

The LIDYL logo features the word 'LIDYL' in a white, bold, sans-serif font. To the right of the text is a stylized graphic of a laser beam or particle path, composed of several overlapping curved lines in shades of yellow and orange, with a small orange dot at the center of the path.


LIDYL

Laboratoire Interactions, Dynamiques et Lasers
EMR9000 CEA, CNRS, Université Paris-Saclay

**Modeling a novel laser-driven electron accelerator concept:
Particle-In-Cell simulations at the exascale**

Neil ZAIM

March 6th, 2023




O. Gobert



T. Ceccotti



A. Panchal



A. Ammar



I. Kara-Mostefa



P. Forestier-Colleoni



S. Dobosz



P. Bartoli



H. Vincenti

(head of numerical division)



P. Martin

Theory/simulations



L. Fedeli



T. Clark




N. Zaim




A. Sainte-Marie

Outline




Laser-plasma interaction & the Particle-In-Cell method



WarpX: a Particle-In-Cell code for the exascale era




WarpX performances




Simulation of a new electron accelerator concept

Outline




Laser-plasma interaction & the Particle-In-Cell method



WarpX: a Particle-In-Cell code for the exascale era




WarpX performances



Simulation of a new electron accelerator concept



High-power Ti:Sapphire femtosecond lasers are the most intense light-sources available on Earth

Ultra-short **20-30 fs** “bullets” of light focused to extreme intensities



High-power Ti:Sapphire femtosecond lasers are the most intense light-sources available on Earth



Ultra-short **20-30 fs** “bullets” of light focused to extreme intensities



From **few TeraWatt** up to **several PetaWatt**
(few PW is power of the sunlight striking Australia!!)

High-power Ti:Sapphire femtosecond lasers are the most intense light-sources available on Earth

Ultra-short **20-30 fs** “bullets” of light focused to extreme intensities

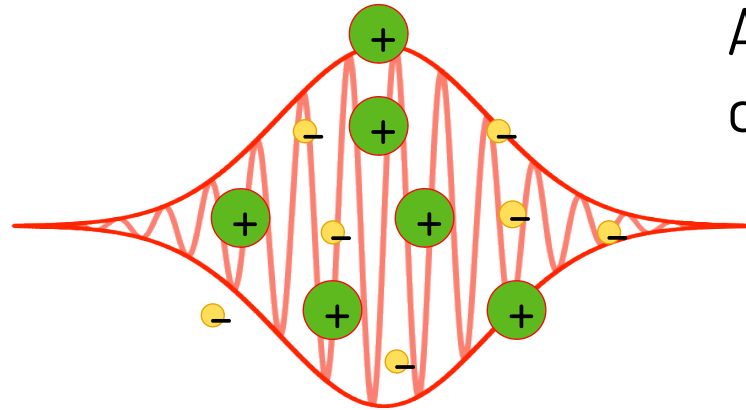


Intensity from **10^{18} W/cm^2** up to **10^{23} W/cm^2**

From **few TeraWatt** up to **several PetaWatt**
(few PW is power of the sunlight striking Australia!!)



Focused down to a **few μm spot**
(wavelength is $\sim 800 \text{ nm}$)

Any target irradiated with these lasers
becomes a plasma almost instantaneously




A “gas” of interacting
charged particles


We are mainly interested in
ultra-intense laser-plasma interaction



We are interested in laser-plasma interaction to study extreme physical regimes & accelerate particles

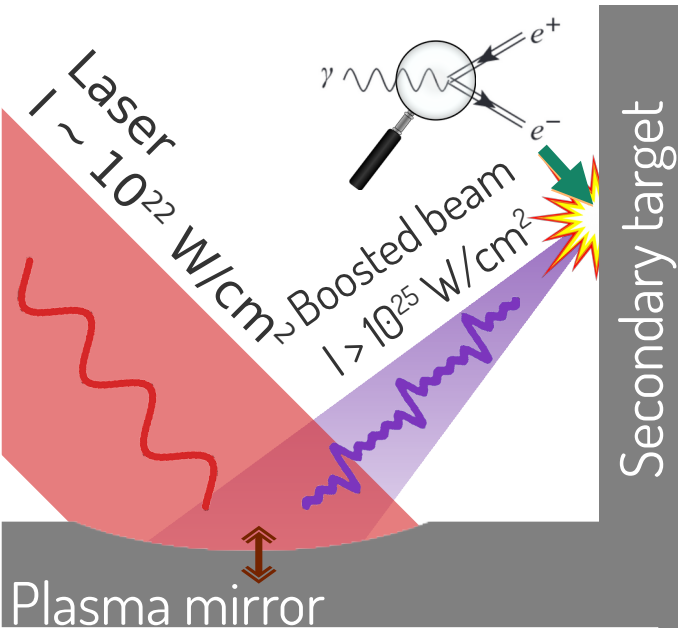


Use ultra-intense lasers to study the strong-field regime of QED




Laser-driven electron acceleration: compact high-energy e^- sources

We are interested in laser-plasma interaction to study extreme physical regimes & accelerate particles





Use ultra-intense lasers to study the strong-field regime of QED




Laser-driven electron acceleration: compact high-energy e^- sources

In laser-driven electron acceleration schemes an ultra-intense laser typically interacts with a low density gas




Reproduced from (CC BY 4.0) :
Kurz, T. et al. Nat. Commun. 12, 2895 (2021),

This generally works well **but...**




We can accelerate electron up to few gigaelectronvolts in few centimeters

This generally works well **but...**
having high charge & high quality at the same time is difficult




This generally works well **but...**
having high charge & high quality at the same time is difficult



**This is not good for many applications,
e.g. FLASH radiobiology studies**




We propose an approach that should give us high-charge, high-quality, ultra-short electron beams



↑ high quality

↑ high charge

We propose an approach that should give us high-charge, high-quality, ultra-short electron beams




↑ high quality

↑ high charge

Does it actually work?

