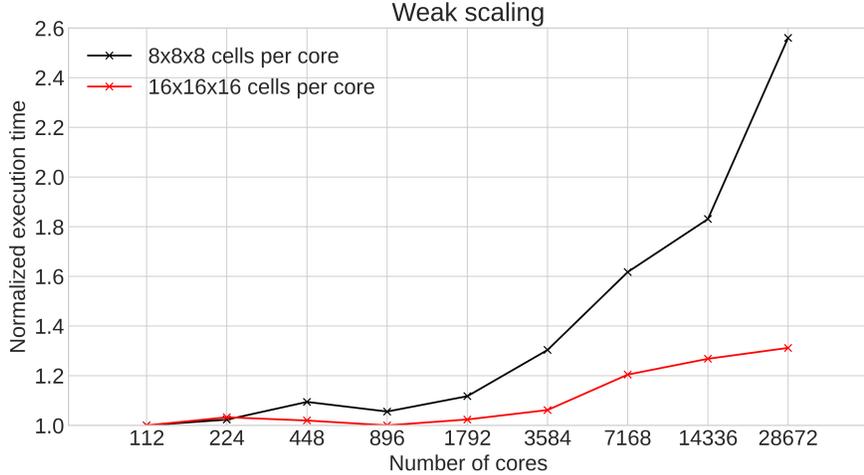


Figure 12: The weak parallel scaling results of FLEKS. The execution time with 112 CPU cores is used as the reference. For perfect scaling the normalized execution time would be 1.0 for all runs.

594 sphere modeling. The global MHD magnetosphere model uses the same
 595 setup as the simulations presented in [5, 24], but the active PIC region is not
 596 limited to be a box anymore. Figure 14 shows how the active PIC region
 597 can efficiently cover the dayside magnetopause, including the dawn-side and
 598 dusk-side flanks, and also the cusp region at the same time.

599 The test particle module enables us to follow the trajectories of particles
 600 in the magnetosphere. Figure 15 shows an example of test particles in Earth’s
 601 magnetosphere.

602 7. Conclusion

603 In this paper, we introduce a new kinetic code FLEKS, which is designed
 604 as the kinetic component of the MHD-AEPIC model [8, 25]. To support
 605 long simulations with varying global configurations, FLEKS allows activating
 606 or deactivating cells dynamically during a simulation to fit the regions of
 607 interest. This feature was introduced by Shou et al. [8] first, but FLEKS
 608 is more flexible since the minimum activation unit is a patch containing N
 609 ($N \geq 2$) cells in each direction instead of a large block used in [8], and FLEKS
 610 supports multiple independent PIC domains in an MHD-AEPIC simulation.

