

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

Yuxi Chen<sup>1a</sup>, Gábor Tóth<sup>a</sup>, Hongyang Zhou<sup>a</sup>, Xiantong Wang<sup>a</sup>

<sup>a</sup> *Center for Space Environment Modeling, University of Michigan, Ann Arbor, Michigan 48109, 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 a box, which is not always feasible for covering the whole physical structure of interest. The FLEKS, which is a new PIC code and allows a PIC region of any shape, is designed to break this restriction. It 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 simulation 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 with the time-dependent electromagnetic that is obtained from a global simulation.

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

---

<sup>1</sup>Corresponding 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, a PIC region is restricted to be a box, which is not always efficient or feasible 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 (MHD) 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 FLEKS, which is designed and implemented as the PIC component of the MHD-AEPIC model. FLEKS shares some similarities with the work by Shou et al. [8], but FLEKS provides a more flexible grid design. FLEKS uses the parallel data structures provided by the AMReX library [9, 10]. The grid of FLEKS has to be uniform and Cartesian so far, but the

38 active PIC region is not limited to be a box anymore since the PIC cells can  
39 be turned off to fit the region of interest. Furthermore, FLEKS also supports  
40 switching on or switching off grid cells dynamically for MHD-AEPIC simu-  
41 lations. The Gauss’s law satisfying energy-conserving semi-implicit method  
42 (GL-ECSIM) [11] is implemented as the base PIC solver. The time step of the  
43 semi-implicit PIC methods is usually limited by a Courant–Friedrichs–Lewy  
44 (CFL) condition in order to be accurate [12]. Since the plasma maximum  
45 characteristic speed may change significantly during a long MHD-AEPIC  
46 simulation, the simulation will be either too slow or inaccurate with a fixed  
47 time step. To keep the simulation efficient and accurate at the same time,  
48 FLEKS uses an adaptive time-stepping algorithm, which still satisfies the  
49 requirement of the energy-conserving semi-implicit method (ECSIM) [13] to  
50 keep energy conservation. Section 2 describes the adaptive grid and temporal  
51 discretization of FLEKS.

52 The statistical noise of macro-particles is a primary numerical error source  
53 of typical PIC simulations. Dozens to hundreds of particles per cell are usu-  
54 ally used to achieve a balance between the accuracy and computational cost.  
55 Since there are much more macro-particles than the grid cells in a kinetic  
56 PIC simulation, particle-related calculations, such as updating particle posi-  
57 tions and velocities, usually dominate the total computational time. What is  
58 worse is that a massive parallel simulation can be further slowed down grad-  
59 ually due to the imbalance of macro-particle numbers among the CPU or  
60 GPU cores. On the other hand, the decreasing of the macro-particle number  
61 in some cells increases the statistical noise and reduces the accuracy. A par-  
62 ticle resampling algorithm that is able to control the macro-particle number  
63 per cell is crucial for improving both the simulation efficiency and accuracy.  
64 More macro-particles should be populated into the cells, which contain fewer  
65 macro-particles than required, to represent the plasma velocity-space distri-  
66 butions more accurately. This goal is usually achieved by splitting particles.  
67 In the cells with more macro-particles than the threshold, a particle merging  
68 algorithm should be applied to reduce the macro-particle number to speed  
69 up simulations. A particle resampling algorithm is even much more crucial  
70 for a PIC code with adaptive mesh refinement, where the motion of macro-  
71 particles between the coarse and fine cells alters the macro-particle number  
72 per cell dramatically [14, 15].

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

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

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

106 When FLEKS is applied to simulate global phenomena as part of the  
107 MHD-AEPIC model, the simulation time is usually long enough so that the  
108 local macro-particle numbers may reduce or increase significantly. Splitting  
109 particles in the low particle number cells improves statistical representation  
110 and reduces noise. Merging particles alleviates load imbalance and speeds  
111 up simulations. Our particle resampling algorithms are described in section  
112 4.

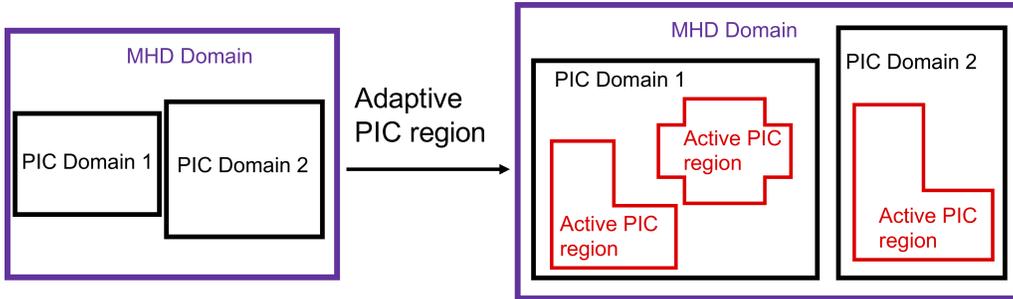


Figure 1: A schematic shows the improvement of the MHD-AEPIC (right) model from the MHD-EPIC (left) model.

113 Tracking the motion of macro-particles is useful for learning the trajec-  
 114 tory and energization of plasma, so FLEKS provides a parallel test particle  
 115 module to follow the motion of macro-particles and save the massive particle  
 116 trajectory data to disk. The test particle module can be used either as a  
 117 module of the PIC code, or as an independent component to directly couple  
 118 to the MHD model. Section 5 describes the implementation details of the  
 119 test 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 shows  
 124 the 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 have  
 130 developed new features to make it more flexible to use. It supports multiple  
 131 independent PIC domains to cover several regions of interest [4], and it also  
 132 allows rotating a PIC box domain to align with the simulation objective  
 133 [24]. However, a box is not always feasible or efficient to cover a 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 a  
137 large portion of cells, where the kinetic effects are not important, to slow  
138 down the simulation. A flexible grid that allows applying an active PIC,  
139 which is in the shape of a paraboloid, to fit the magnetopause would solve  
140 the problems. A dynamically adaptive grid is also useful to improve the  
141 efficiency of some simulations. For instance, the near-Earth X-line may move  
142 from the inner magnetotail to the middle or even far magnetotail [25], and  
143 an adaptive grid that only covers the environment around the X-line is much  
144 more efficient compared to a large PIC box that covers the whole magnetotail.  
145 The MHD-AEPIC model is designed to solve these problems, and FLEKS  
146 is the key component. Figure 1 is a cartoon that shows the concept of the  
147 MHD-AEPIC model.

148 FLEKS still requires the shape of a PIC domain to be a box, and the  
149 grid has to be uniform. But it allows switching off part of the cells to fit a  
150 region of any shape. The most straightforward approach is using a bit-wise  
151 array to switch on/off each cell. However, we make the algorithm a bit more  
152 sophisticated. We divide the whole PIC domain into patches (Figure 2(a)).  
153 Each patch contains  $N$  cells in each direction, and one can turn on or turn  
154 off each patch. The patch size  $N$  is required to be larger or equal to 2.  
155 FLEKS does not allow switching on/off each cell independently ( $N=1$ ) for the  
156 following reason. FLEKS requires two ghost cell layers. If  $N=1$ , the boundary  
157 ghost cells of an active region may overlap with the physical cells of another  
158 active region, and hence introduces more difficulties to handle the boundary  
159 ghost cells. A large patch size also benefits the coupling efficiency. In MHD-  
160 AEPIC simulations, the fluid model controls the status of the patches based  
161 on either geometric or physical criteria [25]. The fluid model passes the bit-  
162 wise patch status array to FLEKS through the Message Passing Interface  
163 (MPI), and the size of this array is reduced significantly with a large patch  
164 size  $N$ . In this paper, we use the word ‘active’ to describe the patches or cells  
165 that are switched on. The active cells do not have to be connected, and the  
166 boundary ghost cells of the active regions are filled in with the information  
167 obtained from the fluid model [1]. Figure 2(a) shows an example that contains  
168 two separated active regions.

169 FLEKS uses the data structures provided by the AMReX library to store  
170 the fields and also the particles. After the patch status array is obtained from  
171 the fluid model, FLEKS uses the functions provided by the AMReX library  
172 to divide the active regions into blocks. AMReX does not require all the  
173 blocks to have the same size. We note that the patch and the block are two

174 independent concepts. The patches are only used to activate or deactivate  
175 cells. For example, the 'L' shape active region in Figure 2(a) consists of 3  
176 patches and it can be divided into 2 blocks (Figure 2(b)).

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

184 FLEKS has two ghost cell layers, but the outer layer is only used to receive  
185 and store the magnetic fields, which are necessary for calculating currents on  
186 the nodes of the inner ghost cell layer from  $\vec{J} = \nabla \times \vec{B}$ . The currents are used  
187 to generate particles with correct velocities in the inner layer ghost cells. To  
188 simplify the description, we ignore the outer layer in Figure 2(c) and also  
189 in the rest of the paper unless otherwise specified. The principle of setting  
190 boundary conditions of the electromagnetic fields and the particles is still the  
191 same as the MHD-EPIC coupling algorithm [1]. However, the non-box shape  
192 of an active region introduces some extra implementation difficulties. They  
193 are three types of ghost cells for a block: the internal ghost cells (blue cells  
194 in Figure 2(c)), the exclusive boundary ghost cells (gray cells in Figure 2(c))  
195 and the shared boundary ghost cells (cyan cells in Figure 2(c)). The internal  
196 ghost cells are not boundary cells, and there is no need to apply boundary  
197 conditions. The exclusive boundary ghost cells are not overlapped with any  
198 cells of the neighboring blocks, and they should be filled in with new macro-  
199 particles as the particle boundary condition. The shared boundary ghost  
200 cells are overlapped with the boundary ghost cells of the neighboring blocks,  
201 and only one of these blocks should generate boundary particles. Here is the  
202 algorithm to choose the block for populating new particles. The first step is  
203 to distinguish the boundary ghost cells from the internal ghost cells. Then,  
204 for each boundary ghost cell, either the exclusive type or the shared type,  
205 we loop through its at most 26 neighboring cells (3D) in a fixed order (we  
206 choose to loop through all the face neighbors first, then the edge neighbors,  
207 and finally the corner neighbors), skip the cells does not exist and find out  
208 the first neighboring cell that is either physical cell or internal ghost cell. If  
209 this neighboring cell is inside the physical domain of this block, this block  
210 should generate particles inside this boundary ghost cell. For example, in  
211 Figure 2(c)), C1 and C3 are overlapped with each other. We loop through

212 the neighboring cells of C1 (C3) and find C2 (C4) is its first neighboring cell  
 213 that is either a physical cell or an internal ghost cell, and block-1 (block-  
 214 2) should (not) generate particles in C1 (C3) since C2 (C4) is (not) inside  
 215 block-1 (block-2).

