

# Towards Kinetic Whole Device Modelling of Low-Temperature Plasma Devices

Algorithms and Computer Science

*Andrew (Tasman) Powis*

# Talk Overview

## Particle-in-Cell Method

*(Another) introduction to PIC and a discussion on why we use it*

## Leveraging Modern Supercomputers

*Some important considerations for programming on CPU+GPU computers*


## GPUs are Not Enough

*Improving hardware performance is not sufficient to achieve kinetic whole-device modeling, we also need to consider new or alternative algorithms*

# Why Simulate?

The goal of numerical simulations is broadly:

1. To improve our understanding of fundamental physics by resolving features which are not measurable or by reaching conditions which are not achievable in experiments




3D gyro-kinetic simulations of ion-temperature gradient instabilities within a Tokamak. Run using the GTS code [Ma *et al.*]

# Why Simulate?

The goal of numerical simulations is broadly:

1. To improve our understanding of fundamental physics by resolving features which are not measurable or by reaching conditions which are not achievable in experiments
2. To reduce the cost of engineering design or prototyping through whole-device modelling (WDM)




3D gyro-kinetic simulations of ion-temperature gradient instabilities within a Tokamak. Run using the GTS code [Ma *et al.*]

# Why Kinetic Simulations?

The global properties of *many* low-temperature plasma devices (particularly at low pressure) are dominated by kinetic processes:

- Interaction between the plasma and the wall through the sheath
- Velocity space instabilities (i.e. bump on tail)
- Collisional processes (i.e. ionization)
- Non-classical or turbulence induced transport



Electron density plots demonstrating formation of a rotating spoke within 2D Penning discharge simulations [Powis *et al.*]

# Particle-in-Cell Method

# Particle-in-Cell


Begin with the **Klimontovich representation** of a plasma (single species for clarity). Here we will consider the **electrostatic approximation**, suitable for many low-temperature plasma systems

$$N(\mathbf{q}, t) = \sum_p \delta(\mathbf{q} - \mathbf{Q}_p)$$

$$\mathbf{Q}_p = (\mathbf{X}_p, \mathbf{V}_p) = (X_p, Y_p, Z_p, VX_p, VY_p, VZ_p)$$

$$\frac{d\mathbf{Q}_p}{dt} = (\dot{\mathbf{X}}_p, \dot{\mathbf{V}}_p) = \left( \mathbf{V}_p, -\frac{q}{m} \nabla \phi \right)$$

$$\nabla^2 \phi = -\frac{q}{\epsilon_0} \sum_p \delta(\mathbf{x} - \mathbf{X}_p)$$



# Particle-in-Cell

Simulating the actual number of particles in almost all plasma systems is prohibitive, therefore we “clump” them into **macro-particles** with weight  $w_p$


$$N(\mathbf{q}, t) = \sum_p w_p \delta(\mathbf{q} - \mathbf{Q}_p)$$

$$\mathbf{Q}_p = (\mathbf{X}_p, \mathbf{V}_p) = (X_p, Y_p, Z_p, VX_p, VY_p, VZ_p)$$

$$\frac{d\mathbf{Q}_p}{dt} = (\dot{\mathbf{X}}_p, \dot{\mathbf{V}}_p) = \left( \mathbf{V}_p, -\frac{q}{m} \nabla \phi \right)$$

$$\nabla^2 \phi = -\frac{q}{\epsilon_0} \sum_p \delta(\mathbf{x} - \mathbf{X}_p)$$

$w_p$  particles





# Particle-in-Cell


We compute the electrostatic potential on a grid. This smooths out the small time scale inter-particles forces and reduces the cost of solving Poisson's equation

$$N(\mathbf{q}, t) = \sum_p w_p \delta(\mathbf{q} - \mathbf{Q}_p)$$

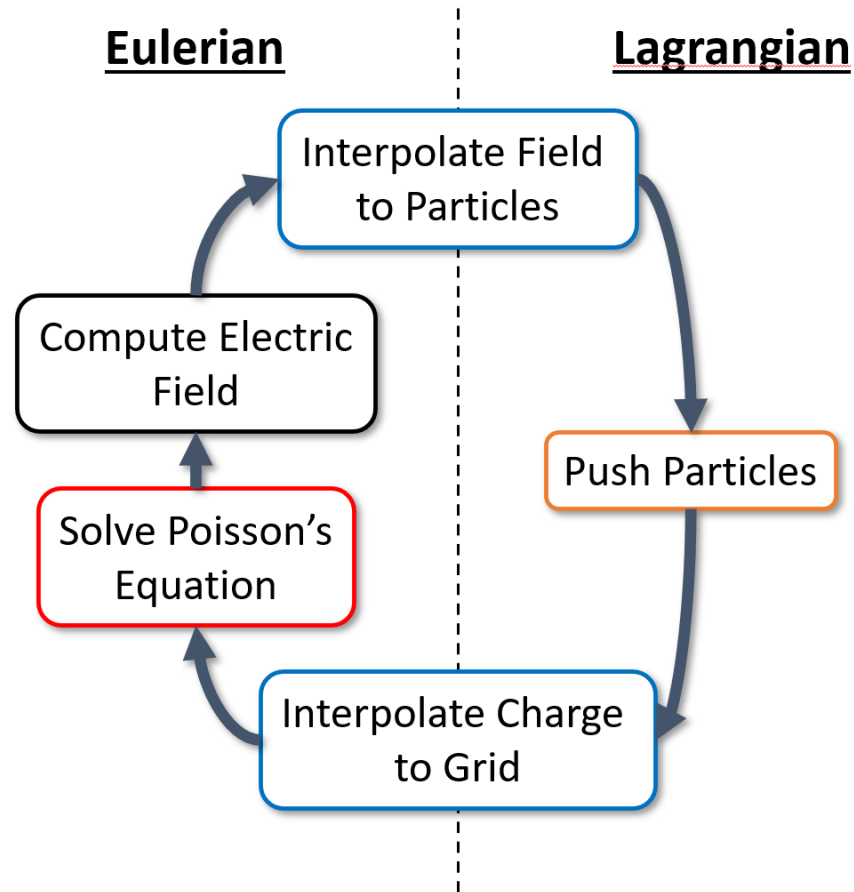
$$\mathbf{Q}_p = (\mathbf{X}_p, \mathbf{V}_p) = (X_p, Y_p, Z_p, VX_p, VY_p, VZ_p)$$

$$\frac{d\mathbf{Q}_p}{dt} = (\dot{\mathbf{X}}_p, \dot{\mathbf{V}}_p) = \left( \mathbf{V}_p, -\frac{q}{m} \sum_{i,j,k} S(\mathbf{x}_{i,j,k} - \mathbf{X}_p) [\nabla\phi]_{i,j,k} \right)$$

$$\nabla^2 \phi_{i,j,k} = -\frac{q}{\epsilon_0} \sum_p S(\mathbf{x}_{i,j,k} - \mathbf{X}_p)$$




# Particle-in-Cell



PIC is therefore a mixed Eulerian/Lagrangian framework for kinetic plasma simulations

# Particle-in-Cell



PIC is therefore a mixed Eulerian/Lagrangian framework for kinetic plasma simulations

Fine-scale physics can be added back in via Monte-Carlo collision algorithms

# What do we want to model?

- Materials processing plasma reactor
- Assume a  $10^{17} \text{ 1/m}^3$  density Argon plasma with  $10 \text{ eV}$  electrons
- Length scale  $\sim 1 \text{ m}$
- Time scale  $\sim 1 \text{ ms}$
- Full 3D kinetic simulation



Materials processing plasma reactor [Booth, Ecole Polytechnique]

# Challenges for Particle-in-Cell

Particle-in-cell (and other kinetic methods) suffer from severe stability constraints:


The time step must suitably resolve the plasma frequency:

$$\Delta t \omega_p \approx 0.2$$

The cell size should resolve the Debye length, to avoid the *finite-grid-instability*

$$\Delta x \approx \lambda_D$$

For plasma problems which have vast time and length scales this results in high computational cost



Plasma wave dispersion relationship for  $\Delta x = 10\lambda_D$  and 0<sup>th</sup> order interpolation (orange) theoretical response (blue) PIC response

# Requirements for Whole Device Modelling

- Materials processing plasma reactor
- Assume a  $10^{17} \text{ 1/m}^3$  density Argon plasma with  $10 \text{ eV}$  electrons
- Length scale  $\sim 1 \text{ m}$
- Time scale  $\sim 1 \text{ ms}$
- Based on PIC requirements we need:
  - $10^8$  time steps
  - $10^{12}$  cells
  - $10^{15}$  particles



Materials processing plasma reactor [Booth, Ecole Polytechnique]

# Time Step Requirements for Practical Simulations

We want to perform simulations in a reasonable time-frame to be useful to industry

**Target:** Simulation time less than 10 days

- For 1 *ms* we need  $\sim 10^8$  time steps, this gives us a performance goal of  $t_{step} \approx 10 \text{ ms}$
- We need high performance super computers to meet these goals

