

A photograph of a modern, multi-story building with a glass and metal facade, likely the Lawrence Berkeley National Laboratory. The building is set against a clear blue sky. In the foreground, there is a paved area and some greenery. A large, tall, thin tree stands to the left of the building. In the background, a large, dome-shaped structure is visible.

Modeling plasma accelerators I: “Full” Electromagnetic Particle-In-Cell codes



Remi Lehe

Lawrence Berkeley National Laboratory, Berkeley, California

February 6, 2023

Numerical simulations of plasmas: motivations

- Typical **needs for modelling**:
 - Designing future experiments:
e.g. quantitative estimate of beam properties
 - Interpreting existing experiments
esp. with limited diagnostics
- **Analytical models** are **key** to our understanding, but cannot always capture all relevant phenomena, e.g.
 - Exact fields in the bubble regime
 - Non-linear focusing of complex lasers in plasmas
 - ...
- Thus, in many cases, **numerical simulations** are needed.



Overview

Modeling plasma accelerators I: “Full” Electromagnetic Particle-In-Cell codes

- Fundamental equations for plasma accelerators
- The electromagnetic Particle-In-Cell (PIC) algorithm
- Discretizing the field and particle equations
- Parallelization of PIC codes

Modeling plasma accelerators II: Making the simulations faster

- The boosted-frame technique
- Cylindrical geometry
- Laser envelope model
- Quasi-static PIC codes

Outline

- **Fundamental equations for plasma accelerators**
- The electromagnetic Particle-In-Cell (PIC) algorithm
- Discretizing the field and particle equations
- Parallelization of PIC codes

Which fundamental equations capture the physics of beam-driven plasma acceleration?

- The driver beam creates strong E and B fields (“space-charge fields”)
- These fields displace the background plasma particles.
- The displaced plasma particles generate E and B fields behind the driver (the “wakefield”)
- The wakefield accelerates the witness beam (and decelerates the driver beam).

Maxwell's equations


$$\nabla \cdot \mathbf{E} = \frac{\rho}{\epsilon_0} \quad \partial_t \mathbf{B} = -\nabla \times \mathbf{E}$$

$$\nabla \cdot \mathbf{B} = 0 \quad \frac{1}{c^2} \partial_t \mathbf{E} = \nabla \times \mathbf{B} - \mu_0 \mathbf{j}$$

Relativistic equations of motion, with Lorentz force

$$\frac{d\mathbf{p}}{dt} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B})$$

$$\frac{d\mathbf{x}}{dt} = \mathbf{v} \quad \left(\mathbf{p} = \frac{m\mathbf{v}}{\sqrt{1 - \mathbf{v}^2/c^2}} \right)$$



Courtesy of
M. Thévenet
and S. Diederichs

Which fundamental equations capture the physics of beam-driven plasma acceleration?

Maxwell's equations

$$\nabla \cdot \mathbf{E} = \frac{\rho}{\epsilon_0} \quad \partial_t \mathbf{B} = -\nabla \times \mathbf{E}$$

$$\nabla \cdot \mathbf{B} = 0 \quad \frac{1}{c^2} \partial_t \mathbf{E} = \nabla \times \mathbf{B} - \mu_0 \mathbf{j}$$

$\rho, \mathbf{j}$: contains the contributions from the **driver and witness beam** and the **plasma particles**

$\mathbf{E}, \mathbf{B}$: contains superposition of the space-charge field from **the driver and witness beam**, and the wakefield from the **plasma particles**.


Relativistic equations of motion, with Lorentz force

$$\frac{d\mathbf{p}}{dt} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B})$$

$$\frac{d\mathbf{x}}{dt} = \mathbf{v} \quad \left(\mathbf{p} = \frac{m\mathbf{v}}{\sqrt{1 - \mathbf{v}^2/c^2}} \right)$$

Is applied to every particles of the **driver and witness beam** as well as the **plasma particles**

Which fundamental equations capture the physics of laser-driven plasma acceleration?