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

### 221 3. Adaptive time-stepping

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

231 Our adaptive time-stepping algorithm is summarized in Figure 3. At the  
 232 end of one cycle, Both the electromagnetic fields and the particle velocities are  
 233 at time stage  $t^n$ , and the particle locations are at the staggered stage  $t^{n+1/2}$ .  
 234 The difference between  $t^{n+1/2}$  and  $t^n$  is  $t^{n+1/2} - t^n = \Delta t^n/2$ . The maximum  
 235 speed  $v_{rms}$  can be obtained with the particle velocities at  $t^n$ , and a new time  
 236 step  $\Delta t^{n+1}$  can be calculated from  $\Delta t^{n+1} = CFL \cdot \Delta x/v_{rms}$ . However, during  
 237 the next cycle to update the electromagnetic fields and particle velocities  
 238 from  $t^n$  to  $t^{n+1}$ , the time step should be  $t^n$  instead of  $t^{n+1}$ , so that the  
 239 particle location  $X^{n+1/2}$  is still at the middle of  $t^n$  and  $t^{n+1}$  and the energy  
 240 conservation property of ECSIM is preserved. In order to adjust the time  
 241 step for the next cycle, we use the time step  $\Delta t^n/2 + \Delta t^{n+1}/2$  for updating  
 242 the particle location from  $X^{n+1/2}$  to  $X^{n+3/2}$ . Since the velocity  $V^{n+1}$  is not  
 243 at the middle of  $X^{n+1/2}$  of  $X^{n+3/2}$  anymore, the second-order accuracy of  
 244 updating particle locations is not satisfied. In practice, the speed  $v_{rms}$  varies  
 245 gradually and the difference between  $\Delta t^n$  and  $\Delta t^{n+1}$  is small, so the loss of  
 246 accuracy is negligible.

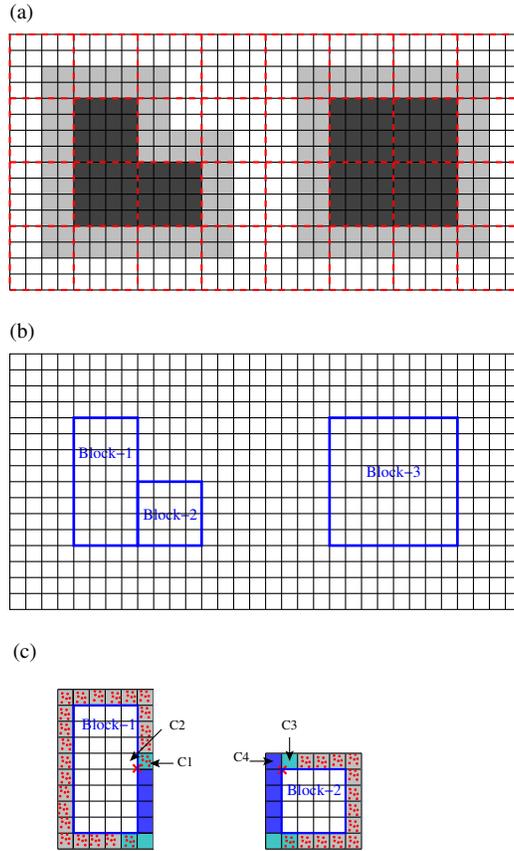


Figure 2: The black lines represent the cells of a PIC domain. The red dashed lines in (a) show the patches, and one patch contains  $4 \times 4$  cells in this example. In (a), the active patches/cells are colored by dark gray, and light gray area represents the ghost cells of the active PIC regions. (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.

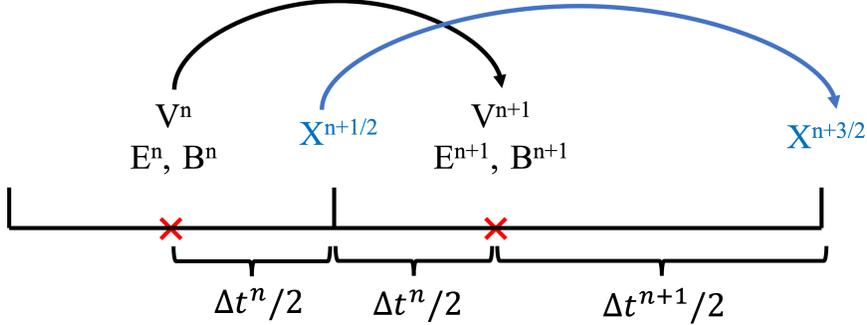


Figure 3: The adaptive temporal discretization.

247 **4. Particle resampling**

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

259 *4.1. Particle splitting*

260 Our particle splitting algorithm is essentially the same as the one in-  
 261 troduced by Lapenta [16], in which one particle is split into two children  
 262 particles. The children particles have the same velocity as their parent par-  
 263 ticle, but their locations are oppositely displaced slightly along the velocity  
 264 direction. By displacing the new particles along the velocity direction, the  
 265 orbits of the new particles are still close to the orbit of the old particle.

266 The particle splitting will be triggered for the cells with ppc less than the  
 267 splitting threshold, which is 80% of the initial ppc by default, and we will

268 use this number for all the simulations presented in section 6. Initially, the  
269 particles that are close to each other have similar weights, but the weights  
270 may become quite different later due to the transport of particles and also the  
271 particle splitting and merging. For each cell, we choose to split the heaviest  
272  $N$  particles to minimize the particle weight variance, where  $N$  is the difference  
273 between the current ppc and the splitting threshold.

#### 274 *4.2. Particle merging*

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

283 Our particle merging algorithm consists of two steps: (1) selecting 6 par-  
284 ticles that are close to each other in the phase space, and (2) merging these  
285 6 particles into 5. In the following subsections, we will describe the merging  
286 step first, and what follows is the selecting step.

##### 287 *4.2.1. Merging particles*

288 Once the 6 old particles for merging have been obtained from the select-  
289 ing step, we use 5 new particles to replace them. The replacement should  
290 not alter the original phase space distribution significantly. However, it is  
291 not trivial to quantitatively measure the change of the phase space. The con-  
292 servation of total mass, momentum, and energy can be used as a guidance  
293 and indicator of preserving phase space structure. However, satisfying the  
294 conservation property is not good enough, it is still possible that the new  
295 particle set occupies a very different velocity space volume than the old set.  
296 Previous methods [26, 17] usually generate new particles with velocities that  
297 are different from the velocities of the old particles, and extra actions are  
298 usually required to ensure the new particles are not too far away from the  
299 old ones in the velocity space. To avoid this difficulty, we choose to delete  
300 one of the 6 old particles and distribute its mass to the rest 5 particles under  
301 the constrain of conserving total mass, momentum, and energy. The weights  
302 of these 5 particles change, but their velocities inherit from the old ones, so

303 the new particle set occupies almost the same phase space volume as the old  
 304 set.

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

$$w_t = \sum_{i=1}^{N_{old}} w_i \quad (1)$$

$$\mathbf{p}_t = \sum_i^{N_{old}} w_i \mathbf{v}_i \quad (2)$$

$$e_t = \sum_{i=1}^{N_{old}} \frac{1}{2} w_i |\mathbf{v}_i|^2, \quad (3)$$

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

$$w_t = \sum_{i=1}^{N_{new}} w_{i,new} \quad (4)$$

$$\mathbf{p}_t = \sum_i^{N_{new}} w_{i,new} \mathbf{v}_i \quad (5)$$

$$e_t = \sum_{i=1}^{N_{new}} \frac{1}{2} w_{i,new} |\mathbf{v}_i|^2 \quad (6)$$

$$w_{i,new} > 0, \quad (7)$$

310 where we choose  $N_{new} = 5$  since there are 5 quantities to conserve, the  
 311 velocities  $\mathbf{v}_i$  are known, and the weights  $w_{i,new}$  are the unknowns of the  
 312 linear equations under the constrain of  $w_{i,new} > 0$ . We note that if the  
 313 solution does not satisfy the constrain, we skip this merging. The procedures  
 314 described above need to calculate the distance in the 6D phase space, and  
 315 its calculation method will be described in the next subsection.

316 By deleting the lighter particle from a pair that is closest to each other in  
 317 the phase space, it is likely its heavier neighbor will gain most of the weight  
 318 and the rest 4 particles adjust their weights relatively slightly. By inheriting  
 319 the velocities and locations from the old particles, the new particles occupy  
 320 almost the same phase space volume as the old particles (Figure 4(d) and  
 321 (e)), so there is no room for the phase space structure to change dramatically.  
 322 Compared to the schemes that allow choosing new particle velocity with fewer  
 323 restrictions [17], our merging algorithm is less efficient to reduce the particle  
 324 number because (1) the new particle set still contains 5 particles, and (2) the  
 325 merging may fail when the constrain  $w_{i,new} > 0$  can not be satisfied. If it  
 326 is required, our algorithm can be easily modified to use a  $N_{old}$  that is larger  
 327 than 6 by deleting  $N_{old} - 5$  particles. However, as it can be seen from the  
 328 following numerical test section, merging 6 particles into 5 is already efficient  
 329 enough for our typical applications.

