

Dynamic Particle Weight Remapping in Hybrid PIC Hall-effect Thruster Simulation

*Presented at Joint Conference of 30th International Symposium on Space Technology and Science,
34th International Electric Propulsion Conference and 6th Nano-satellite Symposium
Hyogo-Kobe, Japan
July 4–10, 2015*

Robert Martin*

ERC Incorporated, Huntsville, AL, 35806-3200, USA

Justin Koo†, David Bilyeu‡

AFRL/RQRS, USA

Control of macroparticle growth and distribution and statistical noise are key challenges for particle kinetic models such as particle-in-cell (PIC). For hybrid fluid-PIC models such as those commonly used in Hall-effect thruster (HET) simulation, the statistical noise adds an additional challenge due to the stiffness imposed by the use of gradients of the mobility tensor in the solution of Ohm's law. Past results indicated the direct hybrid PIC is approach is essentially intractable in more than 1D without advanced filtering for noise reduction. Though increasing the ratio of computationally represented macro particles to real physical particles is one means of variance reduction used to improve the fidelity of particle-based simulations, the high dynamic range necessary for problems involving chain-branching reactions such as ionization tend to require balancing regions of high statistical noise with regions of extreme computational cost. In this work, we use a recently developed method of particle weight remapping that strictly conserves mass, momentum, and energy while simultaneously unlike most prior standard particle merging techniques to adapt macroparticle weights within a simple pseudo-spectral HET simulation. The technique also avoids thermalization by remaining faithful to the original velocity distribution function through the use of octree binning in velocity space. Though demonstrating effective control over macroparticle density in an axial-azimuthal hybrid PIC HET simulation with little impact on solution quality, the method was found to be unnecessary because ions exit the small computational domain too rapidly to impose severe computational challenges in the baseline implementation. However, the method remains a viable candidate to enhance solution tractability as problem complexity increases.

Nomenclature

(a/b)	= a^{th} or b^{th} particle corresponding to $+/-$ if present
c_i	= Thermal (peculiar) velocity
$\hat{R}$	= Unit normal direction
n	= A natural number, $n \in \mathbb{N}$
N	= Number of macroparticles in a subset of macroparticles
(p)	= p^{th} particle of the distribution
v_i	= Velocity in the i^{th} direction
v_i^{th}	= Thermal velocity in the i^{th} direction
w	= Ratio of physical to computational macroparticle weight

*robert.martin.81.ctr@us.af.mil

†justin.koo@us.af.mil

‡david.bilyeu.1@us.af.mil

I. Introduction

Very large dynamic range of plasma densities can result in significant challenges for particle-based methods such as the pseudo-spectral hybrid-PIC model. Reconstruction of the phase space from the stochastic description of PIC macroparticles in order to minimize noise and bound computational cost provides an opportunity for significant gains in problem tractability for particle-based methods.

The standard approach of merging of macroparticles¹ uses pair-wise coalescence (2:1 ratio). New particle with average properties from two initial ones are constructed. Though simple and fast, the method cannot simultaneously conserve momentum and energy due to insufficient degrees of freedom in the resulting particle. This error can be reduced by selecting particles that are near each other in velocity space, but the error remains a fundamental consequence of the reduction of degrees of freedom. Other sophisticated models have been designed to mitigate the error in momentum or energy such as computational particles with internal energy in which the error could be accumulated² and merging values to grid nodes and redistributing the moments to particles.^{3,4} Although exact energy and momentum conservation is possible with these methods, they are considerably more complex than the original naive 2:1 merge and rarely tested in simple reproducible test cases.

To address the insufficient number of degrees of freedom, a simple method which relies on the generation of a *pair* of particles has periodically emerged.^{1,5} The particle pair provides the required freedom to conserve all moments up to 2^{nd} order exactly. The pair of resultant particles have the same mean momentum as the originals, but they also have additional equal and opposite components of velocity in addition to the mean momentum such that energy is conserved. This method would obviously provide no benefits if starting from 2 particles, but an effective 2:1 reduction could be achieved by merging 4 particles into 2. An arbitrary number of initial particles can be considered, and as long as 2 new particles are produced, the method exactly conserves mass, momentum and energy simultaneously. This method can then be called a $N : 2$ merge ($N > 2$).

