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

The WarpX code: Particle-In-Cell simulations at the exascale

Thomas Clark

11/04/2023

PHI
PHYSICS AT HIGH INTENSITY


O. Gobert


T. Ceccotti


A. Panchal


A. Ammar


I. Kara-Mostefa


P. Forestier-Colleoni

Experiments


S. Dobosz Dufrénoy


P. Bartoli


H. Vincenti
(head of numerical division)

Theory/simulations


A. Sainte-Marie


P. Martin


L. Fedeli


T. Clark


N. Zaim

Outline

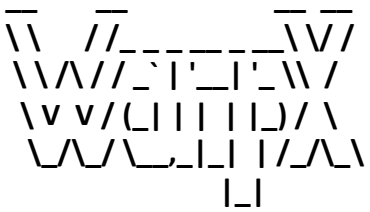

The Particle-In-Cell code WarpX

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

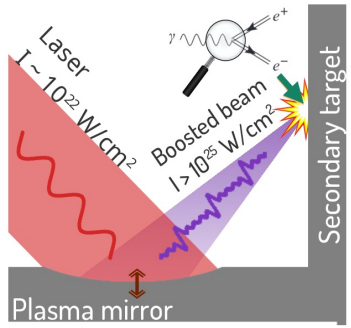

What we study with WarpX

Outline

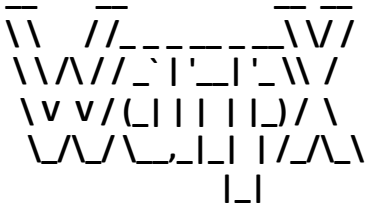

The Particle-In-Cell code WarpX

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

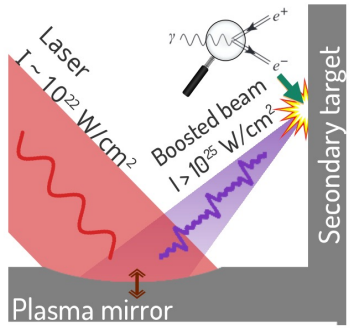

What we study with WarpX

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

30+ contributors

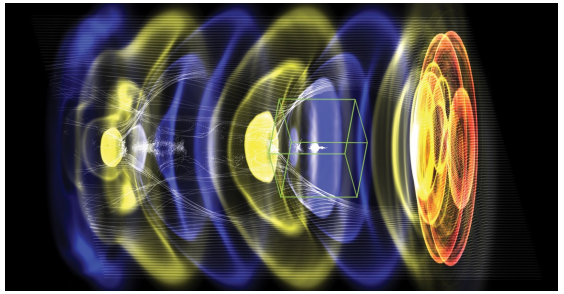

EXASCALE COMPUTING PROJECT

BERKELEY LAB

DE LA RECHERCHE À L'INDUSTRIE

Lawrence Livermore National Laboratory

SLAC NATIONAL ACCELERATOR LABORATORY

DESY.

CERN

MODERN ELECTRON

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

30+ contributors

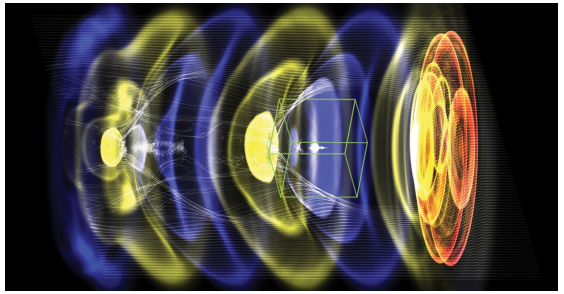

EXASCALE COMPUTING PROJECT

BERKELEY LAB

DE LA RECHERCHE À L'INDUSTRIE

Lawrence Livermore National Laboratory

SLAC NATIONAL ACCELERATOR LABORATORY

DESY.

CERN

MODERN ELECTRON

open source initiative Approved License®

Open-source & available on Github

Documentation: **ecp-warpx.github.io/**

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

30+ contributors

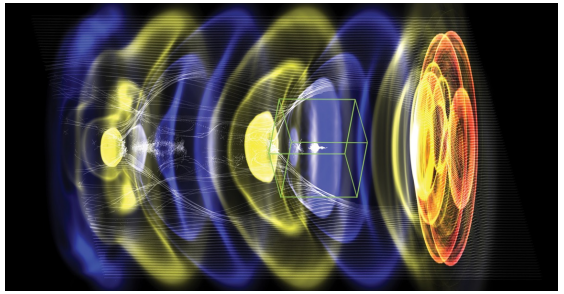

EXASCALE COMPUTING PROJECT

BERKELEY LAB

DE LA RECHERCHE À L'INDUSTRIE

Lawrence Livermore National Laboratory

SLAC NATIONAL ACCELERATOR LABORATORY

DESY.