330 *4.2.2. Selecting particles*

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

348 The particle merging algorithm needs calculating the distance in the 6D  
 349 phase space, and it is defined as:

$$l_{6D} = \Delta l_{3D} / \Delta x + c_1 * \Delta v / v_{th} \quad (8)$$

$$v_{th} = \frac{1}{N_p} \sum_{i=0}^{N_p} |\mathbf{v}_i - \mathbf{v}|^2 \quad (9)$$

350 where  $\Delta l_{3D}$  is the spatial distance,  $\Delta v$  is the distance in the velocity space,  
 351  $\Delta x$  is the simulation cell size,  $N_p$  is the particle number of the corresponding  
 352 cell,  $\mathbf{v}$  it the bulk velocity, and  $v_{th}$  is the thermal velocity. The constant  $c_1$   
 353 determines the relative importance of the spatial distance  $\Delta l_{3D}$  and the veloc-  
 354 ity space distance  $\Delta v$ . We choose  $c_1 = 2$  based on our numerical experiment  
 355 experience.

356 At the end of each computational cycle, the following scheme will be  
 357 executed to select particles for merging if the ppc of a cell is larger than the  
 358 merging threshold, which is 1.5 times of the initial ppc by default.

- 359 • Step-1: binning particles in the velocity space. For each spatial cell  
 360 (Figure 4(a)), we create a grid in the velocity space ranging from  $(v_x -$   
 361  $v_{th}, v_y - v_{th}, v_z - v_{th})$  to  $(v_x + v_{th}, v_y + v_{th}, v_z + v_{th})$ , and assign particles  
 362 to velocity bins (Figure 4(b)). The velocity grid is divided into  $n_{bin} =$

363  $\lceil 0.8 * N_p^{1/3} \rceil$  bins in each direction, where  $N_p$  is the current ppc of the  
 364 cell and the constant 0.8 is chosen based on numerical experiments.  
 365 We note that each bin contains a buffer region (Figure 4(c)), and the  
 366 particles in the buffer region may also belong to other bins. We use  
 367 1/8 of the velocity space bin size as the width of the buffer region  
 368 (Figure 4(c)). Due to the existence of the buffer region, one particle  
 369 may belong to multiple bins, but it can only be selected for merging at  
 370 most once during one cycle.

- 371 • Step-2: selecting particles from a bin. If there are more than 6 particles  
 372 inside a bin, including the buffer region, we choose a cluster of 6 parti-  
 373 cles from them. For each velocity bin, we calculate the velocity center  
 374 of the associated particles (black cross in Figure 4(c)), and find the 6  
 375 particles closest to the center in the 3D velocity space. If a particle in  
 376 the buffer region has been selected for merging by a neighboring bin,  
 377 this particle should not be selected again.

- 378 • Step-3: limiting the 6D distance. The previous step selects particles  
 379 only based on the distance in the velocity space. This step ensures the  
 380 selected particles are also close to each other in the 6D phase space.  
 381 We find the 6D center (blue cross in Figure 4(d)) of these 6 particles,  
 382 and the 6D distance  $l_{6D}$  of all the 6 particles to the center should be  
 383 less than 0.6. Again, the constant 0.6 is chosen based on numerical  
 384 experiments.

- 385 • Step-4: merging 6 particles into 5 with the algorithm described in sec-  
 386 tion 4.2.1.

387 The particle selection method used in step-2 prefers selecting particles in the  
 388 center of a bin. Without applying the buffer region in step-1, the particles  
 389 near the edge of a bin are less likely chosen for merging. On the other hand,  
 390 it is more likely that a bin contains more than 6 particles with a buffer re-  
 391 gion, and hence the merging efficiency is improved. Based on our numerical  
 392 experiments, applying the buffer region does not improve the simulation re-  
 393 sults significantly, but it is still kept by default to avoid the aforementioned  
 394 potential issues.

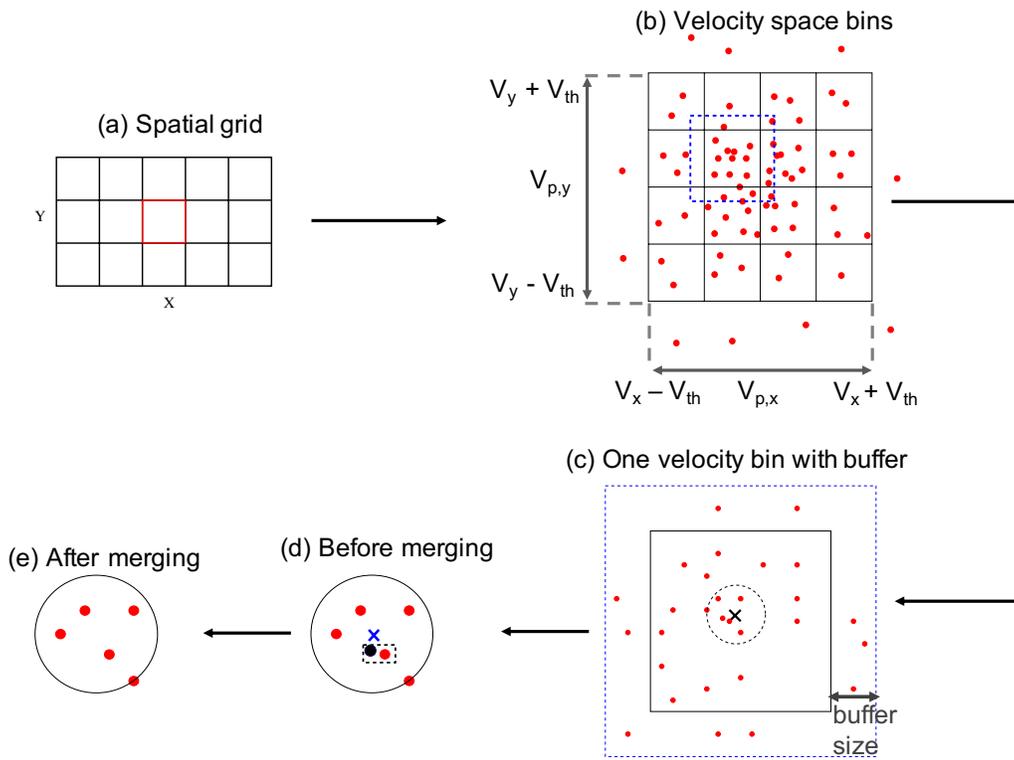


Figure 4: The steps of merging macro-particles.

## 395 5. Test particle module

396 An independent test particle (TP) module is designed to track the mo-  
397 tion of the macro-particles for FLEKS. It can be used either as an auxiliary  
398 component of the PIC algorithm or as an independent component. The TP  
399 module uses the same algorithm to move particles as the GL-ECSIM algo-  
400 rithm. When the TP module is used with the PIC component together, the  
401 TP module shares the same grid layout as the PIC component and uses the  
402 electromagnetic fields calculated by PIC to update test particles. When the  
403 PIC component is turned off, FLEKS becomes a pure test particle code, and  
404 the TP module can directly obtain the grid structure and electromagnetic  
405 fields from the MHD model. Compared to the embedded PIC simulations,  
406 the pure test particle simulations are only one-way coupled, i.e., the MHD  
407 model provides the electromagnetic fields for FLEKS, but there is not any  
408 feedback from FLEKS to the MHD model.

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

## 424 6. Numerical tests

### 425 6.1. *Two-dimensional fast magnetosonic wave propagation with adaptive PIC* 426 *region*

427 We use a two-dimensional (2D) fast magnetosonic wave propagation test  
428 to demonstrate the capability of FLEKS's adaptive grids. The same initial

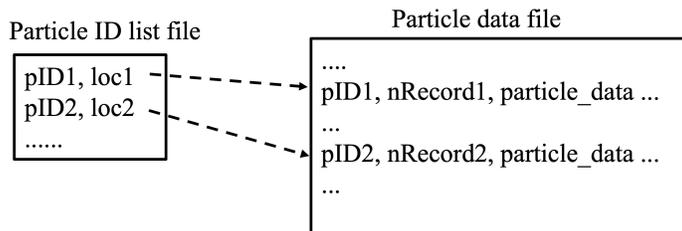


Figure 5: The file structure for storing test particles.

429 condition as what is described in [1] is applied here to produce a propagat-  
 430 ing fast magnetosonic wave. The simulation domain of the MHD code is  
 431  $-160/3 < x < 160/3$  and  $-40 < y < 40$ . Two independent PIC domains  
 432 are used. The left domain in Figure 6 covers the region of  $-40 < x < 0$  and  
 433  $-20 < y < 20$  with a grid resolution of  $\Delta x = \Delta y = 1/16$ . The right domain  
 434 covers the region of  $20 < x < 40$  and  $-10 < y < 10$  with a grid resolution  
 435 of  $\Delta x = \Delta y = 1/8$ . All cells of the right PIC domain are always switched  
 436 on during the simulation. For the left domain, only the cells that satisfy the  
 437 following conditions are switched on:

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

438 where  $r$  is the distance to the center of the PIC domain,  $L_x$  is the length of  
 439 the PIC domain, which is 40 in this case, and  $T$  is the simulate time. The  
 440 central PIC cells ( $r < L_x/10$ ) are always switched on, and the outer shell of  
 441 active PIC cells keep changing during the simulation. A movie that shows  
 442 the adaptation of the active PIC region is provided as an online supplement.  
 443 We note that the simulation parameters for these two PIC domains can be  
 444 specified independently. For example, the cell size is different for these two  
 445 PIC domains as it is described above, and the ion-electron mass ratio  $m_i/m_e$

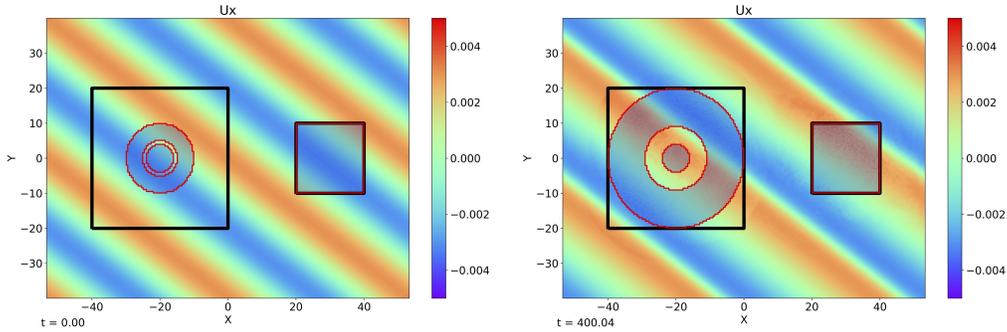


Figure 6: The velocity  $U_x$  of the 2D fast magnetosonic wave test at the beginning (left) and at  $t = 400$  (right). The black rectangles show the area of the PIC grids. Inside each PIC grid, the semi-transparent area, which is enclosed by red lines, represents the active PIC region. Since all PIC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

446 is 25 for the left domain and it is 100 for the right domain. Both PIC domains  
 447 use  $CFL = 0.2$ , and 900 particles per cell (ppc) per species.