Even if a merging algorithm exactly conserves the first three moments (mass, momentum and energy), it can result in artificial thermalization of the velocity distribution. Consider merging three particles selected from a velocity distribution of equal weight particles in one velocity space direction down to two. In general, the mean velocity of the three particles will be on one side of the mean velocity of the distribution. Furthermore, two of the particles are on one side of the mean velocity of the set of three. If randomly selected from the distribution, the two particles are more likely to be on the side closer to the mean of the velocity distribution. If two particles are created with the same mean velocity and RMS velocity as the original set, they tend to be more symmetric about the distribution mean velocity than the original particles. Repeated merging then tends towards a symmetric thermal distribution around the mean velocity, effectively resulting in artificial collision-like thermalization. In multiple spatial dimensions, the problem is further exacerbated if the pair of resultant particles are split along an arbitrarily selected axis in velocity space as this further isotropizes the velocity distribution artificially.

This artificial thermalization is significant for ternary merges and was likely the reason for incorrect wave-speeds in a collisionless shock test case used to test merging schemes.¹ To inhibit thermalization, selection of near neighbors in velocity space can help to mitigate the error introduced by the merge. This is not unlike the near neighbor criteria on the binary merge used to improve energy conservation, but it was concluded that selection of near neighbor particle triplets was an ill-defined and costly problem compared to the neighbor-pair search. To avoid this issue, in this work as with the recent prior work at AFRL,^{6,7} the particle VDF is first binned using an adaptive octree mesh to select near neighbors as described in the following section.

II. Computational Model

The model used in this code mirrors that of the companion paper.⁸ The centerline of the HET is again modeled in the axial-azimuthal plane with PIC using singly charged xenon ions, the electron fluid is solved using a pseudo-spectral generalized Ohm's law solver, and the fluid is advected with a high-order semi-Lagrangian advection scheme and diffusion based on Knudsen using the channel width. To enable direct comparison with the baseline control pseudo-spectral HET model, the particle remapping algorithm directly inserted into the control model using identical boundary conditions.

A. Merging and Splitting

Dynamic phase space remapping was previously performed using the phase space reconstruction model described in Reference⁶ which demonstrate remapping of bi-Maxwellian particle distributions in 0D and 1D potential wells as well as electrons in an unmagnetized DC-breakdown for a 1D spark-gap. This work has been extended and refined significantly for forthcoming publication,⁷ and the description that follows mirrors that work. Further detail and method verification can be found therein.

A conservative $N:2$ merge forms the basic building block for the conservation properties inherent to the particle remapping techniques described here. The formulation focuses on moments of mass, momentum, and energy as they are conserved in elastic binary collisions. The particle representation of these moments are outlined in eq. (1). Here, w represents the statistical “weight” of the particle in question. This is the number of true, physical microscopic particles represented by each computational macroparticle. Particles are assumed to all have the same species type. A merging algorithm would make little sense for particles of different species as the resulting particles would have compound composition and ill-defined trajectories in phase-space. However, for mixtures, the operations can be applied to each species type independently. This also enables independent target statistical weights for different species which can be critical when a trace species has a large impact on results as in for example branch initiation reactions.

The conserved moments are shown in eq. (1), where $N=n+2$ original particles ($\forall n > 0$), denoted with superscript (p) are combined into average quantities $(\bar{\cdot})$.

$$\bar{w} = \sum_p^{n+2} w^{(p)}, \quad \bar{v}_i = \frac{\sum_p^{n+2} w^{(p)} v_i^{(p)}}{\sum_p^{n+2} w^{(p)}}, \quad \bar{c}_i^2 = \frac{\sum_p^{n+2} w^{(p)} (v_i^{(p)} - \bar{v}_i)^2}{\sum_p^{n+2} w^{(p)}}, \quad (1)$$

The total weight is then split into two particles (a, b) , arbitrarily assumed of equal statistical weight to enable closure. Momentum conservation is obtained from the center of mass (CM) velocity. Assigning a mean component of their velocity as the CM velocity, the thermal component can be freely selected to satisfy the conservation of energy, as seen in eq. (2),

$$w^{(a/b)} = \frac{\bar{w}}{2} \quad v_i^{(a/b)} = \bar{v}_i \pm \|c_i\| \hat{R} \quad (2)$$

where $\hat{R}$ denotes a unit vector of random direction and $c_i = \sqrt{\bar{c}_i^2}$ is an RMS thermal velocity vector. Note that the “thermal” component is included symmetrically in the CM frame. The $\pm$ symbol represent either adding or subtracting the same isotropic thermal vector for both the “a” and “b” particles respectively. This procedure is entirely equivalent to computing properties in the CM frame, then adding the CM velocity component (Galilean transformation).