CERN

MODERN ELECTRON

open source initiative Approved License®

Open-source & available on Github

Documentation: **ecp-warpx.github.io/**

mac OS

WarpX offers a comprehensive set of additional physical modules

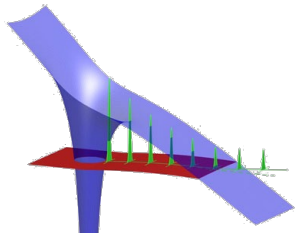

We implement **tunnel ionization** (ADK theory)

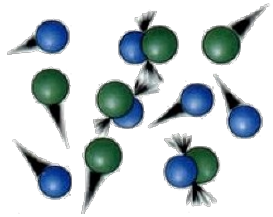

We implement **Coulomb collisions** and collisions with **neutral background**

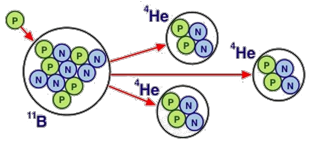

We implement **deuterium-deuterium**, **deuterium-tritium**, **deuterium-helium** and **proton-boron fusion**

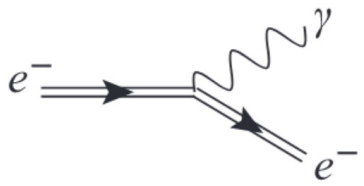

We implement **quantum synchrotron** and **nonlinear Breit-Wheeler pair production**

WarpX provides advanced algorithms

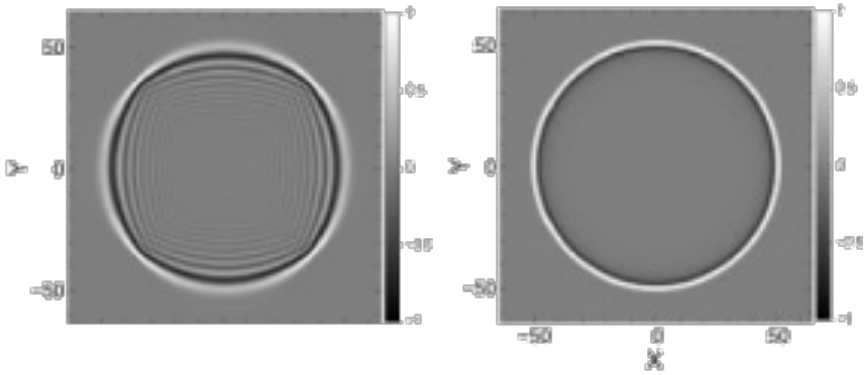

We provide a **pseudo-spectral solver**, that tackles numerical dispersion, avoiding the cost of a global FFT

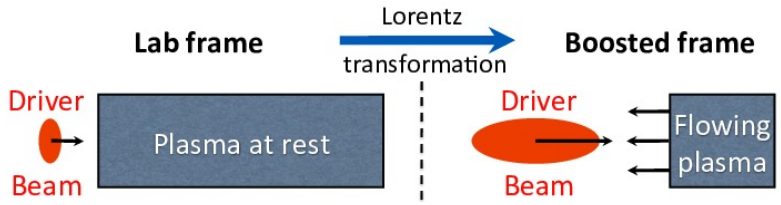

We provide the option of using a **"Boosted frame"**, where the simulation may be orders of magnitude faster

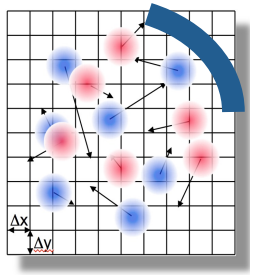

We provide the option of adding **"embedded surfaces"** with complex geometries

WarpX provides advanced algorithms

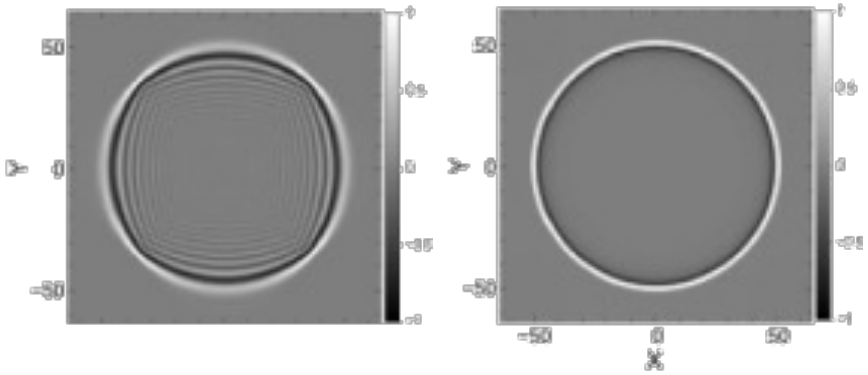

We provide a **pseudo-spectral solver**, that tackles numerical dispersion, avoiding the cost of a global FFT