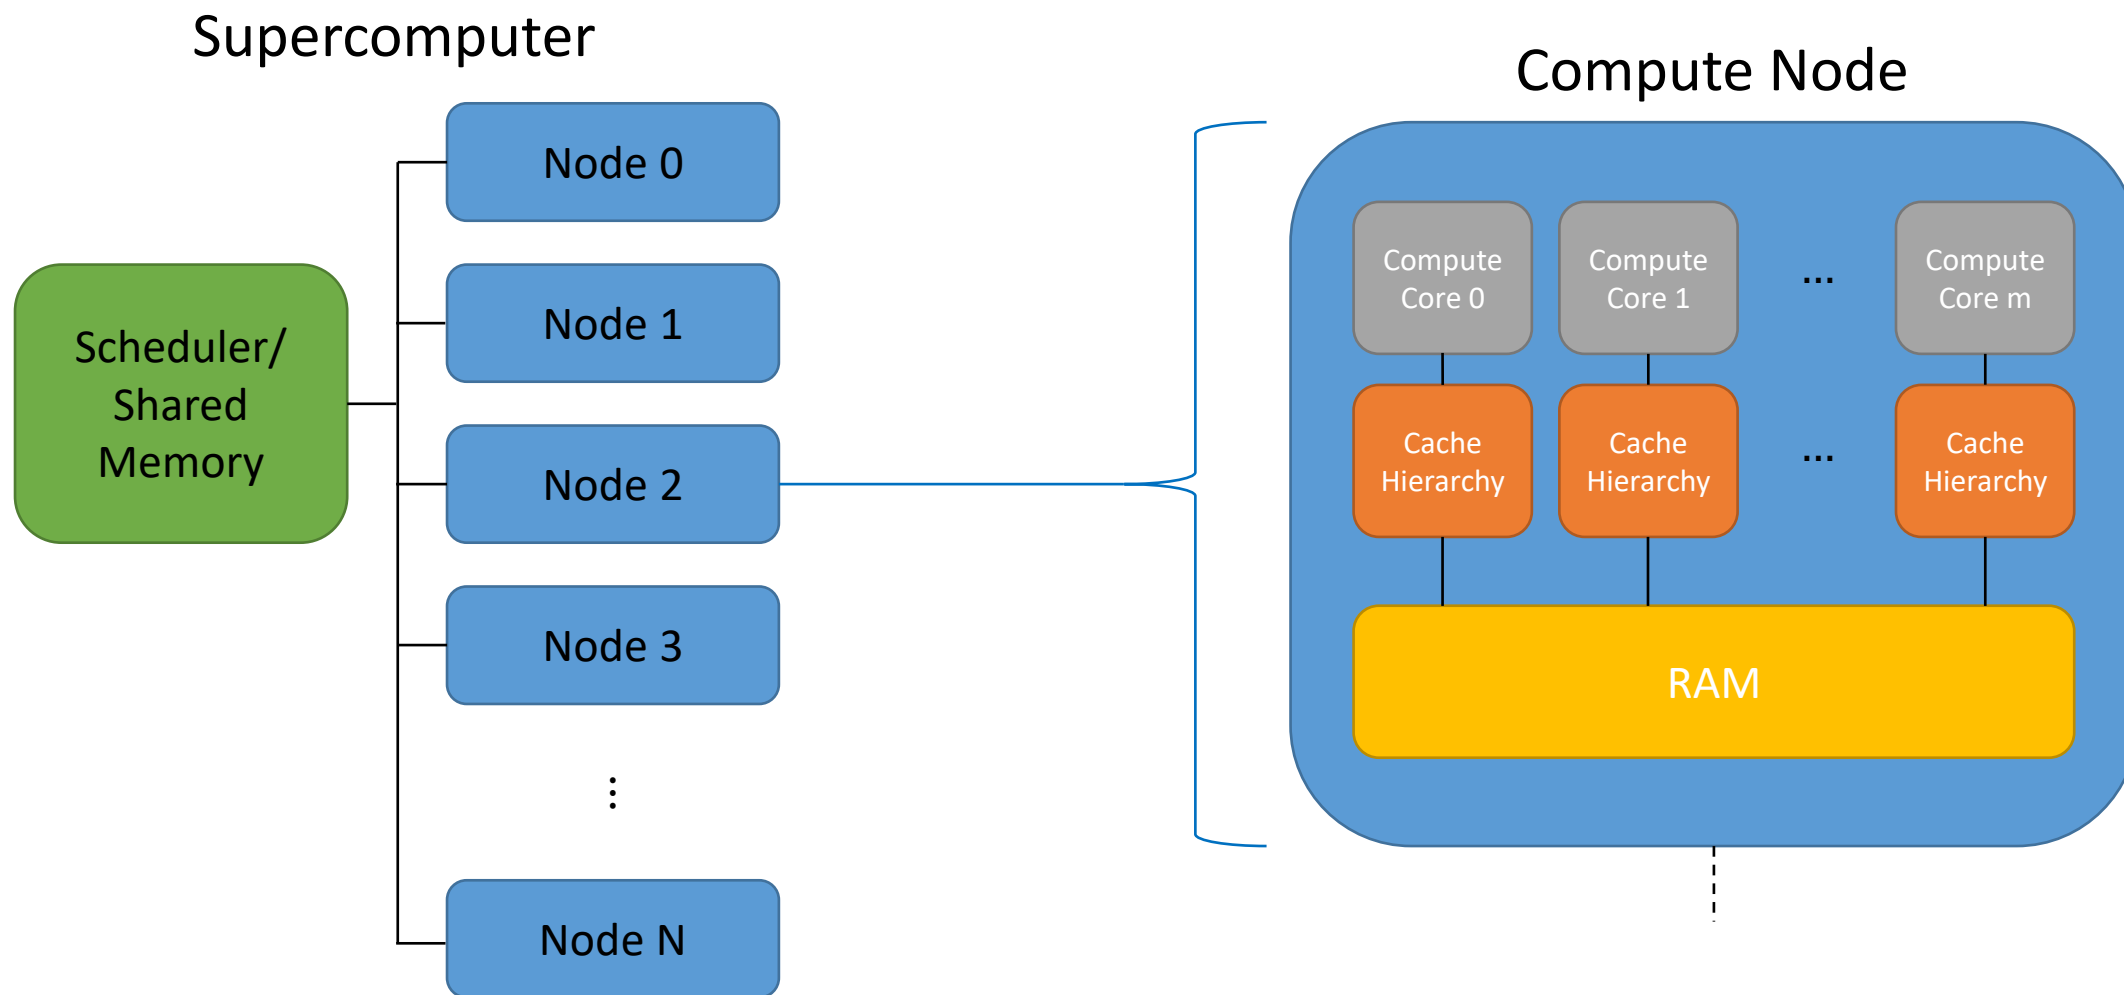


Materials processing plasma reactor [Booth, Ecole Polytechnique]

# Leveraging Modern Supercomputers

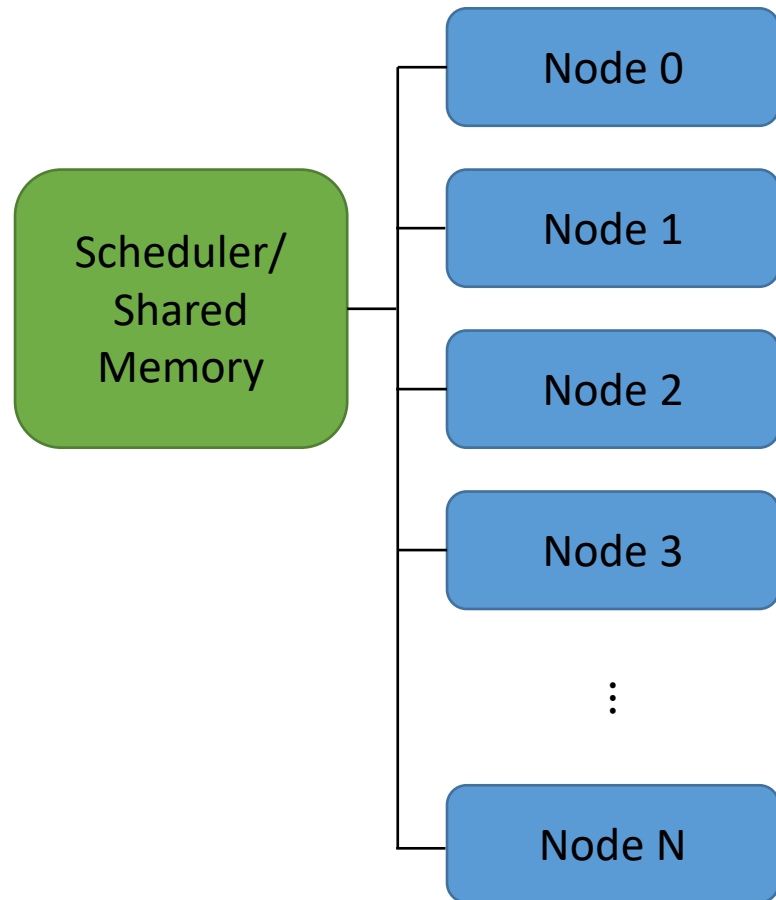


# Anatomy of a Supercomputer



# Coding for Supercomputers


## Supercomputer



- Data sharing between nodes is the slowest form of data transfer on the system
- Communication of data between nodes should be minimized as much as possible
- Data communication is generally handled by a *multiple-data multiple-instruction* (MIMD) paradigm
  - OpenMPI
  - MPICH
  - Intel MPI

# Coding for Supercomputers

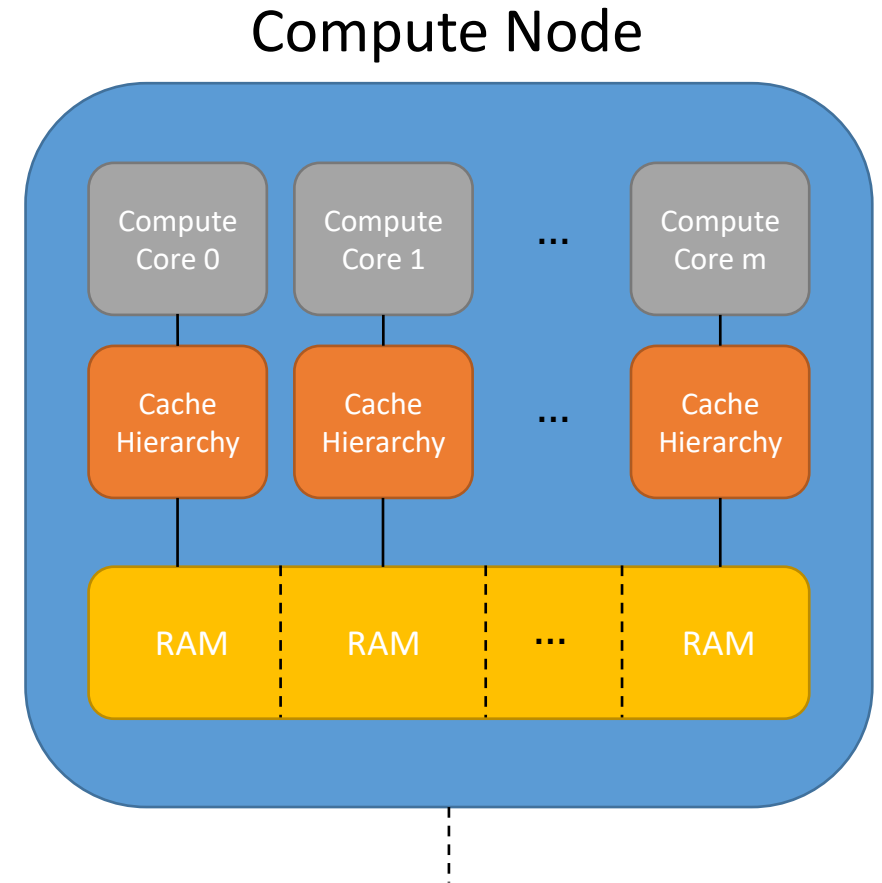
Compute nodes have shared memory



# Coding for Supercomputers

Compute nodes have shared memory

We could allocate different portions of this memory to different compute cores using MIMD instructions