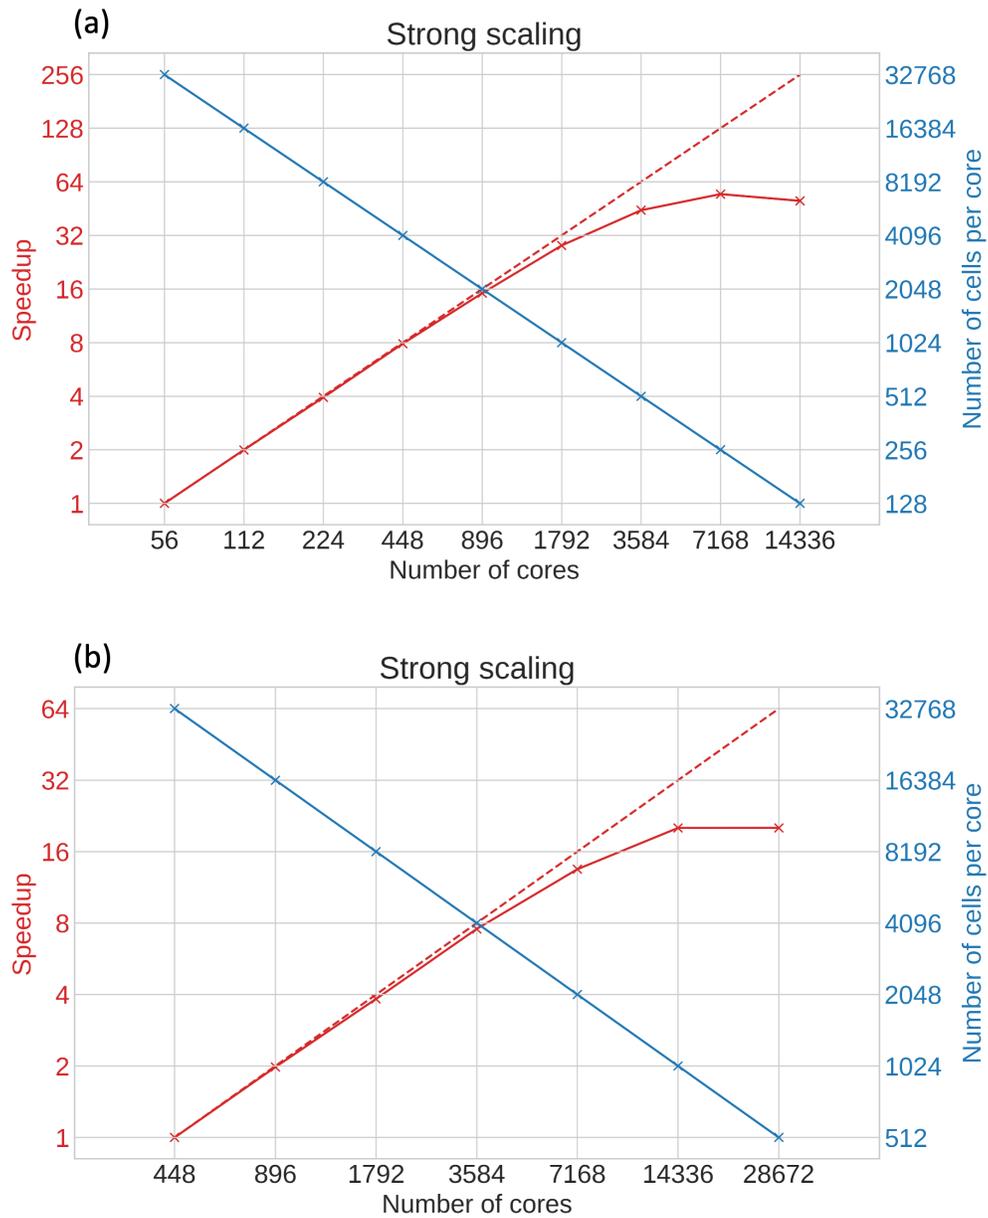


Figure 13: The strong scaling of FLEKS. Panels (a) and (b) show the scaling of simulations with 1.835 million and 14.68 million cells, respectively. The red solid lines represent speedup, and the red dashed lines correspond to perfect speedup. The blue lines show the number of PIC cells per CPU core.

