

Efficient Strict-Binning Particle-in-Cell (PIC) Algorithm for Multi-Core SIMD Processors

Yann Barsamian^{1,2}, Arthur Charguéraud^{2,1}, Sever Hirstoaga^{3,1},
Michel Mehrenberger^{1,3}



1.  **Université** de Strasbourg
2. ICube, CNRS, INRIA Nancy
3. IRMA, CNRS, INRIA Nancy

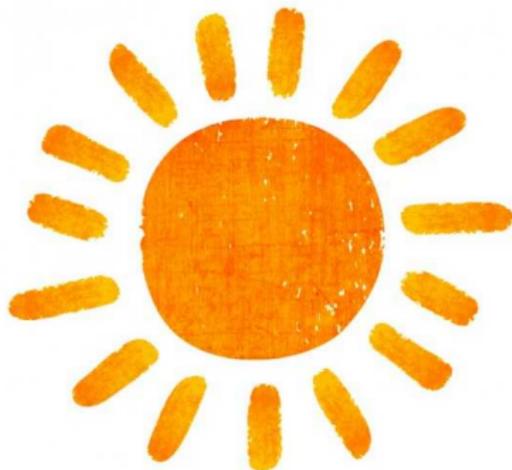


Euro-Par 2018, Torino (Italy)
August 2018



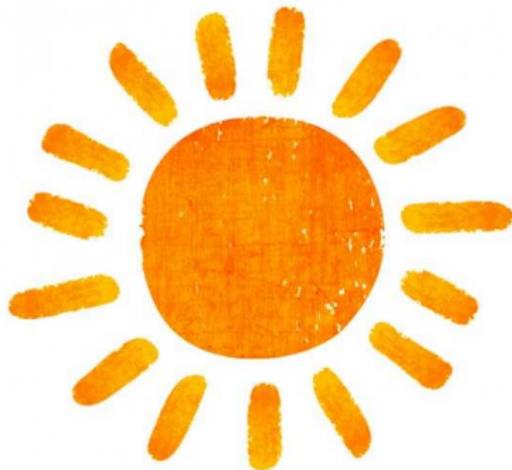


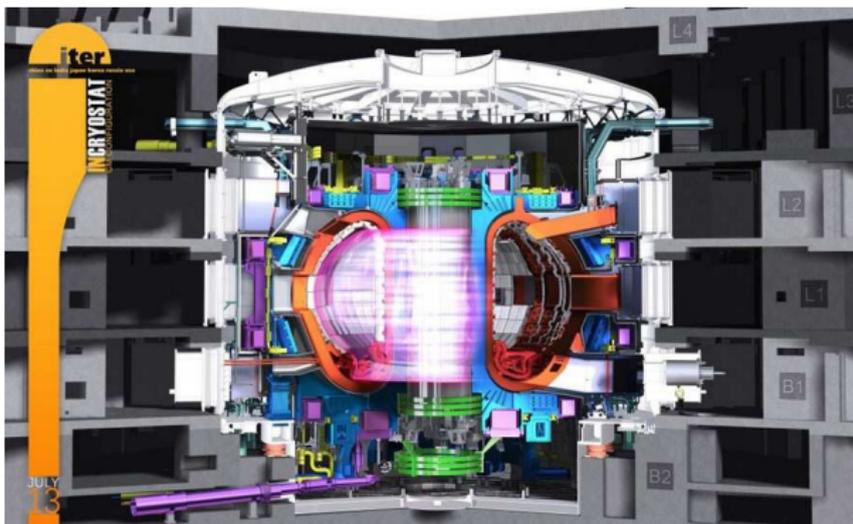
Step 1.





Step 2.





Step 3. ITER¹ tokamak²

(also applicable in other contexts, e.g., astrophysics, where we have to model different particles / planets / ... that interact)

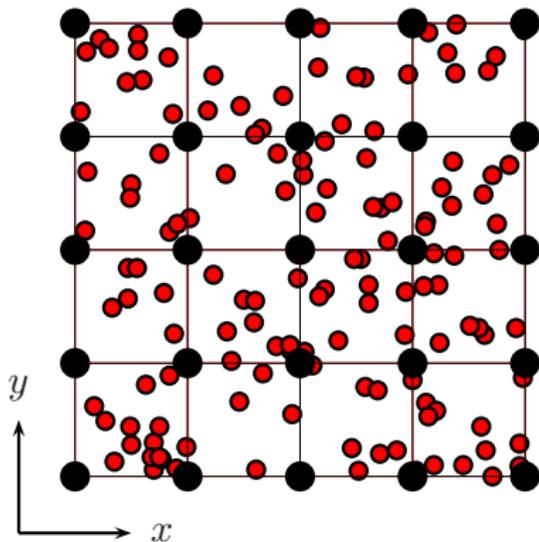
¹“The way” (in Latin) to produce energy (Cadarache, France)

²Токамак: тороидальная камера с магнитными катушками (toroidal chamber with magnetic coils)



$$\begin{cases} \frac{\partial f}{\partial t} + \vec{v} \cdot \nabla_{\vec{x}} f - \vec{E} \cdot \nabla_{\vec{v}} f = 0 & \text{Vlasov} \\ \nabla_{\vec{x}} \cdot \vec{E} = \rho = 1 - \int f d\vec{v} & \text{Poisson} \end{cases}$$