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

### 452 *6.2. One-dimensional non-linear magnetosonic wave evolution*

453 The evolution of the magnetosonic wave is non-linear. The wave may  
 454 finally evolve into a shock, where the plasma phase space distributions may  
 455 become non-Maxwellian. So the non-linear evolution of the magnetosonic  
 456 wave simulation is suitable for testing the particle resampling algorithms.

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

462 initial condition of the 1D magnetosonic wave is:

$$\begin{aligned}
B_x(x) &= B_0 \cos(\theta) \\
B_y(x) &= B_0 [\sin(\theta) + \delta \sin(kx - \omega t)] \\
B_z(x) &= 0 \\
u_x(x) &= \delta \sin(\theta) \frac{v_A^2 v_p}{v_p^2 - v_s^2} \sin(kx - \omega t) \\
u_y(x) &= \delta \cos(\theta) \frac{v_A^2}{v_p} \sin(kx - \omega t) \\
u_z(x) &= 0 \\
\rho(x) &= \rho_0 \left[ 1 + \delta \sin(\theta) \frac{v_A^2}{v_p^2 - v_s^2} \sin(kx - \omega t) \right] \\
p(x) &= p_0 \left[ 1 + \gamma \delta \sin(\theta) \frac{v_A^2}{v_p^2 - v_s^2} \sin(kx - \omega t) \right],
\end{aligned} \tag{11}$$

463 where  $\gamma$  is the specific heat ratio,  $v_A = \frac{B_0}{\sqrt{\rho_0}}$  is the Alfvén speed,  $v_s = \sqrt{\frac{\gamma p_0}{\rho_0}}$   
464 is the sound speed, and  $\theta$  is the angle between the wave vector, which is  
465 the x-direction here, and the background magnetic field. The phase speed  
466  $v_p = \omega/k$  is:

$$v_p^2 = \frac{1}{2} (v_A^2 + v_s^2 + [(v_A^2 + v_s^2)^2 - 4v_s^2 v_A^2 \cos^2 \theta]^{1/2}), \tag{12}$$

467 In this paper, we use  $\gamma = 5/3$ ,  $B_0 = 0.1$ ,  $\theta = 30^\circ$ ,  $\rho_0 = 1$ ,  $p_0 = 1e - 4$ ,  
468  $k = 2\pi/\lambda = 2\pi/64$ , and  $\delta = 0.5$ . We note that the perturbation  $\delta = 0.5$   
469 is not small so that the initial conditions, which are obtained from linear theory,  
470 are not accurate. Since the goal of this test is to compare the simulation  
471 results with and without particle resampling, it is suitable and acceptable to  
472 use such a large perturbation.

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

481 of ppc are about 500 and 3140, respectively, for the simulation without ap-  
 482 plying particle resampling. For the simulation with particle resampling, the  
 483 minimum ppc is about 750 and the maximum ppc is about 1360, and these  
 484 numbers are close to the splitting limit  $0.8 * 900 = 720$  and the merging limit  
 485  $1.5 * 900 = 1350$ . It suggests the particle resampling algorithms are effec-  
 486 tive in controlling the particle number. Except for the particle number, the  
 487 physical quantities of these two simulations are very similar to each other.  
 488 The only noticeable different is that the electric field  $E_y$  of the simulation  
 489 with particle resampling is more noisy near  $x = 20$  due to the reduction  
 490 of particle number. Figure 7(b) shows the ion phase space distribution for  
 491 particles between  $x = 21$  and  $x = 21.2$ . The two mass density distributions  
 492 are comparable even though the particle number density is quite different.

493 Figure 8 shows the simulation speed, which represents the number of PIC  
 494 cells that are updated per second per CPU core. For the first 700 cycles, both  
 495 simulations become slower and slower due to the imbalance of the particle  
 496 number per CPU core. Later, the minimum and maximum ppc reach the  
 497 splitting and merging thresholds and the particle splitting algorithms start  
 498 controlling the further change of the minimum and maximum ppc, so the  
 499 simulation speed stops dropping for the simulation with particle resampling.  
 500 At the end of the simulation, the simulation with particle resampling is almost  
 501 twice faster than the one without particle resampling.

### 502 6.3. Two-dimensional double-current-sheet magnetic reconnection

503 Magnetic reconnection is regarded as one of the most important physical  
 504 processes for energy transfer between magnetic field and plasma in the space  
 505 plasma environment, so it is also widely used to benchmark the performance  
 506 of a kinetic plasma modeling code. Here, we use a two-dimensional (2D)  
 507 asymmetric magnetic reconnection problem to test the particle resampling  
 508 algorithms, because the particles distributions near the reconnection site can  
 509 be non-Maxwellian. It is crucial to demonstrate the particle resampling al-  
 510 gorithms preserve the non-Maxwellian distributions.

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

$$B_x(y) = \left( \frac{B_1 + B_2}{2} \right) \left[ \tanh \left( \frac{y + 0.25L_y}{\delta} \right) - \tanh \left( \frac{y - 0.25L_y}{\delta} \right) \right] - B_2, \quad (13)$$

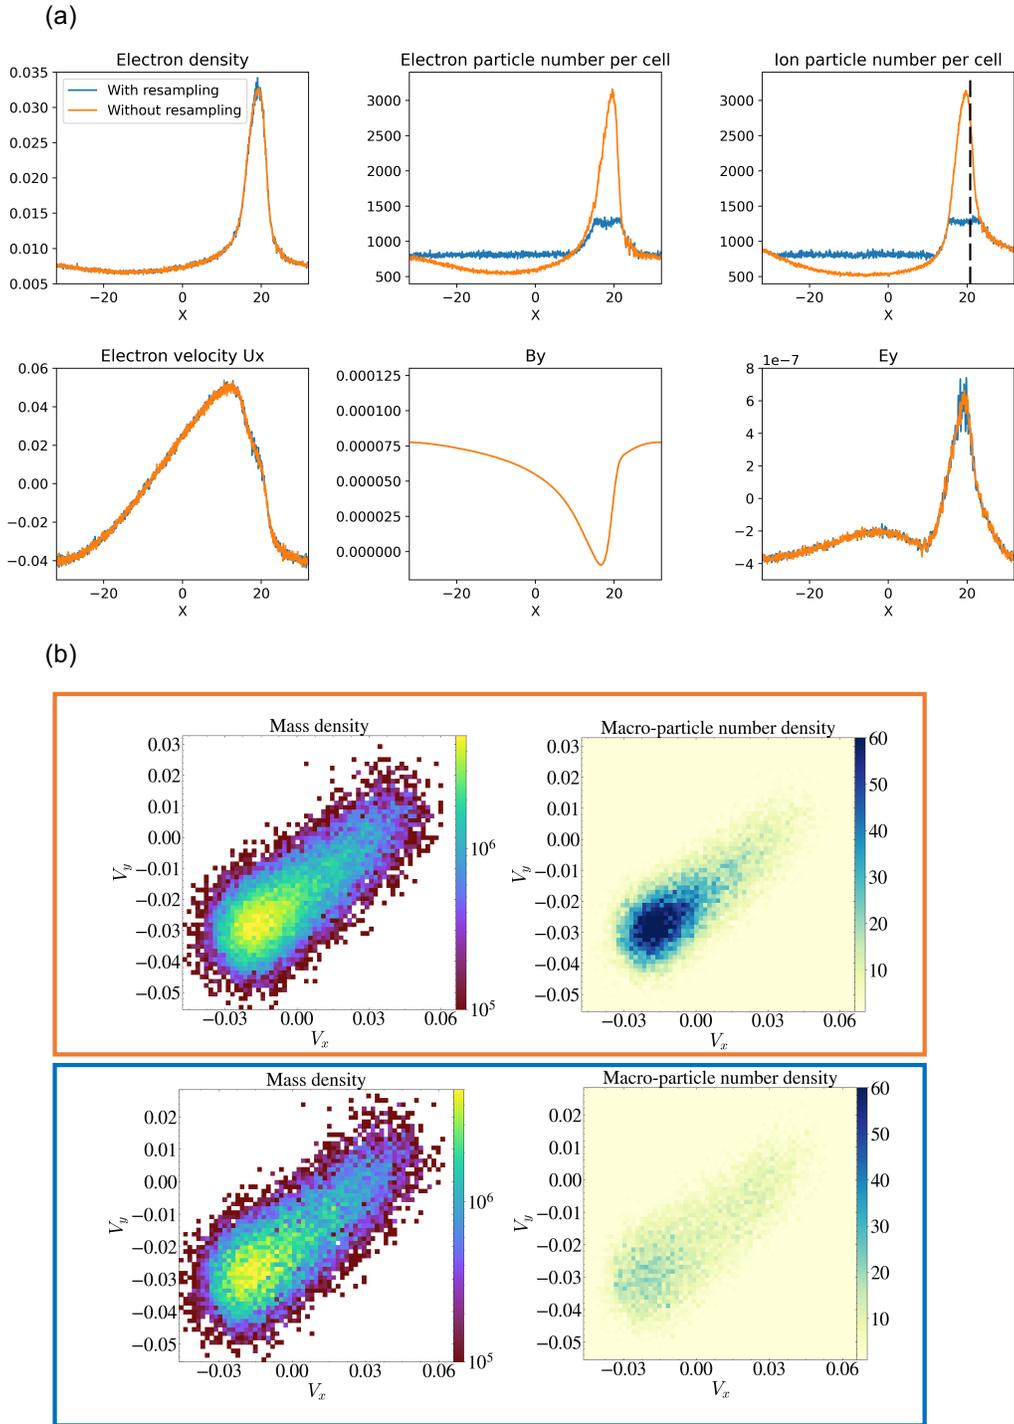


Figure 7: The 1D magnetosonic wave simulation results at  $t = 200$ . (b) shows phase space distributions at  $x = 21$ , where is marked with dashed black line in (a). The upper panel of (b) shows the results without particle resampling, and the lower panel shows the results with particle resampling.

