

# 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, and it is not always feasible to cover the whole physical structure of interest with a box due to the limitation of the computational resources. The FLEKS (Flexible Exascale Kinetic Simulator), which is a new PIC code and allows a PIC region of any shape, is designed to break this restriction and extend the capabilities of the MHD-EPIC model.

FLEKS uses the Gauss's law satisfying energy-conserving semi-implicit method (GL-ECSIM) as the base PIC solver. We have also designed extra numerical techniques, such as the adaptive time stepping and particle resampling algorithms, to further improve the accuracy and flexibility of the PIC solver. The grid of FLEKS has to be Cartesian, but the active PIC region is not necessarily to be a box anymore since any Cartesian cells can be turned off. Furthermore, FLEKS supports switching on or switching off grid cells adaptively during a simulation. The initial conditions and boundary conditions of the active PIC region are provided by the coupled MHD code. FLEKS and the coupled MHD code constitute the MHD with adaptively embedded particle-in-cell (MHD-AEPIC) model.

*Keywords:* particle-in-cell; particle merging

---

<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. Recently, 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 global simulations. The MHD-EPIC model usually covers the dayside or/and the tail 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 a MHD-EPIC simulation, a PIC region is restricted to be a box, which is not always feasible to cover the whole physical structure of interest due to either the limitation of computational resources or the 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 brand new code, the FLEKS (Flexible Exascale Kinetic Simulator), 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 active PIC region is not limited to be a box anymore since the PIC cells can

38 be turned off to fit the region of interest. Furthermore, FLEKS also supports  
39 switching on or switching off grid cells dynamically for MHD-AEPIC simu-  
40 lations. The Gauss’s law satisfying energy-conserving semi-implicit method  
41 (GL-ECSIM) [11] is implemented as the base PIC solver. The time step of the  
42 semi-implicit PIC methods is usually limited by a Courant–Friedrichs–Lewy  
43 (CFL) condition in order to be accurate [12]. Since the plasma maximum  
44 characteristic speed may change significantly during a long MHD-AEPIC  
45 simulation, the simulation will be either too slow or inaccurate with a fixed  
46 time step. To keep the simulation efficient and accurate at the same time,  
47 FLEKS uses an adaptive time-stepping algorithm, which still satisfies the  
48 requirement of the energy-conserving semi-implicit method (ECSIM) [13] to  
49 keep energy conservation. Section 2 describes the adaptive grid and temporal  
50 discretization of FLEKS.

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

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

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

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

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




Figure 1: A schematic shows the grid structure of the MHD-AEPIC model.

112 Tracking the motion of macro-particles is useful for learning the trajectory  
 113 and energization of plasma, so FLEKS provides a parallel test particle module  
 114 to follow the motion of macro-particles and save the massive data to disk.  
 115 The test particle module can be used either as a module of the PIC code, or  
 116 as an independent component to directly couple to the MHD model. Section  
 117 5 describes the implementation details of the test particle module.

118 The paper is organized as follows. Section 2 describes the grid design of  
 119 FLEKS. Section 3 introduces the adaptive time time step. Section 4 focuses  
 120 on the particle merging algorithm. Section 5 presents the implementation of  
 121 the test particle module. Section 6 shows the numerical tests to demonstrate  
 122 the capability of the adaptive active PIC regions, the role of the particle  
 123 resampling algorithms, and the parallel efficiency of FLEKS. Finally, section  
 124 7 presents the conclusions.

## 125 2. Adaptive grid

126 Since the MHD-EPIC model was developed by Daldorff et al. [1], we have  
 127 developed new features to make it more flexible to use. It supports multiple  
 128 independent PIC domains to cover several regions of interest [4], and it also  
 129 allows rotating a PIC box domain to align with the simulation objective  
 130 [24]. However, a box is not always feasible and efficient to cover a region  
 131 of interest. For example, if a PIC box is used to cover the whole dayside  
 132 magnetopause, which is close to a paraboloid, the box will cut through the  
 133 planet and introduce extra difficulties, and the PIC box will also contain  
 134 a large portion of cells where the kinetic effects are not important to slow  
 135 down the simulation. A flexible grid that allows an active PIC region of