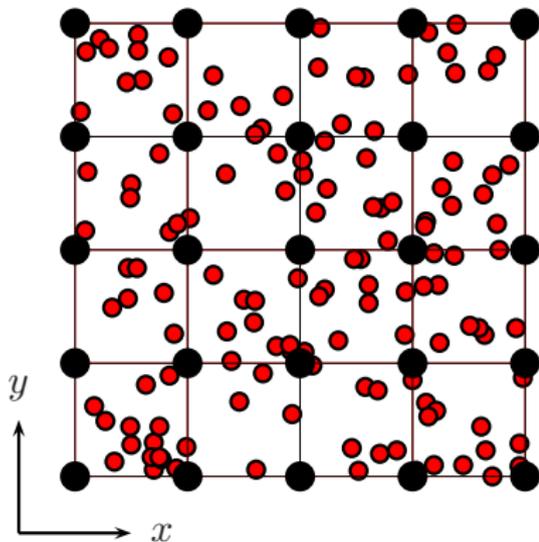
- Distribution function f : N numerical particles (red)
- Electric field \vec{E} and charge density ρ : 3d grids (black)





$$\begin{cases} \frac{\partial f}{\partial t} + \vec{v} \cdot \nabla_{\vec{x}} f - \vec{E} \cdot \nabla_{\vec{v}} f = 0 & \text{Vlasov} \\ \nabla_{\vec{x}} \cdot \vec{E} = \rho = 1 - \int f d\vec{v} & \text{Poisson} \end{cases}$$

- Distribution function f : N numerical particles (red)
- Electric field \vec{E} and charge density ρ : 3d grids (black)

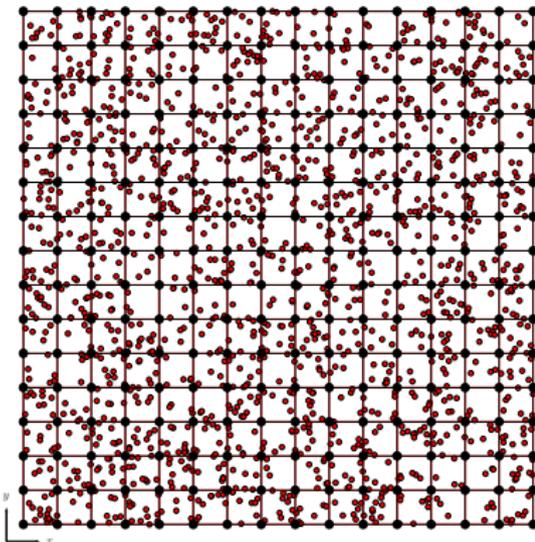


- Physical effects on small scale (+ large scale)
- Noise (numerical errors when N is small)
- Frequent particle motion



$$\begin{cases} \frac{\partial f}{\partial t} + \vec{v} \cdot \nabla_{\vec{x}} f - \vec{E} \cdot \nabla_{\vec{v}} f = 0 & \text{Vlasov} \\ \nabla_{\vec{x}} \cdot \vec{E} = \rho = 1 - \int f d\vec{v} & \text{Poisson} \end{cases}$$

- Distribution function f : N numerical particles (red)
- Electric field \vec{E} and charge density ρ : 3d grids (black)

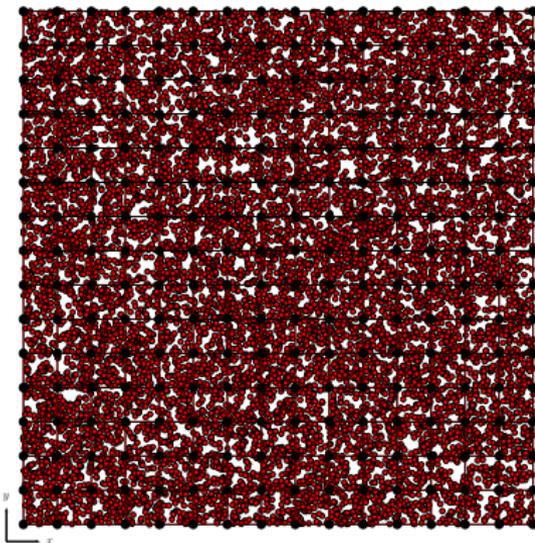


- Physical effects on small scale (+ large scale)
 \Rightarrow increase $ncx \times ncy \times ncz$
 $(1\,000 \times 1\,000 \times 1\,000)$
- Noise (numerical errors when N is small)
- Frequent particle motion



$$\begin{cases} \frac{\partial f}{\partial t} + \vec{v} \cdot \nabla_{\vec{x}} f - \vec{E} \cdot \nabla_{\vec{v}} f = 0 & \text{Vlasov} \\ \nabla_{\vec{x}} \cdot \vec{E} = \rho = 1 - \int f d\vec{v} & \text{Poisson} \end{cases}$$

- Distribution function f : N numerical particles (red)
- Electric field \vec{E} and charge density ρ : 3d grids (black)