Instead of choosing a random isotropic unit vector, a simpler procedure is to add and subtract the thermal velocity in each of the three spatial directions independently augmented by a random sign, i.e.:

$$\bar{v}_i^{th} = \text{sgn}(\text{rand}(-1, 1)) c_i \quad \Rightarrow v_i^{(a/b)} = \bar{v}_i \pm v_i^{th} \quad (3)$$

Plugging these velocities for the “a” and “b” particles into the original moment equations demonstrates the moment conservation. This approach, of conserving each direction’s component of CM kinetic energy independently, involves fewer operations which makes it more direct and less computationally expensive. However, only the total kinetic energy in the CM frame, the sum of the three components, is invariant during a collision just as only the trace of the second moment tensor is an invariant independent of coordinate system, and so independently conserving the direction components is a somewhat arbitrary choice.

Spatial moments are conserved in an analogous manner. Merge result particles conserve spatial center of mass as well as RMS position variance in each direction independently. Future work in progress focuses on investigating conservation of mixed spatial-velocity moments as well as higher order moment conservation via merging to sets larger than two particles as well.

B. Octree Velocity Bins

The binning of particles into the octree structure originally proposed⁶ as an equivalent to a simple recursive sorting with a few constraints on velocity bin subdivision. For the bounding box of particle velocities in a




Figure 1: Three-level octree binning diagram

given spatial cell, every particle is marked as belonging in one of the eight octants. These octants are then selectively refined further. Figure 1 conceptually depicts three levels of such an octree binning procedure.

The number of particles and the sum of the computational weight in each octant is simultaneously accumulated. The particles are then sorted into a second buffer using the running summed number of particles for each octant as the first particle position in each octant and incrementing the particle position in each octant as particles are placed in the octant. Because only velocity cells with 3 or more particles can be merged using the $N : 2$ scheme, only octants with 3 or more particles can be marked for further refinement. Additionally, for a target number of particles per cell and a given space cell density, a target computational weight per particle can be defined. An additional constraint on the refinement of an octant is applied such that the octant is only subdivided further if the total weight of the octant exceeds a threshold based on the target weight. This constraint ensures that even large velocity bins with a few particles of small computational weight are merged into higher computational weight particles preferentially to the merging of particles with large computational weight in small velocity bins and has been verified⁷ through the evolution of particle weight distribution through repeated application of merging followed by the splitting procedure described in the following section. The choice essentially pushes the particle weights towards the target weight and therefore the number of particles per cell towards that target. Though a seemingly trivial point, it is actually an important modification to the algorithm that ensures that the core of the VDF is not represented by only a few super massive particles surrounded by a cloud of very light outliers. Only sufficiently light particles in sufficiently small velocity cells are merged such that the distribution of particle weights fall within a reasonable bound around the target mean. Once octants have been marked for further subdivision, the original and temporary buffers are swapped and each refined octant undergoes the same process copying the resorted particles back into the original buffer. Figure 2 depicts this sorting procedure though for clarity only a binary sort with two colors is shown. The full octree version would use eight colors and produce eight child cells.

C. Particle Splitting

For fully dynamic adaptive particle weights, an analogous particle split method is necessary to re-populate depleted velocity distributions. These may result from the particle merging or from the distribution being forced to occupy a greater volume such as in a plasma plume expansion. The split is achieved by performing similar operations to the merge using the conservative moment sums within velocity bins as well.

The particles in a given velocity cell are first conceptually split while retaining their original spatial and velocity coordinates. As macro-particles already represent many physical particles, this fission process is purely conceptual and has no impact on any conservation property. However, multiple macro-particles with identical positions in phase-space do not provide additional fidelity to the simulation. These newly created




Figure 2: Recursive velocity cell index particle sorting procedure.

fragments are instead merged using the same conservative merge algorithm described previously. The only exception is a slight modification to subdivide particles within a velocity bin further to compensate for the weak growth in the number of computational particles it provided when the number is relatively large.

After merging, the particles created from the fragments no longer reside in identical points in phase space as the original remaining fragments while the moments remain conserved and the velocity distribution is preserved up to the fidelity dictated by the velocity binning process. These new particles repopulate the phase space following slightly different though potentially divergent trajectories than the particles from which they were created. This is particularly important in regions of rapid expansion such as expansion around high angle edges. Though not included in the current model, this is expected to be particularly relevant as particles exit the thruster into the near-plume region once resolved.

