AP-Cloud: Adaptive Particle-in-Cloud Method for Optimal Solutions to Vlasov-Poisson Equation

Xingyu Wang^a, Roman Samulyak^{1a,b}, Xiangmin Jiao^a, Kwangmin Yu^{a,b}

^a*Department of Applied Mathematics and Statistics, Stony Brook University, Stony Brook, NY 11794*

^b*Computational Science Center, Brookhaven National Laboratory, Upton, NY 11973*

Abstract

We propose a new adaptive Particle-in-Cloud (AP-Cloud) method for obtaining optimal numerical solutions to the Vlasov-Poisson equation. Unlike the traditional particle-in-cell (PIC) method, which is commonly used for solving this problem, the AP-Cloud adaptively selects computational nodes or particles to deliver higher accuracy and efficiency when the particle distribution is highly non-uniform. Unlike other adaptive techniques for PIC, our method balances the errors in PDE discretization and Monte Carlo integration, and discretizes the differential operators using a generalized finite difference (GFD) method based on a weighted least square formulation. As a result, AP-Cloud is independent of the geometric shapes of computational domains and is free of artificial parameters. Efficient and robust implementation is achieved through an octree data structure with 2:1 balance. We analyze the accuracy and convergence order of AP-Cloud theoretically, and verify the method using an electrostatic problem of a particle beam with halo. Simulation results show that the AP-Cloud method is substantially more accurate and faster than the traditional PIC, and it is free of artificial forces that are typical for some adaptive PIC techniques.

Keywords: particle method, generalized finite difference, PIC, AMR-PIC
2000 MSC: 65M06, 70F99, 76T10

¹Corresponding author, roman.samulyak@stonybrook.edu

1. Introduction

The Particle-in-Cell (PIC) method [1] is a popular method for solving the Vlasov-Poisson equations for a class of problems in plasma physics, astrophysics, and particle accelerators, for which electrostatic approximation applies, as well as for solving the gravitational problem in cosmology and astrophysics. In such a hybrid particle-mesh method, the distribution function is approximated using particles and the Poisson problem is solved on a rectangular mesh. Charges (or masses) of particles are interpolated onto the mesh, and the Poisson problem is discretized using finite differences or spectral approximations. On simple rectangular domains, FFT methods are most commonly used for solving the Poisson problem. In the presence of irregular boundaries, finite difference approximations are often used, complemented by a cut-cell (a.k.a. embedded boundary) method [2] for computational cells near boundaries, and fast linear solvers (including multigrid iterations) for the corresponding linear system. The computed force (gradient of the potential) on the mesh is then interpolated back to the location of particles. For problems with irregular geometry, unstructured grid with finite element method is often used.

The traditional PIC method has several limitations. It is less straightforward to use for geometrically complex domains. The aforementioned embedded boundary method, while maintaining globally second order accuracy for the second order finite difference approximation, usually results in much larger errors near irregular boundaries [3]. It is also difficult to generalize to higher order accuracy.

Another major drawback of the PIC method is associated with highly non-uniform distribution of particles. As shown in Section 2, the discretization of the differential operator and the right hand side in the PIC method is not balanced in terms of errors. The accuracy is especially degraded in the presence of non-uniform particle distributions. The AMR-PIC [4, 5] improves this problem by performing block-structured adaptive mesh refinement of a rectangular mesh, so that the number of particles per computational cell is approximately the same. However, the original AMR-PIC algorithms suffered from very strong artificial self-forces due to spurious images of particles across boundaries between coarse and refined mesh patches. Analysis of self-force sources and a method for their mitigation was proposed in [6].

In this paper, we propose a new adaptive Particle-in-Cloud (AP-Cloud) method for obtaining optimal numerical solutions to the Vlasov-Poisson

equation. Instead of a Cartesian grid as used in the traditional PIC, the AC-Cloud uses adaptive computational nodes or particles with an octree data structure. The quantity characterizing particles (charge in electrostatic problems or mass in gravitational problems) is assigned to computational nodes by a weighted least squares approximation. The partial differential equation is then discretized using a generalized finite difference (GFD) method and solved with fast linear solvers. The density of nodes is chosen adaptively, so that the error from GFD and that from Monte Carlo integration are balanced, and the total error is approximately minimized. The method is independent of geometric shape of computational domains and free of artificial self-forces.

The remainder of the paper is organized as follows. In Section 2, we analyze numerical errors of the traditional PIC method and formulate optimal refinement strategy. The AP-Cloud method, generalized finite differences, and the relevant error analysis are presented in Section 3. Section 4 describes some implementation details of the method. Section 5 presents numerical verification tests using 2D and 3D problems of particle beams with halo and additional tests demonstrating the absence of artificial self-forces. We conclude this paper with a summary of our results and perspectives for the future work.

2. Error analysis of particle-in-cell method

In a particle-in-cell (PIC) method, the computational objects include a large number of particles and an associated Cartesian grid. These particles are typically randomly sampled and represent an even greater number of physical particles (e.g., protons), so they are also known as “macro-particles”, but conventionally simply referred to as particles. For simplicity of presentation, we will focus on electrostatic problems, for which the states are particle charges. Suppose there are N charged particles at positions $\mathcal{P} = \{\mathbf{p}^i \mid i = 1, 2, \dots, N\}$ in D dimensions, and let q^i denote the charge at $\mathbf{p}^i$. For simplicity, we assume that all the particles carry the same amount of charge, the total charge is 1, i.e., $Nq^i = \int_{\Omega} \rho(\mathbf{x}) d\mathbf{x} = 1$, and the charges of the particles can be represented accurately by a continuous charge distribution function ρ . We assume that ρ is smooth and positive, and its value and all derivatives have comparable magnitude. Let M denote the Cartesian grid, and without loss of generality, suppose its edge length is h along all directions, and let $\mathbf{y}^j$ denote the j th grid point in M . A PIC method estimates

the charge density ρ on M , then solves the Poisson equation

$$\Delta\phi = c\rho \quad (1)$$

on M to obtain the potential ϕ , whose gradient is the electric field $\mathbf{E}$. In this setting, a PIC method consists of the following three steps:

1. Approximate the right-hand side of (1) by interpolating the states from particles to the grid points $\mathbf{y}^j$, i.e.,

$$\tilde{\rho}(\mathbf{y}^j, \mathcal{P}, h) = \frac{1}{h^D} \sum_{i=1}^N q^i \Phi\left(\frac{\mathbf{p}^i - \mathbf{y}^j}{h}\right) \approx \underbrace{\frac{1}{h^D} \int_{\Omega} \rho(\mathbf{x}) \Phi\left(\frac{\mathbf{x} - \mathbf{y}^j}{h}\right) d\mathbf{x}}_{\tilde{\rho}(\mathbf{y}^j, h)} \approx \rho(\mathbf{y}^j), \quad (2)$$

where Φ is the interpolation kernel, a.k.a. the charge assignment scheme.

2. Discretize the left-hand side of (1) on M , typically using the finite difference method, and then solve the resulting linear system.
3. Obtain the electric field $\mathbf{E}$ by computing $\nabla\phi$ using finite difference, and then interpolating $\mathbf{E}$ from the grid points to the particles using the same interpolation kernel Φ as in Step 1, i.e.,

$$\tilde{\mathbf{E}}(\mathbf{p}^i, \mathcal{P}, h) = \frac{1}{h^D} \sum_{\mathbf{y}^j \in M} \tilde{\mathbf{E}}(\mathbf{y}^j, \mathcal{P}, h) \Phi\left(\frac{\mathbf{p}^i - \mathbf{y}^j}{h}\right) \approx \mathbf{E}(\mathbf{p}^i). \quad (3)$$

One of the most commonly used charge assignment schemes Φ is the cloud-in-cell (CIC) scheme

$$\Phi(\mathbf{x}) = \prod_{d=1}^D \max(1 - |x_d|, 0), \quad (4)$$

for which the interpolation in Step 3 corresponds to bilinear and trilinear interpolation in 2-D and 3-D, respectively.