- The laser pulse is an electromagnetic wave (oscillating E and B field)
- The net effect of the oscillating E and B is to displace the plasma particles (“ponderomotive force”)

- The displaced plasma particles generate E and B fields behind the driver (the “wakefield”)
- The wakefield accelerates the witness beam (and decelerates the driver beam).
- The laser pulse evolves as it propagates through the plasma.

Maxwell's equations

$$\nabla \cdot \mathbf{E} = \frac{\rho}{\epsilon_0} \quad \partial_t \mathbf{B} = -\nabla \times \mathbf{E}$$

$$\nabla \cdot \mathbf{B} = 0 \quad \frac{1}{c^2} \partial_t \mathbf{E} = \nabla \times \mathbf{B} - \mu_0 \mathbf{j}$$

Relativistic equations of motion, with Lorentz force

$$\frac{d\mathbf{p}}{dt} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B})$$

$$\frac{d\mathbf{x}}{dt} = \mathbf{v} \quad \left(\mathbf{p} = \frac{m\mathbf{v}}{\sqrt{1 - \mathbf{v}^2/c^2}} \right)$$

Which fundamental equations capture the physics of laser-driven plasma acceleration?

Maxwell's equations

$$\nabla \cdot \mathbf{E} = \frac{\rho}{\epsilon_0} \quad \partial_t \mathbf{B} = -\nabla \times \mathbf{E}$$

$$\nabla \cdot \mathbf{B} = 0 \quad \frac{1}{c^2} \partial_t \mathbf{E} = \nabla \times \mathbf{B} - \mu_0 \mathbf{j}$$

$\rho, \mathbf{j}$: contains the contributions from the **witness beam** and the **plasma particles**

$\mathbf{E}, \mathbf{B}$: contains superposition of the **laser electromagnetic field**, the wakefield from the **plasma particles**, and the space charge field from the **witness beam**

Relativistic equations of motion, with Lorentz force

$$\frac{d\mathbf{p}}{dt} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B})$$


$$\frac{d\mathbf{x}}{dt} = \mathbf{v} \quad \left(\mathbf{p} = \frac{m\mathbf{v}}{\sqrt{1 - \mathbf{v}^2/c^2}} \right)$$

Is applied to every particles of the **witness beam** as well as the **plasma particles**

Outline

- Fundamental equations for plasma accelerators
- **The electromagnetic Particle-In-Cell (PIC) algorithm**
- Discretizing the field and particle equations
- Parallelization of PIC codes

The fields and particles are represented in different ways.




The **E and B fields** are represented on a grid (standard way to solve PDEs).

The **beam and plasma particles** are represented by discrete particles, that move through the grid.


The plasma and beams are actually represented by “macroparticles”.

- Typical beams : $\sim 10^8 - 10^{10}$ particles.
($100\ \mu m$)³ of plasma at $10^{18}\ cm^{-3}$: 10^{12} particles.
 $\Rightarrow$ **It is too computationally expensive to represent and track every single physical particle!**
 - In practice, particles that are **initially close** in 6D phase space follow **similar trajectories**.
 - Simulations use **macroparticles**, which represent several physical particles that are close in phase space.
- Each macroparticle is characterized by its:
- Position **x**
 - Momentum **p**
 - “Weight” **w** (= how many physical particles it represents)



The workflow of Particle-In-Cell simulation: initialization

1. Initialize the plasma and laser at t=0




Note: another popular method is to emit the laser **continuously throughout the simulation**, from one of the boundaries or from an “antenna”