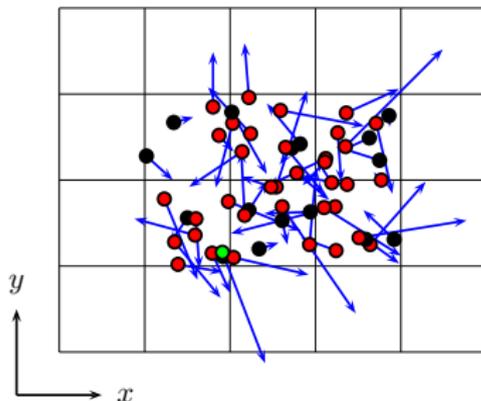


- Physical effects on small scale (+ large scale)
 \Rightarrow increase $ncx \times ncy \times ncz$
 (1 000 \times 1 000 \times 1 000)
- Noise (numerical errors when N is small)
 \Rightarrow increase $\frac{N}{ncx \times ncy \times ncz}$
 (10 000 to 1 000 000)
- Frequent particle motion



- Three levels of parallelism :
 - network (MPI, inter-node),
 - socket (OpenMP, intra-node),
 - instruction (SIMD),
- Maximization of the number of particles that can fit in memory,

- Maximization of the throughput of the simulation which is memory bound,
- Handling particles moving more than 2 cells per time step (“fast-moving particles”), without loss of performance,
- Comparison to other implementations.





Initialization:

- 1 Initialize N particles icell , $d\{x,y,z\}$, $v\{x,y,z\}$ of size $[N]$
- 2 Compute ρ and E rho , $E\{x,y,z\}$ of size $[\text{ncx}] [\text{ncy}] [\text{ncz}]$

Algorithm:

- 3 **Foreach** time iteration **do**
- 4 **If** (*condition*) **then**
- 5 Sort the particles³ $\mathcal{O}(N)$ counting sort
- 6 **End If**
- 7 Set all cells of ρ to 0
- 8 **Foreach** particle **do**
- 9 Update the velocity $v_+ = -E\Delta t$
- 10 Update the position $x_+ = v\Delta t$
- 11 Accumulate the charge on the nearest ρ cells
- 12 **End Foreach**
- 13 Compute E from ρ FFT Poisson solver
- 14 **End Foreach**

³Decyk, Karmesin, de Boer, & Liewer (1996)



Initialization:

- 1 Initialize N particles icell , $d\{x,y,z\}$, $v\{x,y,z\}$ of size $[N]$
- 2 Compute ρ and E rho , $E\{x,y,z\}$ of size $[\text{ncx}] [\text{ncy}] [\text{ncz}]$

Algorithm:

Execution time breakdown

- 3 **Foreach** time iteration **do**
- 4 **If** (*condition*) **then**
- 5 Sort the particles³ 10%⁴
- 6 **End If**
- 7 Set all cells of ρ to 0
- 8 **Foreach** particle **do**
- 9 Update the velocity 50%⁴
- 10 Update the position 25%⁴
- 11 Accumulate the charge on the nearest ρ cells 15%⁴
- 12 **End Foreach**
- 13 Compute E from ρ <1%⁴
- 14 **End Foreach**

³Decyk, Karmesin, de Boer, & Liewer (1996)

⁴Any difference in system hardware or software design or configuration may affect actual performance (-):

To sort or not to sort?



	Sort	Upd. v	Upd. x	Deposit	Total
Do not sort	0.0	98.0	64.6	35.9	199.0
Sort every 100	3.6	78.3	64.4	25.6	177.0
Always sort	209.0	66.3	64.2	13.4	353.0

Execution time (in s). Test case: 200 000 000 particles, 128×128 grid, $\Delta t = 0.1$, 500 iterations. Architecture: Intel Broadwell, 18 cores, 76.8 GB/s.

Periodic sorting: **better data locality**, and **shorter overall time**: find the best frequency⁵.

Sorting at each iteration⁶: **enhancement of the data locality & vectorization of the update velocities loop**, but **too costly**.

Efficient data structure to keep particles sorted⁷: **avoid the sorting step**.