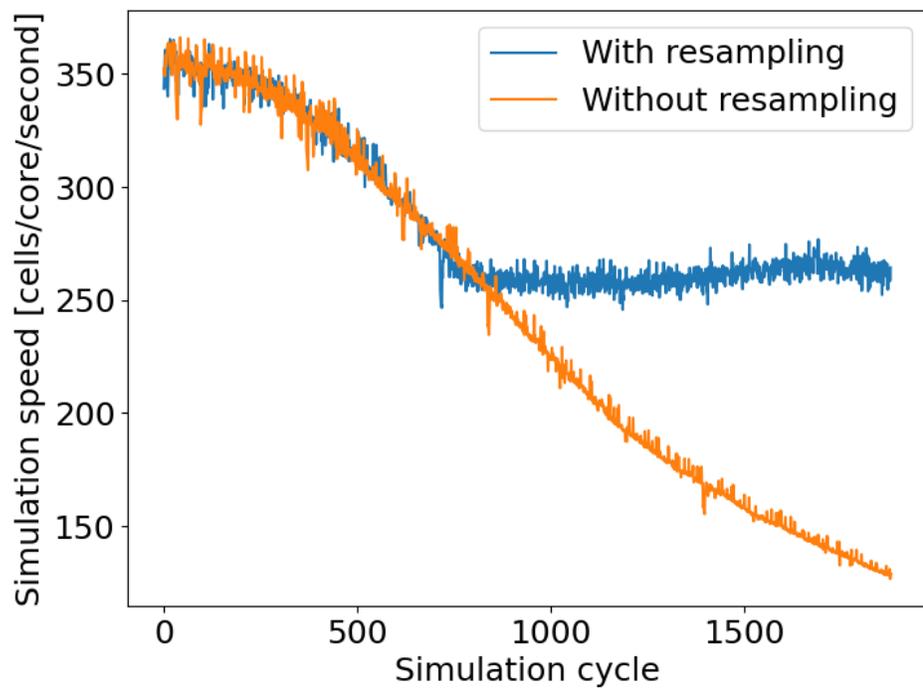


Figure 8: The simulation speed of the 1D magnetosonic wave simulations.

515 where  $B_1 = 1$  and  $B_2 = 2$  are the asymptotic magnetic field amplitudes.  $L_y =$   
516  $32$  is the width of the simulation domain, and the centers of the two current  
517 sheets are at  $y = -8$  and  $y = 8$ , respectively. The plasma pressure is set to  
518 balance the magnetic field pressure  $p_B$ . To mimic the plasma environment  
519 of Earth's magnetopause, the asymptotic plasma beta  $\beta = (p_i + p_e)/p_B$  are  
520  $3.6$  and  $0.15$  on the "1" and "2" side, respectively. The initial pressure ratio  
521 between electrons and ions is  $p_i/p_e = 5$  in the whole simulation domain. The  
522 ion temperature is:

$$T_i(y) = \left( \frac{T_{i,1} + T_{i,2}}{2} \right) \left[ \tanh \left( \frac{y + 0.25L_y}{\delta} \right) - \tanh \left( \frac{y - 0.25L_y}{\delta} \right) \right] + T_{i,2}. \quad (14)$$

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

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

545 Figure 9 and Figure 10 also compare the physical quantities of the sim-  
546 ulations. All simulations show essentially the same structures, including the  
547 off-diagonal electron tensor. It demonstrates the particle resampling algo-  
548 rithms do not introduce any significant artificial effect.

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

#### 570 6.4. Strong and weak scalings

571 3D asymmetric magnetic reconnection simulations are used to test the  
 572 strong and weak scaling of FLEKS. The setup of the 3D test is similar to the  
 573 2D simulation in previous subsection, and it is uniform in the z-direction.  
 574 Since FLEKS uses the parallel field and particle data structures provided by  
 575 AMReX, The scaling results largely depend on the performance of AMReX  
 576 [9, 10]. Figure 12 shows the weak scalings. With  $8^3$  cells per core, the  
 577 performance is still good with up to about 10k cores. With  $16^3$  cells per  
 578 CPU core, it reaches good performance even with 28672 cores. Figure 13  
 579 shows the strong scalings of two problems. The speedup is not too far away  
 580 from the ideal scaling up to about 7k (Figure 13(a)) or 14k (Figure 13(b))  
 581 CPUs for these problems.

#### 582 6.5. Magnetospheric simulations

583 Magnetospheric simulation is the most important application of MHD-  
 584 AEPIC. Here, we show examples of how FLEKS benefits magnetosphere

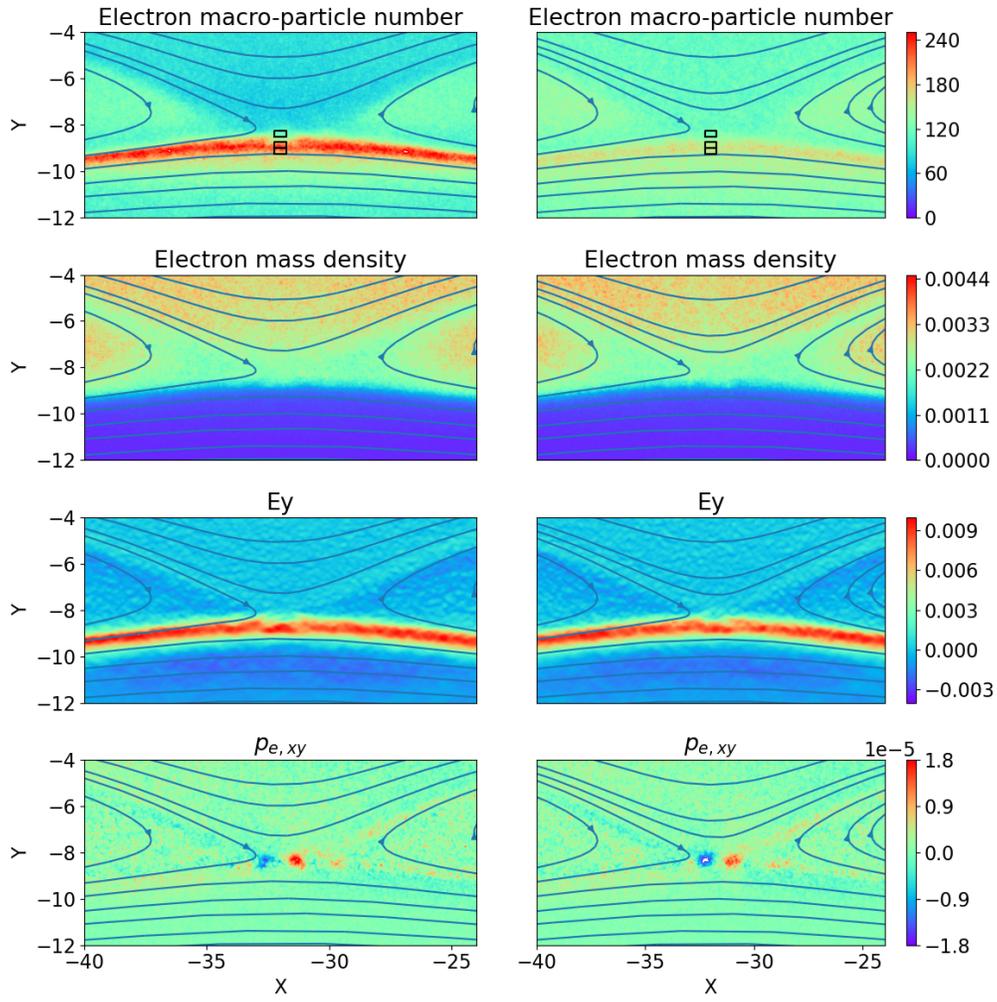


Figure 9: 2D magnetic reconnection results with (right column) or without (left column) particle resampling. The initial particle number per cell is 100.