# Coding for Supercomputers

Compute nodes have shared memory

We could allocate different portions of this memory to different compute cores using MIMD instructions

But the modern trend is to use what are called *single-instruction multiple-data* (SIMD) instruction sets

The standard for this is known as OpenMP and comes with all modern compilers




# Coding for Supercomputers

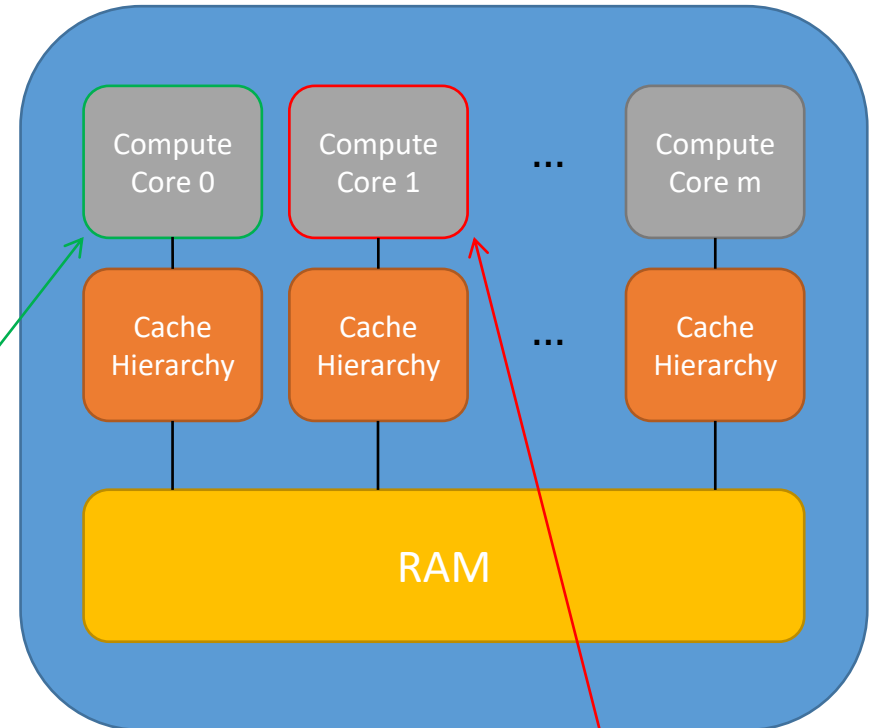
Compute nodes have shared memory so we can rely on *single-instruction multiple-data* (SIMD) standards such as OpenMP

- Avoid editing the same memory from different cores (race condition)

$$\frac{d\mathbf{Q}_p}{dt} = \dot{\mathbf{Q}}_p = \left( v_p, -\frac{q}{m} \nabla \phi \right)$$