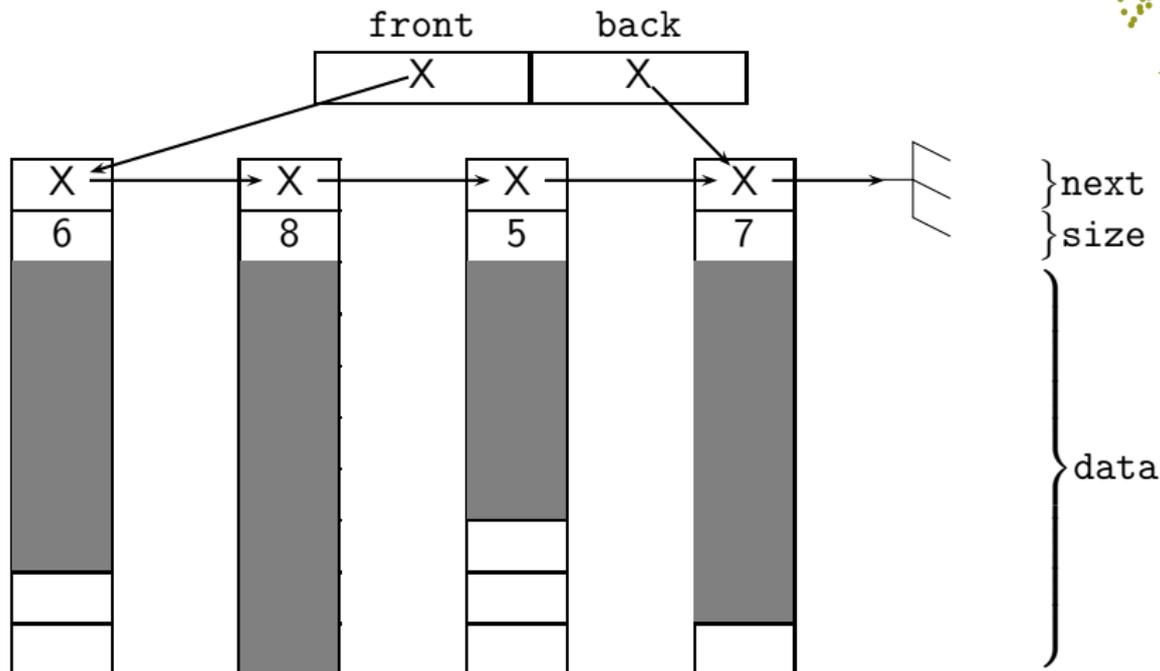


⁵Marin, Jin, & Mellor-Crummey (2008)

⁶Lanti, Tran, Jocksch, Hariri, Brunner, Gheller, & Villard (2016)

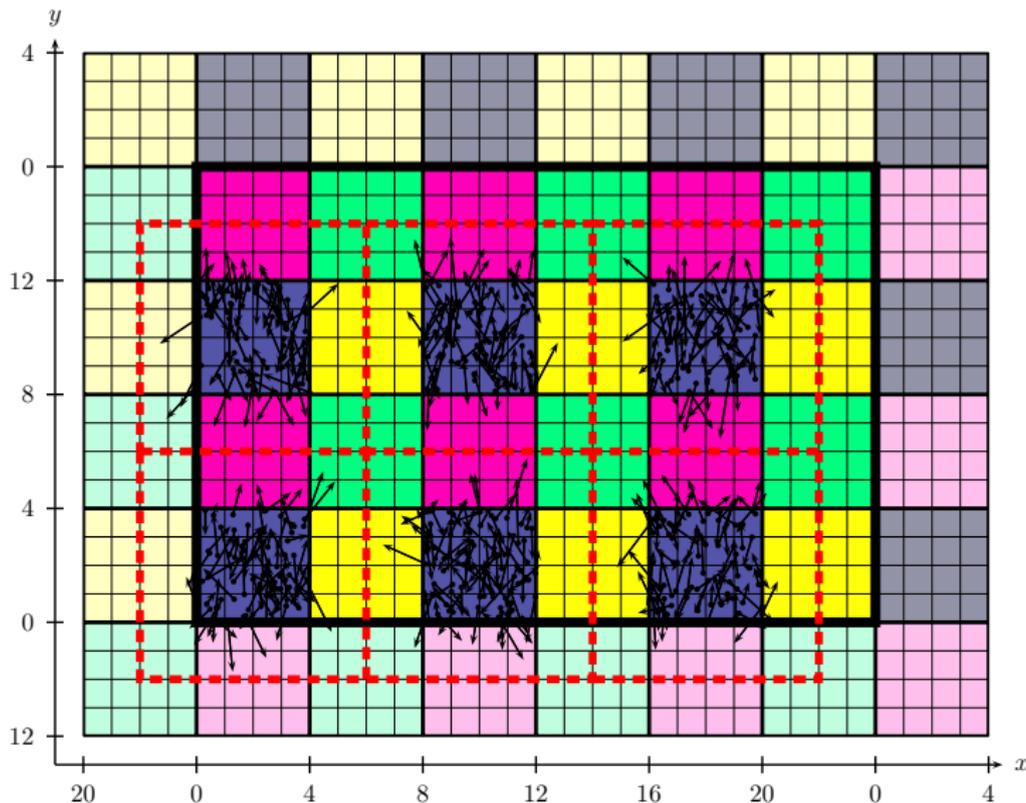
⁷Durand, Raffin, & Faure (2012); Nakashima, Summura, Kikura, & Miyake (2017); Barsamian, Charguéraud, & Ketterlin (2017)

Chunk Bags: Linked Lists of Fixed-Size Arrays



```
struct chunk { struct chunk* next; int size; // 0<=size<=K
               float dx[K], dy[K], dz[K];
               double vx[K], vy[K], vz[K]; } chunk;
struct { chunk* front, back; } bag;
```