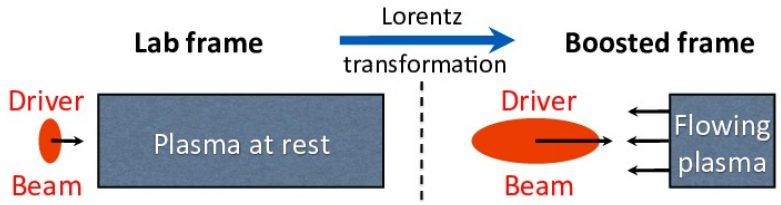

We provide the option of using a **"Boosted frame"**, where the simulation may be orders of magnitude faster

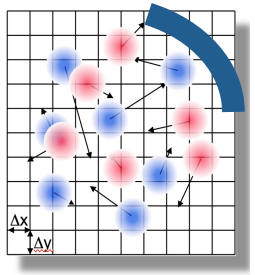

We provide the option of adding **"embedded surfaces"** with complex geometries

(and several others!)

WarpX provides critical features to run efficiently at scale

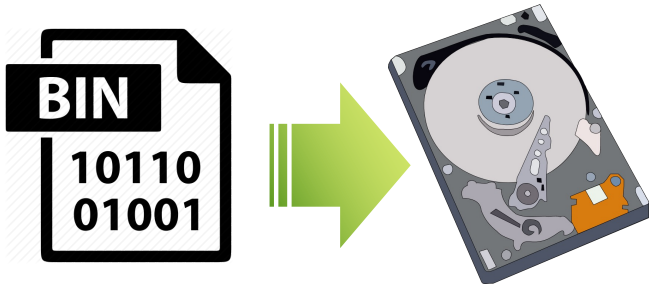

Scalable output based on OpenPMD+ADIOS2 allows writing efficiently 10s Terabytes of data per simulation

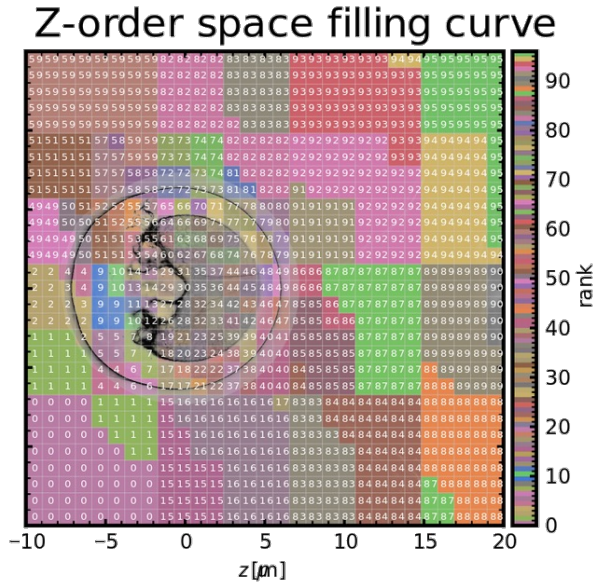

Dynamic load balancing: redistributing "chunks" of the simulation among the nodes to ensure that each one has an approximately equal amount of work

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

30+ contributors

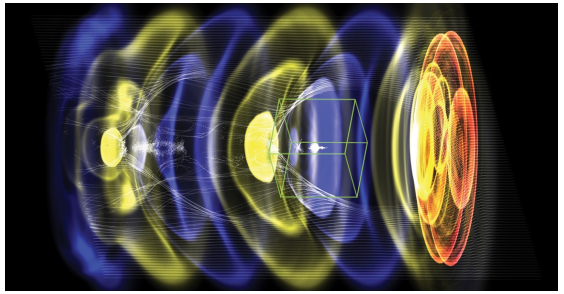

EXASCALE COMPUTING PROJECT

BERKELEY LAB

DE LA RECHERCHE À L'INDUSTRIE

Lawrence Livermore National Laboratory

SLAC NATIONAL ACCELERATOR LABORATORY

DESY.

CERN

MODERN ELECTRON

open source initiative Approved License®

Open-source & available on Github

Documentation: **ecp-warpx.github.io/**

mac OS

From your laptop to the largest supercomputers in the world!