The workflow of Particle-In-Cell simulation: time evolution

1. Initialize the plasma and laser at t=0



2. Repeatedly update the fields and particles, in discrete timesteps using a discretized version of the Maxwell equations and equations of motion




$$\partial_t \mathbf{B} = -\nabla \times \mathbf{E}$$

$$\partial_t \mathbf{E} = c^2 \nabla \times \mathbf{B} - \mu_0 c^2 \mathbf{j}$$



$$\frac{d\mathbf{p}}{dt} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B})$$

$$\frac{d\mathbf{x}}{dt} = \mathbf{v} \quad \left(\mathbf{p} = \frac{m\mathbf{v}}{\sqrt{1 - v^2/c^2}} \right)$$


The time evolution of particles and fields is computed over discrete timesteps, using the PIC loop



At each Δt : one iteration of the following loop



Instead of having the simulation box cover the whole plasma, PIC simulations often use a moving window



Note: the grid cells are not drawn to scale. (They should be much finer, to resolve the laser wavelength.)

Ensuring that $\nabla \cdot \mathbf{B} = 0$


- The PIC loop only uses 2 out of the 4 Maxwell equations.
How do we ensure that $\nabla \cdot \mathbf{B} = 0$?
- For the **continuous** Maxwell equations:

- $\nabla \cdot \mathbf{B} = 0$ at $t = 0$

- $\mathbf{B}$ evolves according to:
 $\partial_t \mathbf{B} = -\nabla \times \mathbf{E}$

$\Rightarrow$

$$\nabla \cdot \mathbf{B} = 0 \text{ at any } t$$



- This remains true for the **discretized** Maxwell equations.
(with the most common discretization schemes)
- It is therefore important to ensure $\nabla \cdot \mathbf{B} = 0$ in the simulation box, when **initializing** the simulation

Ensuring that $\nabla \cdot \mathbf{E} = \rho/\epsilon_0$

- The PIC loop only uses 2 out of the 4 Maxwell equations.
How do we ensure that $\nabla \cdot \mathbf{E} = \rho/\epsilon_0$?

- For the **continuous** Maxwell equations:

- $\nabla \cdot \mathbf{E} = \rho/\epsilon_0$ at $t = 0$

- $\mathbf{E}$ evolves according to:

$$\partial_t \mathbf{E} = c^2 \nabla \times \mathbf{B} - \mu_0 c^2 \mathbf{j}$$


$\Rightarrow$

$$\nabla \cdot \mathbf{E} = \rho/\epsilon_0 \text{ at any } t$$

- $\mathbf{j}, \rho$ satisfy:

$$\partial_t \rho + \nabla \cdot \mathbf{j} = 0$$



- This remains true for the **discretized** Maxwell equations (with the most common discretization schemes).
- However, making sure that the deposited $\mathbf{j}, \rho$ satisfy $\partial_t \rho + \nabla \cdot \mathbf{j} = 0$ is not trivial!



Outline

- Fundamental equations for plasma accelerators
- The electromagnetic Particle-In-Cell (PIC) algorithm
- **Discretizing the field and particle equations**
- Parallelization of PIC codes

Discretizing the field and particle equations



Discretizing the equations of motion: derivatives are replaced by finite-difference approximations

Continuous equation:

$$\frac{dx}{dt} = v$$



Discretized equation:

$$\frac{x^{new} - x^{old}}{\Delta t} = v$$

$$\text{i.e. } x^{new} = x^{old} + v\Delta t$$



Discretizing the equations of motion: derivatives are replaced by finite-difference approximations

Continuous equation:

$$\frac{d\mathbf{p}}{dt} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B})$$





Discretized equation:

$$\frac{\mathbf{p}^{new} - \mathbf{p}^{old}}{\Delta t} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B})$$

$$\left(\mathbf{p} = \frac{m\mathbf{v}}{\sqrt{1 - v^2/c^2}} \right)$$

Different algorithms exist (e.g. Boris, Vay, Higuera-Cary, ...) depending on the details of how the term $\mathbf{v} \times \mathbf{B}$ is treated.