136 a paraboloid to fit the magnetopause solves the problems. A dynamically  
137 adaptive grid is also useful to improve the efficiency of some simulations.  
138 For instance, the near-Earth X-line may move from the inner magnetotail  
139 to the middle or even far magnetotail [25], and an adaptive grid that only  
140 covers the environment around the X-line is much more efficient compared  
141 to a large PIC box that covers the whole magnetotail. The MHD-AEPIC  
142 model is designed to solve these problems, and FLEKS is the key component.  
143 Figure 1 is a carton that shows the concept of the MHD-AEPIC model.

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

165 FLEKS uses the data structures provided by the AMReX library to store  
166 the fields and also the particles. After the patch status array is obtained from  
167 the fluid model, FLEKS uses the functions provided by the AMReX library  
168 to divide the active regions into blocks. AMReX does not require all the  
169 blocks to have the same size. We note that the patch and the block are two  
170 independent concepts. The patches are only used to activate or deactivate  
171 cells. For example, the ‘L’ shape active region in Figure 2(a) consists of 3  
172 patches and it can be divided into 2 blocks (Figure 2(b)).

173 FLEKS allows activating or deactivating patches during a simulation. If

174 the active regions change, FLEKS will produce a new set of blocks to cover the  
 175 new active regions. With the function provided by AMReX, FLEKS copies  
 176 the fields and particles from the old blocks to the new ones for the cells that  
 177 are already active and deletes the information of the newly deactivated cells.  
 178 The newly activated cells are filled in with the information obtained from  
 179 the fluid model as what is done for FLEKS initialization.

180 FLEKS has two ghost cell layers, but the outer layer is only used to  
 181 receive and store the magnetic fields, which are necessary for calculating  
 182 currents on the nodes of the inner ghost cell layer from  $\vec{J} = \nabla \times \vec{B}$ . The  
 183 currents are used to generate particles with correct velocities in the inner  
 184 layer ghost cells. To simplify the description, we ignore the outer layer in  
 185 Figure 2(c) and also in the rest of the paper unless otherwise specified. The  
 186 principle of setting boundary conditions of the electromagnetic fields and the  
 187 particles is still the same as the MHD-EPIC algorithm with a box PIC region  
 188 [1]. However, the non-box shape of an active region introduces some extra  
 189 implementation difficulties. They are three types of ghost cells for a block:  
 190 the internal ghost cells (blue cells in Figure 2(c)), the exclusive boundary  
 191 ghost cells (gray cells in Figure 2(c)) and the shared boundary ghost cells  
 192 (cyan cells in Figure 2(c)). The internal ghost cells are not boundary cells,  
 193 and there is no need to apply boundary conditions. The exclusive boundary  
 194 ghost cells are not overlapped with any cells of the neighbor blocks, and they  
 195 should be filled in with macro-particles as the particle boundary condition.  
 196 The shared boundary ghost cells are overlapped with the boundary ghost cells  
 197 of the neighbor blocks. Only one of these blocks should generate boundary  
 198 particles. Here is the algorithm to choose the block. The first step is to  
 199 distinguish the boundary ghost cells from the internal ghost cells. Then, for  
 200 each boundary ghost cell, either the exclusive type or the shared type, we  
 201 loop through its at most 26 neighbor cells (3D) in a certain order (we choose  
 202 to loop through all the face neighbors first, then the edge neighbors, and  
 203 finally the corner neighbors), skip the cells does not exist and find out the  
 204 first neighbor cell that is either physical cell or internal ghost cell. If the  
 205 location of this neighbor cell is inside the physical domain of this block, this  
 206 block should generate particles inside this boundary ghost cell. For example,  
 207 in Figure 2(c)), C1 and C3 are overlapped with each other. We loop through  
 208 the neighbor cells of C1 (C3) and find C2 (C4) is its first neighbor cell that  
 209 is either a physical cell or an internal ghost cell, and block-1 (block-2) should  
 210 (not) generate particles in C1 (C3) since C2 (C4) is (not) inside block-1  
 211 (block-2).