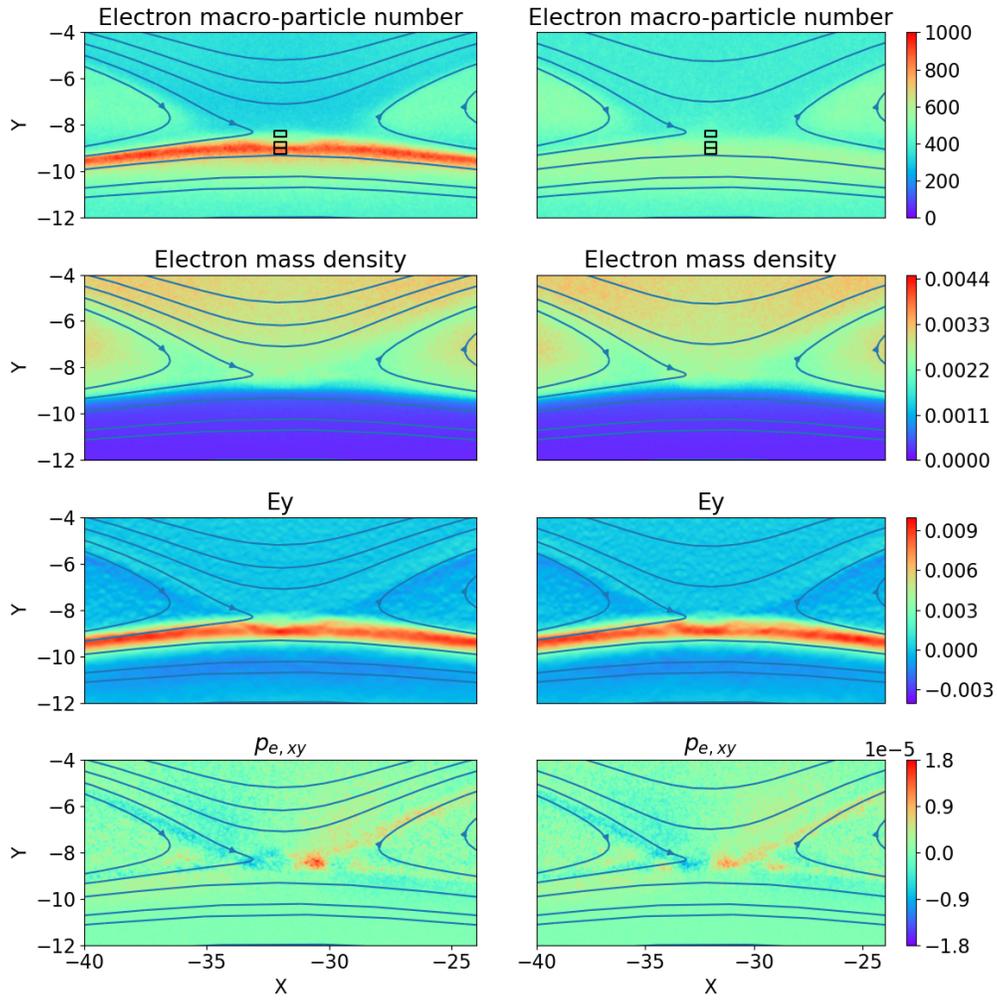


Figure 10: 2D magnetic reconnection results with (right column) or without (left column) particle resampling. The initial particle number per cell is 400.

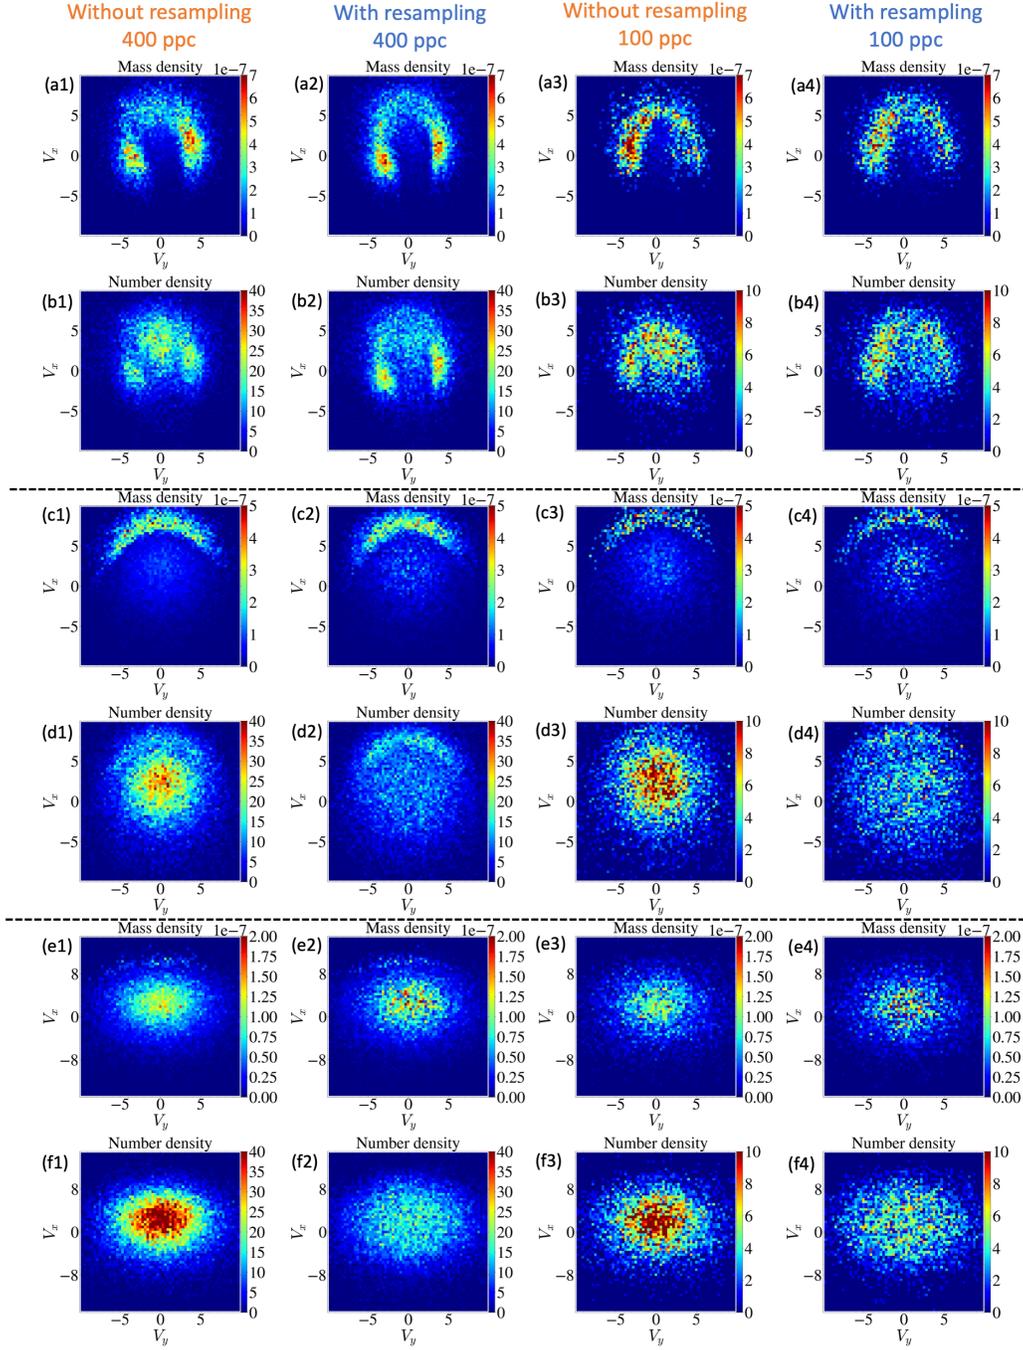


Figure 11: Each column shows the phase space distributions from the same simulations. The first two rows, middle two rows and the last two rows represent the distributions of box-A, box-B and box-C, respectively. From top to bottom, the black rectangles in Figure 9 and Figure 10 show the locations of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

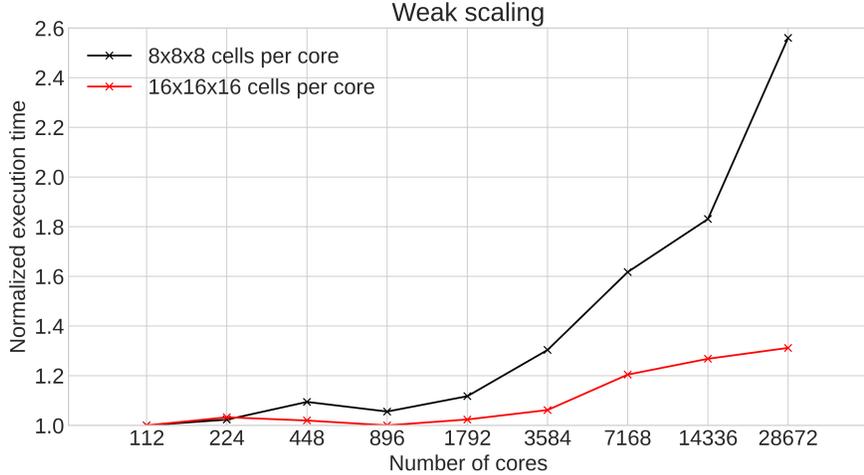


Figure 12: The weak scaling results of FLEKS. The execution time with 112 CPU cores is used as the reference. If the scaling is perfect, the normalized execution time should be 1.0 for all cases.

585 modeling. The global MHD magnetosphere model uses the same setup as  
 586 the simulations presented in [5, 24], but the active PIC region is not limited  
 587 to be a box anymore. In the example of Figure 14, the active PIC region  
 588 can cover the dayside magnetopause, including the dawn-side and dusk-side  
 589 flanks, and also the cusp region at the same time.

590 The test particle module enables us to follow the trajectories of particles  
 591 in the magnetosphere. Figure 15 shows an example of test particles in Earth's  
 592 magnetosphere.