What are the best parameters to use in experiments?

We propose an approach that should give us high-charge, high-quality, ultra-short electron beams




↑ high quality

↑ high charge

Does it actually work?

What are the best parameters to use in experiments?




We need simulations


The Particle-In-Cell method is the standard tool to model relativistic kinetic plasmas

Particle-In-Cell codes are the tool of choice to model kinetic plasma phenomena

Electromagnetic fields




Charged macro-particles




Particle-In-Cell codes are the tool of choice to model kinetic plasma phenomena

Electromagnetic fields

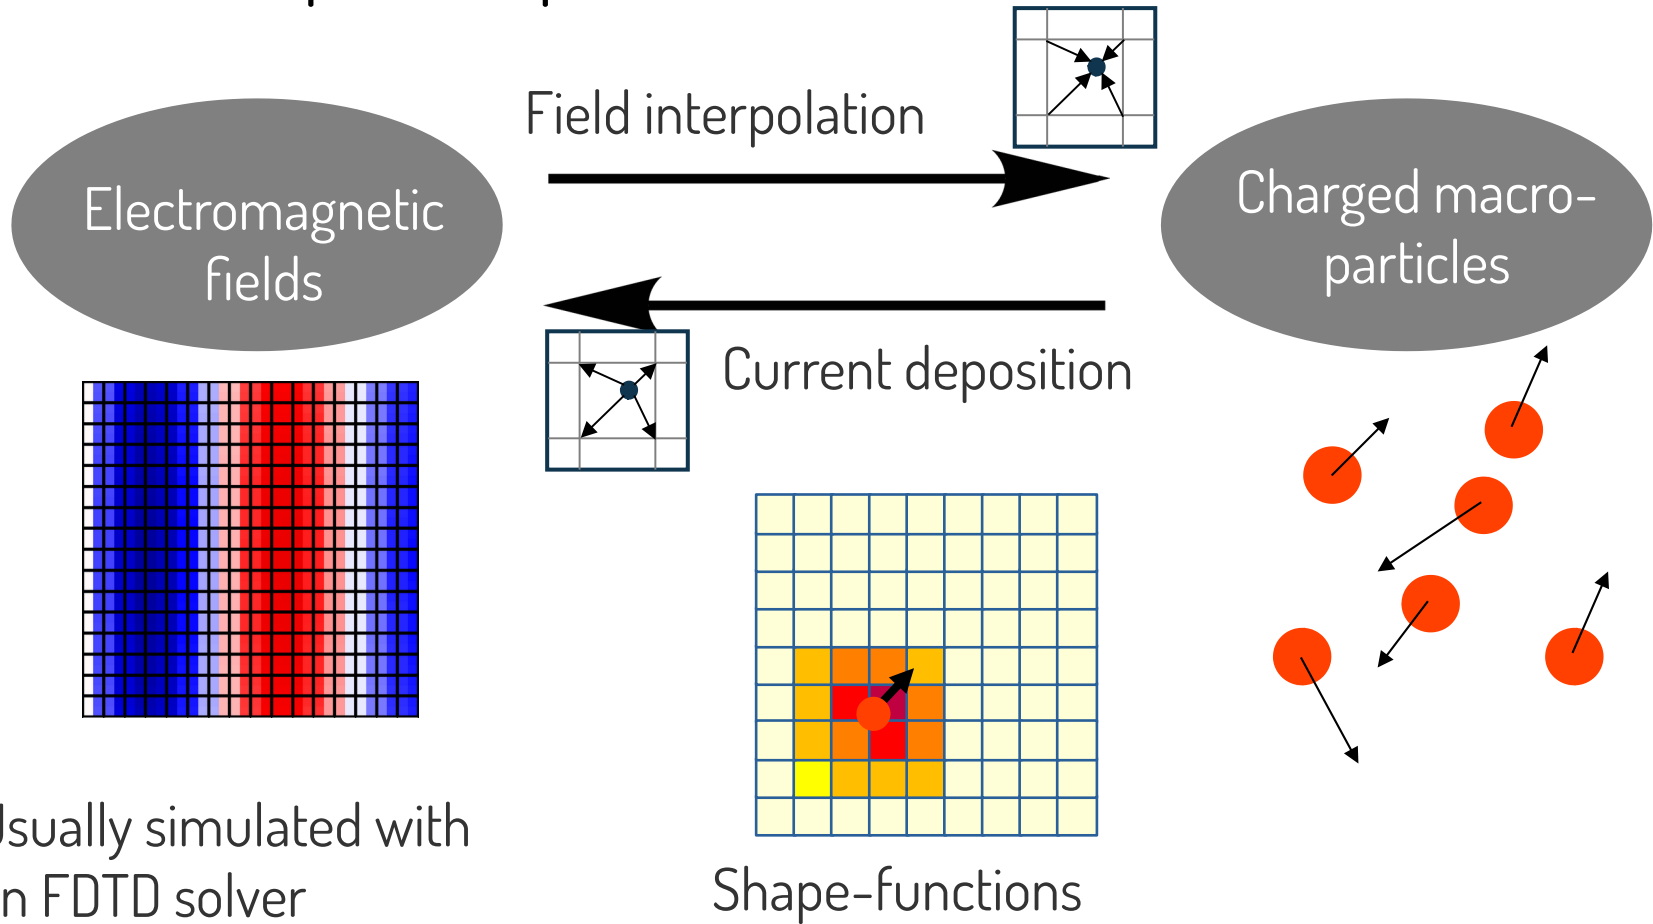


Usually simulated with an FDTD solver

Charged macro-particles




Particle-In-Cell codes are the tool of choice to model kinetic plasma phenomena




Particle-In-Cell codes are the tool of choice to model kinetic plasma phenomena

This is (typically) the most expensive part !