212 The electric fields are node-based in FLEKS. For a node that is shared by  
 213 multiple blocks, such as the one indicated by red-cross in Figure 2(c)), only  
 214 one block should take care of the shared node when solving the linear equation  
 215 of the electric fields. The algorithm described in the previous paragraph is  
 216 also applied to choose the proper block for the shared nodes.

### 217 3. Adaptive time-stepping

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

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

### 243 4. Particle resampling

244 We implemented particle resampling algorithms to control the macro-  
 245 particle number of each cell. Every computational cycle, a particle splitting




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 colors the ghost cells of the active cells. (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 neighbor 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 neighbor blocks. One of the blocks should generate boundary particles.




Figure 3: The adaptive temporal discretization.

246 (merging) algorithm is applied to produce (remove) macro-particles for the  
 247 cells that contain fewer (more) particles than the low (high) threshold. The  
 248 goal of splitting and merging is to keep the number of particles per cell  
 249 (PPC) close to a constant in the whole computational domain. This constant  
 250 number is the initial PPC in FLEKS. Essentially, the particle resampling  
 251 algorithms use a new set of particles to replace the old ones. Our guiding  
 252 principle of designing the algorithms is that the replacement should keep  
 253 the original particle phase-space distribution as much as possible. In order  
 254 to conveniently apply the resampling algorithms, FLEKS stores the particle  
 255 data cell by cell.

#### 256 4.1. Particle splitting

257 Our particle splitting algorithm is essentially the same as the one intro-  
 258 duced by Lapenta[16], in which one particle is split into two children parti-  
 259 cles. The children particles have the same velocity as their parent particle,  
 260 but their location is oppositely displaced slightly along the velocity direction.  
 261 We choose to displace the new particles along the velocity direction so that  
 262 the orbits of the new particles are still close to the orbit of the old particle.

263 Initially, the particles that are close to each other have similar weights, but  
 264 the weights may become quite different later due to the transport of particles  
 265 and the effect of particle resampling. We choose to split the heaviest particles  
 266 to minimize the particle weight variance.




Figure 4: The steps of merging macro-particles.

267 *4.2. Particle merging*

268 The essence of particle merging is replacing a set of particles with a new  
 269 set, and the new set contains fewer particles than the old one. Particle merg-  
 270 ing reduces the particle number in some cells and improves load balancing.  
 271 Particle resampling has a negative impact on the accuracy of a simulation  
 272 because (1) the replacement introduces errors, and (2) fewer particles lead  
 273 to larger statistical noise in the subsequent simulation. The increasing of  
 274 statistical noise is inevitable, but the errors caused by the replacement can  
 275 be minimized with a proper merging algorithm.

## 276 5. Test particle module

277 An independent test particle (TP) module is designed to track the mo-  
278 tion of the macro-particles for FLEKS. It can be used either as an auxiliary  
279 component of the PIC algorithm or as an independent component, which is  
280 directly coupled with the MHD model. Similar to the macro-particles of the  
281 PIC code, the test particles are also generated based on local plasma fluid  
282 quantities that are obtained from the MHD model. The TP module also uses  
283 the same algorithm to move particles as the GL-ECSIM algorithm. When the  
284 TP module is used with the PIC component together, the TP module shares  
285 the same grid layout as the PIC component and uses the electromagnetic  
286 fields calculated by PIC to update test particles. When the PIC component  
287 is turned off, the TP module can directly obtain the grid structure and elec-  
288 tromagnetic fields from the MHD model. FLEKS is a pure test particle code  
289 for this case. Compared to the embedded PIC simulations, the test particle  
290 simulations are only one-way coupled, i.e., the MHD model provides the elec-  
291 tromagnetic fields for FLEKS, but there is not any feedback from FLEKS to  
292 the MHD model.