In PIC, the error in potential ϕ comes from two sources. One is from the first approximation in (2), for which the analysis is similar to Monte Carlo integration within a control volume associated with $\mathbf{y}^j$, under the assumption that (2) is a continuous function. The other source is the discretization error of both the second approximation in (2) in Step 1 and the left-hand side of (1) on M in Step 2. We denote the above two errors from these two sources

as $\mathcal{E}_M$ and $\mathcal{E}_D$, respectively. As shown in Appendix A, under the assumption that the interpolation kernel Φ satisfies the positivity condition, the expected value of the former is

$$\mathbb{E}[\|\mathcal{E}_M\|] = \mathcal{O} \left(\sqrt{\frac{\rho(\mathbf{y})}{Nh^D}} \right),$$

and the discretization error is

$$\mathcal{E}_D = \mathcal{O}(\rho(\mathbf{y})h^2).$$

Let A denote the coefficient matrix of the linear system in step 2, and suppose $\|A^{-1}\|$ is bounded by a constant. The expected total error in the computed potential ϕ is then

$$\mathcal{E} = \mathcal{O}(\mathbb{E}[\|\mathcal{E}_M\|]) + \mathcal{O}(\mathcal{E}_D) = \mathcal{O} \left(\sqrt{\frac{\rho(\mathbf{y})}{Nh^D}} + \rho(\mathbf{y})h^2 \right).$$

In general, $\mathcal{E}_D$ dominates the total error for coarse grids and $\mathcal{E}_M$ dominates for finer grids. The total expected error is approximately minimized if $\mathcal{E}_M$ and $\mathcal{E}_D$ are balanced. If the particles are uniformly distributed, then the errors are balanced when

$$h = \mathcal{O}(N\rho(\mathbf{y}))^{-\frac{1}{4+D}}. \quad (5)$$

In this setting, the discretization error in ϕ is second order in h . The discretization error in numerical differentiation $\mathbf{E} = \nabla\phi$ is also second order, a fact called supraconvergence [12]. Thus, although the optimal mesh size is deduced to minimize the error in ϕ , the error in $\mathbf{E}$ is also minimized.

In many applications, the particle distribution is highly non-uniform, for which the PIC is neither efficient nor accurate. In [5], an adaptive method, called AMR-PIC, was proposed, which fixed the number of particles per cell, and hence

$$h = \mathcal{O}(N\rho(\mathbf{y}))^{-\frac{1}{D}}.$$

The AMR-PIC over-refines the grid compared to the optimal grid resolution in (5). In addition, the original AMR-PIC technique also introduces artificial self forces. New adaptive strategies are needed to resolve both of these issues.

3. Adaptive Particle-in-Cloud method

In this section, we describe a new adaptive method, called Adaptive Particle-in-Cloud or AP-Cloud, which approximately minimizes the error by balancing Monte Carlo noise and discretization error, and at the same time is free of the artificial self forces present in AMR-PIC.

The AP-Cloud method can be viewed as an adaptive version of PIC that replaces the traditional Cartesian mesh of PIC by an octree data structure. We use a set of computational nodes, which are octree cell centres, instead of the Cartesian grid, of which the distribution is derived using an error balance criterion. Computational nodes will be referred to as nodes in the remainder of the paper. Instead of the finite difference discretization of the Laplace operator, we use the method of generalized finite-difference (GFD) [9], based on a weighted least squares formulation. The framework includes interpolation, least squares approximation, and numerical differentiation on a stencil in the form of cloud of nodes in a neighborhood of the point of interest. It is used for the charge assignment scheme, numerical differentiation, and interpolation of solutions. The advantage of GFD is that it can treat coarse regions, refined regions, and refinement boundaries in the same manner, and it is more flexible for problems in complex domain or with irregular refinement area. As a method of integration, GFD will be used in the quadrature rule in charge assignment scheme. The new charge assignment scheme, together with GFD differentiation and interpolation operators from computational nodes to particles is easily generalizable to higher order schemes. We described the key components of the AP-Cloud in this section, and discuss its implementation details in Section 4.

3.1. Generalized finite-difference method

For simplicity of presentation, we consider a second order generalized finite-difference method.

Let $\mathbf{y}^j$, $j \in \{1, 2, \dots, m\}$ be the nodes in a neighborhood of reference node $\mathbf{y}^0$. Given a C^2 function f , by Taylor expansion we have

$$f(\mathbf{y}^j) = f(\mathbf{y}^0) + (\mathbf{y}^j - \mathbf{y}^0)^T \nabla f(\mathbf{y}^0) + (\mathbf{y}^j - \mathbf{y}^0)^T H(\mathbf{y}^0) (\mathbf{y}^j - \mathbf{y}^0) + O(h^3), \quad (6)$$

where h is the characteristic interparticle distance in the neighborhood, for example, $h = \max_{d,j} |y_d^j - y_d^0|$, and H is the Hessian matrix. Putting equations for all neighbors together and omitting higher order term, we obtain

$$V(\mathbf{y}^0) \partial f(\mathbf{y}^0) = \delta f(\mathbf{y}^0), \quad (7)$$

where $V(\mathbf{y}^0)$ is a generalized Vandermonde matrix, $\partial f(\mathbf{y}^0)$ is the first order and second order derivative of f at $\mathbf{y}^0$, and $\delta f(\mathbf{y}^0)$ is the increment of f .

For example, in 2-D, let $m = 5$ and $\delta y_d^j = y_d^j - y_d^0$, then

$$V(\mathbf{y}^0) = \begin{bmatrix} \delta y_1^1 & \delta y_2^1 & \frac{1}{2}(\delta y_1^1)^2 & \delta y_1^1 \delta y_2^1 & \frac{1}{2}(\delta y_2^1)^2 \\ \delta y_1^2 & \delta y_2^2 & \frac{1}{2}(\delta y_1^2)^2 & \delta y_1^2 \delta y_2^2 & \frac{1}{2}(\delta y_2^2)^2 \\ \delta y_1^3 & \delta y_2^3 & \frac{1}{2}(\delta y_1^3)^2 & \delta y_1^3 \delta y_2^3 & \frac{1}{2}(\delta y_2^3)^2 \\ \delta y_1^4 & \delta y_2^4 & \frac{1}{2}(\delta y_1^4)^2 & \delta y_1^4 \delta y_2^4 & \frac{1}{2}(\delta y_2^4)^2 \\ \delta y_1^5 & \delta y_2^5 & \frac{1}{2}(\delta y_1^5)^2 & \delta y_1^5 \delta y_2^5 & \frac{1}{2}(\delta y_2^5)^2 \end{bmatrix}, \quad (8)$$

$$\partial f(\mathbf{y}^0) = [f_{y_1}(\mathbf{y}^0), f_{y_2}(\mathbf{y}^0), f_{y_1 y_1}(\mathbf{y}^0), f_{y_1 y_2}(\mathbf{y}^0), f_{y_2 y_2}(\mathbf{y}^0)]^T, \quad (9)$$

where $f_{y_d}(\mathbf{y}^0)$ denotes the derivative of f with respect to y_d , and

$$\delta f(\mathbf{y}^0) = [f(\mathbf{y}^1) - f(\mathbf{y}^0), f(\mathbf{y}^2) - f(\mathbf{y}^0), \dots, f(\mathbf{y}^5) - f(\mathbf{y}^0)]^T. \quad (10)$$

To analyze the error in GFD, let $\xi_d^j = h^{-1} \delta y_d^j$. Rewrite (7) as

$$V_0(\mathbf{y}^0) \partial f_0(\mathbf{y}^0) = \delta f(\mathbf{y}^0), \quad (11)$$

where

$$V_0(\mathbf{y}^0) = \begin{bmatrix} \xi_1^1 & \xi_2^1 & \frac{1}{2}(\xi_1^1)^2 & \xi_1^1 \xi_2^1 & \frac{1}{2}(\xi_2^1)^2 \\ \xi_1^2 & \xi_2^2 & \frac{1}{2}(\xi_1^2)^2 & \xi_1^2 \xi_2^2 & \frac{1}{2}(\xi_2^2)^2 \\ \xi_1^3 & \xi_2^3 & \frac{1}{2}(\xi_1^3)^2 & \xi_1^3 \xi_2^3 & \frac{1}{2}(\xi_2^3)^2 \\ \xi_1^4 & \xi_2^4 & \frac{1}{2}(\xi_1^4)^2 & \xi_1^4 \xi_2^4 & \frac{1}{2}(\xi_2^4)^2 \\ \xi_1^5 & \xi_2^5 & \frac{1}{2}(\xi_1^5)^2 & \xi_1^5 \xi_2^5 & \frac{1}{2}(\xi_2^5)^2 \end{bmatrix}, \quad (12)$$

$$\partial f_0(\mathbf{y}^0) = [h f_{y_1}(\mathbf{y}^0), h f_{y_2}(\mathbf{y}^0), h^2 f_{y_1 y_1}(\mathbf{y}^0), h^2 f_{y_1 y_2}(\mathbf{y}^0), h^2 f_{y_2 y_2}(\mathbf{y}^0)]^T. \quad (13)$$

