

Array-of-Struct particles for iPic3D on MIC

Alec Johnson and Giovanni Lapenta

Centre for mathematical Plasma Astrophysics
Mathematics Department
KU Leuven, Belgium

EASC2014
Stockholm, Sweden
April 3, 2014

Abstract: We are porting iPic3D to the MIC for particle processing. iPic3D advances both the electromagnetic field and the particles implicitly, requiring typically 100-200 iterations of the field advance and 3-5 iterations of the particle advance for each cycle. We use particle subcycling to limit particle motion to one cell per cycle, which improves accuracy and simplifies sorting. To accelerate sorting, we represent particles in AoS format in double precision so that particle data exactly fits the cache line width. To vectorize particle calculations, we process particles in blocks: a fast 8x8 matrix transpose implemented in intrinsics converts each 8-particle block between SoA and AoS representation.



Goal: efficiently converged multiscale simulation of plasma

Tool: iPic3D, an implicit particle-in-cell code

Task: port to Xeon + Phi (MIC):

- improve MPI
- use OMP threads
- vectorize

Key issue: data layout of particles

- Ordering:
 - SoA for vectorization (push and sum)
 - AoS for localization (sorting)
- Granularity of particles:
 - grouped by cell: vectorization efficiency
 - grouped by thread subdomain: cache efficiency

The purpose of this presentation is to justify four algorithm choices:

Two fundamental determinations:

- 1 **Subcycle particles:**
 - for each particle, break time step into substeps
 - move the particle at most one cell per substep
 - motivation: *accurate simulation of fast-moving particles*
 - benefit: simpler sorting
- 2 Use **Array-of-Structs (AoS)** for particles.
 - motivation: fast sorting
 - can still vectorize via fast transpose/intrinsics

Two secondary determinations:

- 1 Use **double precision** for particles.
 - Vlasov solver via resampling.
 - no mixed precision.
 - particle exactly fits cache line.
- 2 Use **AoS field to push particles.**
 - motivation: better localization of field data access
 - justification: one transpose per cycle is justified by numerous particle iterations and amortized by many iterations of SoA field solver.

1 iPic3D algorithm

2 Algorithm choices

iPic3D simulates charged particles interacting with the electromagnetic field. It solves the following equations:

Fields:

$$\partial_t \mathbf{B}(\mathbf{x}) + \nabla \times \bar{\mathbf{E}}(\mathbf{x}) = 0$$

$$\partial_t \mathbf{E}(\mathbf{x}) - c^2 \nabla \times \bar{\mathbf{B}}(\mathbf{x}) = -\bar{\mathbf{J}}(\mathbf{x})/\epsilon_0,$$

Particles:

$$\partial_t \mathbf{v}_p = \frac{q_p}{m_p} \left(\bar{\mathbf{E}}'(\bar{\mathbf{x}}'_p) + \bar{\mathbf{v}}_p \times \bar{\mathbf{B}}'(\bar{\mathbf{x}}'_p) \right),$$

$$\partial_t \mathbf{x}_p = \bar{\mathbf{v}}_p,$$

Moments (10):

$$(1) \quad \sigma(\mathbf{x}) := \sum_p S(\mathbf{x} - \mathbf{x}_p) q_p$$

$$(3) \quad \mathbf{J}(\mathbf{x}) := \sum_p S(\mathbf{x} - \mathbf{x}_p) q_p \mathbf{v}_p$$

$$(6) \quad \mathcal{P}(\mathbf{x}) := \sum_p S(\mathbf{x} - \mathbf{x}_p) q_p \mathbf{v}_p \mathbf{v}_p$$

The Implicit Moment Method uses these 10 moments (with \mathbf{E} and \mathbf{B}) to estimate $\bar{\mathbf{J}}$.

Discretization:

$$\bullet \quad \partial_t X := \frac{X^{n+1} - X^n}{\Delta t}.$$

$$\bullet \quad \bar{X} = \frac{1}{2} X^{n+1} + \frac{1}{2} X^n.$$

Current Evolution

$$\partial_t \mathbf{J}_s + \nabla \cdot \mathcal{P}_s = \frac{q_s}{m_s} \left(\sigma'_s \bar{\mathbf{E}} + \bar{\mathbf{J}}_s \times \mathbf{B}' \right)$$