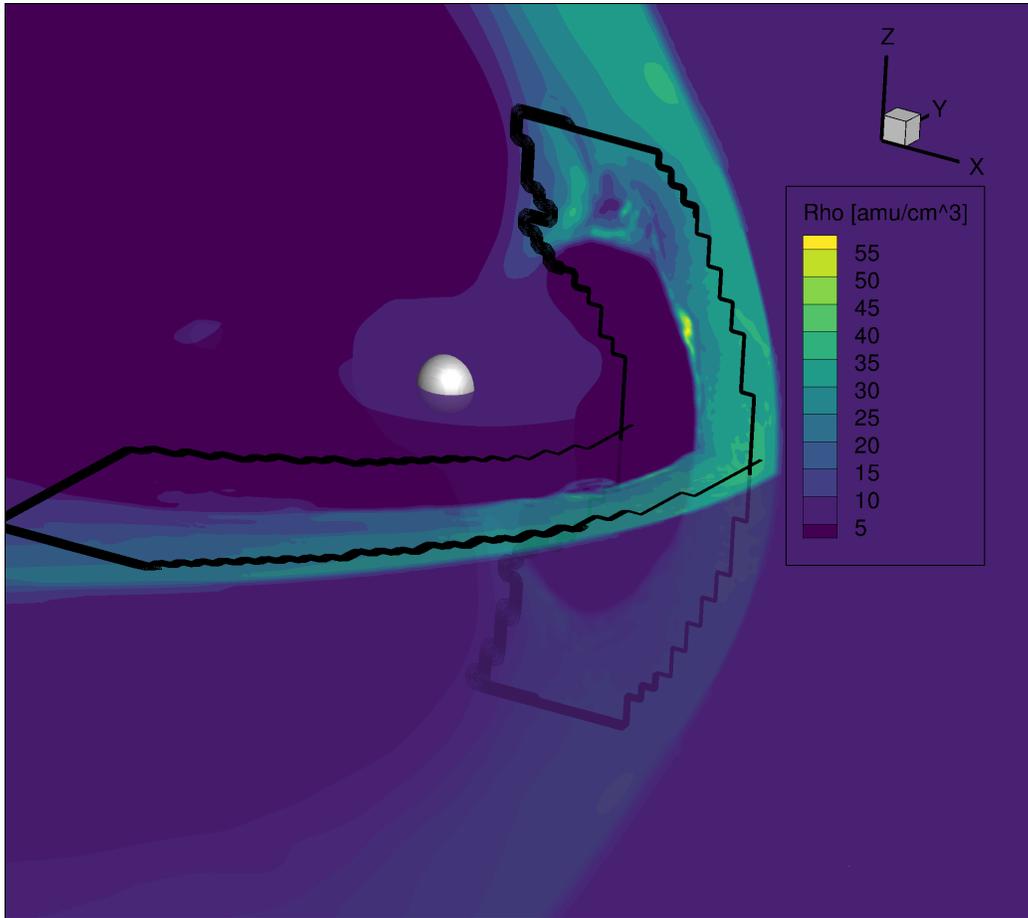


Figure 14: An MHD-AEPIC simulation of Earth's magnetosphere with the dayside magnetopause and the cusps covered by FLEKS. The black lines indicate the edge of the active PIC region.

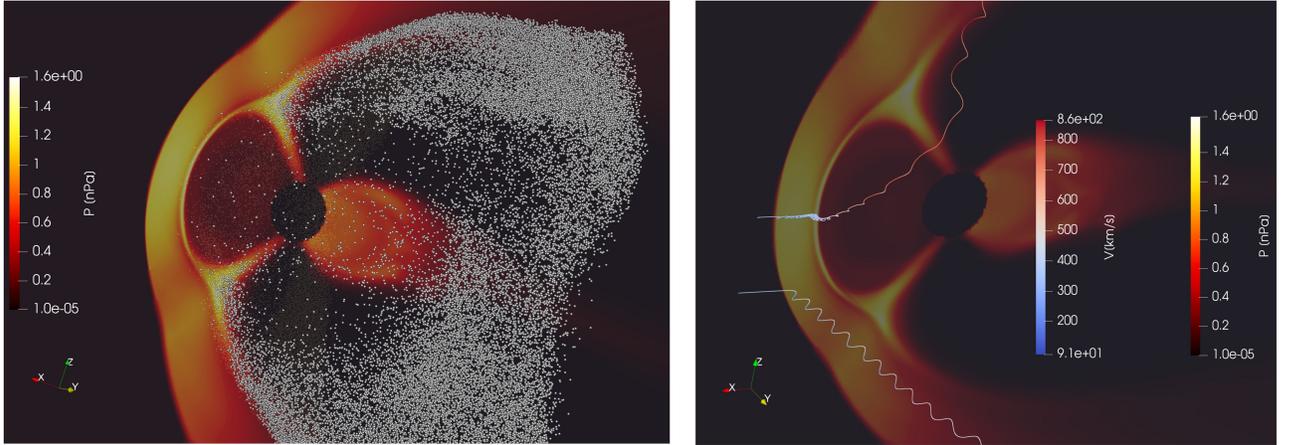


Figure 15: The locations (left) and the example trajectories (right) of test particles.