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

30+ contributors

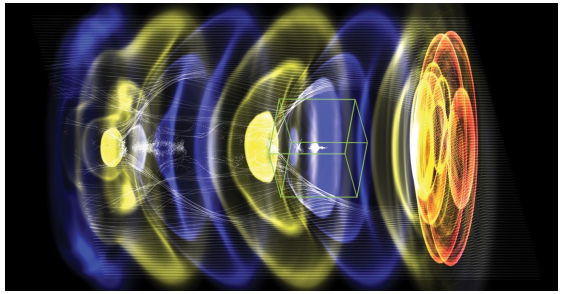

EXASCALE COMPUTING PROJECT

BERKELEY LAB

DE LA RECHERCHE À L'INDUSTRIE

Lawrence Livermore National Laboratory

SLAC NATIONAL ACCELERATOR LABORATORY

DESY.

CERN

MODERN ELECTRON

open source initiative Approved License®

Open-source & available on Github

Documentation: **ecp-warpx.github.io/**

Gordon Bell prize winner @

SC22

From your laptop to the largest supercomputers in the world!

Outline

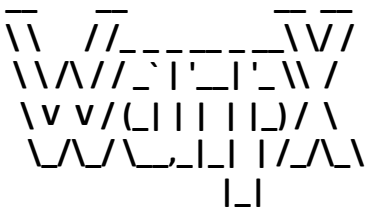

The Particle-In-Cell code WarpX

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

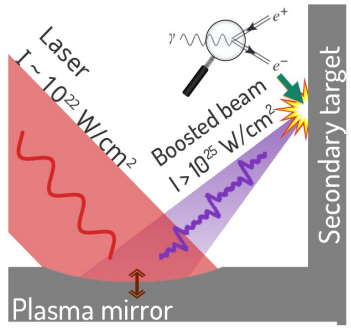

What we study with WarpX

Why do we need portability ?

		Architecture	Rank in TOP500
	Perlmutter	Nvidia A100	8
	Summit	Nvidia V100	5
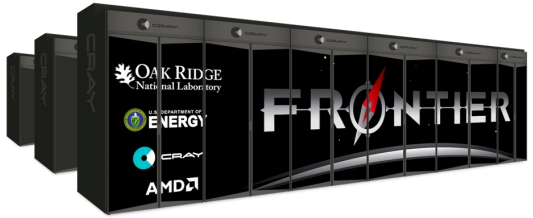	Frontier	AMD MI250X	1
	Fugaku	Fujitsu A64FX	2