Compute Node



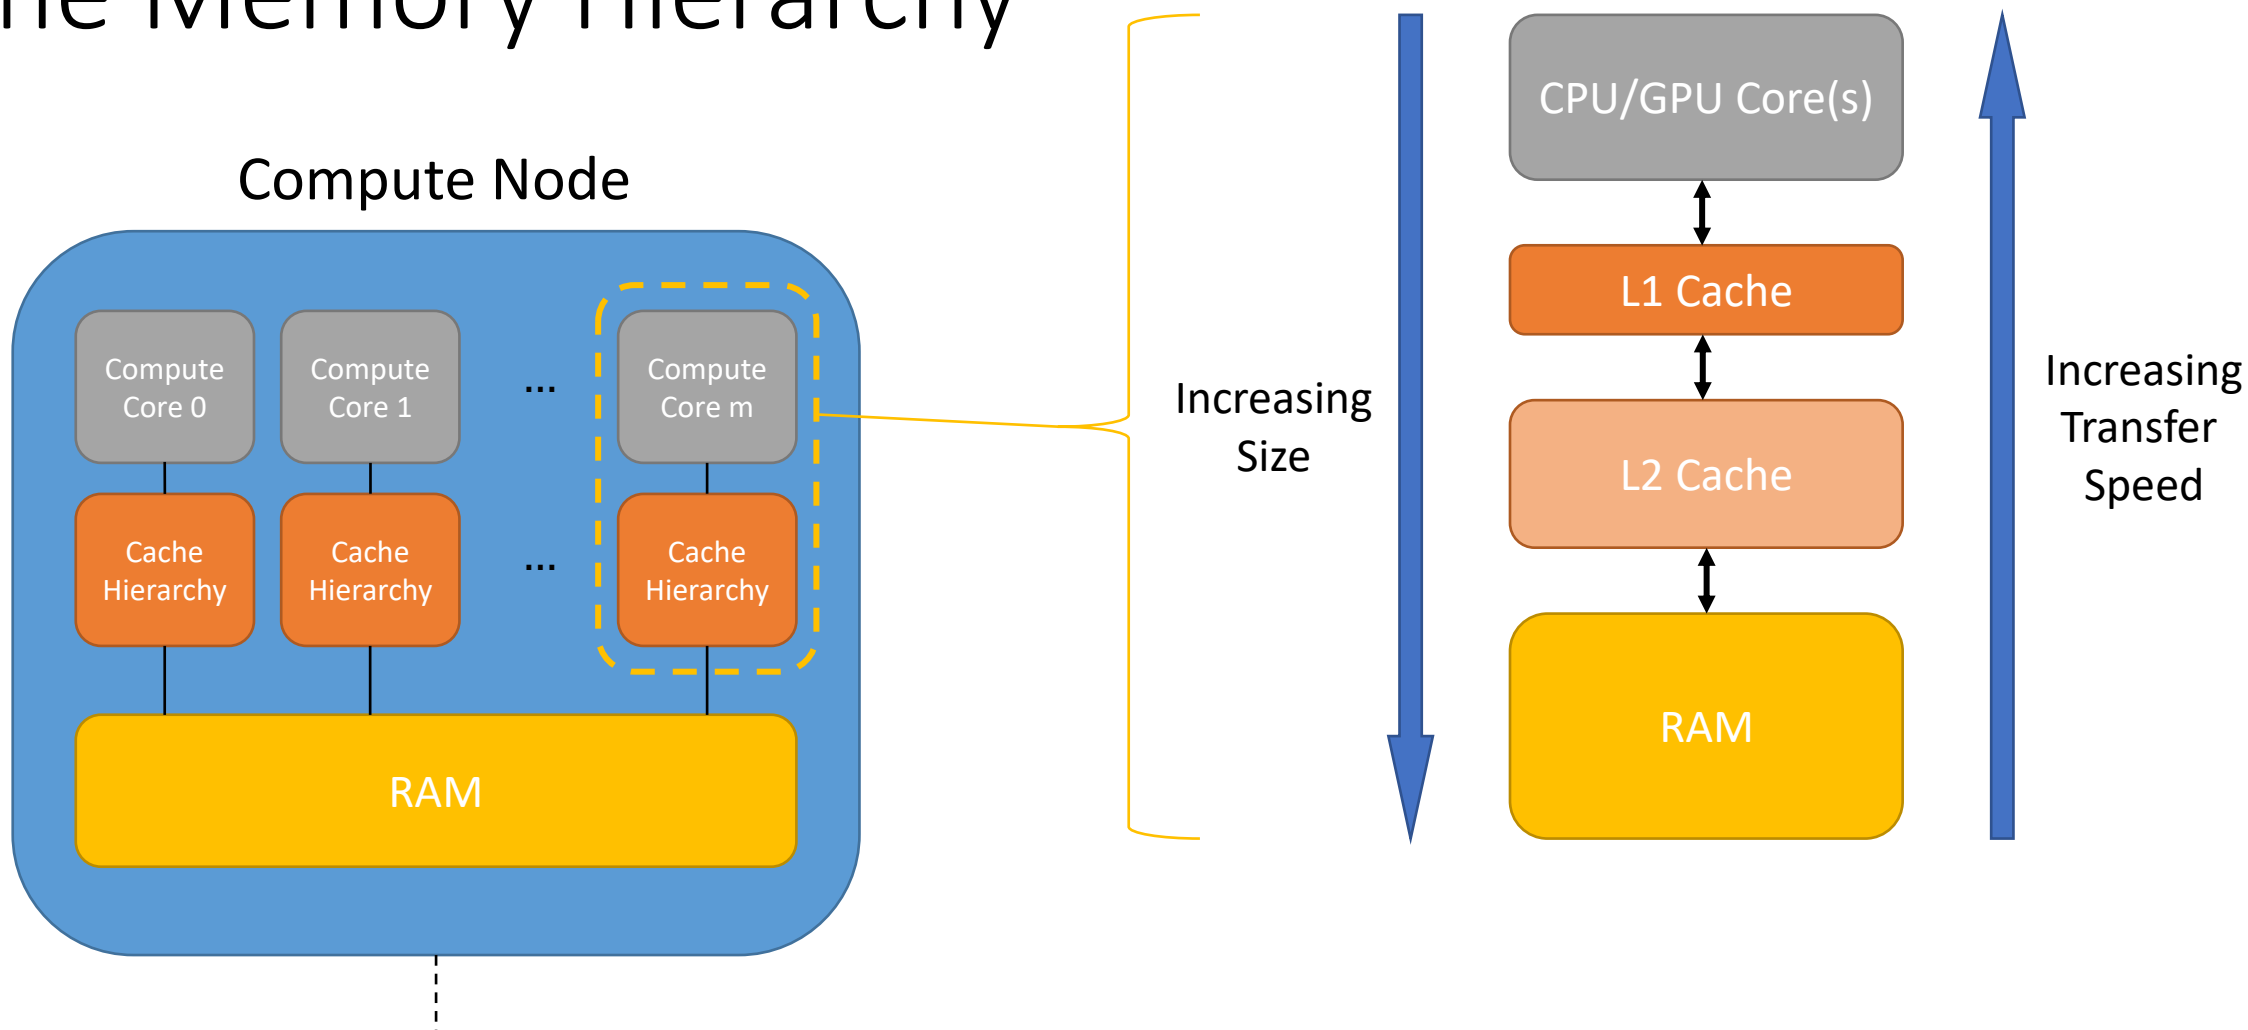
# Coding for Supercomputers

Compute nodes have shared memory so we can rely on *single-instruction multiple-data* (SIMD) standards such as OpenMP

- Avoid editing the same memory from different cores (race condition)
- Avoiding loop branches (i.e. if statements) can also enable acceleration via vectorization



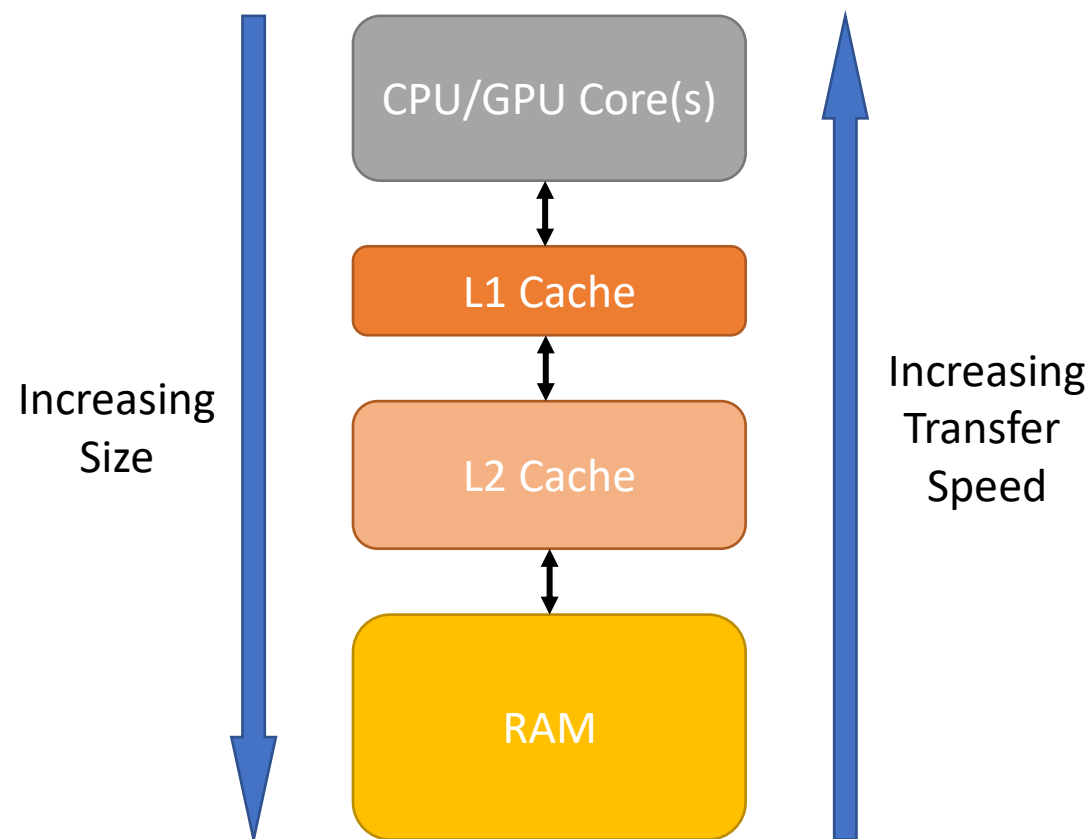
# The Memory Hierarchy






# The Memory Hierarchy

Many algorithms access memory contiguously (i.e. all in order). Knowing this, CPUs will load *chunks* of memory all at once



L1 Cache Chunk



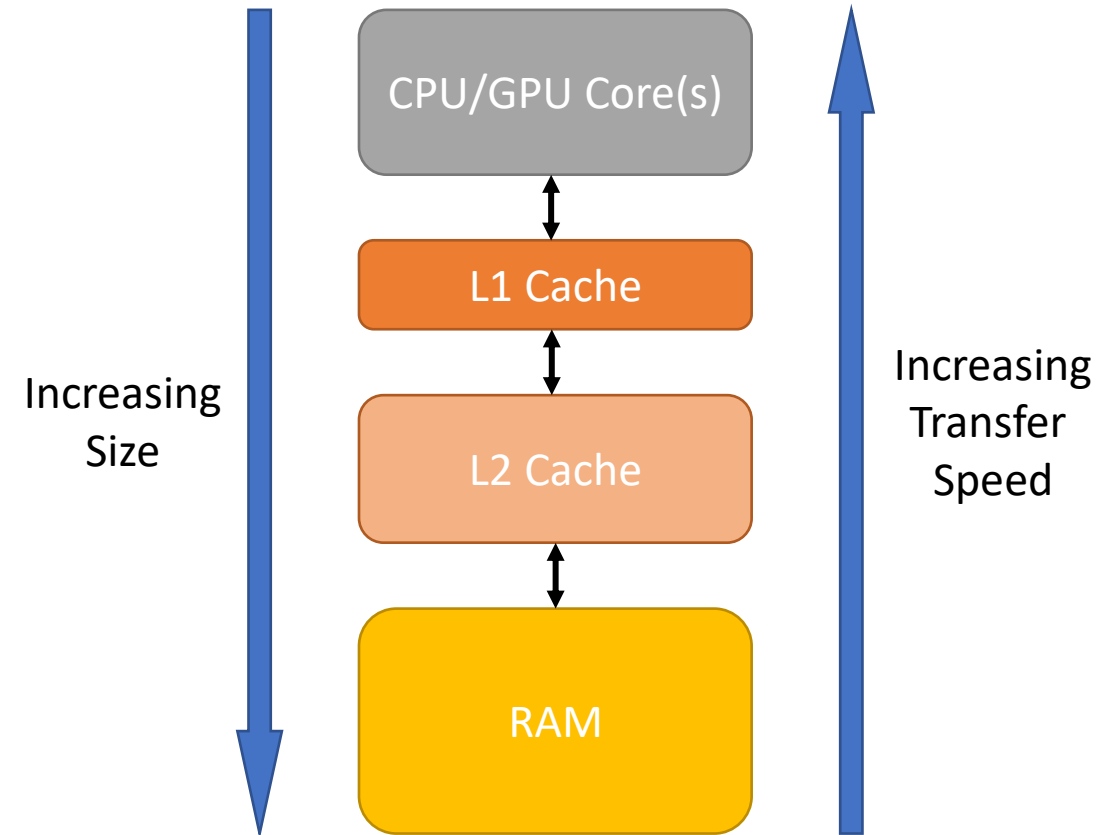

L2 Cache Chunk

# The Memory Hierarchy

## Why PIC is good

- When updating particles we access memory contiguously and using the same instructions. This leads to optimal memory use efficiency

$$\frac{d\mathbf{Q}_p}{dt} = \dot{\mathbf{Q}}_p = \left( \mathbf{v}_p, -\frac{q}{m} \nabla \phi \right)$$




# Strong Scaling in 3D

Modelling a thermal cube of plasma with periodic boundary conditions

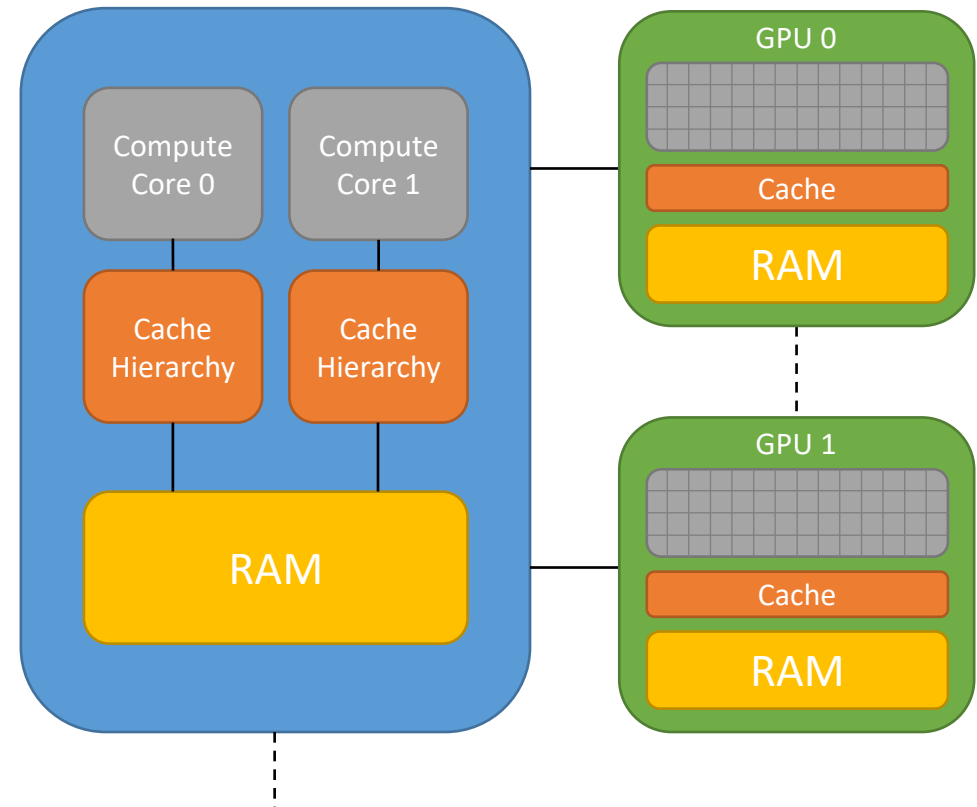
$32 \times 32 \times 32$  cells with 1,000 particles-per-cell