Average current responds linearly to electric field:

$$\bar{\mathbf{J}} = \hat{\mathbf{J}} + \mathbb{A} \cdot \bar{\mathbf{E}},$$

where:

$$\left[\begin{array}{l} \hat{\mathbf{J}} := \sum_s \hat{\mathbf{J}}_s, \\ \hat{\mathbf{J}}_s := \Pi_s \cdot \left(\mathbf{J}_s^n - \frac{\Delta t}{2} \nabla \cdot \mathcal{P}_s \right), \\ \mathbb{A} := \sum_s \beta_s \sigma'_s \Pi_s, \\ \Pi_s := \frac{\mathbb{I} - \tilde{\mathbf{B}}_s \times \mathbb{I} + \tilde{\mathbf{B}}_s \tilde{\mathbf{B}}_s}{1 + |\tilde{\mathbf{B}}_s|^2}, \\ \tilde{\mathbf{B}}_s := \beta_s \mathbf{B}', \\ \beta_s := \frac{q_s \Delta t}{2 m_s}. \end{array} \right]$$

Implicit Particle Advance

$$\bar{\mathbf{v}}_p = \left(\frac{\mathbb{I} - \tilde{\mathbf{B}}_p \times \mathbb{I} + \tilde{\mathbf{B}}_p \tilde{\mathbf{B}}_p}{1 + |\tilde{\mathbf{B}}_p|^2} \right) \cdot \left(\mathbf{v}_p^n + \beta_s \bar{\mathbf{E}}'(\bar{\mathbf{x}}_p) \right), \text{ where}$$

$$\tilde{\mathbf{B}}_p := \frac{q_p \Delta t}{2 m_p} \bar{\mathbf{B}}'(\bar{\mathbf{x}}_p), \text{ and}$$

$$\bar{\mathbf{x}}_p = \mathbf{x}_p^0 + \Delta t \bar{\mathbf{v}}_p.$$

iPic3D cycles through three tasks:

- 1 `fields.advance(moments)`
- 2 `particles[s].move(fields)`
- 3 `moments[s].sum(particles[s])`

Moving particles consists of pushing and sorting, e.g.:

```
foreach subcycle c:  
  foreach particle:  
    particle.push(field(cell(particle)))  
    particle.sort()  
  particles.communicate()
```

- 1 iPic3D algorithm
- 2 Algorithm choices

Balance two goals:

- 1 flexibility where architectures/algorithms differ
- 2 best particulars where architectures/algorithms agree

Architecture key attributes:

- 1 Width of **cache line**: **8 doubles** = 512 bits (fairly universal)
- 2 Width of **vector unit**:
 - **8 doubles** = 512 bits for MIC
 - 4 doubles = 256 bits for Xeon with AVX
 - 2 doubles = 128 bits for SSE2

Data of **algorithm**:

- 1 fields: 6 doubles (two vectors) per mesh cell:
 - 1 B_x magnetic field
 - 2 B_y magnetic field
 - 3 B_z magnetic field
 - 4 ψ (\mathbf{B} correction potential)
 - 5 E_x electric field
 - 6 E_y electric field
 - 7 E_z electric field
 - 8 ϕ (\mathbf{E} correction potential)
- 2 100s of particles per mesh cell; 8 doubles (2 vectors + 2 scalars) per particle:
 - 1 u velocity
 - 2 v velocity
 - 3 w velocity
 - 4 q charge (or particle ID)
 - 5 x position
 - 6 y position
 - 7 z position
 - 8 t subcycle time

(1) Why subcycle?

Traditionally the implicit moment method moves all particles with the same time step.

We are implementing subcycling:

- For each particle, the global time step is partitioned into substeps.
- Substeps stop particles at cell boundaries.

Benefits of subcycling:

- 1 Simplifies sorting:
 - SoA vectorization requires sorting particles by mesh cell.
 - Subcycling guarantees that particles move only one mesh cell per subcycle.
 - Without subcycling, particles can move arbitrarily far between sorts.
 - Without subcycling, particles must be sorted with every iteration of the implicit mover.
 - Without subcycling, sorted particle data must include average position data and no longer fits in a single cache line.
- 2 Subcycling is needed to resolve fast particles accurately.
 - Maxwell's equations need time-averaged current.
 - Subcycling is needed to get correct time-averaged current of fast particles.

(2) AoS particle vectorization

Usually SoA is preferred for vectorization.

But AoS particles can still be vectorized in one of two ways:

Fast matrix transpose

- **8-component particles**

(subcycled case):

- Represent as AoS
- Process in 8-particle blocks
- Convert blocks to/from SoA using fast 8x8 matrix blocked transpose (28-36 8-wide vector instructions)

- **12-component particles**

(non-subcycled case):

- Consider padding extra components to 8 (faster sort); otherwise:
- first 8 components handled like 8-component particles
- last 4 components handled like 4-component particles using fast 4x8 \leftrightarrow 8x4 transpose (16 8-wide vector instructions).