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

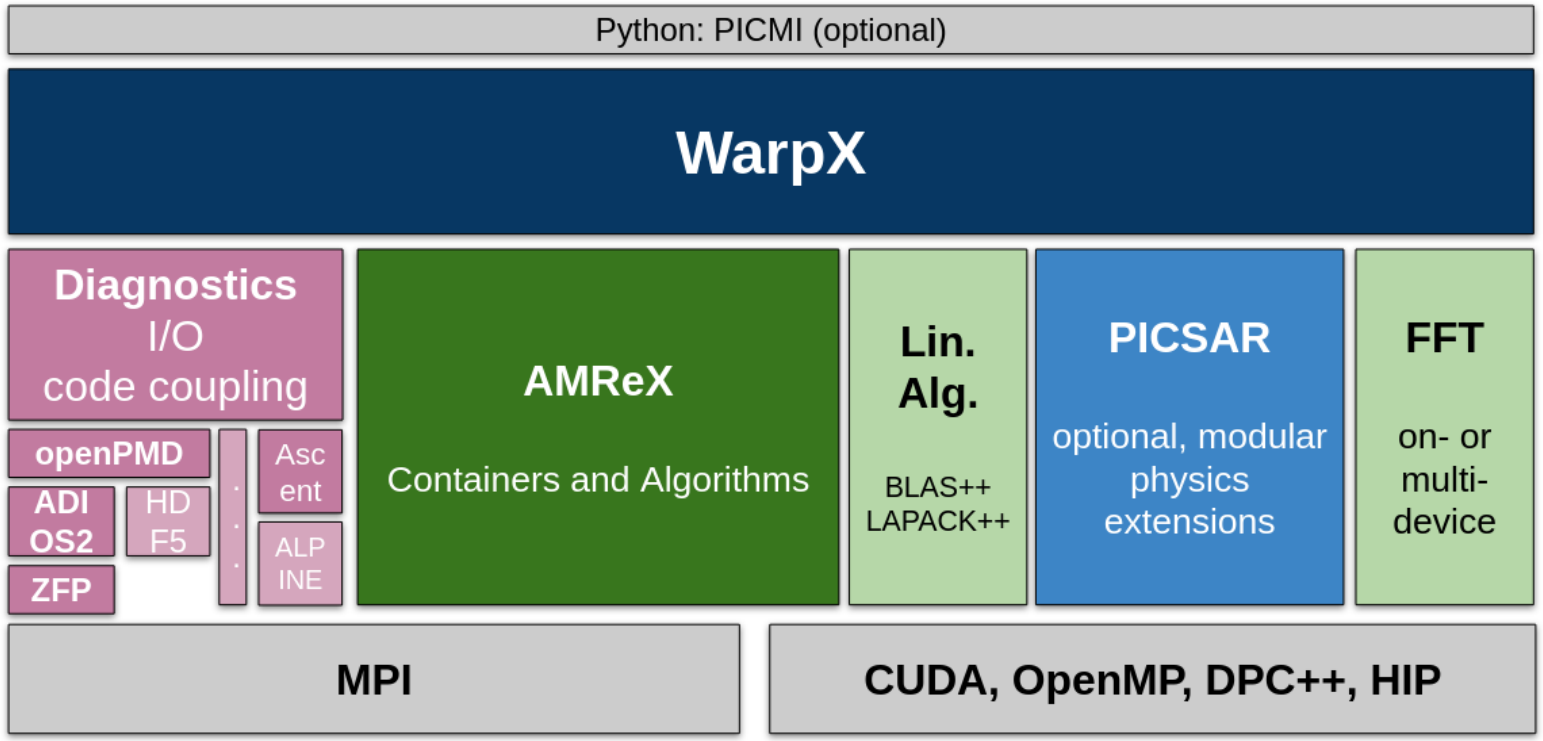

WarpX is built on top of the AMReX library, which provides performance portability