611 During a long simulation, since the plasma properties inside the active PIC
 612 region may change greatly, we design an adaptive time-stepping algorithm
 613 to adjust PIC time step accordingly. The adaptive time-stepping scheme
 614 preserves the energy conservation property of the ECSIM algorithm.

615 Since the number of particles per cell may change dramatically during a
 616 long simulation and leads to load imbalance and loss of accuracy in the cells
 617 with fewer particles, a particle splitting and a particle merging algorithms are
 618 designed to control the change of ppc. The particle merging algorithm selects
 619 6 particles that are close to each other in the phase space and combines them
 620 into 5 new particles. The merging conserves the total mass, momentum, and
 621 energy, and it also preserves the phase space structure as much as possible
 622 by inheriting velocities from the old particles. We have presented several
 623 non-trivial tests showing that the particle splitting and merging algorithm
 624 does not introduce any spurious features.

625 The particle resampling improves the efficiency of FLEKS substantially by
 626 not allowing the ppc to drop to very small values or increase to unnecessarily
 627 high values. In addition, load balancing the PIC domain becomes much easier
 628 with roughly the same number of particles in each grid cell. Indeed, FLEKS
 629 shows excellent weak and strong parallel scaling. Finally, the test-particle

630 module expands the capability of FLEKS, and provides a useful tool for
631 investigating the transport and energization of particles in magnetospheres.

632 FLEKS improves the quality and efficiency of MHD-AEPIC simulation
633 results significantly. For example, Wang et. al. [27] use FLEKS inside the
634 Space Weather Modeling Framework to model a complete magnetospheric
635 storm with kinetic reconnection in the tail.

636 **Acknowledgments:** This work was primarily supported by the NSF PREEVENTS
637 grant 1663800. Y. Chen was also supported by NASA under grant 80NSSC19K0161
638 and 80NSSC20K1416. Computational resources supporting this work were
639 provided by an NSF LRAC allocation at the Texas Advanced Computing
640 Center (TACC) at The University of Texas at Austin.

641 **References**

- 642 [1] L. K. S. Daldorff, G. Tóth, T. I. Gombosi, G. Lapenta, J. Amaya,
643 S. Markidis, J. U. Brackbill, Two-way coupling of a global Hall mag-
644 netohydrodynamics model with a local implicit Particle-in-Cell model,
645 J. Comput. Phys. 268 (2014) 236. doi:10.1016/j.jcp.2014.03.009.
- 646 [2] M. Rieke, T. Trost, R. Grauer, Coupled Vlasov and two-fluid codes on
647 GPUs, Journal of Computational Physics 283 (2015) 436–452. doi:
648 10.1016/j.jcp.2014.12.016.
649 URL <http://dx.doi.org/10.1016/j.jcp.2014.12.016>
- 650 [3] T. Sugiyama, K. Kusano, Multi-scale plasma simulation by the interlock-
651 ing of magnetohydrodynamic model and particle-in-cell kinetic model,
652 J. Comput. Phys. 227 (2007) 1340–1352. doi:10.1016/j.jcp.2007.
653 09.011.
- 654 [4] G. Tóth, X. Jia, S. Markidis, B. Peng, Y. Chen, L. Daldorff, V. Ten-
655 ishev, D. Borovikov, J. Haiducek, T. Gombosi, A. Glocer, J. Dorelli,
656 Extended magnetohydrodynamics with embedded particle-in-cell sim-
657 ulation of ganymede’s magnetosphere, J. Geophys. Res. 121. doi:
658 10.1002/2015JA021997.
- 659 [5] Y. Chen, G. Tóth, P. Cassak, X. Jia, T. I. Gombosi, J. Slavin,
660 S. Markidis, B. Peng, Global three-dimensional simulation of earth’s
661 dayside reconnection using a two-way coupled magnetohydrodynamics
662 with embedded particle-in-cell model: initial results, J. Geophys. Res.
663 122 (2017) 10318. doi:10.1002/2017JA024186.
- 664 [6] Y. Chen, G. Tóth, X. Jia, J. A. Slavin, W. Sun, S. Markidis, T. I.
665 Gombosi, J. M. Raines, Studying Dawn-Dusk Asymmetries of Mercury’s
666 Magnetotail Using MHD-EPIC Simulations, Journal of Geophysical Re-
667 search: Space Physics 124 (11) (2019) 8954–8973. arXiv:1904.06753,
668 doi:10.1029/2019JA026840.
- 669 [7] H. Zhou, G. Tóth, X. Jia, Y. Chen, S. Markidis, Embedded Kinetic Sim-
670 ulation of Ganymede’s Magnetosphere: Improvements and Inferences,
671 Journal of Geophysical Research: Space Physics 124 (7) (2019) 5441–
672 5460. doi:10.1029/2019JA026643.