Now $V_0(\mathbf{y}^0)$ depends on the shape but not the diameter of the GFD stencil. The error in solving linear system is

$$\|\text{Error}(\partial f_0(\mathbf{y}^0))\| \leq \|V_0^{-1}(\mathbf{y}^0)\| \|\text{Error}(\delta f(\mathbf{y}^0))\|. \quad (14)$$

The error in right hand size comes from the omitted term in the Taylor expansion $\mathcal{O}(f_{y_i y_j y_k} h^3) = \mathcal{O}(f h^3)$ for $i, j, k \in 1, 2$, and $\|V_0^{-1}(\mathbf{y}^0)\|$ is a constant independent of h , so the error in $\partial f_0(\mathbf{y}^0)$ is also $\mathcal{O}(f h^3)$. Because the coefficient before the l th order derivative in $\partial f_0(\mathbf{y}^0)$ is h^l , the error for l th order derivative is $\mathcal{O}(f h^{3-l})$.

In this example, the number of neighbors is equal to the number of unknowns, and it is quite likely for $V_0(\mathbf{y}^0)$ to be nearly singular. In practice,

we use more neighbors in the stencil than the number of coefficients in the Taylor series to improve the stability. In AP-Cloud method, 8 neighbors instead of 5 are used for the second order GFD in two dimensions, and 17 neighbors instead of 9 are used in three dimensions. In this case, the linear system is a least square problem. It is often helpful to assign more weights to closer neighbors to improve the accuracy, which is called the weighted least square method [8]. AP-Cloud method uses a normalized Gaussian weight function [7]:

$$W(r) = \frac{e^{-r^2/r_{max}^2} - e^{-c}}{1 - e^{-c}}, \quad (15)$$

where W is the weight, r is the distance of the neighbor from the reference particle, r_{max} is the maximum distance of all neighbors in the stencil from the reference particle, $c = 4$.

By solving the linear system or least square problem (7), we can express the gradient ∂f as linear combinations of δf . For example, once the potential ϕ is computed at nodes, we can find its gradient by generalized finite-difference, and then interpolate it to particles by Taylor expansion. Generally, the error of the k th order GFD interpolation is $\mathcal{O}(f(\mathbf{y})h^{k+1})$, and its approximation of the l th order derivative is $\mathcal{O}(f(\mathbf{y})h^{k-l+1})$.

Given a set of nodes, the selection for GFD neighbors, or the shape of the GFD stencil, is important for both accuracy and stability. Simply choosing the nearest nodes to be neighbors may lead to an imbalanced stencil. We follow the quadrant criterion in [10] and select two nearest nodes from each quadrant to be neighbors.

3.2. Algorithm of AP-Cloud method

AP-Cloud also has three steps to calculate the electric field given by a particle distribution: a density estimator, a Poisson solver, and an interpolation step, but each step is different from its counterpart in PIC. Let $\mathcal{C}$ be the set of all computational nodes, and $f(\mathcal{C}) = (f(\mathbf{y}^1), f(\mathbf{y}^2), \dots, f(\mathbf{y}^n))$, where n is the total number of nodes. Below is a detailed description of the three steps.

1. Approximate density by interpolating states from particles $\mathcal{P}$ to com-

putational nodes $\mathcal{C}$ by

$$\begin{aligned}
& a(\mathbf{y}^j)^T \partial \tilde{\rho}(\mathbf{y}^j) \\
: &= \sum_{l=0}^k \sum_{d_1, d_2, \dots, d_l=1}^D \underbrace{\frac{1}{h^D} \int_{\Omega} \prod_{i=1}^l (x_{d_i} - y_{d_i}^j) \Phi\left(\frac{\mathbf{x} - \mathbf{y}^j}{h}\right) d\mathbf{x}}_{a_{d_1, d_2, \dots, d_l}(\mathbf{y}^j)} \tilde{\rho}_{y_{d_1} y_{d_2} \dots y_{d_l}}(\mathbf{y}^j) \\
&= \underbrace{\frac{1}{h^D} \sum_{i=1}^N q^i \Phi\left(\frac{\mathbf{p}^i - \mathbf{y}^j}{h}\right)}_{\rho_M(\mathbf{y}^j, h)}.
\end{aligned} \tag{16}$$

The right hand side of (16) is identical to the Monte Carlo integration in PIC, but the left hand side is a linear combination of derivatives, instead of the simple $\tilde{\rho}(\mathbf{y}^j, \mathcal{P}, h)$ in the PIC method. Because the coefficients in the linear combination, $a(\mathbf{y}^j) = (a_{d_1, d_2, \dots, d_l}(\mathbf{y}^j))$, depend only on h and the interpolation kernel, they can be easily pre-calculated and tabulated in a lookup table. The derivatives $\tilde{\rho}_{y_{d_1} y_{d_2} \dots y_{d_l}}(\mathbf{y}^j)$ are in turn linear combinations of density values $\rho(\mathbf{y}^j)$ given by least square solution of (7):

$$\partial \tilde{\rho}(\mathbf{y}^j) = V(\mathbf{y}^j)^+ \delta \tilde{\rho}(\mathbf{y}^j), \tag{17}$$

where $V(\mathbf{y}^j)^+$ is the pseudo-inverse of the Vandermonde matrix.

Let $C(\mathbf{y}^j)$ be a matrix such that $\delta \tilde{\rho}(\mathbf{y}^j) = C(\mathbf{y}^j) \tilde{\rho}(\mathcal{C})$. Substituting (17) into (16), we get a linear equation for density values at the reference node $\mathbf{y}^j$ and its neighbors

$$\underbrace{a(\mathbf{y}^j)^T V(\mathbf{y}^j)^+ C(\mathbf{y}^j)}_{b(\mathbf{y}^j)} \tilde{\rho}(\mathcal{C}) = \rho_M(\mathbf{y}^j, h). \tag{18}$$

Putting equations for all nodes together, we obtain a global linear system for density values

$$B(\mathcal{C}) \tilde{\rho}(\mathcal{C}) = \rho_M(\mathcal{C}, h). \tag{19}$$

where $B(\mathcal{C}) = [b(\mathbf{y}^1); b(\mathbf{y}^2); \dots; b(\mathbf{y}^n)]$. Solution of (19) is the estimated density in AP-Cloud method.

2. Discretize the left-hand side of (1) on $\mathcal{C}$ using GFD method. Solve the resulting linear system for $\phi(\mathcal{C})$.
3. Obtain the electric field $\mathbf{E}$ by computing $\nabla \phi$ using GFD method, and then interpolating ϕ and $\mathbf{E}$ from the $\mathcal{C}$ to $\mathcal{P}$ using a Taylor expansion.

3.3. Error analysis for AP-Cloud

Similar to PIC, the error in AP-Cloud also contains the Monte Carlo noise $\mathcal{E}_M$ and the discretization error $\mathcal{E}_D$. The Monte Carlo noise in replacing $\bar{\rho}(\mathbf{y}^j, h)$ by $\rho_M(\mathbf{y}^j, h)$ is identical to the Monte Carlo noise in PIC, that is, $\mathcal{O}(\sqrt{\frac{\rho(\mathbf{y}^j)}{Nh^D}})$. The discretization error in the first step has two sources: the Taylor expansion in (16) and the GFD approximation of the gradient in (17). The difference between the average of the k th order Taylor expansion of ρ and ρ itself is $\mathcal{O}(\rho(\mathbf{y}^j)h^{2\lfloor \frac{k}{2} \rfloor + 2})$, where we obtain an additional order for even k due to the symmetry of the kernel Φ . Because the error of l th order derivative is $\mathcal{O}(\rho(\mathbf{y}^j)h^{k-l+1})$, and the coefficient for l th order derivative $a_{d_1, d_2, \dots, d_l}(\mathbf{y}^j) = \mathcal{O}(h^l)$, the discretization error given by the GFD derivative approximation is $\mathcal{O}(\rho(\mathbf{y}^j)h^{k+1}) \geq \mathcal{O}(\rho(\mathbf{y}^j)h^{2\lfloor \frac{k}{2} \rfloor + 2})$. Thus the total discretization error in step 1 is $\mathcal{O}(\rho(\mathbf{y}^j)h^{k+1})$.