Physical vectors (intrinsic-heavy):

- **MIC:**

- process 2 particles at a time
- concatenate velocity vectors:
[u1, v1, w2, q1, u2, v2, w2, q2]
- concatenate position vectors:
[x1, y1, z1, t1, x2, y2, z2, t2]
- Use physical vector operations (use swizzle for cross-product)

- **Xeon**

- process 1 particle at a time (or 2 at a time for single precision)

Pusher times in iPic3D:

time	pusher
=====	=====
0.102 sec	SoA (but also need to sort each iteration)
0.202 sec	AoS_intr (no sort required, but helps cache)
0.259 sec	AoS (no sort required, but helps cache)

Pusher times for a single iteration:

time	pusher
=====	=====
.07 Mcycles	SoA
.13 Mcycles	AoS_tran (8-pcl blocks via fast 8x8 transpose)
.21 Mcycles	AoS_intr (2-pcl blocks with intrinsics mover)

Pusher times for 4 iterations stopping at cell boundary:

time	pusher
=====	=====
.36 Mcycles	SoA
.40 Mcycles	AoS_tran (8-pcl blocks via fast 8x8 transpose)
[unimplemented]	AoS_intr (no need to sort with each subcycle)

Cache-line-sized particles facilitate sorting:

- can transfer particles directly to memory destination with no-read writes
- no cache contention
- vector unit divides cache line size, so fully utilized

Sort particles by:

- 1 process subdomain (for MPI),
- 2 thread subdomain, and
- 3 mesh cell (for vectorization)

To hide communication latency, overlap process-level communication with thread-level sorting.

- General sort:

```
send exiting particles
sort particles in process
wait on incoming particles
sort incoming particles
```

- Subcycle sort (moving ≤ 1 cell per subcycle):

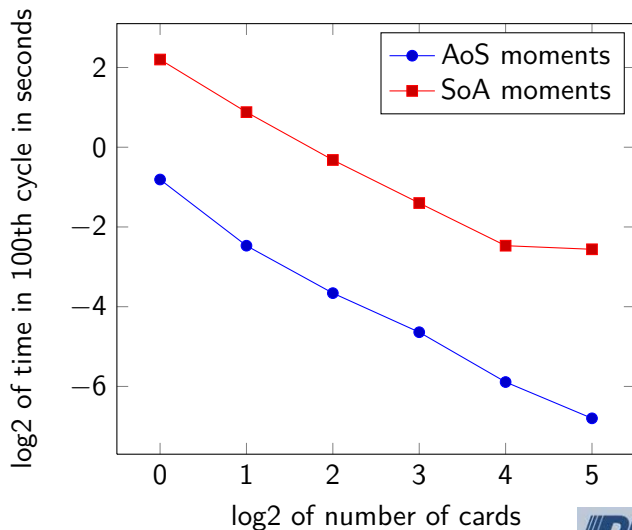
```
move particles in boundary cells
send particles in ghost cells
move particles in interior cells
move incoming particles
```

In the field solver we represent fields and moments in SoA format. This allows better vectorization of the implicit solver.

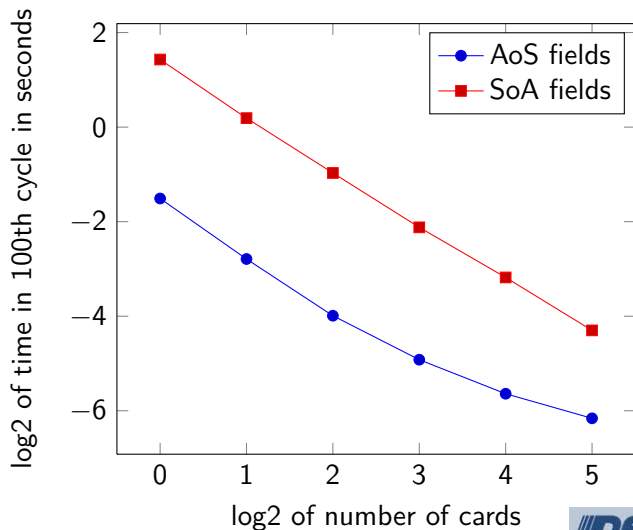
In the particle solver, we represent fields and moments with AoS format:

- AoS gives better localization of random access.
- SoA fields and moments offer no benefit to vectorization of particle processing.
- The transpose is done only once per cycle.

summing moments, computation, Xeon Phi [strong scaling]



moving particles, computation, Xeon Phi [strong scaling]



Trade-offs in SoA vs. AoS particles:

- SoA particles better for vectorization.
- SoA particles require sorting with each iteration or subcycle.
- AoS particles better for sorting.
- AoS allows infrequent sorting.
- fast 8x8 transpose converts between AoS and SoA.

Conclusions:

- Use AoS for basic particle representation.
- Convert to SoA in blocks when beneficial for vectorization.