- 673 [8] Y. Shou, V. Tenishev, Y. Chen, G. Toth, N. Ganushkina, Magne-
674 tohydrodynamic with Adaptively Embedded Particle-in-Cell model:
675 MHD-AEPIC, *Journal of Computational Physics* 446 (2021) 110656.
676 doi:<https://doi.org/10.1016/j.jcp.2021.110656>.
677 URL [https://www.sciencedirect.com/science/article/pii/
678 S0021999121005519](https://www.sciencedirect.com/science/article/pii/S0021999121005519)
- 679 [9] W. Zhang, A. Almgren, V. Beckner, J. Bell, J. Blaschke, C. Chan,
680 M. Day, B. Friesen, K. Gott, D. Graves, M. Katz, A. Myers, T. Nguyen,
681 A. Nonaka, M. Rosso, S. Williams, M. Zingale, AMReX: a framework
682 for block-structured adaptive mesh refinement, *Journal of Open Source
683 Software* 4 (37) (2019) 1370. doi:[10.21105/joss.01370](https://doi.org/10.21105/joss.01370).
684 URL <https://doi.org/10.21105/joss.01370>
- 685 [10] W. Zhang, A. Myers, K. Gott, A. Almgren, J. Bell, AMReX: Block-
686 structured adaptive mesh refinement for multiphysics applications, *Inter-
687 national Journal of High Performance Computing Applications* 0 (0)
688 (2021) 1–19. doi:[10.1177/10943420211022811](https://doi.org/10.1177/10943420211022811).
- 689 [11] Y. Chen, G. Tóth, Gauss’s law satisfying energy-conserving semi-
690 implicit particle-in-cell method, *J. Comput. Phys.* 386 (2019) 632.
691 doi:[10.1016/j.jcp.2019.02.032](https://doi.org/10.1016/j.jcp.2019.02.032).
- 692 [12] J. U. Brackbill, D. W. Forslund, An implicit method for electromagnetic
693 plasma simulation in two dimensions, *Journal of Computational Physics*
694 46 (2) (1982) 271–308. doi:[10.1016/0021-9991\(82\)90016-X](https://doi.org/10.1016/0021-9991(82)90016-X).
- 695 [13] G. Lapenta, Exactly energy conserving semi-implicit particle in cell for-
696 mulation, *J. Comput. Phys.* 334 (2017) 349. doi:[10.1016/j.jcp.2017.
697 01.002](https://doi.org/10.1016/j.jcp.2017.01.002).
- 698 [14] K. Fujimoto, S. Machida, Electromagnetic full particle code with adap-
699 tive mesh refinement technique: Application to the current sheet evolu-
700 tion, *Journal of Computational Physics* 214 (2) (2006) 550–566. doi:
701 [10.1016/j.jcp.2005.10.003](https://doi.org/10.1016/j.jcp.2005.10.003).
- 702 [15] K. Fujimoto, A new electromagnetic particle-in-cell model with adaptive
703 mesh refinement for high-performance parallel computation, *Journal of
704 Computational Physics* 230 (23) (2011) 8508–8526. doi:[10.1016/j.jcp.2011.08.003](https://doi.org/10.1016/j.jcp.2011.08.003).