Discretizing the field and particle equations



Discretizing the Maxwell equations: derivatives are replaced by finite-difference approximations

Continuous equation:

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




Discretized equation:

$$\frac{\mathbf{B}^{new} - \mathbf{B}^{old}}{\Delta t} = -\hat{\nabla} \times \mathbf{E}$$

i.e. $\mathbf{B}^{new} = \mathbf{B}^{old} - \Delta t \hat{\nabla} \times \mathbf{E}$

$\hat{\nabla}$: discretized spatial derivative

e.g. $\hat{\partial}_x \mathbf{E} = \frac{\mathbf{E}_{i+1,j,k} - \mathbf{E}_{i,j,k}}{\Delta x}$



(In practice, the components of $\mathbf{E}$ and $\mathbf{B}$ are staggered in time and space in order to increase the accuracy of this discretization.)

$\Delta x, \Delta y, \Delta z$: cell size

Discretizing the Maxwell equations: derivatives are replaced by finite-difference approximations

Continuous equation:


$$\frac{\partial \mathbf{E}}{\partial t} = c^2 \nabla \times \mathbf{B} - \mu_0 c^2 \mathbf{j}$$

Discretized equation:

$$\frac{\mathbf{E}^{new} - \mathbf{E}^{old}}{\Delta t} = c^2 \hat{\nabla} \times \mathbf{B} - \mu_0 c^2 \mathbf{j}$$


$$\text{i.e.} \quad \mathbf{E}^{new} = \mathbf{E}^{old} + \Delta t c^2 \hat{\nabla} \times \mathbf{B} - \Delta t \mu_0 c^2 \mathbf{j}$$

- This discretization scheme is called the **Yee scheme**.
- Other discretization schemes are also commonly used (e.g. using more points on the grid to evaluate the finite-difference spatial derivative)



Resolution requirement: beam-driven acceleration

- **Linear regime:**
 $\Delta x, \Delta y, \Delta z$ should be much **smaller** than λ_p and the beam(s) size
- **Nonlinear regime:**
 $\Delta x, \Delta y, \Delta z$ may need to be **smaller** than the thin sheath along the bubble
- If in doubt, studying **convergence** with respect to grid size is always a good idea.




*Courtesy of
M. Thévenet
and S. Diederichs*

Resolution requirement: laser-driven acceleration

Same requirements as for beam-driven, but **in addition**, Δz needs to be **much smaller** than the **laser wavelength**.

e.g.
$$\Delta z \sim \frac{\lambda}{40}$$



Imposes that the grid be very fine in z !

The CFL limit constrains the timestep Δt .

Courant-Friedrichs-Lewy (CFL) limit

In order for the Yee discretization scheme to be **numerically stable**:


$$\Delta t < \frac{1}{c} \frac{1}{\sqrt{1/\Delta x^2 + 1/\Delta y^2 + 1/\Delta z^2}}$$

(Most other discretization schemes also have a CFL limit.)

For instance:

- For square cells ($\Delta x = \Delta y = \Delta z$): $\Delta t < \frac{\Delta z}{c\sqrt{3}}$
- If $\Delta z \ll \Delta x, \Delta y$ (e.g. laser-driven): $\Delta t < \frac{\Delta z}{c}$

Computational cost rapidly increases with grid resolution.



The number of computational operations scales like:

$$N_{comp} \propto \underbrace{\left(\frac{L_x}{\Delta x} \right) \times \left(\frac{L_y}{\Delta y} \right) \times \left(\frac{L_z}{\Delta z} \right)}_{\text{Number of grid points}} \times \underbrace{\left(\frac{T_{interaction}}{\Delta t} \right)}_{\text{Number of timesteps (i.e. number of iterations of the PIC loop)}}$$

Number of grid points


Number of timesteps
(i.e. number of iterations
of the PIC loop)

$T_{interaction}$ = physical duration of the simulated phenomenon
(e.g. time it takes for the driver to propagate through the whole plasma)


Reducing $\Delta x, \Delta y, \Delta z$ by a factor 2 increases the computational cost by a factor 16!