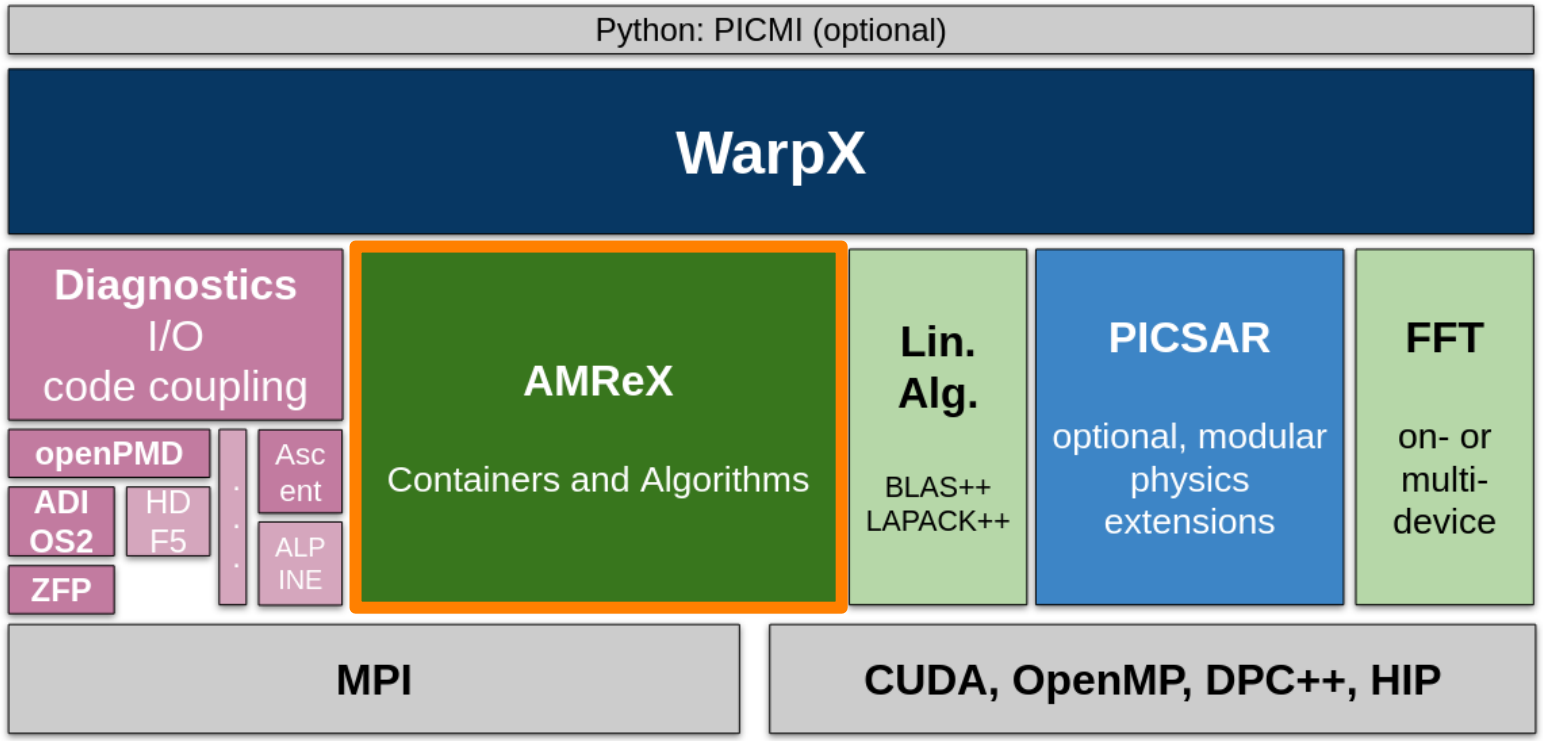

WarpX is built on top of the AMReX library, which provides performance portability

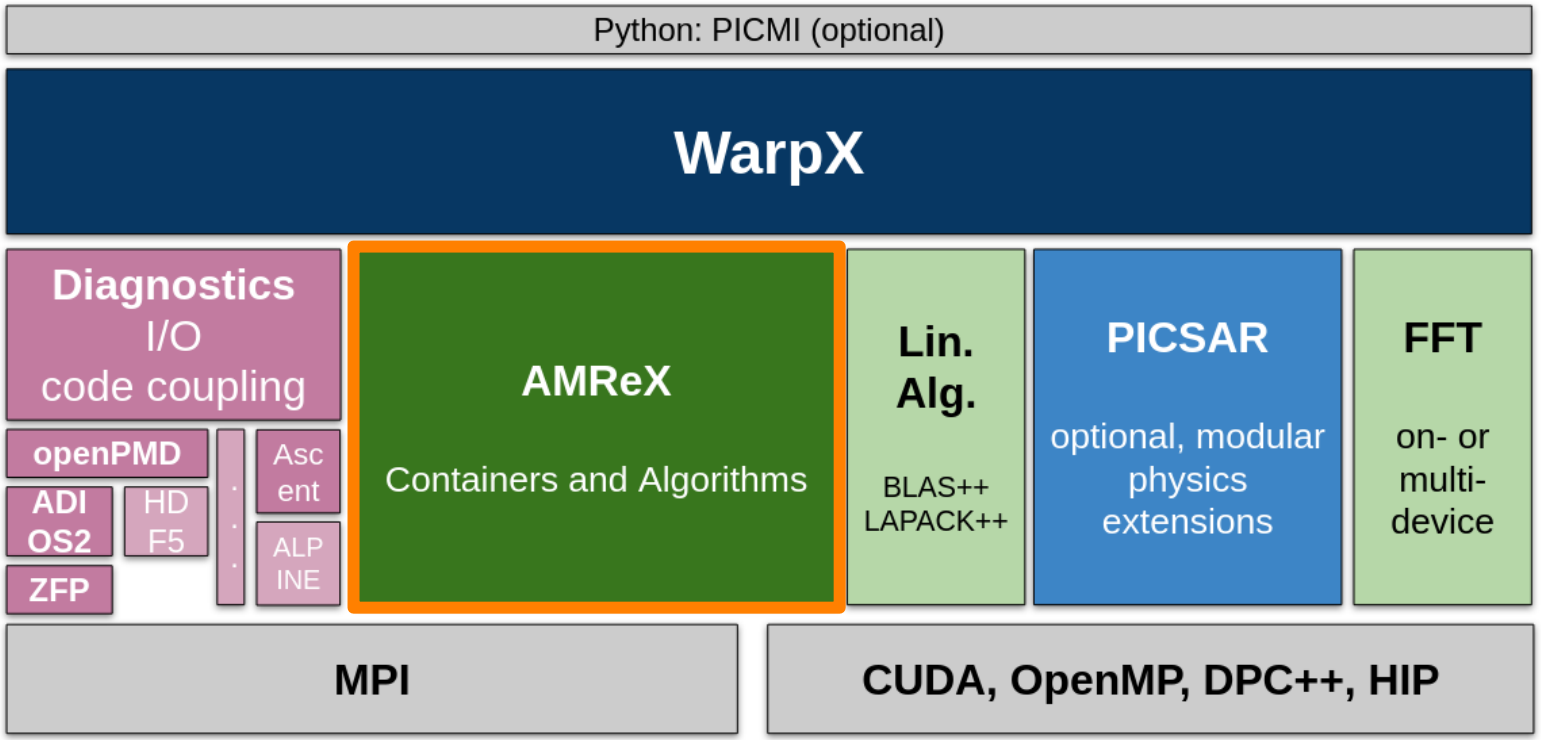

Single source approach

```cpp
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

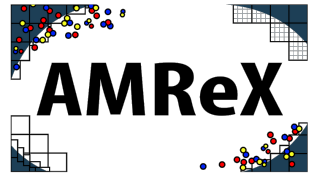

We express our algorithms as lambdas fed to "ParallelFor" functions