Electromagnetic fields



Usually simulated with an FDTD solver




Charged macro-particles




Shape-functions


We can add more physical processes to the “core” Particle-In-Cell algorithm




Ionization



Collisions




Nuclear fusion

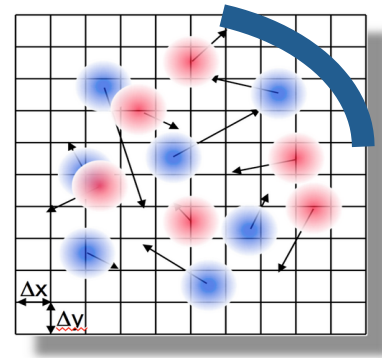


Strong-field QED



And we can also add sophisticated numerical methods



Mesh refinement




Complex surfaces



Advanced solvers ...


If we want to perform 3D simulations,
we often end up needing a lot of computing power

Physical
problem



200 fs to
simulate

If we want to perform 3D simulations,
we often end up needing a lot of computing power

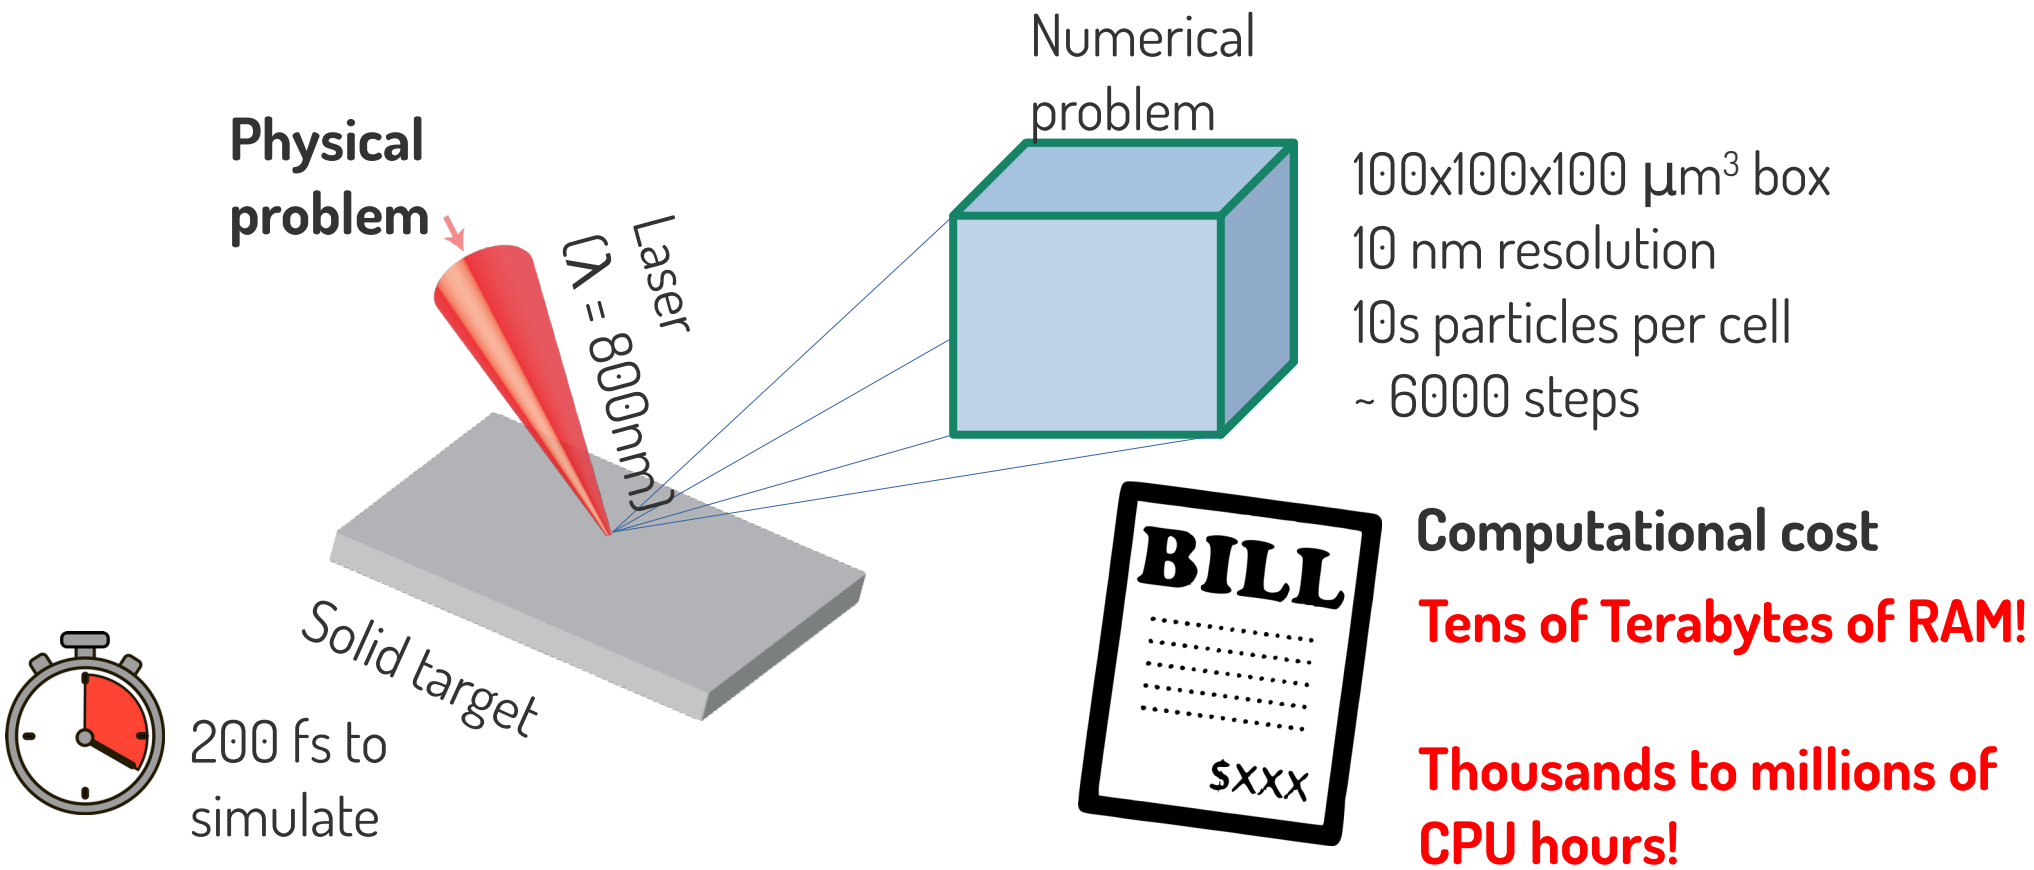


100x100x100 μm^3 box
10 nm resolution
10s particles per cell
~ 6000 steps





200 fs to simulate


If we want to perform 3D simulations,
we often end up needing a lot of computing power




In order to perform large-scale Particle-In-Cell simulations we need supercomputers



Outline




Laser-plasma interaction & the Particle-In-Cell method



WarpX: a Particle-In-Cell code for the exascale era




WarpX performances



Simulation of a new electron accelerator concept

WarpX is an open-source Particle-In-Cell code for the exascale era.

30+ contributors




Open-source & available on Github
Documentation: ecp-warpX.github.io/

Gordon Bell prize
winner @




From your laptop to the largest supercomputers in the world!




WarpX offers a very rich set of features


Comprehensive additional physics modules set →




Ionization



Collisions




Nuclear fusion




Strong-field QED

Cartesian & cylindrical geometries →




3D Cartesian grid




Cylindrical grid (schematic)

Scalable output →





Advanced solvers →




Advanced methods →

Z-order space filling curve




WarpX is used for many different applications!





Plasma accelerators (LBNL, DESY, SLAC)

Laser-ion acceleration - advanced mechanisms (LBNL)




Plasma mirrors and high-field physics + QED (CEA Saclay/LBNL)




Laser-ion acceleration - laser pulse shaping (LLNL)


Fusion devices (Zap Energy, Avalanche Energy)




Thermionic converter (Modern Electron)