Generalized finite-difference Poisson solver has the same accuracy with its estimation of $\Delta\phi$, i.e., $\mathcal{O}(\rho(\mathbf{y})h^{k-1})$. However, for the second order GFD we observe a supraconvergence when the GFD stencil is well-balanced due to the error cancellation similar to that in the standard five point finite difference stencil. In this case, the error for the solution and its gradient of the GFD Poisson solver are both $\mathcal{O}(f(\mathbf{y})h^2)$, as observed in our numerical experiments. The interpolation from $\mathcal{C}$ to $\mathcal{P}$, based on the k th order Taylor expansion of ϕ where the derivatives are given by GFD, is $(k+1)$ th order accurate for ϕ and k th order accurate for $\mathbf{E}$.

3.4. Refinement strategy for AP-Cloud

Generally, when k th order GFD is used in the charge assignment scheme, the Poisson solver, and the differentiation and interpolation routines, the total error for both ϕ and $\mathbf{E}$ are

$$\mathcal{E} = \mathcal{O}(\mathbb{E}[|\mathcal{E}_M|]) + \mathcal{O}(\mathcal{E}_D) = \mathcal{O}\left(\sqrt{\frac{\rho(\mathbf{y})}{Nh^D}} + \rho(\mathbf{y})h^{k-1}\right), \quad (20)$$

where the $\mathcal{O}(\rho(\mathbf{y})h^{k-1})$ leading term in discretization error is from GFD Poisson solver. To minimize the error, the optimal mesh size is

$$h = \mathcal{O}\left(\frac{1}{N\rho(\mathbf{y})}\right)^{\frac{1}{2k+D-2}}, \quad (21)$$

and the minimized error is

$$\mathcal{E}_{total} = \rho(\mathbf{y})^{\frac{k+D-2}{2k+D-2}} N^{-\frac{k-1}{2k+D-2}}. \quad (22)$$

For the second order GFD in particular, we have better error bound due to the symmetry of interpolation kernel and stencil and supraconvergence

$$\mathcal{E} = \mathbb{E}[\mathcal{E}_M] + \mathcal{E}_D = \mathcal{O} \left(\sqrt{\frac{\rho(\mathbf{y})}{Nh^D}} + \rho(\mathbf{y})h^2 \right), \quad (23)$$

and the optimal mesh size is the same as in (5)

$$h = \mathcal{O} (N\rho(\mathbf{y}))^{-\frac{1}{4+D}}. \quad (24)$$

4. Implementation

We use a 2^D -tree data structure to store particles, and select some of its cell centres as computational nodes. The 2^D -tree data structure is a tree data structure in a D -dimensional space in which each cell has at most 2^D children. Quadtree and octree are standard terms in 2D and 3D spaces, respectively.

The algorithm in [13] is used in the 2^D -tree construction. The first step is to sort the particles by their Morton key, so particles in the same cell are contiguous in the sorted array. Then leaf cells are constructed by an array traversal, during which we record the number of particles and the index of the first particle in each cell. At last the interior cells are constructed in a depth decreasing order by a traversal of cells of the deeper level. The overall time complexity, dominated by the Morton key sorting, is $\mathcal{O}(N \log N)$, where N is the total number of particles. This parallel 2^D -tree construction algorithm, together with parallel linear solver, enables efficient parallel implementation of AP-Cloud.

Because all computational nodes are cell centres of a 2^D -tree, their distribution will be similar to an AMR-PIC mesh. This is a result of our implementation method and not an internal property of AP-Cloud.

4.1. Error balance criterion

The optimal interparticle distance h given in (21) depends on the charge density ρ . In most applications, we do not know ρ in advance; otherwise, we

do not need the charge assignment scheme to estimate it. We use a Monte Carlo method to obtain a rough estimation of ρ :

$$\rho(\mathbf{y}) = \frac{N'}{NV}, \quad (25)$$

where V is the volume of a neighborhood of $\mathbf{y}$, and N' is the number of particles in the neighborhood.

If the neighborhood is the box with the edge length h centred at $\mathbf{y}$, $V = h^D$, we substitute (25) into (21) and obtain

$$h = \mathcal{O}(N'^{-\frac{1}{2k-2}}). \quad (26)$$

4.2. 2:1 mesh balance

If the charge density undergoes rapid changes, as is typical for certain applications (such as particle accelerators and cosmology), the optimal h (26) also changes rapidly. This causes two potential problems. First, when the difference between levels of refinement on two sides of a cell is too large, that cell cannot find a balanced GFD stencil. If no particle in the coarse side is chosen to be its neighbor, the information on that side is missing. If we force the algorithm to choose a particle on the coarse side as a neighbor, the truncation error from this particle is much larger than that from the others. Second, in some cases, there are almost no particles in the region near the boundary. In order to enforce the boundary condition, interior nodes need to use far away nodes located on the boundary as their neighbors.

To avoid these two problems, we enforce a 2:1 mesh balance. The 2:1 mesh balance requires that the difference between the levels of refinement of two neighbors is at most one. Because the mesh size changes smoothly, both imbalanced GFD stencils and empty regions are avoided.

4.3. Algorithm to select computational nodes and search GFD neighbors

For clarity, we will focus on the selection of nodes in 3D in this subsection. The selection of nodes in 2D is similar and easier. An octree cell and the centre of a cell will be used interchangeably in this subsection.

We say an octree cell $\mathbf{z}$ is a neighbor of another octree cell $\mathbf{y}$ in a set of octree cells S , if

- $\mathbf{y}, \mathbf{z} \in S$;
- $Cell(\mathbf{y}) \cap Cell(\mathbf{z})$ is a face.

- $level(\mathbf{z}) \leq level(\mathbf{y})$;
- No ancestor of $\mathbf{z}$ satisfies the previous three conditions.

The neighbors defined here are different from the neighbors in [13] or the neighbors in GFD stencil. Generally, for any cell $\mathbf{y}$ in any set of octree cells S , it has at most 6 neighbors, each corresponding to one of its 6 faces. It is possible for a cell to have less than 6 neighbors. For example, root cell has no neighbor in any S in a non-periodic region.

During the selection of nodes, we will keep a queue of octree cells, Q , and a list $L(\mathbf{y})$ for each cell $\mathbf{y}$ in Q containing its neighbors in Q , which we call neighbor list. The basic operation is to open a cell $\mathbf{y} \in Q$:

1. Mark $\mathbf{y}$ as non-node candidate.
2. Add all children cells of $\mathbf{y}$ at the end of Q , mark them as node candidates.
3. Initialize the neighbor lists of the new added cells. Some of the neighbors are their siblings, while others are the neighbors of $\mathbf{y}$ or children of the neighbors of $\mathbf{y}$.
4. Update the neighbor lists of the descendent of the neighbors of $\mathbf{y}$.

The algorithm for the selection of nodes is as follows.

1. Initialize a queue Q , which contains only the root cell.
2. Traverse Q . For each cell $\mathbf{y}$ in Q , test if it satisfies

$$h < cN'^{-\frac{1}{2k-2}}, \quad (27)$$

where c is a tuning parameter, h is the diameter of the subtree, n is the number of particles in the subtree. If the condition is not satisfied, open $\mathbf{y}$. Let l be the deepest level in Q at the end of this traverse.

3. Traverse Q . For each leaf cell $\mathbf{y}$ at level l , check if the neighbors of $\mathbf{y}$ satisfy the 2:1 mesh balance. Open each neighbor $\mathbf{z}$ that does not satisfy 2:1 mesh balance.
4. If $l > 0$, let $l \leftarrow l - 1$, then repeat step 3. If $l = 0$, output all node candidates as computational nodes. If non-periodic boundary condition is used, add additional nodes on the boundary.