```
#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)

← 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

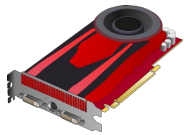

WarpX scales **very well** over 4-5 orders of magnitude

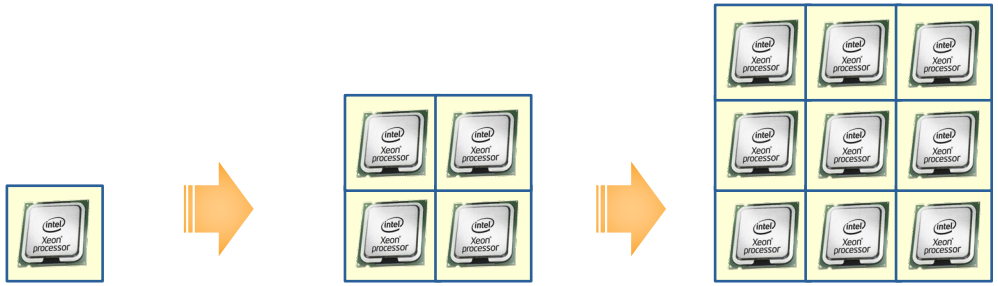

L.Fedeli et al. SC22 (2022)

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

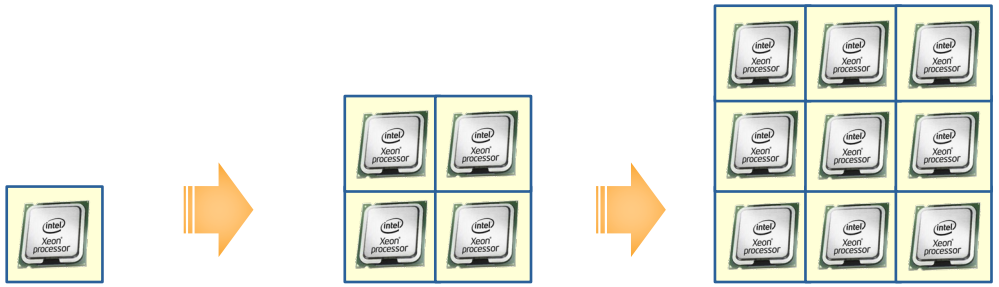

Nodes

Frontier: 1 – 8,576 (pre-acceptance)