## 593 7. Conclusion

594 In this paper, we introduce a new kinetic code FLEKS, which is designed  
 595 as the kinetic component of the MHD-AEPIC model [8, 25]. To support  
 596 long simulations with varying global configurations, FLEKS allows activating  
 597 or deactivating cells dynamically during a simulation to fit the regions of  
 598 interest. This feature is introduced by Shou et al. [8] first, but FLEKS  
 599 is more flexible that the minimum activation unit is a patch contains  $N$   
 600 ( $N \geq 2$ ) cells in each direction instead of a large block, and FLEKS supports  
 601 multiple independent PIC domains in an MHD-AEPIC simulation. During

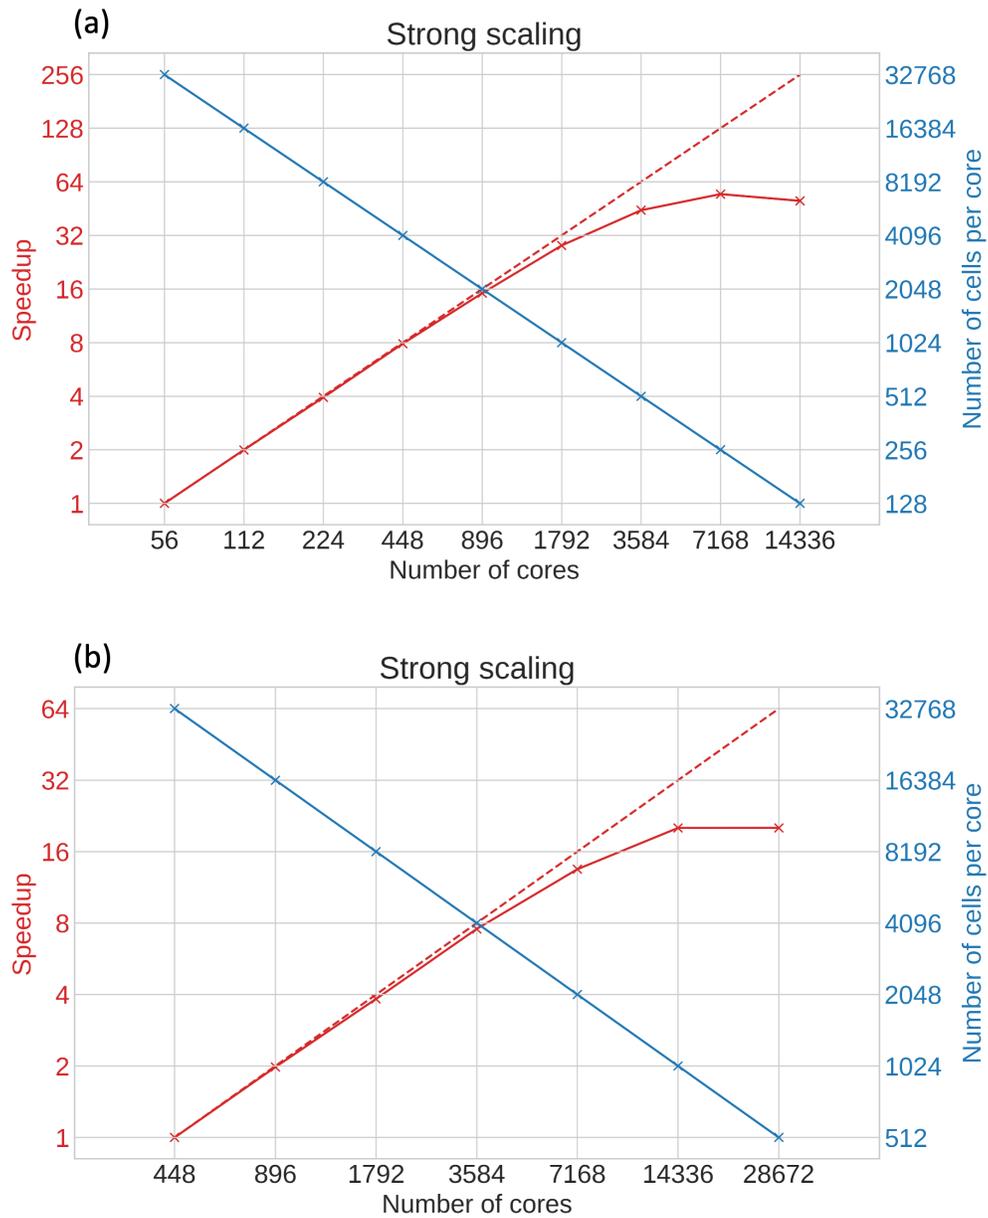


Figure 13: The strong scaling of FLEKS. (a) and (b) show the scalings of simulations with 1.835 million and 14.68 million cells, respectively. The red solid lines represent speedup, and the red dashed lines are perfect speedup. The blue lines show the number of PIC cells per CPU core.

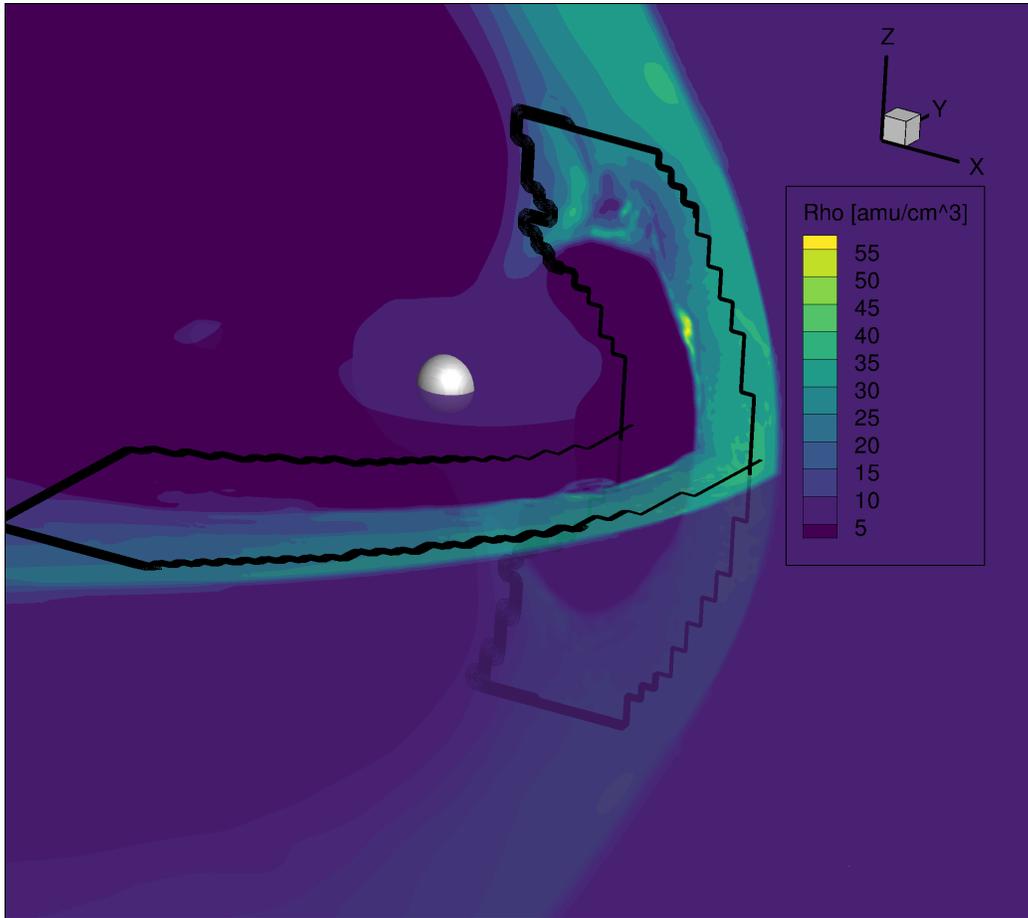


Figure 14: An MHD-AEPIC simulation of Earth's magnetosphere with the dayside magnetopause and the cusps covered by FLEKS. The black lines indicate the edge of the active PIC region.

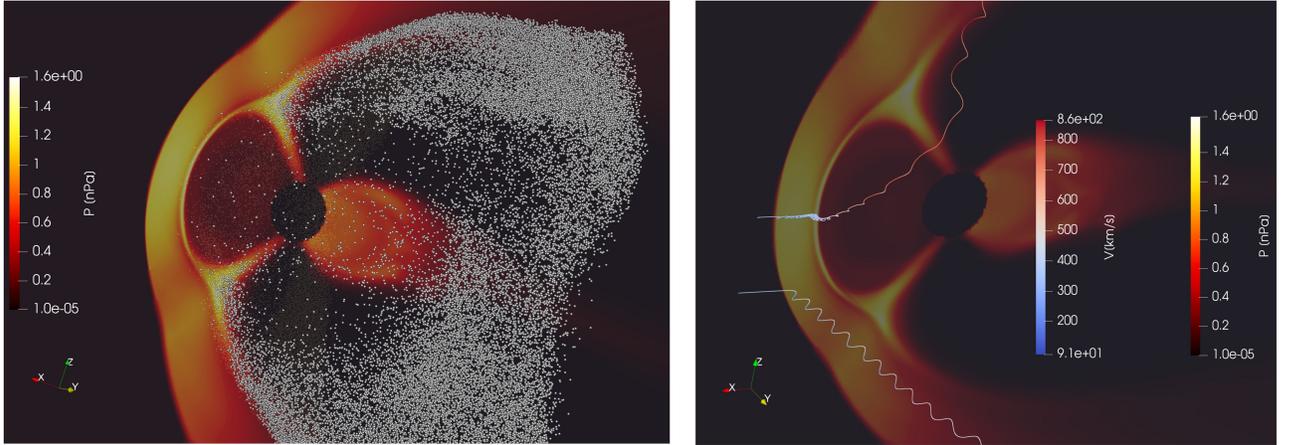


Figure 15: The locations (left) and the example trajectories (right) of test particles.

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

606 Since the number of particle per cell may change dramatically during a  
 607 long simulation and leads to load imbalance and loose of accuracy in the cells  
 608 with fewer particles, a particle splitting and a particle merging algorithms are  
 609 designed to control the change of ppc. The particle merging algorithm selects  
 610 6 particles that are close to each other in the phase space and combines them  
 611 into 5 new particles. The merging conserves the total mass, momentum and  
 612 energy, and it also preserves the phase space structure as much as possible  
 613 by inheriting velocities from the old particles.

614 The test-particle module expands the capability of FLEKS, and provides  
 615 a useful tool for investigating the transport and energization of particles in  
 616 magnetospheres.

617 **Acknowledgments:** This work was primarily supported the NSF PREEVENTS  
 618 grant 1663800. Computational resources supporting this work were pro-  
 619 vided by an NSF LRAC allocation at the Texas Advanced Computing Center

620 (TACC) at The University of Texas at Austin.