293 In a three-dimensional (3D) simulation, it is common to track the motion  
294 of millions of test particles, and a few thousand steps of the update will  
295 easily produce a few hundred Gigabytes of particle trajectory data. The test  
296 particle module should organize the data properly to improve both the IO  
297 performance of writing data to disk and also the efficiency of finding all the  
298 data of a certain particle for data analysis. To reduce the IO frequency, the  
299 TP module of FLEKS saves the particle trajectory data every 100 cycles,  
300 and all the processors write to the same file with MPI-IO APIs. We note  
301 that if a test particle moves from one processor to another in the middle  
302 of two IO operations, its trajectory data should also be transferred to the  
303 destination processor. Besides the particle trajectory data file, a particle ID  
304 list file, which maps a particle ID to its data location in the particle data file,  
305 is also created. An example of these two files is shown in Figure 5. With this  
306 file structure, it is efficient to find all the trajectory data of a certain particle  
307 for data analysis.

## 308 6. Numerical tests

### 309 6.1. Two-dimensional double-current-sheet magnetic reconnection

310 The two-dimensional magnetic reconnection problem is widely used to  
311 test plasma simulation codes. The double-current-sheet setup allows periodic




Figure 5: The file structures for storing test particles.




Figure 6: 2D fast wave test.





Figure 7: 1D fast wave.




Figure 8: Simulation speed.

312 boundary conditions for both directions. Here we use a setup based on the  
 313 GEM-challenge [26].

314 The initial condition is set to satisfy the fluid force balance for both  
 315 electrons and ions [27]. The simulation domain is  $-12.8 < x < 12.8$  and  
 316  $-6.4 < y < 6.4$  in normalized CGS unit. The speed of light is set to be  $c = 1$ .  
 317 The ion density is uniform and  $n_i = 0.0975$ . The ion plasma frequency is  
 318  $\omega_{pi} = \sqrt{\frac{4\pi n_i e^2}{m_i}} = 1.107$  and the ion inertial length  $d_i = c/\omega_{pi} = 0.903$  since  
 319  $m_i = 1$  and  $q_i = -q_e = 1$ . A reduced ion-electron mass ratio  $m_i/m_e = 25$  is  
 320 used, so the electron skin depth is about  $d_e = d_i/5 = 0.18$ . Initially, there is  
 321 no charge separation,  $n_e = n_i$ , and the electric field is  $\mathbf{E} = 0$ .

322 The background magnetic field is

$$B_x = B_0 \left( -1 + \tanh \frac{y - y_B}{\delta} + \tanh \frac{y_T - y}{\delta} \right) \quad (1)$$

323 where  $B_0 = 0.07$ , the positions of the two current sheets are  $y_B = -3.2$  and  
 324  $y_T = 3.2$ , respectively, and the width of the current sheets are controlled by  
 325  $\delta = 0.5$ . The electrons have a velocity in the z-direction to generate current  
 326 equal to the curl of the magnetic field, i.e.,  $J_z = n_e q_e u_{e,z} = -\partial B_x / \partial y$ . The  
 327 ion pressure  $p_i$  is uniform in the whole domain. Far away from the current  
 328 sheets, the ion plasma beta is 1, and the electron pressure is 1/5 of the  
 329 ion pressure. Near the current sheet, the electrons are heated to balance  
 330 the magnetic field gradient force, which is the same as the Lorentz force  
 331  $-n_e q_e u_{e,z} B_x$ . This unperturbed initial condition is in fluid force balance [27].

332 A perturbation is added to excite the reconnection. The magnetic field  
 333 perturbation vector potential is  $A_x = 0$ ,  $A_y = 0$  and:

$$A_z = A_0 B_0 \left\{ - e^{-\frac{(x-x_T)^2}{G_x^2} - \frac{(y-y_T)^2}{G_y^2}} \cos [k_x(x - x_T)] \cos [k_y(y - y_T)] \right. \\ \left. + e^{-\frac{(x-x_B)^2}{G_x^2} - \frac{(y-y_B)^2}{G_y^2}} \cos [k_x(x - x_B)] \cos [k_y(y - y_B)] \right\} \quad (2)$$

334 where the perturbation amplitude is set by  $A_0 = 0.1$ , the locations along the  
 335 top and bottom current sheets are  $x_T = 6.4$  and  $x_B = -6.4$ , respectively,  
 336 the width of Gaussian profiles are  $G_x = G_y = 0.5$ , and the wave vectors are  
 337  $k_x = 2\pi/25.6$  and  $k_y = 2\pi/12.8$ . Since these two reconnection sites, i.e., the  
 338 bottom left one at  $(x_B, y_B)$  and the top right one at  $(x_T, y_T)$ , produce the  
 339 same signatures, we only plot and discuss the bottom left reconnection site  
 340 for simplicity.