- 705 jcp.2011.08.002.
706 URL <http://dx.doi.org/10.1016/j.jcp.2011.08.002>
- 707 [16] G. Lapenta, Particle rezoning for multidimensional kinetic particle-in-
708 cell simulations, *Journal of Computational Physics* 181 (1) (2002) 317–
709 337. doi:10.1006/jcph.2002.7126.
- 710 [17] M. Vranic, T. Grismayer, J. L. Martins, R. A. Fonseca, L. O. Silva,
711 Particle merging algorithm for PIC codes, *Computer Physics Commu-
712 nications* 191 (1) (2015) 65–73. arXiv:1411.2248, doi:10.1016/j.
713 cpc.2015.01.020.
714 URL <http://dx.doi.org/10.1016/j.cpc.2015.01.020>
- 715 [18] F. Assous, T. Pougeard Dulimbert, J. Segré, A new method for coalesc-
716 ing particles in PIC codes, *Journal of Computational Physics* 187 (2)
717 (2003) 550–571. doi:10.1016/S0021-9991(03)00124-4.
- 718 [19] D. R. Welch, T. C. Genoni, R. E. Clark, D. V. Rose, Adaptive particle
719 management in a particle-in-cell code, *Journal of Computational Physics*
720 227 (1) (2007) 143–155. doi:10.1016/j.jcp.2007.07.015.
- 721 [20] M. Pfeiffer, A. Mirza, C. D. Munz, S. Fasoulas, Two statistical particle
722 split and merge methods for Particle-in-Cell codes, *Computer Physics
723 Communications* 191 (1) (2015) 9–24. doi:10.1016/j.cpc.2015.01.
724 010.
725 URL <http://dx.doi.org/10.1016/j.cpc.2015.01.010>
- 726 [21] D. Faghihi, V. Carey, C. Michoski, R. Hager, S. Janhunen, C. S. Chang,
727 R. D. Moser, Moment preserving constrained resampling with applica-
728 tions to particle-in-cell methods, *Journal of Computational Physics* 409
729 (2020) 109317. doi:10.1016/j.jcp.2020.109317.
730 URL <https://doi.org/10.1016/j.jcp.2020.109317>
- 731 [22] J. Teunissen, U. Ebert, Controlling the weights of simulation particles:
732 Adaptive particle management using k-d trees, *Journal of Computa-
733 tional Physics* 259 (2014) 318–330. arXiv:1301.1552, doi:10.1016/j.
734 jcp.2013.12.005.
735 URL <http://dx.doi.org/10.1016/j.jcp.2013.12.005>
- 736 [23] P. T. Luu, T. Tückmantel, A. Pukhov, Voronoi particle merging al-
737 gorithm for PIC codes, *Computer Physics Communications* 202 (2016)

- 738 165–174. arXiv:1504.00636, doi:10.1016/j.cpc.2016.01.009.
739 URL <http://dx.doi.org/10.1016/j.cpc.2016.01.009>
- 740 [24] Y. Chen, G. Tóth, H. Hietala, S. K. Vines, Y. Zou, Y. Nishimura, M. V.
741 Silveira, Z. Guo, Y. Lin, S. Markidis, Magnetohydrodynamic with em-
742 bedded particle-in-cell simulation of the Geospace Environment Mod-
743 eling dayside kinetic processes challenge event, Earth and Space Sci-
744 encedoi:10.1029/2020ea001331.
- 745 [25] X. Wang, Y. Chen, G. Tóth, Global magnetohydrodynamic magneto-
746 sphere simulation with an adaptively embedded particle-in-cell model,
747 Earth and Space Science Open Archive (2021) 21doi:10.1002/essoar.
748 10508044.1.
749 URL <https://doi.org/10.1002/essoar.10508044.1>
- 750 [26] G. Lapenta, J. U. Brackbill, Nonlinear evolution of the lower hybrid
751 drift instability: Current sheet thinning and kinking, Phys. Plasmas 9
752 (2002) 1544.
- 753 [27] H. Wang, A. Ridley, et al., Substorm onset dynamics in the magneto-
754 tail as derived from joint TC-1 and Cluster data analysis, J. Geophys.
755 Res.Submitted.