For the merged and original fragments to tend towards approximately the same computational weight, the original weights of particles is split by the factor $f = N_s / (N_s + 2)$, where N_s is the number of particles in the original set of particles to be split. If, for example, the original set has 4 particles, $f = 2/3$. Assuming the original particles were the same weight, w_0 , each of the original particles' weight is scaled down by f to weigh fw_0 or $2w_0/3$. The remainder pieces to be merged would have a total weight of $4(1 - f)w_0$ or again $2w_0/3$ per merge product particle when divided between the two resulting particles. Similarly, $f = 3/5$ for $N_s = 3$. If the original set of particles have all different computational weights, the particles resulting from the merge portion have the same weight as the mean of the original fragments as well. This again helps to push the particle weight towards a mean value rather than extrema and is important when merging and splitting is applied continuously.

As demonstrated for collisional simulations,⁹ another method of repopulating the phase space is to modify the collision models to a fractional form rather than the random all-or-nothing form characteristic of DSMC and MCC algorithms, but this is beyond the scope of the current investigation.

III. Results

Figure 3 shows traces of the macroparticle count and wall-clock per iteration for the HET discharge both without (control) and with (dynamic) adaptive phase space remapping. The unmerged solution stabilizes at approximately 4.3 million macroparticles and the dynamic remapping case at about 1.7 million. The unmerged iteration time similarly stabilizes around 13s per iteration whereas the remapping case is around 9s. This corresponds to about a 1.4x speedup with a 2.5x reduction in particle count. Considering the additional cost of the merging and splitting operations must be performed, these results appear to be reasonable. Though the time spent in the added merge and split operations along with their corresponding sort routines represent only about 8.4% of the total runtime for the merge case, the other components such as the 18% time spent just in the spectral potential solve also prevent speedups proportional to the

relative particle count reduction. The key observation is that despite the relatively complex recursive nature of the merge and split algorithm, the merging and splitting operations represent only a small manageable computational cost compared to the rest of the simulation.




Figure 3: Comparison of original and dynamically remapped total macroparticle counts (millions) and wall clock per iteration. Because macroparticle count saturates quickly for the control case, dynamic remapping demonstrates little gain for this problem.

Figure 4 compares the ion density and macroparticle counts averaged over the Θ -direction over time. It is clear that the dynamic merge and split control the particle count within a relatively narrow band in macroparticles per cell. However, because the ions exit the computational domain in relatively few iterations due to the strong imposed field, the number of macroparticles remains tractable for the entire simulation window in the control case. The results show that the dynamic remapping does not impair the solution structure, but is unnecessary for the specified inputs.

Similar to the time-history plot, the θ - z -plane density results appear similar up to the statistical noise as demonstrated in Figure 5. The primary perceptible difference is a slight increase in this noise downstream in the dynamically controlled case. Again, the macroparticle count piece of the plot demonstrates the tight control over particle count the method affords with minimal impact on solution density results.

Finally, the axial-velocity distributions are shown in Figure 6. As in the prior plots, the only apparent difference in the density plots is a slight increase in noise in the dynamically controlled case. In the mean macroparticle weight plots, the dynamically controlled case shows similar structure to the control case. The primary difference is increased peak average velocity along the arc from the strong ionization front (violet extension to colormap) and increased noise. Much of this weight structure is the result of the variable ionization rate with a fixed number of added ionization macroparticles per cell per timestep. The merging and splitting retains most of the VDF structure despite enforcing a nearly constant particle count per spatial cell.




Figure 4: Comparison of PIC ion densities for original and dynamically remapped particle populations. Data is averaged over θ -direction and plotted over time. The figure demonstrates little impact on number density structure while the dynamically controlled case maintains nearly uniform numbers of macroparticles per cell.

IV. Conclusion and Future Work

Because the ions rapidly stream out of the solution domain for the baseline HET discharge setup, there is little need for dynamic particle remapping in this problem. The ionization model used in this work samples new ions from a thermal population and weights them according to the ionization reaction rate using the ion and electron densities. Unlike the chain branching $e^- + Ar^0 \rightarrow e^- + e^- + Ar^+$ reaction used in ionizing breakdown cases in prior work, this model already inherently adjusts macroparticle weight as ions are created to avoid runaway macroparticle counts. The dynamic remapping appears to provide similar results to the original solution, but not a compelling reason to include it with this model.