Pulsars, magnetic reconnection (LBNL)



Magnetic fusion sheaths (LLNL)



Microelectronics (LBNL) - ARTEMIS




WarpX runs on GPUs (AMD, NVIDIA) and on CPUs (AMD, Intel, ARM...)



WarpX runs on GPUs (AMD, NVIDIA) and on CPUs (AMD, Intel, ARM...)



We achieve performance portability across different architectures thanks to **AMReX**



WarpX is built on top of the AMReX library, which provides performance portability

Single source approach



```
using namespace amrex;



int N = 1'000'000;

Gpu::ManagedVector<double> a(N);
Gpu::ManagedVector<double> b(N);
Gpu::ManagedVector<double> c(N);
Gpu::ManagedVector<double> result(N);

/* OTHER CODE*/

auto d_a = a.data();
auto d_b = b.data();
auto d_c = c.data();
auto d_result = c.data();

ParallelFor(N,
 [=] AMREX_GPU_DEVICE (int i){
     d_result[i] = d_a[i]*d_b[i] + d_c[i];
 });
```



We express our algorithms as lambdas
fed to “ParallelFor” functions

```
#ifdef AMREX_USE_OMP
    #pragma omp parallel for
#endif
for (WarpXParIter pti(*this, lev); pti.isValid(); ++pti)
{
    // ..
    amrex::ParallelFor(number_of_particles,
        [=] AMREX_GPU_DEVICE (int i)
        {
            // ..
        }
    // ..
}
```

We express our algorithms as lambdas
fed to “ParallelFor” functions

```
#ifdef AMREX_USE_OMP
    #pragma omp parallel for
#endif
for (WarpXParIter pti(*this, lev); pti.isValid(); ++pti)
{
    // ..
    amrex::ParallelFor(number_of_particles,
        [=] AMREX_GPU_DEVICE (int i)
        {
            // ..
        }
    // ..
}
```

← On GPUs, this is a
CUDA/HIP/DPC++ kernel call

On CPUs this is just a loop
(possibly SIMD)



We express our algorithms as lambdas fed to “ParallelFor” functions

```
#ifdef AMREX_USE_OMP
    #pragma omp parallel for
#endif
for (WarpXParIter pti(*this, lev); pti.isValid(); ++pti)
{
    // ..
    amrex::ParallelFor(number_of_particles,
        [=] AMREX_GPU_DEVICE (int i)
        {
            // ..
        })
    // ..
}
```

← On GPUs, this is a CUDA/HIP/DPC++ kernel call

On CPUs this is just a loop (possibly SIMD)



← AMReX also provides GPU-friendly containers, drop-in replacement for some STL features, parallel reductions...