Given the number of particles and the order of GFD is fixed, the tuning parameter c in (27) determines the number of nodes. Ideally, c can be

computed from the constant in the proportional relationship in (21), which in turn depends on the order of GFD, the kernel function, and the relative magnitude of ρ and its gradients. However, in most applications, the relative magnitude of ρ and its gradients is unknown, so we try different values of c and compare their results to estimate its optimal value in numerical tests.

Checking error balance criterion and 2:1 mesh balance takes only constant number of operations per cell. Except the part to update the neighbor lists, each open operation takes constant number of elementary operations as well, which can be charged on the 8 new added cells, so the time complexity is $\mathcal{O}(|Q|)$, where $|Q|$ is the number of cells in Q in the end of the selection. To analyze the complexity to update the neighbor lists, we note that each time we update the neighbor of a cell $\mathbf{z}$ in Q , the level of its neighbor increases. Since the level of its neighbor is bounded by the height of the octree l_{max} , the total running time to update neighbors of all particles is $\mathcal{O}(|Q|l_{max})$. Because each interior cell of Q has 8 children, Q is a complete octree, we have $|Q| < 8/7n$, where n is the number of not opened cells, i.e., the number of computational nodes. In conclusion, the complexity to select nodes is $\mathcal{O}(|Q|l_{max})$.

After selecting the nodes, the neighbor list can be used to search GFD neighbors. If $\mathbf{z}$ is a neighbor of node $\mathbf{y}$, and $\mathbf{z}$ is not a node itself, then the 4 children of $\mathbf{z}$ that share a face with $\mathbf{y}$ must be nodes because of 2:1 mesh balance. The nodes among the neighbors of $\mathbf{y}$ and the children of these neighbors that share a face with $\mathbf{y}$ are called 1-ring. The union of the k -ring and the nodes among the neighbors of k -ring nodes and the children of these neighbors that share a face with k -ring nodes are called $(k+1)$ -ring. GFD neighbors are chosen from 2-ring by the quadrant criterion if there are enough number of nodes in 2-ring. If there are not enough number of nodes in k -ring, we will try to select GFD neighbors from $(k+1)$ -ring. In our simulations, 5-ring always contains enough neighbors. Because this neighbor searching algorithm only depends on the local nodes information and local data structures, we claim the complexity to search neighbors of a node is independent from the total number of nodes, and this algorithm takes $\mathcal{O}(n)$ time to find neighbors of all computational nodes.

One problem is related to the fact that we do not know in advance how deep the nodes are in the octree while we build the octree. In other words, it is possible that during the algorithm to select nodes, a leaf cell in the octree needs to be opened. However, this is very unlikely to happen in practice, if we always use the maximum depth supported in the implementation. For

example, in the 2D Gaussian beam with halo test in Section 5, the order of GFD is 2, the region is $[-1, 1] \times [-1, 1]$, and the minimum tuning parameter c used in simulations is 0.01. If we use a 64 bits Morton key, the maximum depth of the octree is 21, the cell size of the leaf cell is $(1 - (-1)) \times 2^{-21} = 2^{-20}$. According to (27), $0.01^2 \times 2^{20 \times 2} = 1.10 \times 10^8$ particles must be in the same leaf cell in order to open it, which is more than the total number of particles in the whole domain.

5. Numerical results

The AP-Cloud method approximately minimizes the error. However, it does incur additional cost for the construction and search of the octree data structure. In addition, the fast Fourier transform can no longer be used for solving the resulting linear system, so we must replace it with a sparse linear solver. Therefore, the practical advantage of AP-Cloud is by no means obvious. In this section, we present some numerical results for problems with non-uniform distribution, and demonstrate the advantages of AP-Cloud in terms of both accuracy and efficiency compared to PIC. We also discuss potential advantages over AMR-PIC.

5.1. 2D Gaussian beam with halo

We have performed verification of the Adaptive Particle-in-Cloud method using examples of highly non-uniform distributions of particles typical for accelerator beams with halos. In such problems, a high-intensity, small-sigma particle beam is surrounded by a larger radius halo containing from 3 to 6 orders of magnitude smaller number of particles compared to the main beam. As accurate modeling of realistic accelerator beam and halo distributions is unnecessary for the numerical verification, we represent the system by axially symmetric Gaussian distributions. This also allows us to obtain a benchmark solution.

Consider the following 2D electrostatic problem

$$\Delta\phi = \rho, \tag{28}$$

where charge density ρ is given by two overlapping Gaussian distribution:

$$\rho(\mathbf{x}) = a_1 \left[\exp\left(-\frac{|\mathbf{x}|^2}{\tau_1^2}\right) + a_2 \exp\left(-\frac{|\mathbf{x}|^2}{\tau_2^2}\right) \right], \tag{29}$$

in the domain $\Omega = [-1, 1] \times [-1, 1]$. We use the following values for the coefficients: the radius of the main beam $\tau_1 = 0.02$, the halo intensity $a_2 = 10^{-5}$, and the width of the halo $\tau_2 = 0.3$. Coefficient $a_1 = 396.1$ is a normalization parameter to ensure $\int_{\Omega} \rho(\mathbf{x}) d\mathbf{x} = 1$. The model is consistent (in terms of the order of magnitude for the beam versus halo ratio) with real particles beams in accelerators.

While the AP-Cloud method is independent of the geometric shape of the computational domain, we solve the problem in a square domain to enable the comparison with the traditional PIC method. The benchmark solution is obtained in the following way. The problem is embedded in a larger domain, a radius 2 disk, using the same charge density function ρ and the homogeneous Dirichlet boundary condition. A solution, obtained by a highly refined 1D solver in cylindrically symmetric coordinates, is considered as the benchmark solution. The Dirichlet boundary condition for the two-dimensional problem is computed by interpolating the 1D solution at the location of the 2D boundary. This boundary condition function is then used for both the second-order AP-Cloud and PIC methods.

In our numerical simulations, CIC scheme (4) is used in charge assignment and interpolation in PIC method. Theoretically, there are more accurate schemes available, such as triangular shaped cloud with reshaping step. However, these higher order schemes are very computationally intensive, and are not able to give better result than CIC with the same CPU time in our numerical tests. The order of accuracy of AP-Cloud method does not depend on the particular kernel function Φ , so we choose the nearest grid point scheme for its simplicity, that is, Φ in (16) is set to be the characteristic function of the corresponding octree cell.

Figure 1: Distribution of 10^6 particles and the electric potential.

Figure 2: Distribution of nodes. Green nodes are given by error balance criterion, blue nodes are given by 2:1 mesh balance, and red nodes are on the boundary.

Figure 3: Distribution of error of potential ϕ computed by AP-Cloud. ($\|\phi\|_{\infty} = 0.7226$.)

Figure 4: Distribution of error of gradient of potential ϕ_x computed by AP-Cloud. ($\|\phi_x\|_{\infty} = 3.581$.)

n	Running time	Error of ϕ	Error of ϕ_x
121	0.341	0.118	2.05
441	0.327	0.0648	1.87
1681	0.353	0.0347	1.41
6561	0.490	0.0139	0.674
25921	1.60	0.00371	0.214

Table 1: CPU time and accuracy of traditional PIC with 10^6 particles.

Table 2: CPU time and accuracy of AP-Cloud with 2:1 mesh balance with 10^6 particles.

n	Running time	Error of ϕ	Error of ϕ_x
256	0.929	0.00289	0.0515
428	0.933	0.0183	0.0218
1156	0.970	0.00886	0.00927
3652	1.14	0.00365	0.00750
7559	1.45	0.000233	0.00725
19077	2.80	0.000145	0.00724

Figure 5: Dependence of error on the number of nodes for PIC and AP-Cloud.

Figure 7: Stacked plot of running times of AP-Cloud.

Table 3: Breakdown of running times of AP-Cloud.