One consequence of using low resolution: spurious numerical dispersion

Yee scheme with low
resolution (unphysical)





Expected result




If Δz is not **small enough** compared to the **laser wavelength**, the **group velocity** of the simulated laser pulse may be erroneous. (“spurious numerical dispersion”)

(Picking a different discretization scheme than the Yee scheme can mitigate spurious numerical dispersion.)

Discretizing the field and particle equations



The gathered field is a weighted sum of the field on the grid.



The field felt by a macroparticle is a **weighted sum** of the field at **neighboring grid points**:

$$\mathbf{E}(\mathbf{x}_p) = \sum_{i,j,k \in \mathcal{V}(\mathbf{x}_p)} S(\mathbf{x}_p - \mathbf{x}_{i,j,k}) \mathbf{E}_{i,j,k}$$


Diagram illustrating the equation for the gathered field $\mathbf{E}(\mathbf{x}_p)$ at macroparticle position $\mathbf{x}_p$:

- $\mathbf{E}(\mathbf{x}_p)$: Macroparticle position
- $\sum_{i,j,k \in \mathcal{V}(\mathbf{x}_p)}$: Sum over grid points in the vicinity of the macroparticle
- $S(\mathbf{x}_p - \mathbf{x}_{i,j,k})$: "Shape factor" (Grid point position)
- $\mathbf{E}_{i,j,k}$: Field value at grid point


Different shape factors are commonly used.

$$\mathbf{E}(\mathbf{x}_p) = \sum_{i,j,k \in \mathcal{V}(\mathbf{x}_p)} S(\mathbf{x}_p - \mathbf{x}_{i,j,k}) \mathbf{E}_{i,j,k}$$


$$S(\mathbf{x}_p - \mathbf{x}_{i,j,k}) = \hat{S}\left(\frac{x_p - x_i}{\Delta x}\right) \hat{S}\left(\frac{y_p - y_i}{\Delta y}\right) \hat{S}\left(\frac{z_p - z_i}{\Delta z}\right)$$




- **Linear shape factor** (a.k.a. “1st order”, “CIC”) involves **2 grid points** along each direction



- **Quadratic shape factor** (a.k.a. “2nd order”, “TSC”) involves **3 grid points** along each direction




- **Cubic shape factor** (a.k.a. “3rd order”, “PQS”) involves **4 grid points** along each direction



High-order particle shapes are computationally more **expensive** but often lead to **smoother** results.

Macroparticles distribute charge/current to neighboring points



Macroparticles distribute their **charge** to neighboring grid points:

$$\rho_{i,j,k} = \frac{\sum_{p \in \mathcal{V}(\mathbf{x}_{i,j,k})} S(\mathbf{x}_{i,j,k} - \mathbf{x}_p) q_p}{\Delta x \Delta y \Delta z}$$

Field value at grid point


Sum over macroparticles in the vicinity of the grid point

Cell volume

Total charge of the macroparticle

(takes into account weight)

Macroparticles distribute charge/current to neighboring points



Macroparticles distribute their **current** to neighboring grid points:

$$\mathbf{J}_{i,j,k} = \frac{\sum_{p \in \mathcal{V}(\mathbf{x}_{i,j,k})} S(\mathbf{x}_{i,j,k} - \mathbf{x}_p) q_p \mathbf{v}_p}{\Delta x \Delta y \Delta z}$$

Note: that are also more advanced deposition schemes for $\mathbf{J}$ (e.g. Esirkepov scheme), that automatically ensure

$$\partial_t \rho + \nabla \cdot \mathbf{J} = 0$$

Outline

- Fundamental equations for plasma accelerators
- The electromagnetic Particle-In-Cell (PIC) algorithm
- Discretizing the field and particle equations
- **Parallelization of PIC codes**

Full PIC simulations can rapidly be very computationally expensive.

$$N_{comp} \propto \underbrace{\left(\frac{L_x}{\Delta x}\right) \times \left(\frac{L_y}{\Delta y}\right) \times \left(\frac{L_z}{\Delta z}\right)}_{\text{Number of grid points}} \times \underbrace{\left(\frac{T_{interaction}}{\Delta t}\right)}_{\text{Number of timesteps (i.e. number of iterations of the PIC loop)}}$$

Example: typical 3D laser-wakefield simulation


3D grid with 200 x 200 x 2500 grid points

140,000 timesteps (CFL limit!)