$$\Delta x = \lambda_D, \quad \Delta t \omega_p = 0.2$$





# Acceleration

- GPUs are like a compute node but with thousands of low performance cores
- They perform well on algorithms which repeat the same thing on huge amounts of data
- This makes PIC particle routines (update, collisions, interpolation) ideal for acceleration!
- Data must be transferred from the node to CPU to the GPU which can be time consuming
- Highly recommend 2021 paper by *Juhasz, Durian, Derzsi, Matejcek, Donko and Hartmann* for far more details!



# Improved Strong Scaling in 3D


Brings us closer to the goal of 10ms time step for 3D simulations



# What about the Field Solver?

Can we achieve a speedup with the field solver on GPUs?

Unfortunately in the region where GPUs could help, the time step is too slow for whole device modelling




# GPUs are Not Enough

Development of new algorithms for kinetic WDM

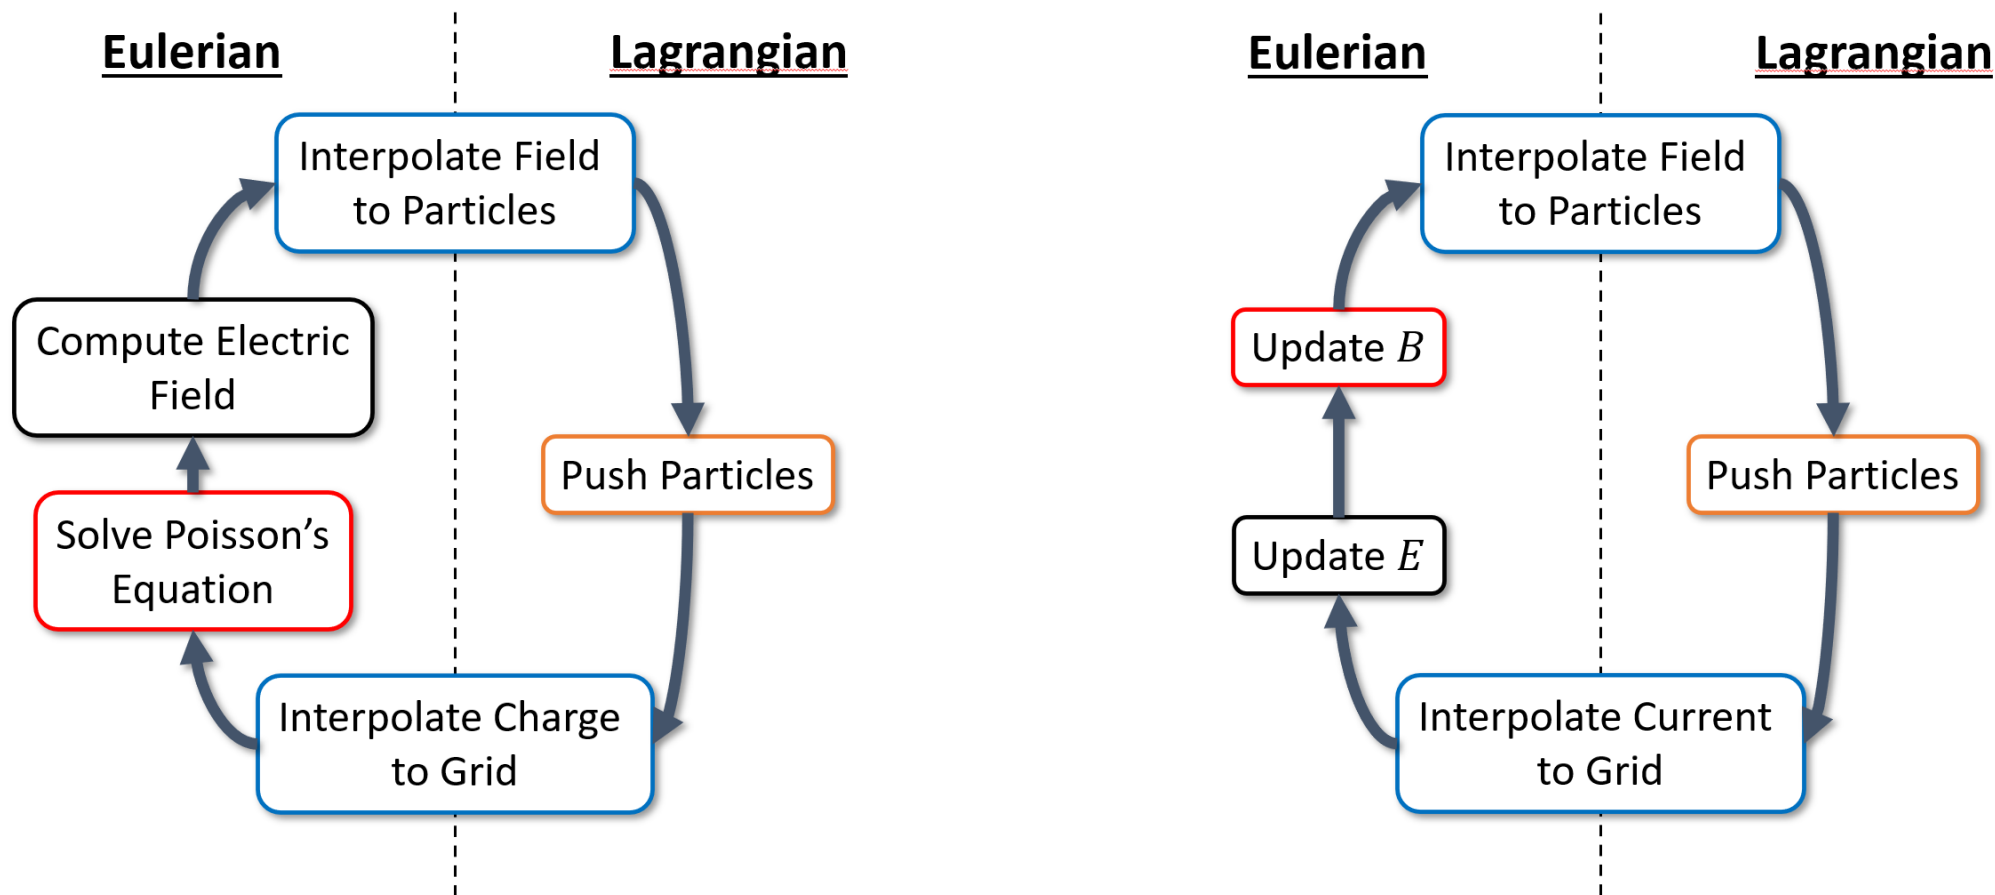
# Re-considering the Electrostatic Approximation