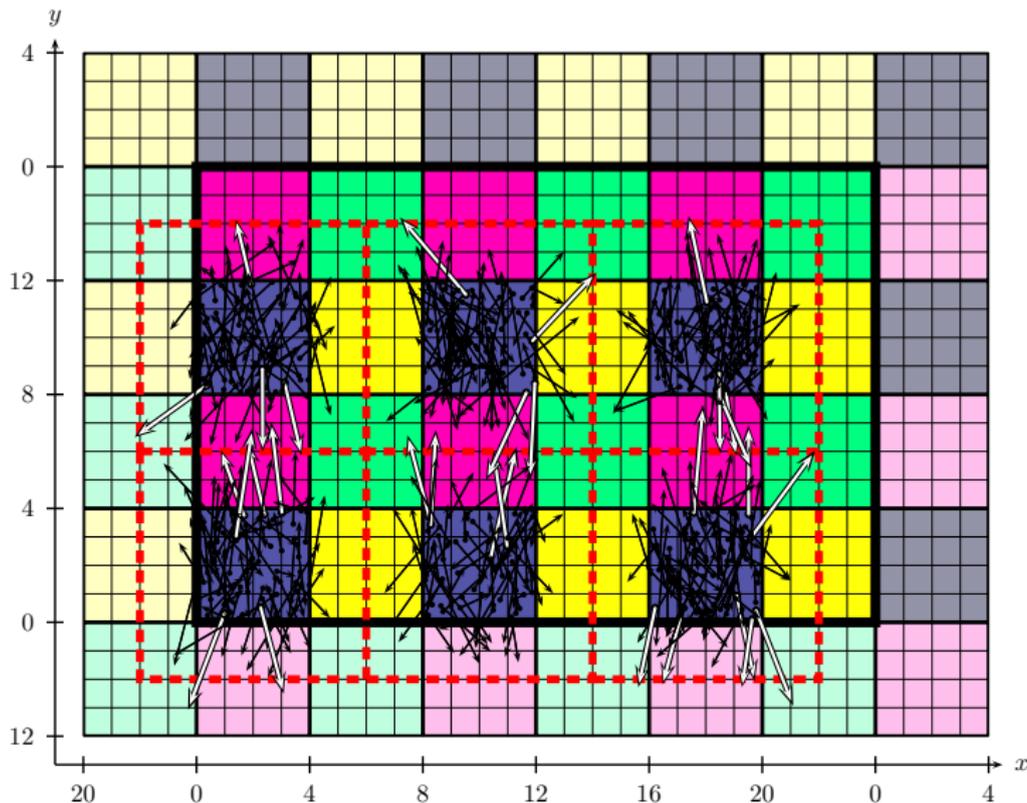
The Eight-Colors Algorithm⁸



8 phases to tame the number of data races when moving particles.

⁸Kong, Huang, Ren, & Decyk (2011)

The Eight-Colors Algorithm⁸

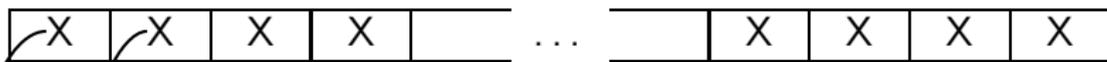


Particles moving more than half a tile away require special care.

⁸Kong, Huang, Ren, & Decyk (2011)



```
chunkbag particles[nbCells] // nbCells = ncx*ncy*ncz
```