N	10^4	10^5	10^6	10^6	10^6
n	828	1900	1888	4420	11190
Build quadtree	7.67e-03	5.48e-02	4.85e-01	4.85e-01	4.89e-01
Search nodes	3.11e-04	6.58e-04	6.46e-04	1.45e-03	3.70e-03
Build linear systems	1.11e-02	2.44e-02	2.39e-02	5.62e-02	1.41e-01
Solve linear system for ρ	1.64e-01	1.68e-01	1.68e-01	1.95e-01	2.76e-01
Solve linear system for ϕ	1.81e-01	1.95e-01	1.95e-01	2.99e-01	7.93e-01
Find interpolation coefficient	1.41e-02	3.20e-02	3.15e-02	7.46e-02	1.88e-01
Interpolate	8.18e-04	9.11e-03	1.40e-01	1.40e-01	1.45e-01
Total running time	3.82e-01	4.89e-01	1.05e+00	1.26e+00	2.05e+00

From theoretical complexity analysis in Section 4 and experimental results in Table 5.1 and Figure 7, the steps of AP-Cloud can be divided into 3 main groups:

- Quadtree construction and interpolation, which time complexity is $\mathcal{O}(N \log N)$. The running time for this group is $\mathcal{O}(N \log N)$, which dominates when $N \gg n$.
- Searching for nodes, building linear systems and finding interpolation coefficients. The running time for this group is $\mathcal{O}(nl_{max})$, which is small compared to the running time of other two groups.
- Solving the linear system for ρ and ϕ . CPU time depends on both the linear solver and n , and dominates for small ratios of N/n .

In this test, we did not obtain the second order convergence due to the Monte Carlo noise. Table 4 shows the result of another test, where $\bar{\rho}(\mathbf{y}^j, h)$ in (2) is given by the integral of the exact density function instead of Monte Carlo integration. The convergence is second order for both the potential and gradient, as expected.

Table 4: Convergence of AP-Cloud without Monte Carlo noise.

n	Error of ϕ	Error of ϕ_x	Order of ϕ	Order of ϕ_x
240	0.00181	0.00123	-	-
863	0.000564	0.000315	1.82	2.13
3336	0.000150	7.62e-05	1.95	2.09
13043	3.80e-05	2.006e-05	2.01	1.96

5.2. 3D Gaussian beam with halo

In this Section, we investigate the accuracy of the AP-Cloud method in 3D. To enable comparison with a simple benchmark solution, we study a spherically symmetric extension of the beam-with-halo problem. Despite the loss of physics relevance, it is a useful problem that tests adaptive capabilities of the method. Consider the Poisson equation with the charge density ρ given by two overlapping Gaussian distributions (29) in the domain $\Omega = [-1, 1]^3$. The radius of center beam is $\tau_1 = 0.02$, the strength of the halo is $a_2 = 10^{-5}$, and the width of the halo is $\tau_1 = 0.3$. The coefficient $a_1 = 7677$ provides the normalization $\int_{\Omega} \rho(\mathbf{x}) d\mathbf{x} = 1$. The benchmark solution and the boundary condition function were obtained similarly to the 2D case. The distribution of 10^6 particles is shown in Figure 8. The AP-Cloud computation, performed using 4067 nodes (Figure 9), gives a solution with the normalized norm of

ϕ on particles is 0.0352 ($\|\phi\|_\infty = 3.038$), and the normalized norm of ϕ_x on particles is 0.578 ($\|\phi_x\|_\infty = 41.18$).

Figure 8: Distribution of 10^6 particles.

Figure 9: Distribution of 4067 nodes. Green nodes are given by error balance criterion, blue nodes are given by 2:1 mesh balance, and red nodes are on the boundary.

Figure 10: Distribution of error of potential ϕ by AP-Cloud ($\|\phi\|_\infty = 3.038$).

Figure 11: Distribution of error of gradient of potential ϕ_x by AP-Cloud ($\|\phi_x\|_\infty = 41.18$).

n	Running time	Error of ϕ	Error of ϕ_x
8000	0.443	1.40	19.6
64000	1.48	0.726	16.1
512000	24.1	0.219	8.60
4096000	361	0.0606	2.92

Table 5: CPU time and accuracy of traditional PIC with 10^6 particles in 3D.

n	Running time	Error of ϕ	Error of ϕ_x
1546	0.921	0.0402	1.17
4067	1.14	0.0352	0.578
13687	2.10	0.0183	0.329
59349	7.22	0.00443	0.244

Table 6: CPU time and accuracy of AP-Cloud with 2:1 mesh balance with 10^6 particles in 3D.

Figure 12: Dependence of error on the number of nodes for PIC and AP-Cloud in 3D.

Figure 13: Dependence of error on CPU time for PIC and AP-Cloud in 3D.

Results presented in Tables 5 and 6 and Figures 12 and 13 show that the advantage of AP-Cloud is more evident in 3D problem. With only 1546 nodes and less than 1 second running time, AP-Cloud has more accurate result than PIC with 4096000 cells and more than 361 seconds running time.

5.3. Test for self-force effect with single particle

As mentioned in the introduction, Vlasov-Poisson problems with highly non-uniform distributions of matter can be solved using the adaptive mesh refinement technique for PIC [4, 5]. However, it is well known that AMR-PIC introduces significant artifacts in the form of artificial image particles across boundaries between coarse and fine meshes. These images introduce spurious forces that may potentially alter the particle motion to an unacceptable level [4, 5]. Methods for the mitigation of the spurious forces have been designed in [6]. The traditional PIC on a uniform mesh is free of such artifacts.

The convergence of Adaptive Particle-in-Cloud solutions to benchmark solutions, discussed in the previous Section, already indicates the absence of artifacts. To further verify that AP-Cloud is free of artificial forces present in the original AMR-PIC, we have performed an additional test similar to the one in [4], which involved the motion of a single particle across the coarse and fine mesh interface. For AP-Cloud, we studied the motion of a single test particle represented by a moving cloud of nodes with refined distances towards the test particle. The test particle contained a smooth,

sharp, Gaussian-type charge distribution to satisfy the requirements of the GFD method.

The forces and motion of a single test particle obtained with PIC and AP-Cloud methods are shown in Figure 14. We observe that the electric forces computed by the AP-Cloud method are more accurate and smoother compared to even the traditional PIC. But the oscillatory deviation of forces in PIC from the correct direction does not cause accumulation of the total error due to conservative properties of PIC. The trajectories of the particle obtained by both methods are close. The test provides an additional assurance that artificial images are not present in the AP-Cloud method.

Figure 14: Forces (left) and motion (right) of a single test particle obtained with PIC and AP-Cloud methods demonstrating the absence of artifacts in the AP-Cloud method.

6. Summary and Conclusions

We have developed an Adaptive Particle-in-Cloud (AP-Cloud) method that replaces the Cartesian grid in the traditional PIC with adaptive computational nodes. Adaptive particle placement balances the errors of the differential operator discretization and the source computation (analogous to the error of the Monte Carlo integration) to minimize the total error.

AP-Cloud uses GFD based on weighted least squares (WLS) approximations on a stencil of irregularly placed nodes. The framework includes interpolation, least squares approximation, and numerical differentiation capable of high order convergence.

The adaptive nature of AP-Cloud gives it significant advantages over the traditional PIC for non-uniform distributions of particles and complex

boundaries. It achieves significantly better accuracy in the gradient of the potential compared to the traditional PIC for the problem of particle beam with halo. The method is independent of the geometric shape of the computational domain, and can achieve highly accurate solutions in geometrically complex domains. The optimal mesh size based on error-balance criterion gives AP-Cloud a potential advantage over AMR-PIC in terms of accuracy, and specially designed tests showed that the AP-Cloud method is free of artificial images and spurious forces typical for the original AMR-PIC without special mitigation techniques. Another advantage of AP-Cloud over AMR-PIC is the ease of implementation, as AP-Cloud does not require special remapping routines between different meshes. Our future work will focus on higher convergence rates of the method, performance optimization, parallel implementation using hybrid technologies, as well as applications to practical problems with non-uniform distribution of matter. A direct comparison of AP-Cloud with AMR-PIC in terms of accuracy and efficiency will also be addressed in the future work.

Acknowledgement

This work was supported in part by the U.S. Department of Energy, Contract No. DE-AC02-98CH10886.