Figure 9: 2D magnetic reconnection results with (right column) or without (left column) particle resampling.




Figure 10: 2D magnetic reconnection results with (right column) or without (left column) particle resampling.




Figure 11: Phase space distributions.




Figure 12: Weak scaling results.

## 341 7. Conclusion

342 In this paper, we present the algorithms and implementation of FLEKS.  
 343 The adaptive grid allows FLEKS covering a region of any shape. The adap-  
 344 tive temporal discretization and particle resampling algorithms improves the  
 345 accuracy and efficiency of simulations.

346 **Acknowledgments:** This work was supported by the INSPIRE NSF grant  
 347 PHY-1513379 and the NSF PREEVENTS grant 1663800. Computational  
 348 resources supporting this work were provided on the Blue Waters super com-  
 349 puter by the NSF PRAC grant ACI-1640510, on the Pleiades computer by  
 350 NASA High-End Computing (HEC) Program through the NASA Advanced  
 351 Supercomputing (NAS) Division at Ames Research Center, and from Yel-  
 352 lowstone (ark:/85065/d7wd3xhc) provided by NCAR’s Computational and  
 353 Information Systems Laboratory, sponsored by the National Science Foun-  
 354 dation.





Figure 13: Strong scaling results.




Figure 14: A MHD-AEPIC simulation of Earth's magnetosphere with the dayside magnetopause covered by FLEKS.




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

## 355 References

- 356 [1] L. K. S. Daldorff, G. Tóth, T. I. Gombosi, G. Lapenta, J. Amaya,  
 357 S. Markidis, J. U. Brackbill, Two-way coupling of a global Hall mag-  
 358 netohydrodynamics model with a local implicit Particle-in-Cell model,  
 359 J. Comput. Phys. 268 (2014) 236. doi:10.1016/j.jcp.2014.03.009.
- 360 [2] M. Rieke, T. Trost, R. Grauer, Coupled Vlasov and two-fluid codes on  
 361 GPUs, Journal of Computational Physics 283 (2015) 436–452. doi:  
 362 10.1016/j.jcp.2014.12.016.  
 363 URL <http://dx.doi.org/10.1016/j.jcp.2014.12.016>
- 364 [3] T. Sugiyama, K. Kusano, Multi-scale plasma simulation by the interlock-  
 365 ing of magnetohydrodynamic model and particle-in-cell kinetic model,  
 366 J. Comput. Phys. 227 (2007) 1340–1352. doi:10.1016/j.jcp.2007.  
 367 09.011.
- 368 [4] G. Tóth, X. Jia, S. Markidis, B. Peng, Y. Chen, L. Daldorff, V. Ten-  
 369 ishev, D. Borovikov, J. Haiducek, T. Gombosi, A. Glocer, J. Dorelli,  
 370 Extended magnetohydrodynamics with embedded particle-in-cell sim-