“ParallelFor” now supports also compile-time optimization for runtime parameters

```
amrex::ParallelFor(TypeList<CompileTimeOptions<A0,A1,A2,A3>>{},
{runtime_option},
box, [=] AMREX_GPU_DEVICE (int i, int j, int k, auto control)
{
    //...
    if constexpr (control.value == A0) {
        //...
    } else if constexpr (control.value == A1) {
        //...
    } else if constexpr (control.value == A2) {
        //...
    } else {
        //...
    }
    //...
});
```

← Thanks to template programming, under the hood, it generates all the possible combinations

“ParallelFor” now supports also compile-time optimization for runtime parameters



```

amrex::ParallelFor( TypeList<CompileTimeOptions<A0,A1,A2,A3>> {},
{runtime_option},
box, [=] AMREX_GPU_DEVICE (int i, int j, int k, auto control)
{
    //...
    if constexpr (control.value == A0) {
        //...
    } else if constexpr (control.value == A1) {
        //...
    } else if constexpr (control.value == A2) {
        //...
    } else {
        //...
    }
    //...
});


```

← Thanks to template programming, under the hood, it generates all the possible combinations


Helpful to reduce registry pressure on GPUs and for vectorization on CPUs



Outline




Laser-plasma interaction & the Particle-In-Cell method



WarpX: a Particle-In-Cell code for the exascale era



WarpX performances




Simulation of a new electron accelerator concept


Let's have a look at WarpX performances




1st



Frontier
AMD MI250X



2nd



Fugaku
Fujitsu A64FX



5th




Summit
NVIDIA V100



8th



Perlmutter
NVIDIA A100



A Particle-In-Cell code is **memory bound**:
we expect **only few % peak FLOP/s efficiency**



Perlmutter A100

DP PFlop/s

3.38

Linpack
Benchmark

12.9%

HPCG
Benchmark

223%





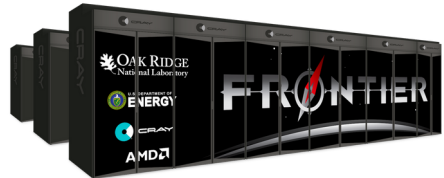
Summit V100

11.79





8.3%

435%

A Particle-In-Cell code is **memory bound**:
we expect **only few % peak FLOP/s efficiency**

		DP PFlop/s	Linpack Benchmark	HPCG Benchmark
	Perlmutter A100	3.38	12.9%	223%
	Summit V100	11.79	8.3%	435%
	Frontier MI250X	43.45	3.3%	310%

A Particle-In-Cell code is **memory bound**:
we expect **only few % peak FLOP/s efficiency**

		DP PFlop/s	Linpack Benchmark	HPCG Benchmark
	Perlmutter A100	3.38	12.9%	223%
	Summit V100	11.79	8.3%	435%
	Frontier MI250X	43.45	3.3%	310%
	Fugaku A64FX	5.31	1.1%	35%

A Particle-In-Cell code is **memory bound**:
 we expect **only few % peak FLOP/s efficiency**



Perlmutter A100

DP PFlop/s

3.38

Linpack
Benchmark

12.9%

HPCG
Benchmark

223%



Summit V100

11.79

8.3%

435%




Frontier MI250X

43.45

3.3%

310%



Fugaku A64FX


5.31

1.1%

35%

→ **Specific tuning for Fugaku (3.3X perf. in SP)** **Atos**

WarpX **very well** over 4-5 orders of magnitude




Nodes

Frontier: 1 – 8,576 (pre-acceptance)


Fugaku: 1 – 152,064

Summit: 2 – 4,263

Perlmutter: 1 – 1,088 (pre-acceptance)



WarpX can be **strong-scaled by an order of magnitude** when needed




Nodes

Frontier: 512 – 8,192 (pre-acceptance)


Fugaku: 6,144 – 152,064

Summit: 512 – 4,096


Perlmutter: 15 – 480 (pre-acceptance)