References

- [1] R.W. Hockney, J.W. Eastwood, Computer simulation using particles, (CRC Press, 1988), 21.
- [2] H. Johansen, P. Colella, A cartesian grid embedding boundary method for poisson's equation on irregular domains, J Comput Phys, 1998 (147), 60?85.
- [3] S. Wang, R. Samulyak, T. Guo, An embedded boundary method for parabolic problems with interfaces and application to multi-material systems with phase transitions, Acta Mathematica Scientia, 30B (2010), No. 2, 499 - 521.
- [4] Vay, J-L., et al. "Mesh refinement for particle-in-cell plasma simulations: applications to and benefits for heavy ion fusion." Laser and Particle Beams 20.04 (2002): 569-575.

- [5] Vay, J-L., et al. "Application of adaptive mesh refinement to particle-in-cell simulations of plasmas and beams." *Physics of Plasmas* (1994-present) 11.5 (2004): 2928-2934.
- [6] Colella, Phillip, and Peter C. Norgaard. "Controlling self-force errors at refinement boundaries for AMR-PIC." *Journal of Computational Physics* 229.4 (2010): 947-957.
- [7] Onate, E., et al. "A finite point method in computational mechanics. Applications to convective transport and fluid flow." *International journal for numerical methods in engineering* 39.22 (1996): 3839-3866.
- [8] Benito, J. J., F. Urena, and L. Gavete. "Influence of several factors in the generalized finite difference method." *Applied Mathematical Modelling* 25.12 (2001): 1039-1053.
- [9] Benito, J. J., et al. "An h-adaptive method in the generalized finite differences." *Computer methods in applied mechanics and engineering* 192.5 (2003): 735-759.
- [10] Liszka, Tadeusz, and Janusz Orkisz. "The finite difference method at arbitrary irregular grids and its application in applied mechanics." *Computers & Structures* 11.1 (1980): 83-95.
- [11] Golub, Gene H., and Charles F. Van Loan, *Matrix Computations*, 4th edition, Johns Hopkins University Press, 2012.
- [12] Barbeiro, S., J. A. Ferreira, and R. D. Grigorieff. "Supraconvergence of a finite difference scheme for solutions in $H^s(0, L)$." *IMA journal of numerical analysis* 25.4 (2005): 797-811.
- [13] Zhou, Kun, et al. "Data-parallel octrees for surface reconstruction." *Visualization and Computer Graphics, IEEE Transactions on* 17.5 (2011): 669-681.
- [14] Cottet, G-H., and P-A. Raviart. "Particle methods for the one-dimensional Vlasov-Poisson equations." *SIAM journal on numerical analysis* 21.1 (1984): 52-76.
- [15] Wang, Bei, Gregory H. Miller, and Phillip Colella. "A particle-in-cell method with adaptive phase-space remapping for kinetic plasmas." *SIAM Journal on Scientific Computing* 33.6 (2011): 3509-3537.

Appendix A. Error analysis for PIC

Assume the interpolation kernel Φ is symmetric, non-negative, bounded by 1, integrable and normalized over a compact support. That is, $0 \leq \Phi(\mathbf{x}) = \Phi(-\mathbf{x}) \leq \|\Phi\|_\infty \leq 1$ for all $\mathbf{x} \in \mathbb{R}^D$, $\Phi(\mathbf{x}) = 0$ for $\|\mathbf{x}\| > r$ for some r , and $\int_{\mathbb{R}^D} \Phi(\mathbf{x}) d\mathbf{x} = 1$. These properties hold for all commonly used charge assignment schemes, including the nearest grid point, cloud-in-cell, and triangular shaped cloud schemes without reshaping step. In the following subsections, we will analyze the errors in the three steps of PIC, respectively.

Appendix A.1. Error in Step 1

To analyze the error in (2), let us first define an average quantity at a grid point $\mathbf{y}^j$ as

$$\bar{\rho}(\mathbf{y}^j, h) := \frac{1}{h^D} \int_{\Omega} \rho(\mathbf{x}) \Phi\left(\frac{\mathbf{x} - \mathbf{y}^j}{h}\right) d\mathbf{x}.$$

Then,

$$|\tilde{\rho}(\mathbf{y}^j, \mathcal{P}, h) - \rho(\mathbf{y}^j)| \leq \underbrace{|\tilde{\rho}(\mathbf{y}^j, \mathcal{P}, h) - \bar{\rho}(\mathbf{y}^j, h)|}_{\mathcal{E}_M} + \underbrace{|\bar{\rho}(\mathbf{y}^j, h) - \rho(\mathbf{y}^j)|}_{\mathcal{E}_{D_1}},$$

where $\mathcal{E}_M$ is analogous to the error in Monte Carlo integration of a continuous function, and $\mathcal{E}_{D_1}$ is known as moment error or discretization error in step 1, which depends on the interpolation kernel Φ .

Appendix A.1.1. Monte Carlo noise

To bound the error $\mathcal{E}_M$, note that the expected value of $\tilde{\rho}$ is

$$\mathbb{E}[\tilde{\rho}(\mathbf{y}^j, \mathcal{P}, h)] = \bar{\rho}(\mathbf{y}^j, h). \quad (\text{A.1})$$

Therefore,

$$\text{Var}(\tilde{\rho}(\mathbf{y}^j, \mathcal{P}, h)) = \mathbb{E}[|\mathcal{E}_M|^2].$$

By the definition of variance,

$$\text{Var}(\tilde{\rho}(\mathbf{y}^j, \mathcal{P}, h)) = \frac{1}{N} \text{Var}(\tilde{\rho}(\mathbf{y}^j, \mathbf{p}^i, h)), \quad (\text{A.2})$$

where

$$\tilde{\rho}(\mathbf{y}^j, \mathbf{p}^i, h) = \frac{1}{h^D} \Phi\left(\frac{\mathbf{p}^i - \mathbf{y}^j}{h}\right)$$

is the estimation of ρ from a single particle $\mathbf{p}^i$, and

$$\mathbb{E}[\tilde{\rho}(\mathbf{y}^j, \mathbf{p}^i, h)] = \bar{\rho}(\mathbf{y}^j, h). \quad (\text{A.3})$$

Let $\tau(\mathbf{y}^j, \mathbf{p}^i, h) = \tilde{\rho}(\mathbf{y}^j, \mathbf{p}^i, h) - \bar{\rho}(\mathbf{y}^j, h)$. Then,

$$\text{Var}(\tilde{\rho}(\mathbf{y}^j, \mathbf{p}^i, h)) = \mathbb{E}[\tau^2(\mathbf{y}^j, \mathbf{p}^i, h)] \quad (\text{A.4})$$

$$= \int_{\Omega} \tau^2(\mathbf{y}^j, \mathbf{x}, h) \rho(\mathbf{x}) \mathbf{d}\mathbf{x}. \quad (\text{A.5})$$

Consider the small neighborhood about $\mathbf{y}^j$ of radius rh , i.e., $\mathcal{A}^j = \{\mathbf{x} : \|\mathbf{x} - \mathbf{y}^j\| < rh\}$, so that $\tilde{\rho}(\mathbf{y}^j, \mathbf{x}, h) = 0$ in $\Omega \setminus \mathcal{A}^j$. Then,

$$\text{Var}(\tilde{\rho}(\mathbf{y}^j, \mathbf{p}^i, h)) = \underbrace{\int_{\mathcal{A}^j} (\tau(\mathbf{y}^j, \mathbf{x}, h))^2 \rho(\mathbf{x}) \mathbf{d}\mathbf{x}}_{I_1} + \underbrace{\int_{\Omega \setminus \mathcal{A}^j} (\tau(\mathbf{y}^j, \mathbf{x}, h))^2 \rho(\mathbf{x}) \mathbf{d}\mathbf{x}}_{I_2}.$$