If the model were switched back to using neutral macroparticles to accurately simulate the kinetic features of neutrals as well, the need for merging and splitting in that population would be likely much more relevant. The ion neutralizing channel wall collisions would inject a significantly different neutral population with vastly different macroparticle weight. These slow cold neutrals would have a much longer residence time in the simulation. In addition with fully kinetic neutrals, the injection of ionization particles ought to be sampled from the neutral VDF rather than an assumed thermal profile as the ionization reaction would have minor impact on the neutral momentum. Because the velocities imparted by the potential are much larger than the neutral kinetic velocities, this distributions may play little role in the axial results. However, in directions transverse to the field, more accurate distribution would potentially impact wall neutralization rates and plume shape.

This topics will be explored further in future work. However, as in the companion paper,⁸ the primary concern for the development of these models is incorporation of an electron energy equation. As in experimental work, breathing modes were observed in both the work of Coche¹⁰ (full-PIC) as well as Lam and Fernandez¹¹ (hybrid-PIC). In particular, the breathing modes observed in Coche's results were particularly large implying a larger potential benefit to tight control of macroparticle count. The addition of dynamic particle weight remapping shows minimal detrimental impact to solution quality and improves computational performance, demonstrating its potential as an useful additional tool for more challenging computational problems to be addressed in the future.




Figure 5: Ion density (top) and macroparticle count (bottom) at $t = 15\mu\text{s}$ for control case (left) and dynamic remapping (right). Mean flow is from bottom (anode) to top (cathode).

References

- ¹Lapenta, G. and Brackbill, J. U., “Dynamic and Selective Control of the Number of Particles in Kinetic Plasma Simulations,” *J. Comp. Phys.*, Vol. 115, No. 1, 1994, pp. 213 – 227.
- ²Hewett, D. W., “Fragmentation, merging, and internal dynamics for PIC simulation with finite size particles,” *J. Comp. Phys.*, Vol. 189, No. 2, 2003, pp. 390 – 426.
- ³Assous, F., Dulimbert, T. P., and Segré, J., “A new method for coalescing particles in PIC codes,” *J. Comp. Phys.*, Vol. 187, No. 2, 2003, pp. 550 – 571.
- ⁴Welch, D., Genoni, T., Clark, R., and Rose, D., “Adaptive particle management in a particle-in-cell code,” *J. Comp. Phys.*, Vol. 227, No. 1, 2007, pp. 143 – 155.
- ⁵Cambier, J.-L., “AFOSR Computational Mathematics Program Contractor’s Review,” .
- ⁶Martin, R. and Cambier, J.-L., “Moment Preserving Adaptive Particle Weights using Octree Velocity Distributions for PIC Simulations,” *Proceedings of the 28th International Symposium on Rarefied Gas Dynamics*, 2012, pp. 872–879.
- ⁷Martin, R. and Cambier, J.-L., “Octree Particle Management for DSMC and PIC Simulations. Part I: Basic Algorithms,” .
- ⁸Koo, J. and Martin, R., “Pseudospectral model for hybrid PIC Hall-effect thruster simulation,” *34th Int. Electric Propulsion Conf. (IEPC)*, Kobe, Japan, July 2015, pp. IEPC-2015-XXX.
- ⁹Martin, R. and Cambier, J.-L., “Low Noise Fractional NTC Collisions for DSMC,” *Proceedings of the 28th International Symposium on Rarefied Gas Dynamics*, 2012, pp. 615–620.
- ¹⁰Coche, P. and Garrigues, L., “A two-dimensional (azimuthal-axial) particle-in-cell model of a Hall thruster,” *Physics of Plasmas*, Vol. 21, 2014, pp. 023503–1–023503–10.
- ¹¹Lam, C. M., Fernandez, E., and Cappelli, M. A., “A 2-D Hybrid Hall Thruster Simulation That Resolves the ExB Electron Drift Direction,” *IEEE Transactions on Plasma Science*, Vol. 43, No. 1, 2015, pp. 86–94.




Figure 6: Axial velocity distribution (VDF) for control case (left) and dynamic remapping (right). These figures plot particle statistical weight within $z - v_z$ bins. Top shows ion VDF summed bin weight and bottom shows bin averaged macroparticle physical to computational weight ratio, $\bar{w}$ for all particles within the axial-velocity bin.