Outline




Laser-plasma interaction & the Particle-In-Cell method



WarpX: a Particle-In-Cell code for the exascale era

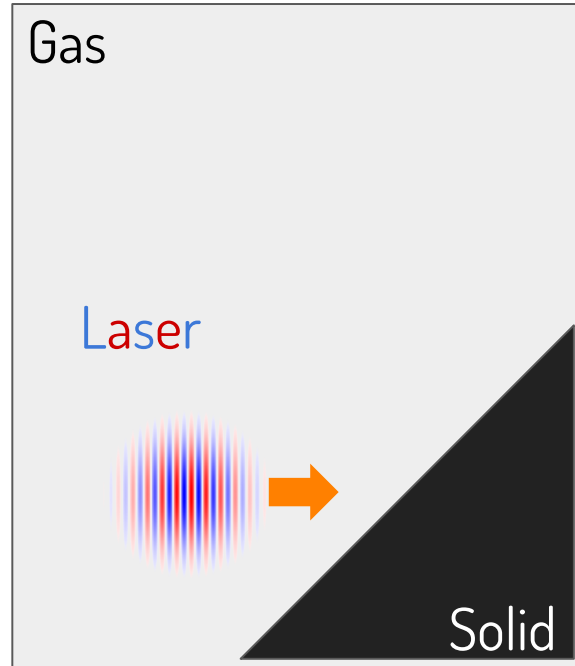



WarpX performances

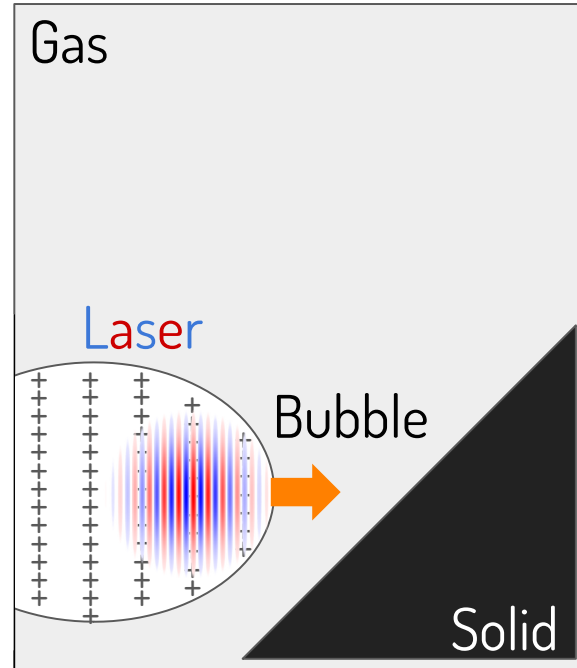



Simulation of a new electron accelerator concept


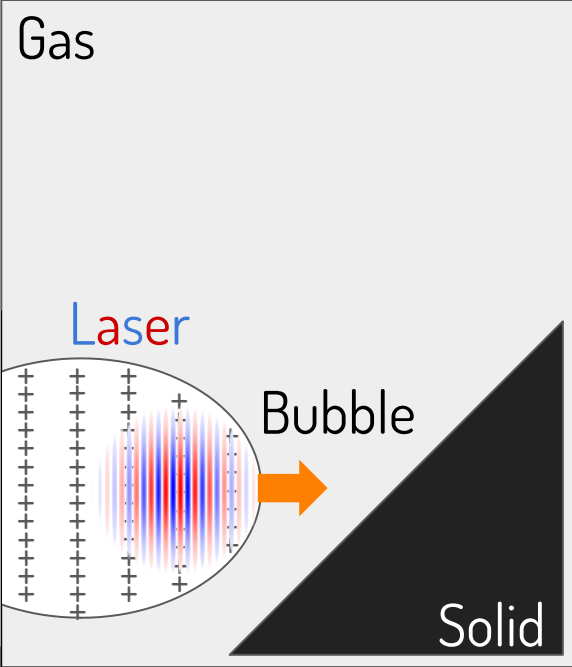

An ultra-short laser beam propagates in a low density gas






The laser pushes electrons away and generates a positively charged “bubble”




The laser is reflected by the high-density plasma and the bubble traps some of its electrons





The bubble accelerates electrons over few millimeters



We can have smaller simulation boxes with a "moving window"




We only need to simulate
 $\sim 100 \times 100 \times 100 \mu\text{m}^3 \rightarrow$





The simulation box follows the laser

The main challenge concerns laser-solid interaction




We need a resolution of **-10 nanometers** for laser-solid interaction →

We only need to simulate **$\sim 100 \times 100 \times 100 \mu\text{m}^3$** →





The main challenge concerns laser-solid interaction




We only need to simulate
 $\sim 100 \times 100 \times 100 \mu\text{m}^3 \rightarrow$

Enabled by **very good weak scaling** \rightarrow





The main challenge concerns laser-solid interaction




We only need to simulate
 $\sim 100 \times 100 \times 100 \mu\text{m}^3 \rightarrow$


Enabled by **very good weak scaling** \rightarrow




The main challenge concerns laser-solid interaction



Enabled by **very good strong scaling** →



Enabled by **very good weak scaling** →



The simulation box follows the laser

The main challenge concerns laser-solid interaction




Enabled by **very good strong scaling** →



How do we switch resolution in the middle of the simulation?



Enabled by **very good weak scaling** →



Mesh refinement, one of the most advanced WarpX features, comes to help

Mesh refinement in a Particle-In-Cell code is **a nightmare!**


Electromagnetic waves have different dispersion relations in the two areas!
(spurious reflections, unphysical effects...)





High resolution
here!

J.-L. Vay et al, Phys. Plasmas 11, 2928 (2004)
R. Lehe et al, Phys. Rev. E 106, 045306 (2022)