The Poisson equation is now the bottleneck to performance



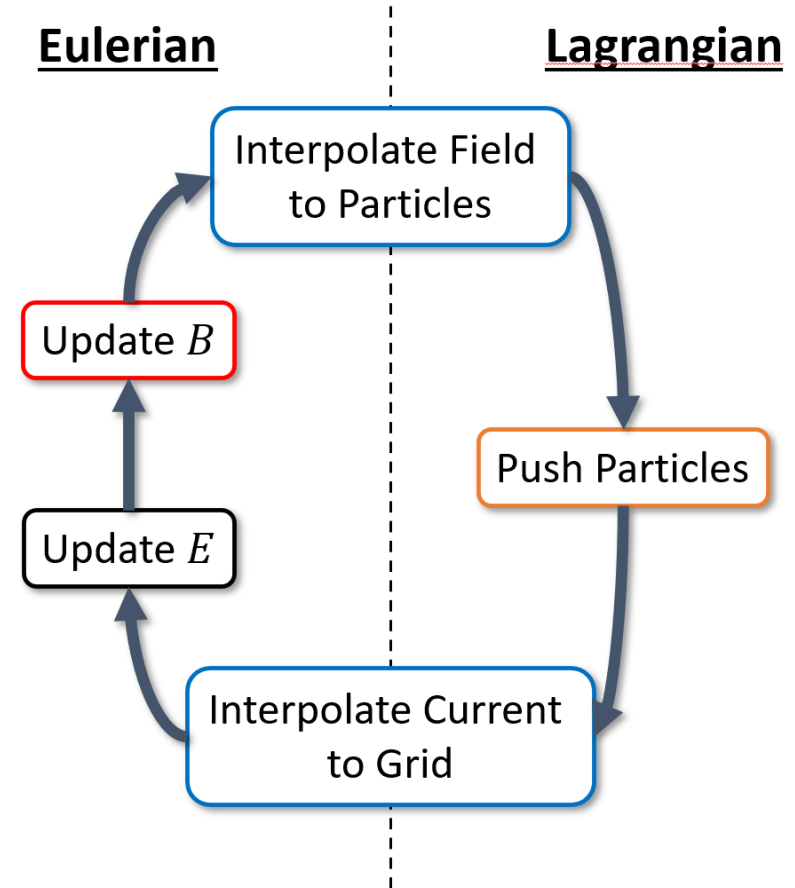


# Poisson vs Maxwell Field Solver



# Poisson vs Maxwell Field Solver

For our particles, we need  
 $\Delta t_{part} < 0.2/\omega_p$



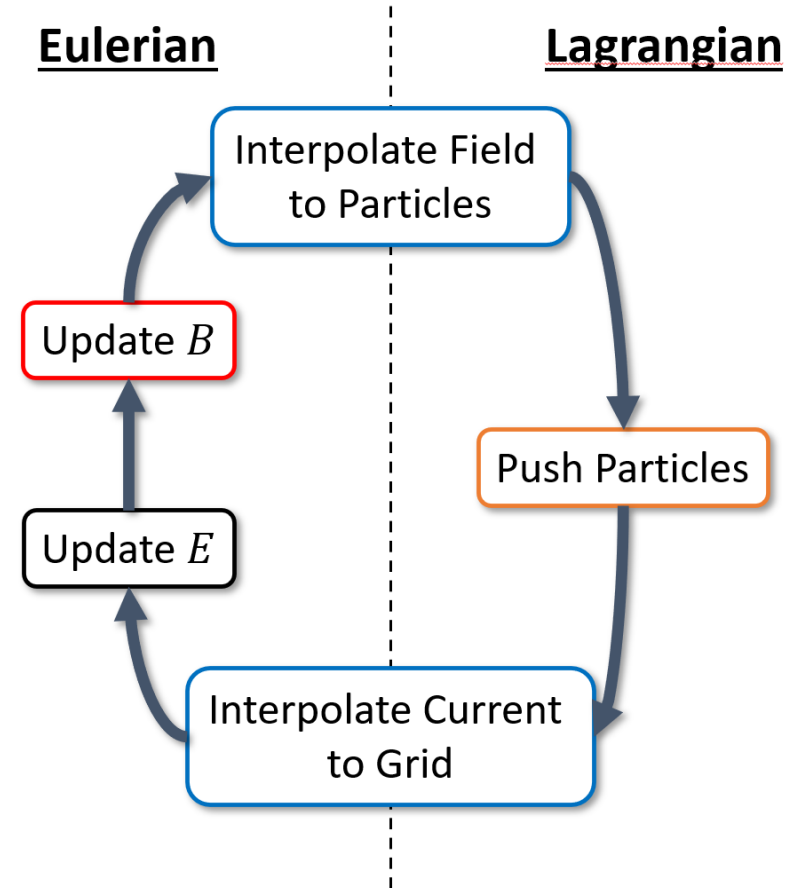
# Poisson vs Maxwell Field Solver

For our particles, we

$$\Delta t_{part} < 0.2/\omega_p$$

But, for an explicit electromagnetic field solver we require:

$$\Delta t_{field} < \Delta x/c$$



# Poisson vs Maxwell Field Solver

For our particles, we need

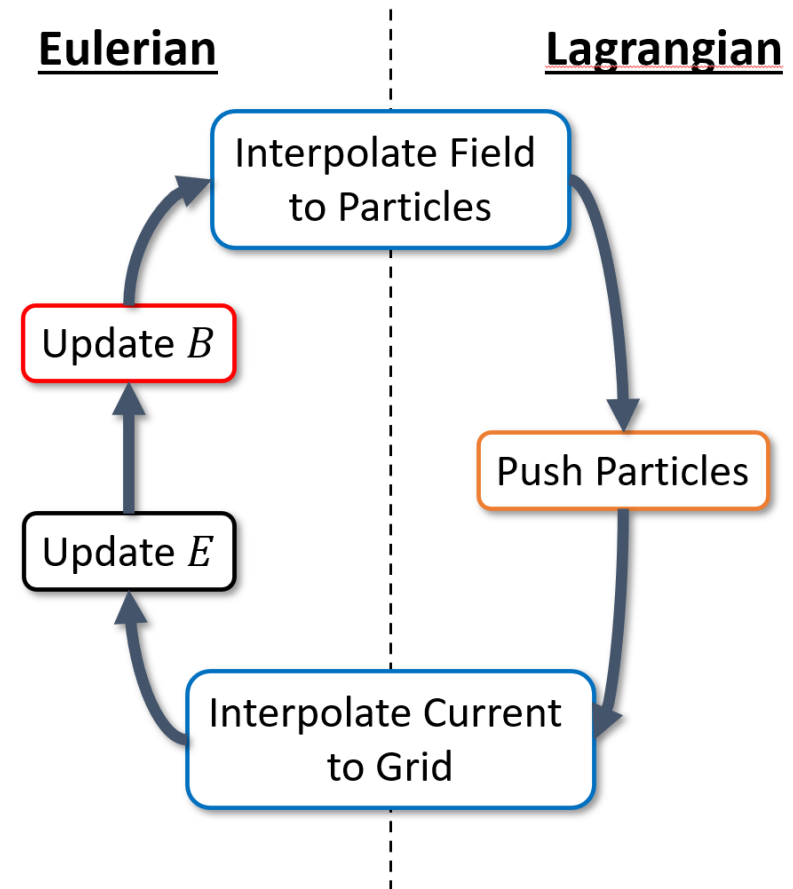
$$\Delta t_{part} < 0.2/\omega_p$$

But, for an explicit electromagnetic field solver we require:

$$\Delta t_{field} < \Delta x/c$$

For a 10 eV plasma with  $\Delta x = \lambda_D$  we have

$$n_c = \frac{\Delta t_{part}}{\Delta t_{field}} = \frac{0.2c}{v_{th}} \approx 50$$



# Sub-cycled Maxwell Solver

For our particles, we need

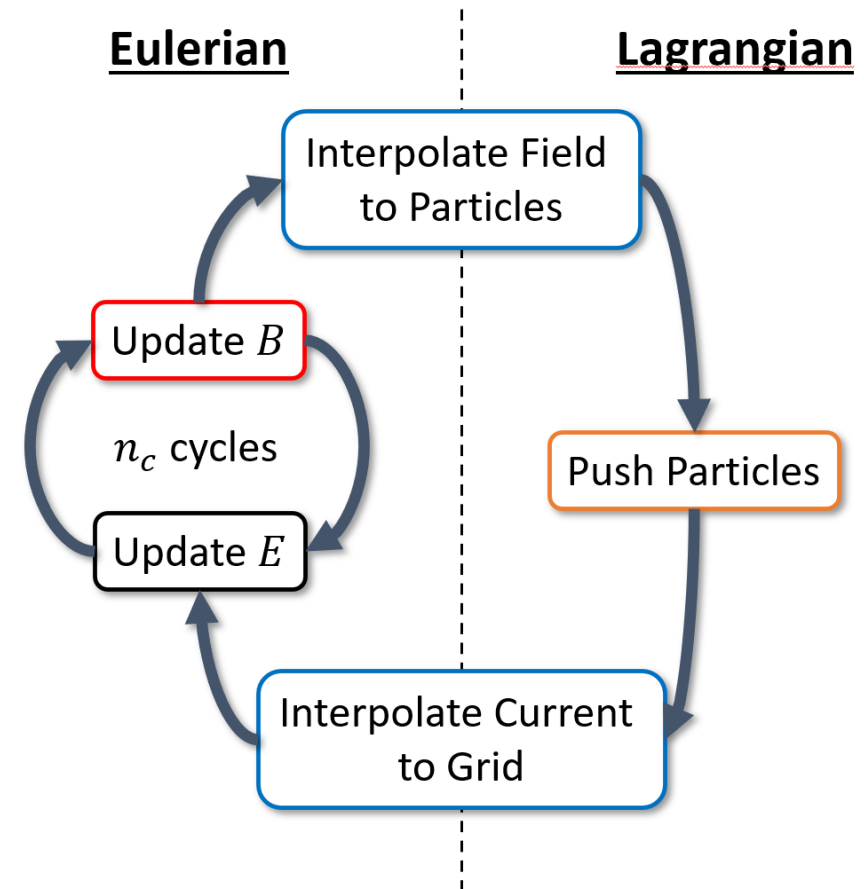
$$\Delta t_{part} < 0.2/\omega_p$$

But, for an explicit electromagnetic field solver we require:

$$\Delta t_{field} < \Delta x/c$$

For a 10 eV plasma with  $\Delta x = \lambda_D$  we have

$$n_c = \frac{\Delta t_{part}}{\Delta t_{field}} = \frac{0.2c}{v_{th}} \approx 50$$




# Sub-cycling a Maxwell Field Solver

Strong scaling tests of the Hypre-PFMG solver vs the sub-cycled Maxwell solver



2D Simulation:  $256 \times 256$  cells

We see a 7 – 12  $\times$  speedup with the sub-cycled Maxwell solver



# Improved Strong Scaling in 3D

Allows us to achieve the goal of 10ms time step size for 3D simulations



# Sub-cycling a Maxwell Field Solver

## Advantages:

- Faster!
- Perfectly scalable, communication time equivalent to domain boundary size
- Less dependent on external software
- We can model EM effects in low-temperature devices!

## Disadvantages:

- **Stability issues**
- Might be noisy?