⇒ ~60,000 hours on 1 core (7 years!)

We need either **faster cores**
or **more cores in parallel**.


Single-core performance is saturating nowadays.



Nowadays, individual cores do not get faster.
We need to use **many cores in parallel**.



Modern supercomputers are organized in nodes, with multiples cores.

Computational work is parallelized across multiple **cores** and across multiples **nodes**.



↔ "Fast" network (~10 Gb/s)

CPU: tens of cores per node




GPUs: thousands of (slow) cores per node



Advantages of GPUs: high memory bandwidth and energy efficiency

Memory bandwidth of different CPU/GPU models


(= how fast can the compute cores fetch data in memory)



Data from github.com/karlrupp/cpu-gpu-mic-comparison


Power efficiency of different supercomputers

(= how much computation per unit of energy consumption)





Data from [top500.org](https://www.top500.org)

Domain-decomposition: groups of cores handle a fixed chunk of space




- Groups of cores (e.g. part of a CPU, one full GPU) work together on the same chunk of space (memory is shared among this group)
- Information is exchanged between groups of core over the network (using MPI)

Domain-decomposition: groups of cores handle a fixed chunk of space



Load imbalance can significantly slow down the simulation.




Load imbalance:

Some groups of cores have many more macroparticles (i.e. more work) than others.



The simulation will always progress at the pace of the slowest group.

Dynamic load-balancing changes the shape of subdomains in order to even out the work load.



¹J. Derouillat *et al*, *Comput. Phys. Comm.* **222**, 351-373 (2018).

²J.-L. Vay *et al*, *Proc. AAC* 2018.




Conclusion

- **Full electromagnetic PIC codes** model the **Maxwell equations**, along with **the equations of motion** for plasma and beam particles
- **Very generic**, but the **computational cost** can be very high. (esp. for laser-driven acceleration)
- Therefore, simulations often **require massively parallel** computing architectures.
- Examples of common full electromagnetic PIC codes in our community:

EPOCH OSIRIS Smilei PConGPU WarpX

Quick announcement: practical exercises

Monday, 6.2.	Tuesday, 7.2.	Wednesday, 8.2.	Thursday, 9.2.	Friday, 10.2.
Breakfast				
Laser-driven plasma acceleration I	Diagnostic techniques I	Laser-driven plasma acceleration II	Diagnostic techniques II	Plasma accelerator applications I
Discussion	Discussion	Discussion	Discussion	Discussion
Coffee & Tea				
Beam-driven plasma acceleration I	Modeling plasma accelerators I	Beam-driven plasma acceleration II	Modeling plasma accelerators II	Plasma accelerator applications II
Discussion	Discussion	Discussion	Discussion	Discussion
Lunch				
High-intensity lasers	Artificial intelligence and controlling acc.	Excursion	High-average power lasers	 Departure
Discussion	Discussion		Discussion	
Coffee & Tea			Coffee & Tea	
Poster session I	Poster session II		Practical exercises	
Dinner				
	Special evening talk			

Please make sure you bring your **laptop**, and have an **active Google account**.


We will be using Google Colab. (no need to install anything)



Thank you for your attention

There are multiple open post-doctoral positions at the BELLA Center (theory & experiments).
If interested, please visit jobs.lbl.gov and search for “BELLA”.

Staggering in time



Staggering in space