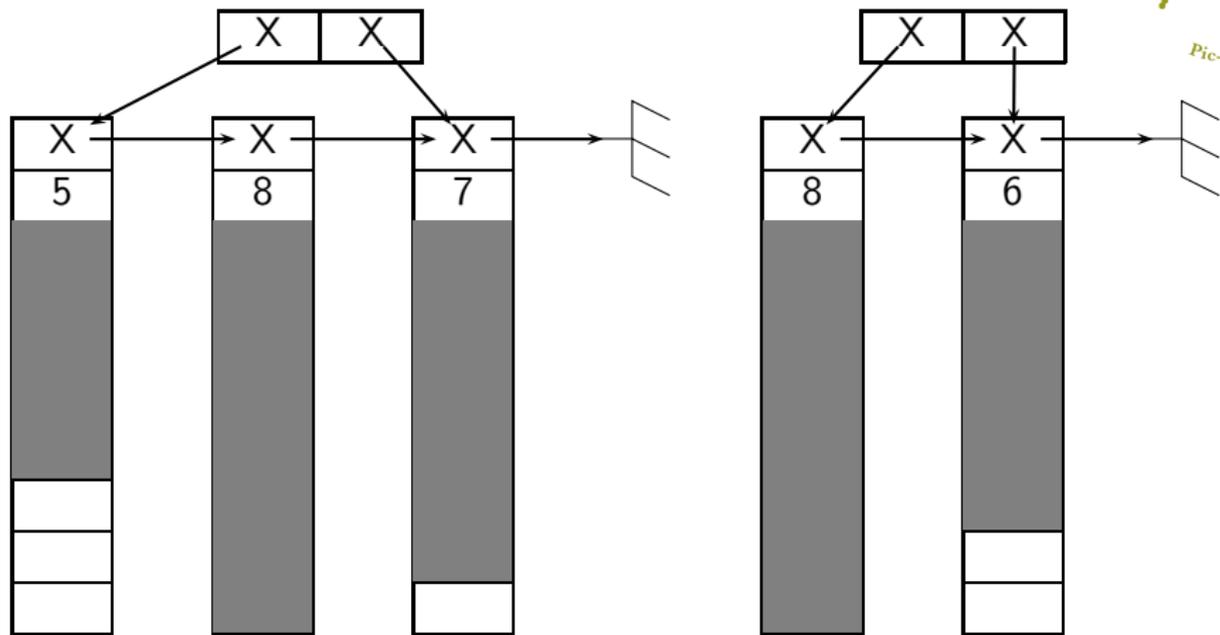
particles with cell identifier 1

particles with cell identifier 0

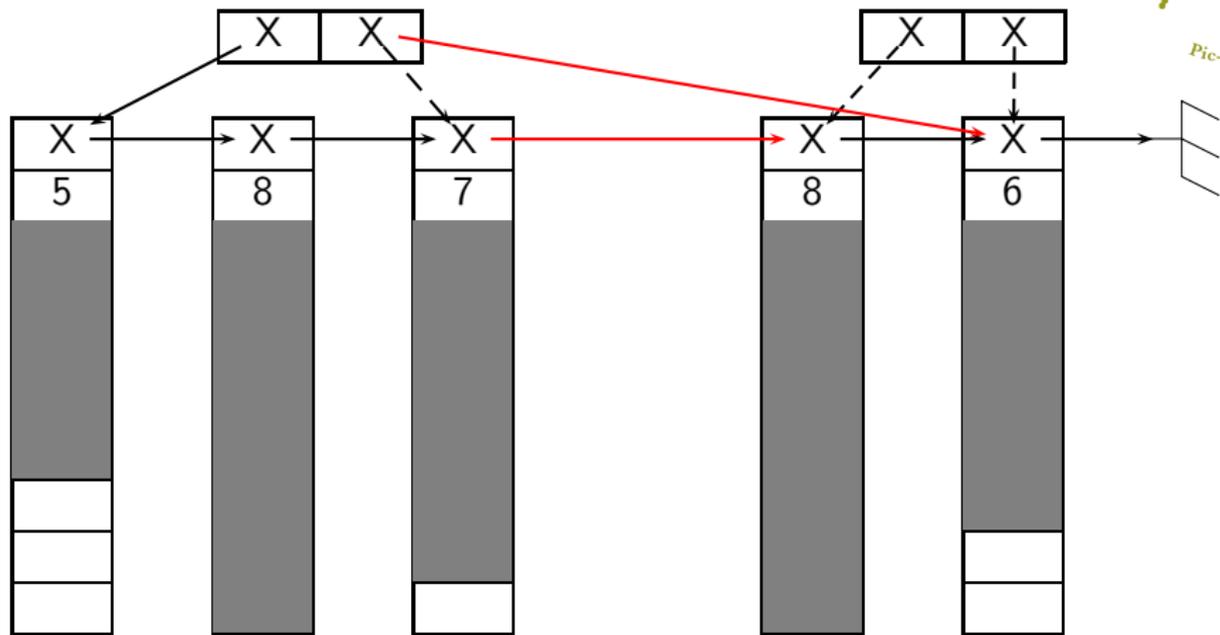
```
chunkbag particlesNextPrivate[nbCells],  
      particlesNextShared[nbCells]
```

- `particlesNextPrivate[i]` receives particles moving to a nearby cell i : no atomic operation required.
- `particlesNextShared[i]` receives particles moving to a remote cell i : atomic push used.
- `particles[i]` at the next time step is obtained by merging the two.