# Requirements for Whole Device Modelling

- Materials processing plasma reactor
- Assume a  $10^{17} \text{ 1/m}^3$  density Argon plasma with  $10 \text{ eV}$  electrons
- Length scale  $\sim 1 \text{ m}$
- Time scale  $\sim 1 \text{ ms}$
- Based on PIC requirements we need:
  - $10^8$  time steps
  - $10^{12}$  cells
  - $10^{15}$  particles
- If we give each node  $32^3$  cells with 1,000 particle-per-cell we need **30 million compute nodes!**
- This is not even possible, let alone practical for industrial design



Materials processing plasma reactor [Booth, Ecole Polytechnique]

# Requirements for Whole Device Modelling

- Materials processing plasma reactor
- Assume a  $10^{17} \text{ 1/m}^3$  density Argon plasma with  $10 \text{ eV}$  electrons
- Length scale  $\sim 1 \text{ m}$
- Time scale  $\sim 1 \text{ ms}$
- Based on PIC requirements we need:
  - $10^8$  time steps
  - $10^{12}$  cells
  - $10^{15}$  particles
- If we give each node  $32^3$  cells with 1,000 particle-per-cell we need **30 million compute nodes!**
- **We must relax the constraint on the cell size and allow for  $\Delta x \gg \lambda_D$**



Materials processing plasma reactor [Booth, Ecole Polytechnique]


# Energy Conserving PIC

This can be achieved via energy-conserving PIC methods [Lewis 1970, Eremin 2022]

Seems to have been rejected by the community due to non-physical cold beam plasma instabilities

Recent work by Barnes & Chacon (2021) has demonstrated that the energy-conserving PIC method is appropriate for modelling low-temperature plasma devices dominated by ambipolar electric fields

Importantly, does not suffer from the *finite-grid-instability* allowing  $\Delta x \gg \lambda_D$




Stability regions for a warm beam moving through a background plasma with the energy-conserving PIC scheme [Barnes & Chacon 2021]

$$u = U_{beam} / \omega_p \Delta x$$


$$M_{\#} = U_{beam} / v_{th,e}$$

# Benchmark Results with Standard PIC





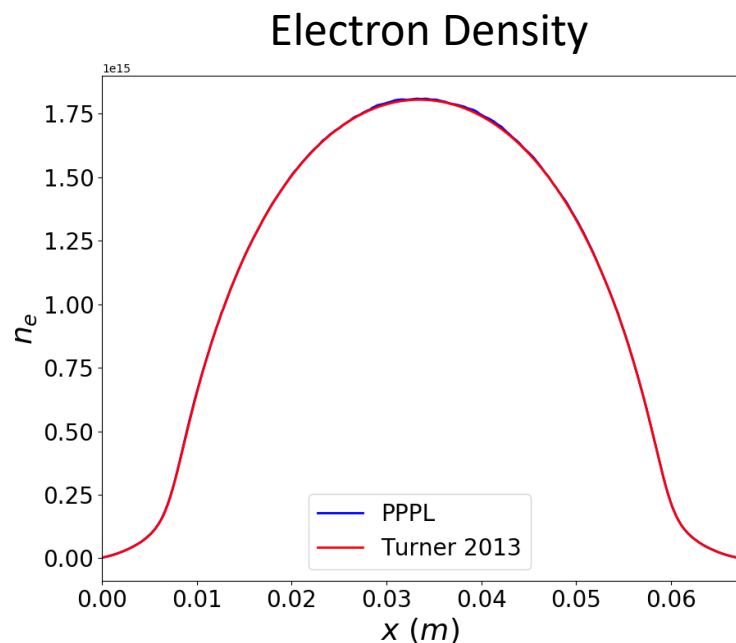
Consider the RF discharge benchmarking case from Turner *et al.* 2013. We begin by studying Case 2 from this paper with the standard electrostatic momentum conserving PIC algorithm

# Results with Energy-Conserving PIC



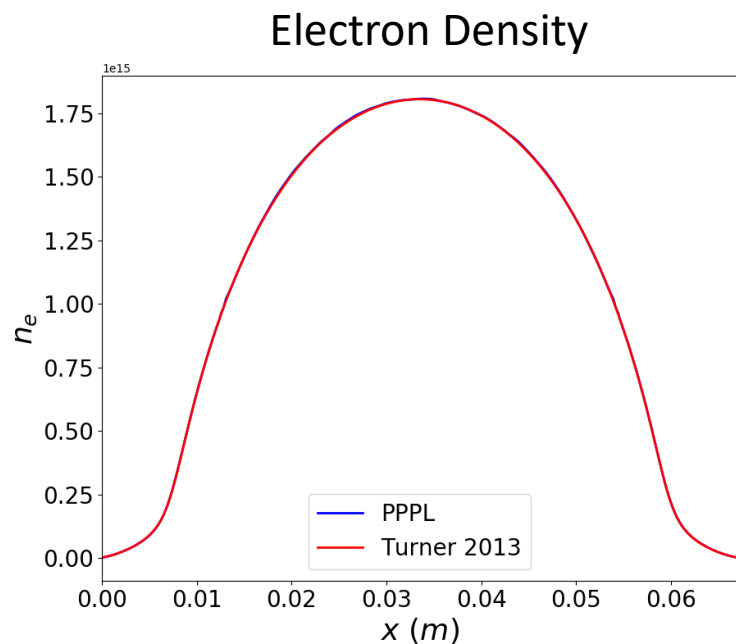


Now compare the results with energy-conserving PIC




# Energy Conserving PIC & Non-Uniform Cells



# Energy Conserving PIC & Non-Uniform Cells

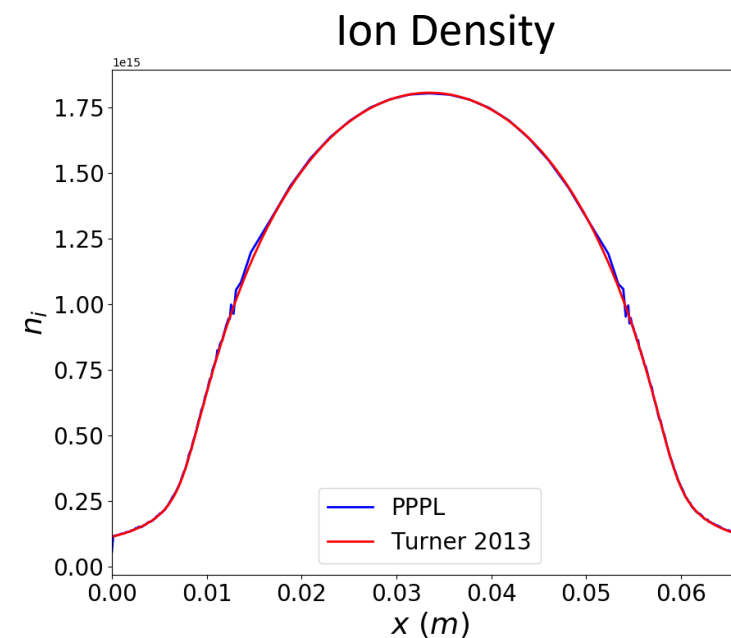



# Energy Conserving PIC & Non-Uniform Cells






# Energy Conserving PIC & Non-Uniform Cells



$$\Delta x_{max} = 16\Delta x_{min}$$



Non-uniform grid

Uniform grid

# Requirements for Whole Device Modelling

- Materials processing plasma reactor
- Assume a  $10^{17} \text{ 1/m}^3$  density Argon plasma with  $10 \text{ eV}$  electrons
- Length scale  $\sim 1 \text{ m}$
- Time scale  $\sim 1 \text{ ms}$
- Based on PIC requirements we need:
  - $10^8$  time steps
  - $10^{12}$  cells
  - $10^{15}$  particles
- Let's assume we can take  $\Delta x \approx 100\lambda_D$



Materials processing plasma reactor [Booth, Ecole Polytechnique]

# Requirements for Whole Device Modelling

- Materials processing plasma reactor
- Assume a  $10^{17} \text{ 1/m}^3$  density Argon plasma with  $10 \text{ eV}$  electrons
- Length scale  $\sim 1 \text{ m}$
- Time scale  $\sim 1 \text{ ms}$
- Based on PIC requirements we need:
  - $10^8$  time steps
  - $10^{12}$  cells  $\rightarrow 10^6$  cells
  - $10^{15}$  particles  $\rightarrow 10^9$  cells
- If we give each node  $32^3$  cells with 1,000 particle-per-cell we need **only need 30 compute nodes**



Materials processing plasma reactor [Booth, Ecole Polytechnique]

# Conclusions

Practical whole-device modelling of low-temperature plasma devices is not yet achievable. To get there we need:

- To fully leverage modern computing systems
- Apply new algorithms to reduce the immense computational cost

## Things I didn't get to (but am happy to discuss)

- Geometric PIC methods (combining energy conservation and sub-cycled Maxwell solvers)
- Memory intensity vs arithmetic intensity
- Comparing PIC and Vlasov solvers

# Backup Slides

# Geometric PIC Methods

Energy conserving PIC methods can be derived directly from the underlying differential geometric structure [Xiao *et al.* 2015, Glasser & Qin, 2020]

Formally, the Vlasov-Maxwell system can be described by Poisson bracket:

$$\begin{aligned} \{F, G\}^{\text{red}}[\bar{f}, \mathbf{B}, \mathbf{E}] = & \int d\mathbf{x} d\mathbf{v} \left[ \bar{f} \left\{ \frac{\delta F}{\delta \bar{f}}, \frac{\delta G}{\delta \bar{f}} \right\}_{xv} \right. \\ & \left. + \bar{f} \mathbf{B} \cdot \left( \frac{\partial}{\partial \mathbf{v}} \frac{\delta F}{\delta \bar{f}} \times \frac{\partial}{\partial \mathbf{v}} \frac{\delta G}{\delta \bar{f}} \right) + \left( \frac{\delta F}{\delta \mathbf{E}} \cdot \frac{\partial \bar{f}}{\partial \mathbf{v}} \frac{\delta G}{\delta \bar{f}} - \frac{\delta G}{\delta \mathbf{E}} \cdot \frac{\partial \bar{f}}{\partial \mathbf{v}} \frac{\delta F}{\delta \bar{f}} \right) \right] \\ & + \int d\mathbf{x} \left( \frac{\delta F}{\delta \mathbf{E}} \cdot \nabla \times \frac{\delta G}{\delta \mathbf{B}} - \frac{\delta G}{\delta \mathbf{E}} \cdot \nabla \times \frac{\delta F}{\delta \mathbf{B}} \right). \end{aligned}$$

With Hamiltonian:

$$H(f, \mathbf{E}, \mathbf{B}) = \int \mathbf{v} \cdot \mathbf{v} f d\mathbf{x} d\mathbf{v} + \frac{1}{2} \int \mathbf{E}^2 d\mathbf{x} + \frac{1}{2} \int \mathbf{B}^2 d\mathbf{x}$$