WarpX allows us to define a region with twice the resolution of the main grid





2D slices of our 3D simulations highlight the acceleration process





← 3D simulation on 4096 Summit nodes

2D slices of our 3D simulations highlight the acceleration process



← 3D simulation
on 4096 Summit
nodes


2D slices of our 3D simulations highlight the acceleration process




← 3D simulation
on 4096 Summit
nodes

2D slices of our 3D simulations highlight the acceleration process

Side view





Front view





← 3D simulation on 4096 Summit nodes

2D slices of our 3D simulations highlight the acceleration process





← 3D simulation
on 4096 Summit
nodes

2D slices of our 3D simulations highlight the acceleration process





← 3D simulation
on 4096 Summit
nodes

2D slices of our 3D simulations highlight the acceleration process





← 3D simulation
on 4096 Summit
nodes

2D slices of our 3D simulations highlight the acceleration process





← 3D simulation on 4096 Summit nodes

2D slices of our 3D simulations highlight the acceleration process





← 3D simulation
on 4096 Summit
nodes

2D slices of our 3D simulations highlight the acceleration process





← 3D simulation
on 4096 Summit
nodes

2D slices of our 3D simulations highlight the acceleration process





← 3D simulation
on 4096 Summit
nodes

2D slices of our 3D simulations highlight the acceleration process





← 3D simulation
on 4096 Summit
nodes

2D slices of our 3D simulations highlight the acceleration process




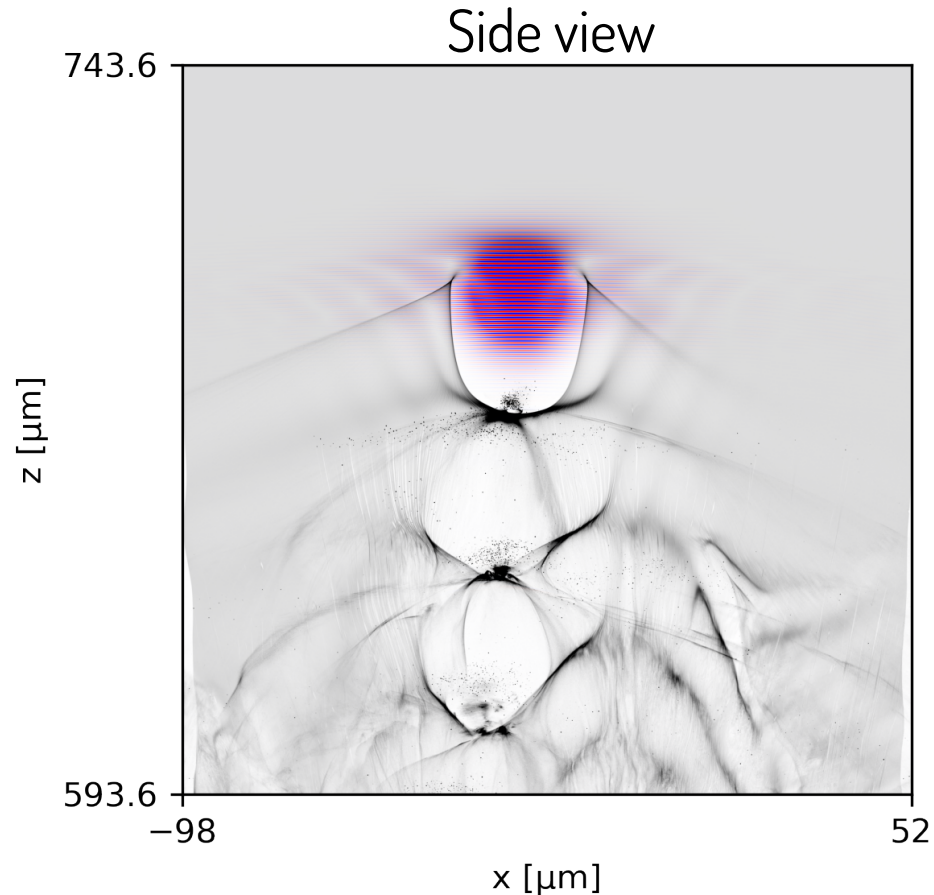
← 3D simulation
on 4096 Summit
nodes

2D slices of our 3D simulations highlight the acceleration process





← 3D simulation on 4096 Summit nodes

2D slices of our 3D simulations highlight the acceleration process





← 3D simulation
on 4096 Summit
nodes

2D slices of our 3D simulations highlight the acceleration process





← 3D simulation
on 4096 Summit
nodes

2D slices of our 3D simulations highlight the acceleration process





← 3D simulation on 4096 Summit nodes

2D slices of our 3D simulations highlight the acceleration process



← 3D simulation
on 4096 Summit
nodes


2D slices of our 3D simulations highlight the acceleration process



← We are mainly concerned with the properties of these electrons

Our simulations with a PW-class laser show that we can accelerate a substantial amount of charge with high quality

After ~ 1mm
(acceleration still in progress)




Production runs on Frontier, Fugaku and Summit



Exascale simulations informed the design of the first experimental validation of our concept (at LOA, France)



Exascale simulations informed the design of the first experimental validation of our concept (at LOA, France)



Laser parameters

$E = 400 \text{ mJ}$

waist = $15 \mu\text{m}$

$P_{\text{peak}} = 10 \text{ TW}$


Results

$Q = 17 \text{ pC}$ (at least)

$dE/E_{\text{peak}} = 8\%$

4X increase of accelerated charge with respect to conventional techniques for the same laser energy!


Conclusions and perspectives



- WarpX is a state-of-the-art open-source Particle-In-Cell code implementing sophisticated numerical algorithms




- WarpX is portable across different architectures and scales well on top machines, including the first exascale supercomputer



- WarpX allows us to simulate novel electron acceleration strategies that could find application for e.g. FLASH radiobiology experiments

Modeling a novel laser-driven electron accelerator concept: Particle-In-Cell simulations at the exascale.



Luca Fedeli, Axel Huebl, France Boillod-Cerneux, Thomas Clark, Kevin Gott, Conrad Hillairet, Stephan Jaure, Adrien Leblanc, Rémi Lehe, Andrew Myers, Christelle Piechurski, Mitsuhsa Sato, Neil Zaim, Weiqun Zhang, Jean-Luc Vay, Henri Vincenti



The End

FLASH effect : cells react differently
to irradiation if it is administered very fast




promising for **cancer treatment** , but still
not well understood


Favaudon et al, Science. Trans. Med., 2014

Bourhis et al, Radiotherapy and Oncology, 2019

With the help of **Atos**
we optimized the most expensive
kernels for A64FX (single precision only)




Current deposition and field gather require to sample a lot of points per particle




Particles have a
shape function

Current deposition and field gather require to sample a lot of points per particle




Particles have a shape function



64 points per particle!