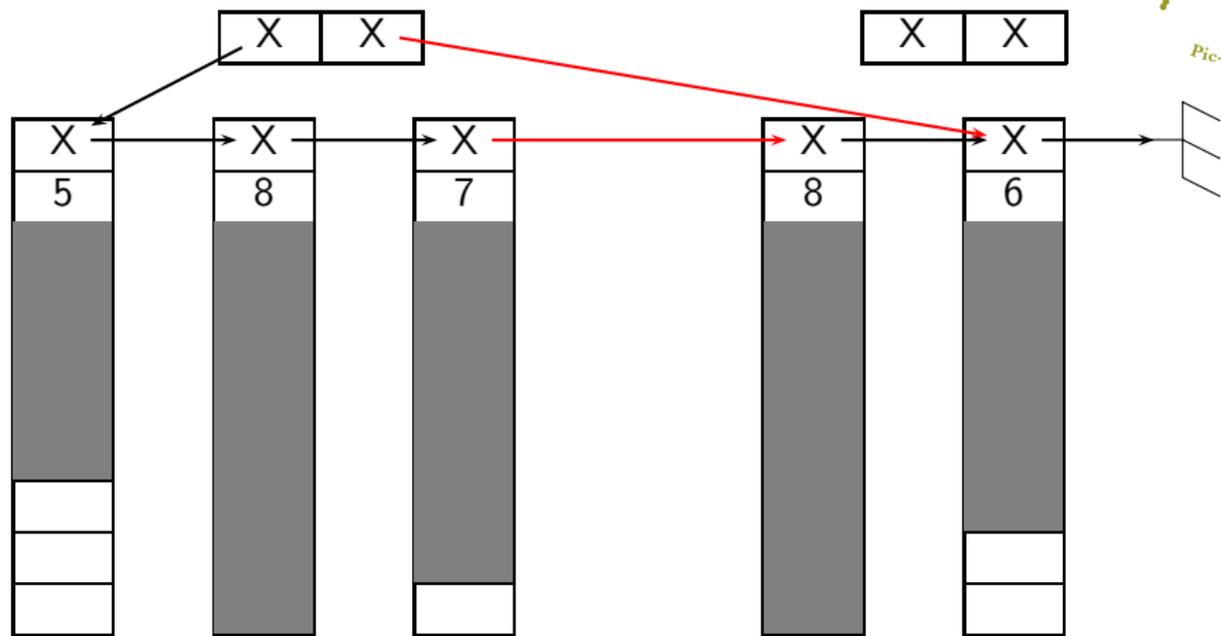
Chunk Bags: Merge Operation



Chunk Bags: Merge Operation



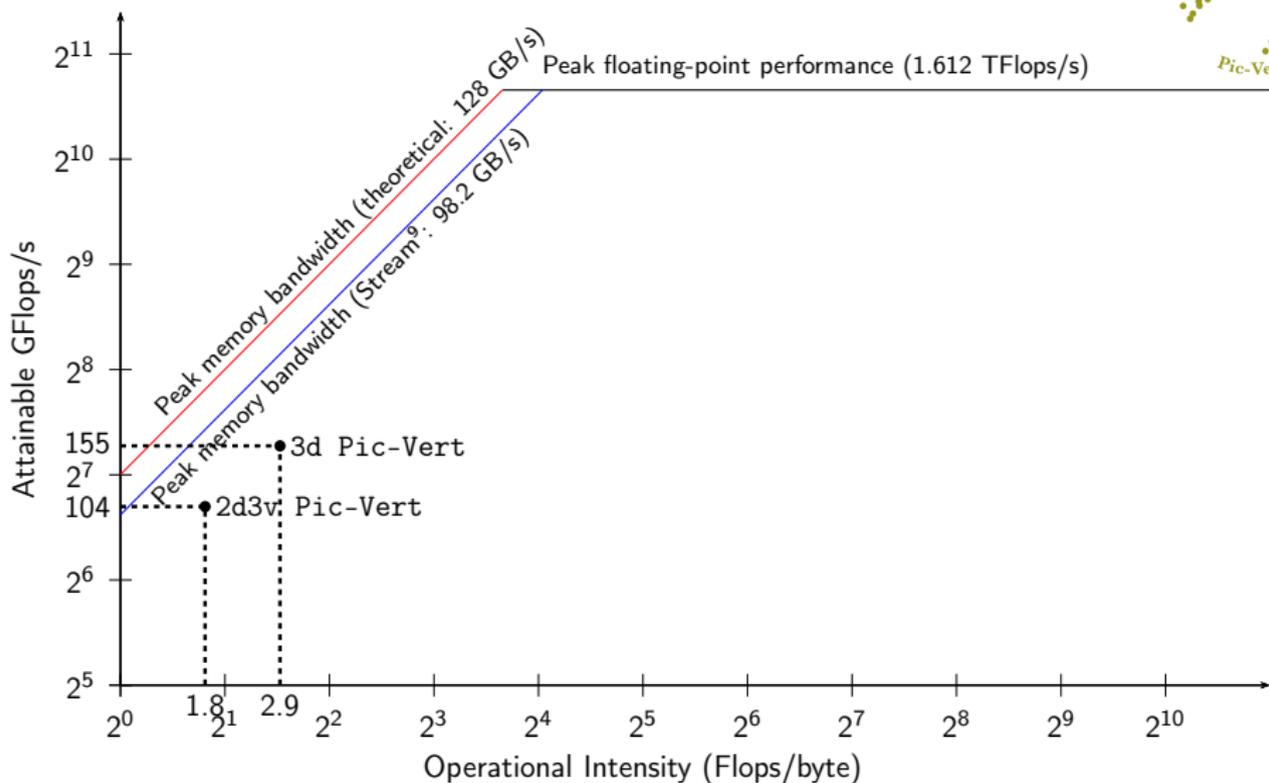
Chunk Bags: Merge Operation



Upper bound on the number of chunks: $\lceil N/K \rceil + 4 \cdot \text{nbCells}$.

All chunks allocated at initialization (no dynamic `malloc/free`).

Roofline Model¹⁰ on Intel Skylake (24 Cores)



⁹McCalpin (1995) - Code v5.10 (2013)

¹⁰Williams, Waterman, & Patterson (2009)



Experiments with different particle velocities:

- up to 4.4% of “fast-moving particles” (more than 2 cells away),
- up to 3.7% of possible conflicts¹¹,
- if processed sequentially: 85% slowdown on 24 cores¹²,
- when processed with our shared bags: only 7.0% slowdown.

¹¹Not all fast-moving particles go out of the “extended tile” — consider a particle on the far left of the tile moving 3 cells to the right.