- 371       ulation of ganymede’s magnetosphere, *J. Geophys. Res.* 121. doi:  
372       10.1002/2015JA021997.
- 373 [5] Y. Chen, G. Tóth, P. Cassak, X. Jia, T. I. Gombosi, J. Slavin,  
374       S. Markidis, B. Peng, Global three-dimensional simulation of earth’s  
375       dayside reconnection using a two-way coupled magnetohydrodynamics  
376       with embedded particle-in-cell model: initial results, *J. Geophys. Res.*  
377       122 (2017) 10318. doi:10.1002/2017JA024186.
- 378 [6] Y. Chen, G. Tóth, X. Jia, J. A. Slavin, W. Sun, S. Markidis, T. I.  
379       Gombosi, J. M. Raines, Studying Dawn-Dusk Asymmetries of Mercury’s  
380       Magnetotail Using MHD-EPIC Simulations, *Journal of Geophysical Re-*  
381       search: *Space Physics* 124 (11) (2019) 8954–8973. arXiv:1904.06753,  
382       doi:10.1029/2019JA026840.
- 383 [7] H. Zhou, G. Tóth, X. Jia, Y. Chen, S. Markidis, Embedded Kinetic Sim-  
384       ulation of Ganymede’s Magnetosphere: Improvements and Inferences,  
385       *Journal of Geophysical Research: Space Physics* 124 (7) (2019) 5441–  
386       5460. doi:10.1029/2019JA026643.
- 387 [8] Y. Shou, V. Tenishev, Y. Chen, G. Toth, N. Ganushkina, Magne-  
388       tohydrodynamic with Adaptively Embedded Particle-in-Cell model:  
389       MHD-AEPIC, *Journal of Computational Physics* 446 (2021) 110656.  
390       doi:<https://doi.org/10.1016/j.jcp.2021.110656>.  
391       URL       [https://www.sciencedirect.com/science/article/pii/](https://www.sciencedirect.com/science/article/pii/S0021999121005519)  
392       S0021999121005519
- 393 [9] W. Zhang, A. Almgren, V. Beckner, J. Bell, J. Blaschke, C. Chan,  
394       M. Day, B. Friesen, K. Gott, D. Graves, M. Katz, A. Myers, T. Nguyen,  
395       A. Nonaka, M. Rosso, S. Williams, M. Zingale, AMReX: a framework  
396       for block-structured adaptive mesh refinement, *Journal of Open Source*  
397       Software 4 (37) (2019) 1370. doi:10.21105/joss.01370.  
398       URL <https://doi.org/10.21105/joss.01370>
- 399 [10] W. Zhang, A. Myers, K. Gott, A. Almgren, J. Bell, AMReX: Block-  
400       structured adaptive mesh refinement for multiphysics applications, *In-*  
401       ternational Journal of High Performance Computing Applications 0 (0)  
402       (2021) 1–19. doi:10.1177/10943420211022811.

- 403 [11] Y. Chen, G. Tóth, Gauss's law satisfying energy-conserving semi-  
404 implicit particle-in-cell method, *J. Comput. Phys.* 386 (2019) 632.  
405 doi:10.1016/j.jcp.2019.02.032.
- 406 [12] J. U. Brackbill, D. W. Forslund, An implicit method for electromagnetic  
407 plasma simulation in two dimensions, *Journal of Computational Physics*  
408 46 (2) (1982) 271–308. doi:10.1016/0021-9991(82)90016-X.
- 409 [13] G. Lapenta, Exactly energy conserving semi-implicit particle in cell for-  
410 mulation, *J. Comput. Phys.* 334 (2017) 349. doi:10.1016/j.jcp.2017.  
411 01.002.
- 412 [14] K. Fujimoto, S. Machida, Electromagnetic full particle code with adap-  
413 tive mesh refinement technique: Application to the current sheet evolu-  
414 tion, *Journal of Computational Physics* 214 (2) (2006) 550–566. doi:  
415 10.1016/j.jcp.2005.10.003.
- 416 [15] K. Fujimoto, A new electromagnetic particle-in-cell model with adaptive  
417 mesh refinement for high-performance parallel computation, *Journal of*  
418 *Computational Physics* 230 (23) (2011) 8508–8526. doi:10.1016/j.  
419 jcp.2011.08.002.  
420 URL <http://dx.doi.org/10.1016/j.jcp.2011.08.002>
- 421 [16] G. Lapenta, Particle rezoning for multidimensional kinetic particle-in-  
422 cell simulations, *Journal of Computational Physics* 181 (1) (2002) 317–  
423 337. doi:10.1006/jcph.2002.7126.
- 424 [17] M. Vranic, T. Grismayer, J. L. Martins, R. A. Fonseca, L. O. Silva,  
425 Particle merging algorithm for PIC codes, *Computer Physics Commu-  
426 nications* 191 (1) (2015) 65–73. arXiv:1411.2248, doi:10.1016/j.  
427 cpc.2015.01.020.  
428 URL <http://dx.doi.org/10.1016/j.cpc.2015.01.020>
- 429 [18] F. Assous, T. Pougeard Dulimbert, J. Segré, A new method for coalesc-  
430 ing particles in PIC codes, *Journal of Computational Physics* 187 (2)  
431 (2003) 550–571. doi:10.1016/S0021-9991(03)00124-4.
- 432 [19] D. R. Welch, T. C. Genoni, R. E. Clark, D. V. Rose, Adaptive particle  
433 management in a particle-in-cell code, *Journal of Computational Physics*  
434 227 (1) (2007) 143–155. doi:10.1016/j.jcp.2007.07.015.