621 **References**

- 622 [1] L. K. S. Daldorff, G. Tóth, T. I. Gombosi, G. Lapenta, J. Amaya,  
623 S. Markidis, J. U. Brackbill, Two-way coupling of a global Hall mag-  
624 netohydrodynamics model with a local implicit Particle-in-Cell model,  
625 J. Comput. Phys. 268 (2014) 236. doi:10.1016/j.jcp.2014.03.009.
- 626 [2] M. Rieke, T. Trost, R. Grauer, Coupled Vlasov and two-fluid codes on  
627 GPUs, Journal of Computational Physics 283 (2015) 436–452. doi:  
628 10.1016/j.jcp.2014.12.016.  
629 URL <http://dx.doi.org/10.1016/j.jcp.2014.12.016>
- 630 [3] T. Sugiyama, K. Kusano, Multi-scale plasma simulation by the interlock-  
631 ing of magnetohydrodynamic model and particle-in-cell kinetic model,  
632 J. Comput. Phys. 227 (2007) 1340–1352. doi:10.1016/j.jcp.2007.  
633 09.011.
- 634 [4] G. Tóth, X. Jia, S. Markidis, B. Peng, Y. Chen, L. Daldorff, V. Ten-  
635 ishev, D. Borovikov, J. Haiducek, T. Gombosi, A. Glocer, J. Dorelli,  
636 Extended magnetohydrodynamics with embedded particle-in-cell sim-  
637 ulation of ganymede’s magnetosphere, J. Geophys. Res. 121 (2016).  
638 doi:10.1002/2015JA021997.
- 639 [5] Y. Chen, G. Tóth, P. Cassak, X. Jia, T. I. Gombosi, J. Slavin,  
640 S. Markidis, B. Peng, Global three-dimensional simulation of earth’s  
641 dayside reconnection using a two-way coupled magnetohydrodynamics  
642 with embedded particle-in-cell model: initial results, J. Geophys. Res.  
643 122 (2017) 10318. doi:10.1002/2017JA024186.
- 644 [6] Y. Chen, G. Tóth, X. Jia, J. A. Slavin, W. Sun, S. Markidis, T. I.  
645 Gombosi, J. M. Raines, Studying Dawn-Dusk Asymmetries of Mercury’s  
646 Magnetotail Using MHD-EPIC Simulations, Journal of Geophysical Re-  
647 search: Space Physics 124 (11) (2019) 8954–8973. arXiv:1904.06753,  
648 doi:10.1029/2019JA026840.
- 649 [7] H. Zhou, G. Tóth, X. Jia, Y. Chen, S. Markidis, Embedded Kinetic Sim-  
650 ulation of Ganymede’s Magnetosphere: Improvements and Inferences,  
651 Journal of Geophysical Research: Space Physics 124 (7) (2019) 5441–  
652 5460. doi:10.1029/2019JA026643.

- 653 [8] Y. Shou, V. Tenishev, Y. Chen, G. Toth, N. Ganushkina, Magne-  
654 tohydrodynamic with Adaptively Embedded Particle-in-Cell model:  
655 MHD-AEPIC, *Journal of Computational Physics* 446 (2021) 110656.  
656 doi:<https://doi.org/10.1016/j.jcp.2021.110656>.  
657 URL [https://www.sciencedirect.com/science/article/pii/  
658 S0021999121005519](https://www.sciencedirect.com/science/article/pii/S0021999121005519)
- 659 [9] W. Zhang, A. Almgren, V. Beckner, J. Bell, J. Blaschke, C. Chan,  
660 M. Day, B. Friesen, K. Gott, D. Graves, M. Katz, A. Myers, T. Nguyen,  
661 A. Nonaka, M. Rosso, S. Williams, M. Zingale, AMReX: a framework  
662 for block-structured adaptive mesh refinement, *Journal of Open Source  
663 Software* 4 (37) (2019) 1370. doi:[10.21105/joss.01370](https://doi.org/10.21105/joss.01370).  
664 URL <https://doi.org/10.21105/joss.01370>
- 665 [10] W. Zhang, A. Myers, K. Gott, A. Almgren, J. Bell, AMReX: Block-  
666 structured adaptive mesh refinement for multiphysics applications, *Inter-  
667 national Journal of High Performance Computing Applications* 0 (0)  
668 (2021) 1–19. doi:[10.1177/10943420211022811](https://doi.org/10.1177/10943420211022811).
- 669 [11] Y. Chen, G. Tóth, Gauss’s law satisfying energy-conserving semi-  
670 implicit particle-in-cell method, *J. Comput. Phys.* 386 (2019) 632.  
671 doi:[10.1016/j.jcp.2019.02.032](https://doi.org/10.1016/j.jcp.2019.02.032).
- 672 [12] J. U. Brackbill, D. W. Forslund, An implicit method for electromagnetic  
673 plasma simulation in two dimensions, *Journal of Computational Physics*  
674 46 (2) (1982) 271–308. doi:[10.1016/0021-9991\(82\)90016-X](https://doi.org/10.1016/0021-9991(82)90016-X).
- 675 [13] G. Lapenta, Exactly energy conserving semi-implicit particle in cell for-  
676 mulation, *J. Comput. Phys.* 334 (2017) 349. doi:[10.1016/j.jcp.2017.  
677 01.002](https://doi.org/10.1016/j.jcp.2017.01.002).
- 678 [14] K. Fujimoto, S. Machida, Electromagnetic full particle code with adap-  
679 tive mesh refinement technique: Application to the current sheet evolu-  
680 tion, *Journal of Computational Physics* 214 (2) (2006) 550–566. doi:  
681 [10.1016/j.jcp.2005.10.003](https://doi.org/10.1016/j.jcp.2005.10.003).
- 682 [15] K. Fujimoto, A new electromagnetic particle-in-cell model with adaptive  
683 mesh refinement for high-performance parallel computation, *Journal of  
684 Computational Physics* 230 (23) (2011) 8508–8526. doi:[10.1016/j.jcp.2011.07.001](https://doi.org/10.1016/j.jcp.2011.07.001).

- 685       jcp.2011.08.002.  
686       URL <http://dx.doi.org/10.1016/j.jcp.2011.08.002>
- 687 [16] G. Lapenta, Particle rezoning for multidimensional kinetic particle-in-  
688 cell simulations, *Journal of Computational Physics* 181 (1) (2002) 317–  
689 337. doi:10.1006/jcph.2002.7126.
- 690 [17] M. Vranic, T. Grismayer, J. L. Martins, R. A. Fonseca, L. O. Silva,  
691 Particle merging algorithm for PIC codes, *Computer Physics Commu-  
692 nications* 191 (1) (2015) 65–73. arXiv:1411.2248, doi:10.1016/j.  
693 cpc.2015.01.020.  
694       URL <http://dx.doi.org/10.1016/j.cpc.2015.01.020>
- 695 [18] F. Assous, T. Pougeard Dulimbert, J. Segré, A new method for coalesc-  
696 ing particles in PIC codes, *Journal of Computational Physics* 187 (2)  
697 (2003) 550–571. doi:10.1016/S0021-9991(03)00124-4.
- 698 [19] D. R. Welch, T. C. Genoni, R. E. Clark, D. V. Rose, Adaptive particle  
699 management in a particle-in-cell code, *Journal of Computational Physics*  
700 227 (1) (2007) 143–155. doi:10.1016/j.jcp.2007.07.015.
- 701 [20] M. Pfeiffer, A. Mirza, C. D. Munz, S. Fasoulas, Two statistical particle  
702 split and merge methods for Particle-in-Cell codes, *Computer Physics  
703 Communications* 191 (1) (2015) 9–24. doi:10.1016/j.cpc.2015.01.  
704 010.  
705       URL <http://dx.doi.org/10.1016/j.cpc.2015.01.010>
- 706 [21] D. Faghihi, V. Carey, C. Michoski, R. Hager, S. Janhunen, C. S. Chang,  
707 R. D. Moser, Moment preserving constrained resampling with applica-  
708 tions to particle-in-cell methods, *Journal of Computational Physics* 409  
709 (2020) 109317. doi:10.1016/j.jcp.2020.109317.  
710       URL <https://doi.org/10.1016/j.jcp.2020.109317>
- 711 [22] J. Teunissen, U. Ebert, Controlling the weights of simulation particles:  
712 Adaptive particle management using k-d trees, *Journal of Computa-  
713 tional Physics* 259 (2014) 318–330. arXiv:1301.1552, doi:10.1016/j.  
714 jcp.2013.12.005.  
715       URL <http://dx.doi.org/10.1016/j.jcp.2013.12.005>
- 716 [23] P. T. Luu, T. Tückmantel, A. Pukhov, Voronoi particle merging al-  
717 gorithm for PIC codes, *Computer Physics Communications* 202 (2016)

- 718 165–174. [arXiv:1504.00636](https://arxiv.org/abs/1504.00636), doi:10.1016/j.cpc.2016.01.009.  
719 URL <http://dx.doi.org/10.1016/j.cpc.2016.01.009>
- 720 [24] Y. Chen, G. Tóth, H. Hietala, S. K. Vines, Y. Zou, Y. Nishimura, M. V.  
721 Silveira, Z. Guo, Y. Lin, S. Markidis, Magnetohydrodynamic with em-  
722 bedded particle-in-cell simulation of the Geospace Environment Model-  
723 ing dayside kinetic processes challenge event, Earth and Space Science  
724 (2020). doi:10.1029/2020ea001331.
- 725 [25] X. Wang, Y. Chen, G. Tóth, Global magnetohydrodynamic magneto-  
726 sphere simulation with an adaptively embedded particle-in-cell model,  
727 Earth and Space Science Open Archive (2021) 21doi:10.1002/essoar.  
728 10508044.1.  
729 URL <https://doi.org/10.1002/essoar.10508044.1>
- 730 [26] G. Lapenta, J. U. Brackbill, Nonlinear evolution of the lower hybrid  
731 drift instability: Current sheet thinning and kinking, Phys. Plasmas 9  
732 (2002) 1544.