# Geometric PIC Methods

The dynamical equations are then recovered using:

$$\frac{\partial f}{\partial t} = \{f, H\} = -\mathbf{v} \cdot \frac{\partial f}{\partial \mathbf{x}} - (\mathbf{E} + \mathbf{v} \times \mathbf{B}) \cdot \frac{\partial f}{\partial \mathbf{v}}$$

$$\frac{\partial \mathbf{E}}{\partial t} = \{\mathbf{E}, H\} = - \int \mathbf{v} f d\mathbf{v} + \nabla \times \mathbf{B}$$

$$\frac{\partial \mathbf{B}}{\partial t} = \{\mathbf{B}, H\} = -\nabla \times \mathbf{E}$$

# Geometric PIC Methods

We can discretise the Poisson bracket as follows:

- Use a Klimontovich particle distribution
- Discretise field quantities over a finite grid
- Interpolations via structure preserving Whitney forms
- Differential operators are chosen via Discrete Exterior Calculus

$$\begin{aligned} \{\{F, G\}\}^{\text{red}}[\bar{f}, \mathbf{B}, \mathbf{E}] &= \int d\mathbf{x} d\mathbf{v} \left[ \bar{f} \left\{ \frac{\delta F}{\delta \bar{f}}, \frac{\delta G}{\delta \bar{f}} \right\}_{xv} \right. \\ &\quad \left. + \bar{f} \mathbf{B} \cdot \left( \frac{\partial}{\partial \mathbf{v}} \frac{\delta F}{\delta \bar{f}} \times \frac{\partial}{\partial \mathbf{v}} \frac{\delta G}{\delta \bar{f}} \right) + \left( \frac{\delta F}{\delta \mathbf{E}} \cdot \frac{\partial \bar{f}}{\partial \mathbf{v}} \frac{\delta G}{\delta \bar{f}} - \frac{\delta G}{\delta \mathbf{E}} \cdot \frac{\partial \bar{f}}{\partial \mathbf{v}} \frac{\delta F}{\delta \bar{f}} \right) \right] \\ &\quad + \int d\mathbf{x} \left( \frac{\delta F}{\delta \mathbf{E}} \cdot \nabla \times \frac{\delta G}{\delta \mathbf{B}} - \frac{\delta G}{\delta \mathbf{E}} \cdot \nabla \times \frac{\delta F}{\delta \mathbf{B}} \right). \end{aligned}$$

$$H(f, \mathbf{E}, \mathbf{B}) = \int \mathbf{v} \cdot \mathbf{v} f d\mathbf{x} d\mathbf{v} + \frac{1}{2} \int \mathbf{E}^2 d\mathbf{x} + \frac{1}{2} \int \mathbf{B}^2 d\mathbf{x}$$



# Geometric PIC Methods

We can discretise the Poisson bracket as follows:

- Use a Klimontovich particle distribution
- Discretise field quantities over a finite grid
- Interpolations via structure preserving Whitney forms
- Differential operators are chosen via Discrete Exterior Calculus

$$\begin{aligned} \{\{F, G\}\}_d^{\text{red}}[X_i, V_i, \mathbf{B}_n, \mathbf{E}_n] &= \sum_{i=1}^L \left( \frac{\partial F}{\partial X_i} \cdot \frac{\partial G}{\partial V_i} - \frac{\partial G}{\partial X_i} \cdot \frac{\partial F}{\partial V_i} + \left[ \frac{\partial F}{\partial V_i} \times \frac{\partial G}{\partial V_i} \right] \cdot \sum_{n=1}^N \mathbf{B}_n W_{\sigma_2}(X_i - \mathbf{x}_n) \right) \\ &+ \sum_{n=1}^N \left( \left[ \sum_{i=1}^L \frac{\partial F}{\partial V_i} W_{\sigma_1}(X_i - \mathbf{x}_n) - \left( \nabla_d^- \times \frac{\partial F}{\partial \mathbf{B}} \right)_n \right] \cdot \frac{\partial G}{\partial \mathbf{E}_n} \right. \\ &\left. - \left[ \sum_{i=1}^L \frac{\partial G}{\partial V_i} W_{\sigma_1}(X_i - \mathbf{x}_n) - \left( \nabla_d^- \times \frac{\partial G}{\partial \mathbf{B}} \right)_n \right] \cdot \frac{\partial F}{\partial \mathbf{E}_n} \right). \end{aligned}$$

$$H_d^{\text{red}}[X_i, V_i, \mathbf{B}_n, \mathbf{E}_n] = \frac{1}{2} \sum_{i=1}^L V_i^2 + \frac{1}{2} \sum_{n=1}^N (\mathbf{E}_n^2 + \mathbf{B}_n^2) = \sum_{\alpha=1}^3 H_V^\alpha + H_B + H_E$$

# Geometric PIC Methods

Lie splitting of the Hamiltonian leads to the following sets of semi-discrete dynamical equations

$$H_d^{\text{red}}[\mathbf{X}_i, \mathbf{V}_i, \mathbf{B}_n, \mathbf{E}_n] = \frac{1}{2} \sum_{i=1}^L V_i^2 + \frac{1}{2} \sum_{n=1}^N (\mathbf{E}_n^2 + \mathbf{B}_n^2) = \sum_{\alpha=1}^3 H_V^\alpha + H_B + H_E$$

$$H_V^\alpha \begin{cases} \dot{X}_i^\beta = \delta_\alpha^\beta V_i^\alpha, \\ \dot{V}_i^\beta = \epsilon_{\beta\alpha\gamma} V_i^\alpha \sum_{n=1}^N B_n^\gamma W_{\sigma_2}(\mathbf{X}_i - \mathbf{x}_n), \\ \dot{B}_n^\beta = 0, \\ \dot{E}_n^\beta = -\delta_\alpha^\beta \sum_{i=1}^L V_i^\alpha W_{\sigma_1}(\mathbf{X}_i - \mathbf{x}_n), \end{cases}$$

$$H_B \begin{cases} \dot{X}_i = 0, \\ \dot{V}_i = 0, \\ \dot{B}_n = 0, \\ \dot{E}_n = (\nabla_d^- \times \mathbf{B})_n, \end{cases}$$

$$H_E \begin{cases} \dot{X}_i = 0, \\ \dot{V}_i = \sum_{n=1}^N \mathbf{E}_n W_{\sigma_1}(\mathbf{X}_i - \mathbf{x}_n), \\ \dot{B}_n = -(\nabla_d^+ \times \mathbf{E})_n, \\ \dot{E}_n = 0. \end{cases}$$

# Geometric PIC Methods

These equations can be integrated exactly, providing phase space maps

$$\Theta_x(\Delta t) = \begin{cases} X_i^\beta(t + \Delta t) &= X_i^\beta(t) + \Delta t \delta_\alpha^\beta V_i^\alpha(t) \\ V_i^\beta(t + \Delta t) &= V_i^\beta(t) + \epsilon_{\beta\alpha\gamma} V_i^\alpha \int_0^{\Delta t} \sum_n W_{\sigma 2} B_n^\gamma(t) (\mathbf{x}_n - \mathbf{X}_i(s)) ds \\ B_n^\beta(t + \Delta t) &= B_n^\beta(t) \\ E_n^\beta(t + \Delta t) &= E_n^\beta(t) - \delta_\alpha^\beta \int_0^{\Delta t} \sum_i V_i^\alpha W_{\sigma 1} (\mathbf{x}_n - \mathbf{X}_i(s)) ds \end{cases}$$

$$\Theta_B(\Delta t) = \begin{cases} \mathbf{X}_i(t + \Delta t) &= \mathbf{X}_i(t) \\ \mathbf{V}_i(t + \Delta t) &= \mathbf{V}_i(t) \\ \mathbf{B}_n(t + \Delta t) &= \mathbf{B}_n(t) \\ \mathbf{E}_n(t + \Delta t) &= \mathbf{E}_n(t) + (\nabla_d^- \times \mathbf{B})_n \end{cases}$$

$$\Theta_E(\Delta t) = \begin{cases} \mathbf{X}_i(t + \Delta t) &= \mathbf{X}_i(t) \\ \mathbf{V}_i(t + \Delta t) &= \mathbf{V}_i(t) + \sum_n \mathbf{E}_n W_{\sigma 1} (\mathbf{x}_n - \mathbf{X}_i(t)) \\ \mathbf{B}_n(t + \Delta t) &= \mathbf{B}_n(t) - (\nabla_d^- \times \mathbf{E})_n \\ \mathbf{E}_n(t + \Delta t) &= \mathbf{E}_n(t) \end{cases}$$

# Geometric PIC Methods

We construct numerical methods by combining these maps:

- First order scheme:

$$\Theta_1(\Delta t) = \Theta_E(\Delta t) \Theta_B(\Delta t) \Theta_z(\Delta t) \Theta_y(\Delta t) \Theta_x(\Delta t)$$

- Second order scheme:

$$\Theta_1(\Delta t) = \Theta_x(\Delta t/2) \Theta_y(\Delta t/2) \Theta_z(\Delta t/2) \Theta_B(\Delta t/2) \Theta_E(\Delta t) \Theta_B(\Delta t/2) \Theta_z(\Delta t/2) \Theta_y(\Delta t/2) \Theta_x(\Delta t/2)$$

- Plus there are simple recipes for higher order schemes


# Geometric PIC Methods

## **Advantages** of these methods:

- They conserve charge exactly
- They narrowly bound energy growth (no finite grid instability)
- There is a clear recipe for building the schemes (no ad-hoc selection of numerical methods)
- The algorithms are explicit (great for GPUs)

## **Disadvantages** of these methods:

- Not momentum conserving (although fairly well bounded)
- Still limited by the standard explicit algorithm constraints on time steps



3D collisionless simulation of the Penning discharge with density contours and electron current streamlines