Vectorization is key to achieve good performances




```
for p : particle
{
  for i : x_indices
    for j : y_indices
      for k : z_indices
        {
          compute n_ijk
        }
      }
    }
  }
}
```

Very small loops (4x4x4)

→ **Inefficient vectorization**

Vectorization is key to achieve good performances



```
for p : particle
{
  for i : x_indices
    for j : y_indices
      for k : z_indices
        {
          compute n_ijk
        }
      }
    }
  }
}
```


Very small loops (4x4x4)
→ **Inefficient vectorization**



```
for i : x_indices
  for j : y_indices
    for k : z_indices
      {
        for p : particle
          {
            compute n_ijk
          }
        }
      }
    }
  }
}
```

Longer inner loop
→ **Efficient vectorization**
→ **Data reorganization**

Vectorization is key to achieve good performances



```
for p : particle
{
  for i : x_indices
    for j : y_indices
      for k : z_indices
      {
        compute n_ijk
      }
    }
  }
}
```

Very small loops (4x4x4)
→ **Inefficient vectorization**

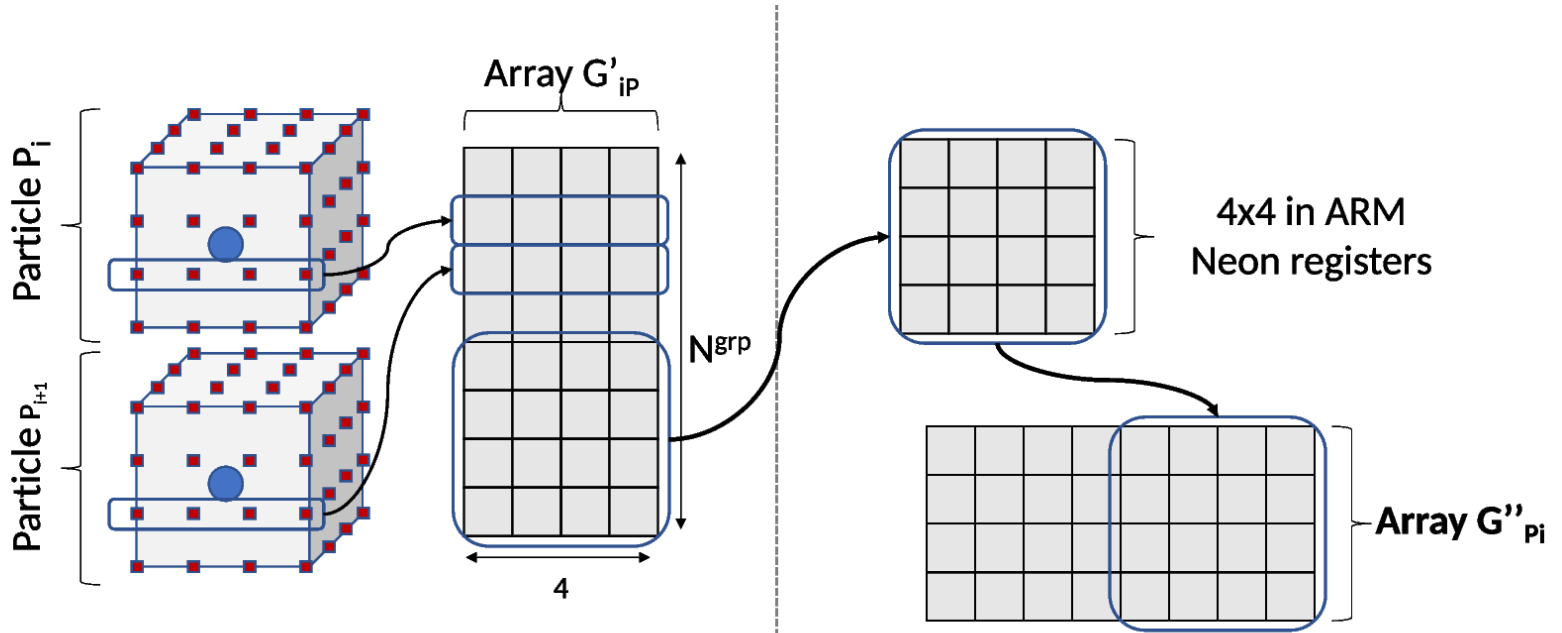


* Particles are grouped (N = 32-128) to keep data in CPU cache.

```
for i : x_indices
  for j : y_indices
    for k : z_indices
      {
        for p : particle
        {
          compute n_ijk
        }
      }
    }
  }
}
```

Longer inner loop
→ **Efficient vectorization**
→ **Data reorganization**

Data re-organization makes heavy use of NEON intrinsics





For each j,k :

For each particle in batch:

Gather 4 contiguous sampling points using ARM Neon intrinsics

Transpose array $4 \times N_{grp} \Rightarrow N_{grp} \times 4$ using ARM Neon intrinsics

Optimized field gather and current deposition lead to very significant speed-ups!



Routine	Speed up
Gather	2.63X
Deposition	4.60X

fipp profiling (whole code):

SIMD inst. rate: 2.2% → 27%

SVE operation rate: 3.6% → 70%