¹²Let t denote the single-core execution time. Assume 3.7% of sequential execution, and 96.3% using 24 cores. The parallel execution time is:

$$0.037t + 0.963t/24 = 1.85t/24.$$

Comparison of Pic-Vert to Other Implementations



Different implementations on different architectures: cores, memory bandwidth in GB/s, number of particles processed by second (absolute and normalized w.r.t. memory bandwidth).
Top: CPUs. Bottom: accelerators (GPUs, MICs).

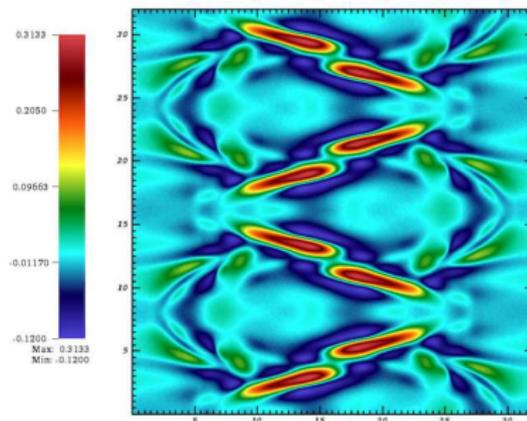
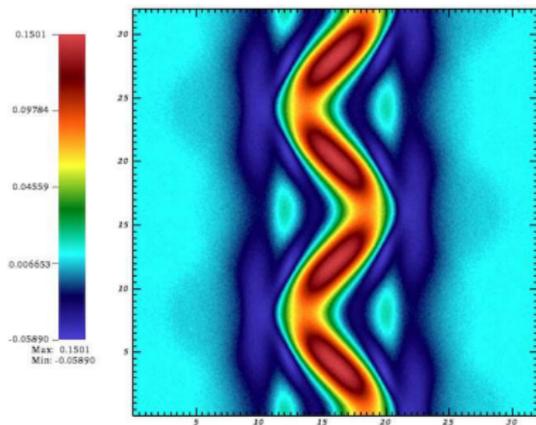
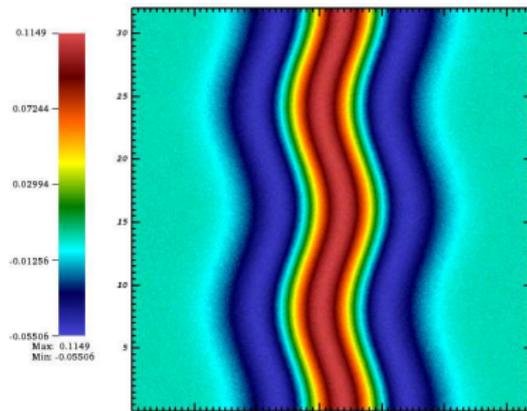
Implem.	Architecture	Cores	M.B.	Part./s	Norm.
VPIC	IBM PowerXCell 8i	9	204.8	$173 \cdot 10^6$	0.85
OSIRIS	Intel Xeon E5-2680	8	51.2	$134 \cdot 10^6$	2.62
ORB5	Intel Xeon E5-2670	8	51.2	$69 \cdot 10^6$	1.35
PICADOR	Intel Xeon E5-2697 v3	14	68	$127 \cdot 10^6$	1.87
GTC-P	Intel Xeon E5 2692 v2	12	59.7	$100 \cdot 10^6$	1.68
PIConGPU	Intel Xeon E5-2698 v3	16	68	$111 \cdot 10^6$	1.63
Pic-Vert	Intel Xeon Platinum 8160	24	128	$740 \cdot 10^6$	5.78
Pic-Vert	Intel Xeon E5-2690 v3	12	68	$374 \cdot 10^6$	5.49
PIConGPU	NVIDIA Tesla GK210	2496	480	$336 \cdot 10^6$	0.70
ORB5	NVIDIA Tesla K20X	2688	250.0	$177 \cdot 10^6$	0.71
PICADOR	Intel Xeon Phi 7250 (KNL)	68	115.2	$298 \cdot 10^6$	2.59
EMSES	Intel Xeon Phi 7250 (KNL)	68	115.2	$1300 \cdot 10^6$	11.3

Validation: 2d3v Electron Hole Test Case¹³



- 64 billion particles,
- grid of size 512×512 ,
- time step 0.1,
- spatial domain $[0, 32]^2$,
- magnetic field 0.2.

Snapshots of ρ at $t=0$ (top right), 20 (bottom left), and 40 (bottom right).



¹²Muschiatti, Roth, Carlson, & Ergun (2000)



- Contributions
 - Particles sorted at all time with low memory overhead (4 · nbCells extra chunks, lowest in the state-of-the-art)
 - Optimal memory bandwidth usage: each particle is loaded from/written to main memory only once per iteration
 - Full advantage of SIMD
 - Efficient handling of fast particles
- Results on Intel Skylake, 24 cores, 128 GB/s
 - 740 million particles / second in 3d
 - 55% of the maximum bandwidth
- Comparison to state-of-the-art implementations
- Future outlooks
 - Solve the Maxwell equations
 - Test Pic-Vert on Many Integrated Core (MIC) architecture
 - Investigate the use of chunks in domain decomposition



That's all Folks!

ybarsamian@unistra.fr