Note that

$$\lim_{h \rightarrow 0} \bar{\rho}(\mathbf{y}^j, h) = \rho(\mathbf{y}^j) \int_{\mathbb{R}^D} \Phi(\mathbf{x}) \mathbf{d}\mathbf{x} = \rho(\mathbf{y}^j) \quad (\text{A.6})$$

and

$$\lim_{h \rightarrow 0} \frac{1}{h^D} \int_{\mathcal{A}^j} \rho(\mathbf{x}) \mathbf{d}\mathbf{x} = r^D V_D \rho(\mathbf{y}^j), \quad (\text{A.7})$$

where V_D is the volume of the unit ball in D dimensions. Assume h is small enough, so that $\bar{\rho}(\mathbf{y}^j, h) < 2\rho(\mathbf{y}^j)$ and $\frac{1}{h^D} \int_{\mathcal{A}^j} \rho(\mathbf{x}) \mathbf{d}\mathbf{x} < 2r^D V_D \rho(\mathbf{y}^j)$. Then,

$$I_1 \leq \int_{\mathcal{A}^j} \left(\frac{1}{h^D} \right)^2 \rho(\mathbf{x}) \mathbf{d}\mathbf{x} < 2 \frac{r^D V_D}{h^D} \rho(\mathbf{y}^j) = \mathcal{O}(h^{-D}) \rho(\mathbf{y}^j), \quad (\text{A.8})$$

and

$$I_2 = \int_{\Omega \setminus \mathcal{A}^j} \bar{\rho}(\mathbf{y}^j, h)^2 \rho(\mathbf{x}) \mathbf{d}\mathbf{x} < \bar{\rho}(\mathbf{y}^j, h)^2 < 4\rho(\mathbf{y}^j)^2 = \mathcal{O}(1) \rho(\mathbf{y}^j). \quad (\text{A.9})$$

Therefore,

$$\begin{aligned} \text{Var}(\tilde{\rho}(\mathbf{y}^j, \mathcal{P}, h)) &= \frac{1}{N} (I_1 + I_2) \\ &= \frac{1}{N} (\rho(\mathbf{y}^j) \mathcal{O}(h^{-D}) + \rho(\mathbf{y}^j)^2 \mathcal{O}(1)) \\ &= \frac{\rho(\mathbf{y}^j)}{N} \mathcal{O}(h^{-D}), \end{aligned} \quad (\text{A.10})$$

and

$$\mathbb{E}[|\mathcal{E}_M|] = \sqrt{\mathbb{E}[|\mathcal{E}_M|^2] - \text{Var}(|\mathcal{E}_M|)} \leq \sqrt{\mathbb{E}[|\mathcal{E}_M|^2]} = \mathcal{O}\left(\sqrt{\frac{\rho(\mathbf{y}^j)}{Nh^D}}\right). \quad (\text{A.11})$$

Note that above bound for $\sqrt{\mathbb{E}[|\mathcal{E}_M|^2]}$ is asymptotically tight. To see this, assume $4h^D\bar{\rho}(\mathbf{y}^j, h) < \|\Phi\|_\infty$, and let $A_1 = \{\mathbf{x} : \Phi((\mathbf{x} - \mathbf{y}^j)/h) > \|\Phi\|_\infty/2\}$. Then,

$$I_1 \geq \int_{A_1} \tau^2(\mathbf{y}^j, \mathbf{x}, h) \rho(\mathbf{x}) d\mathbf{x} > \int_{A_1} \left(\frac{\|\Phi\|_\infty}{4h^D}\right)^2 \rho(\mathbf{x}) d\mathbf{x}.$$

The volume of A_1 is $\mathcal{O}(h^D)$, so I_1 is no smaller than $\mathcal{O}(h^{-D})\rho(\mathbf{y}^j)$, and $\sqrt{\mathbb{E}[|\mathcal{E}_M|^2]}$ is no smaller than $\mathcal{O}\left(\sqrt{\rho(\mathbf{y}^j)h^{-D}N^{-1}}\right)$.

Appendix A.1.2. Moment error

The moment error in charge assignment scheme, $\mathcal{E}_{D_1}$, can be viewed as the error in a numerical quadrature rule. Here we follow the approach taken in [14] and [15]. Let $B(\mathbf{y}^j, a) = \{\mathbf{x} : \|\mathbf{x} - \mathbf{y}^j\| < a\}$, $H(\mathbf{y}^j)$ be the Hessian

matrix for ρ at $\mathbf{y}^j$. The moment error is

$$\begin{aligned}
\mathcal{E}_{D_1} &= |\bar{\rho}(\mathbf{y}^j, h) - \rho(\mathbf{y}^j)| \\
&= \frac{1}{h^D} \int_{\Omega} (\rho(\mathbf{x}) - \rho(\mathbf{y}^j)) \Phi\left(\frac{\mathbf{x} - \mathbf{y}^j}{h}\right) d\mathbf{x} \\
&= \frac{1}{h^D} \int_{B(\mathbf{y}^j, rh)} ((\mathbf{x} - \mathbf{y}^j)^T \nabla \rho(\mathbf{y}^j) + (\mathbf{x} - \mathbf{y}^j)^T H(\mathbf{y}^j)(\mathbf{x} - \mathbf{y}^j) \\
&\quad + \mathcal{O}(\|\mathbf{x} - \mathbf{y}^j\|^3)) \Phi\left(\frac{\mathbf{x} - \mathbf{y}^j}{h}\right) d\mathbf{x} \\
&= \frac{1}{h^D} \int_{B(\mathbf{y}^j, rh)} (\mathbf{x} - \mathbf{y}^j)^T \nabla \rho(\mathbf{y}^j) \Phi\left(\frac{\mathbf{x} - \mathbf{y}^j}{h}\right) d\mathbf{x} \\
&\quad + \frac{1}{h^D} \int_{B(\mathbf{y}^j, rh)} (\mathbf{x} - \mathbf{y}^j)^T H(\mathbf{y}^j)(\mathbf{x} - \mathbf{y}^j) \Phi\left(\frac{\mathbf{x} - \mathbf{y}^j}{h}\right) d\mathbf{x} + \mathcal{O}(h^3) \\
&= 0 + \frac{1}{h^D} \int_{B(\mathbf{y}^j, rh)} \sum_{d=1}^D (x_d - y_d^j)^2 \rho_{y_d y_d}(\mathbf{y}^j) \Phi\left(\frac{\mathbf{x} - \mathbf{y}^j}{h}\right) d\mathbf{x} + \mathcal{O}(h^3) \\
&= h^2 \sum_{d=1}^D \int_{B(\mathbf{y}^j, r)} (z_d - y_d^j)^2 \Phi(\mathbf{z} - \mathbf{y}^j) d\mathbf{z} \rho_{y_d y_d}(\mathbf{y}^j) + \mathcal{O}(h^3) \\
&= \sum_{d=1}^D \rho_{y_d y_d}(\mathbf{y}^j) \mathcal{O}(h^2), \tag{A.12}
\end{aligned}$$

where the fifth equal sign is due to the symmetry of interpolation kernel.

Appendix A.2. Error in Step 2

It is well known that the error in finite difference method for elliptic equation in step 2 is second order

$$\mathcal{E}_{D_2} = \sum_{d=1}^D \rho_{y_d y_d}(\mathbf{y}^j) \mathcal{O}(h^2). \tag{A.13}$$

So the total discretization error in ϕ is

$$\mathcal{E}_D = \mathcal{E}_{D_1} + \mathcal{E}_{D_2} = \sum_{d=1}^D \rho_{y_d y_d}(\mathbf{y}^j) \mathcal{O}(h^2) = \rho(\mathbf{y}^j) \mathcal{O}(h^2), \tag{A.14}$$

where the last equal sign is due to the assumption that ρ and its derivatives have comparable magnitude.

Appendix A.3. Error in Step 3

Step 3 contains a numerical differentiation and a linear interpolation, both of which has second order accuracy. One problem is that numerical differentiation is unstable. Generally, if ϕ is already polluted with a k th order error, then its gradient given by numerical differentiation has at most $(k - 1)$ th order accuracy, because of the h factor in denominator. For finite difference scheme for elliptic equation, it has been proved that, nevertheless, if the solution is smooth enough, then both the solution and its gradient are second order convergent, even if the mesh is non-uniform, which is called supraconvergence [12]. More precisely, let P_h be the interpolation operator from grid function to piecewise linear function, R_h be the restriction operator from continuous function to grid function, $u \in H^3(\Omega)$ be the exact solution of Poisson equation, u_h be the numerical solution, then we have

$$\|P_h(R_h u - u_h)\|_{H^1} \leq C h_{max}^2 \|u\|_{H^3}, \quad (\text{A.15})$$

where h_{max} is the maximal mesh size. Thus we claim the discretization error for $\mathbf{E}$ is also second order.