- 435 [20] M. Pfeiffer, A. Mirza, C. D. Munz, S. Fasoulas, Two statistical particle  
436 split and merge methods for Particle-in-Cell codes, *Computer Physics*  
437 *Communications* 191 (1) (2015) 9–24. doi:10.1016/j.cpc.2015.01.  
438 010.  
439 URL <http://dx.doi.org/10.1016/j.cpc.2015.01.010>
- 440 [21] D. Faghihi, V. Carey, C. Michoski, R. Hager, S. Janhunen, C. S. Chang,  
441 R. D. Moser, Moment preserving constrained resampling with applica-  
442 tions to particle-in-cell methods, *Journal of Computational Physics* 409  
443 (2020) 109317. doi:10.1016/j.jcp.2020.109317.  
444 URL <https://doi.org/10.1016/j.jcp.2020.109317>
- 445 [22] J. Teunissen, U. Ebert, Controlling the weights of simulation particles:  
446 Adaptive particle management using k-d trees, *Journal of Computa-*  
447 *tional Physics* 259 (2014) 318–330. arXiv:1301.1552, doi:10.1016/j.  
448 jcp.2013.12.005.  
449 URL <http://dx.doi.org/10.1016/j.jcp.2013.12.005>
- 450 [23] P. T. Luu, T. Tückmantel, A. Pukhov, Voronoi particle merging al-  
451 gorithm for PIC codes, *Computer Physics Communications* 202 (2016)  
452 165–174. arXiv:1504.00636, doi:10.1016/j.cpc.2016.01.009.  
453 URL <http://dx.doi.org/10.1016/j.cpc.2016.01.009>
- 454 [24] Y. Chen, G. Tóth, H. Hietala, S. K. Vines, Y. Zou, Y. Nishimura, M. V.  
455 Silveira, Z. Guo, Y. Lin, S. Markidis, Magnetohydrodynamic with em-  
456 bedded particle-in-cell simulation of the Geospace Environment Mod-  
457 eling dayside kinetic processes challenge event, *Earth and Space Sci-*  
458 *encedoi:10.1029/2020ea001331*.
- 459 [25] X. Wang, Y. Chen, G. Tóth, Global magnetohydrodynamic magneto-  
460 sphere simulation with an adaptively embedded particle-in-cell model,  
461 *Earth and Space Science Open Archive* (2021) 21doi:10.1002/essoar.  
462 10508044.1.  
463 URL <https://doi.org/10.1002/essoar.10508044.1>
- 464 [26] J. Birn, J. F. Drake, M. A. Shay, B. N. Rogers, R. E. Denton,  
465 M. Hesse, M. Kuznetsova, Z. W. Ma, A. Bhattacharjee, A. Otto,  
466 P. L. Pritchett, Geospace Environmental Modeling (GEM) magnetic  
467 reconnection challenge, *J. Geophys. Res.* 106 (A3) (2001) 3715–3720.  
468 doi:10.1029/1999JA900449.

- 469 [27] Z. Huang, G. Tóth, B. van der Holst, Y. Chen, T. Gombosi, A six-  
470 moment multi-fluid plasma model, *J. Comput. Phys.* 387 (2019) 134.  
471 doi:10.1016/j.jcp.2019.02.023.