Fugaku: 1 – 152,064

Summit: 2 – 4,263

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

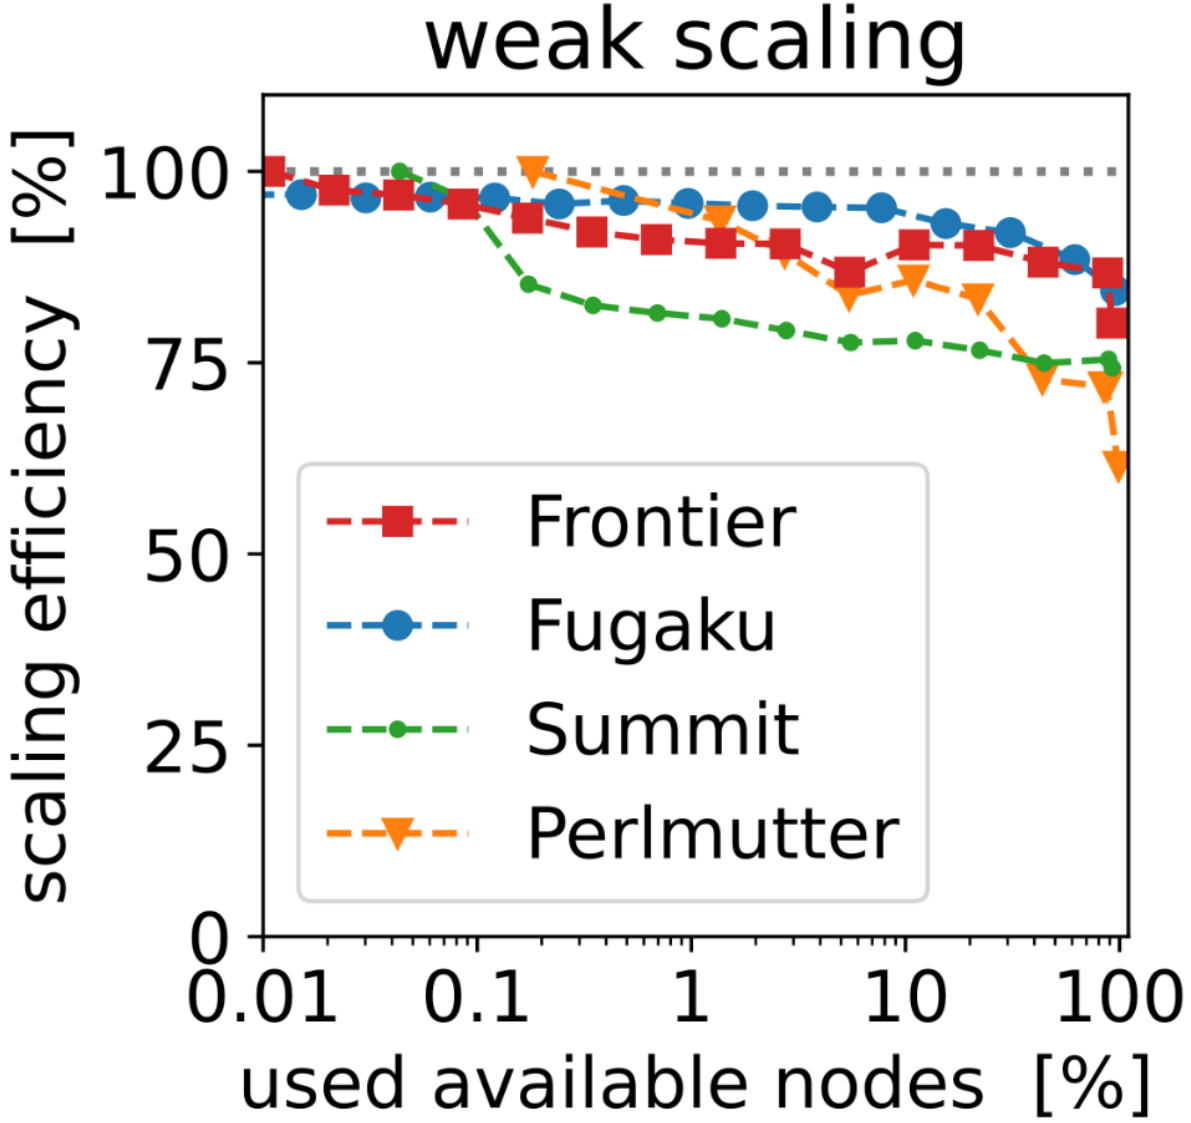

L.Fedeli et al. SC22 (2022)

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

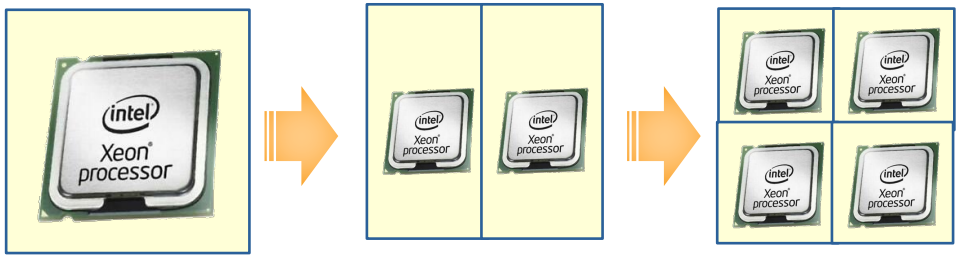

L.Fedeli et al. SC22 (2022)

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

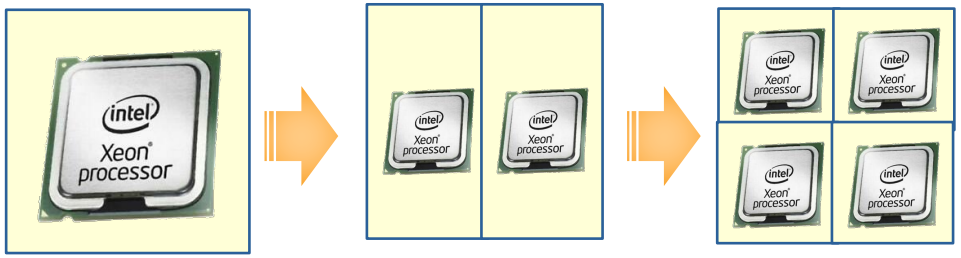

Nodes

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

Fugaku: 6,144 - 152,064

Summit: 512 - 4,096

Perlmutter: 15 – 480 (pre-acceptance)

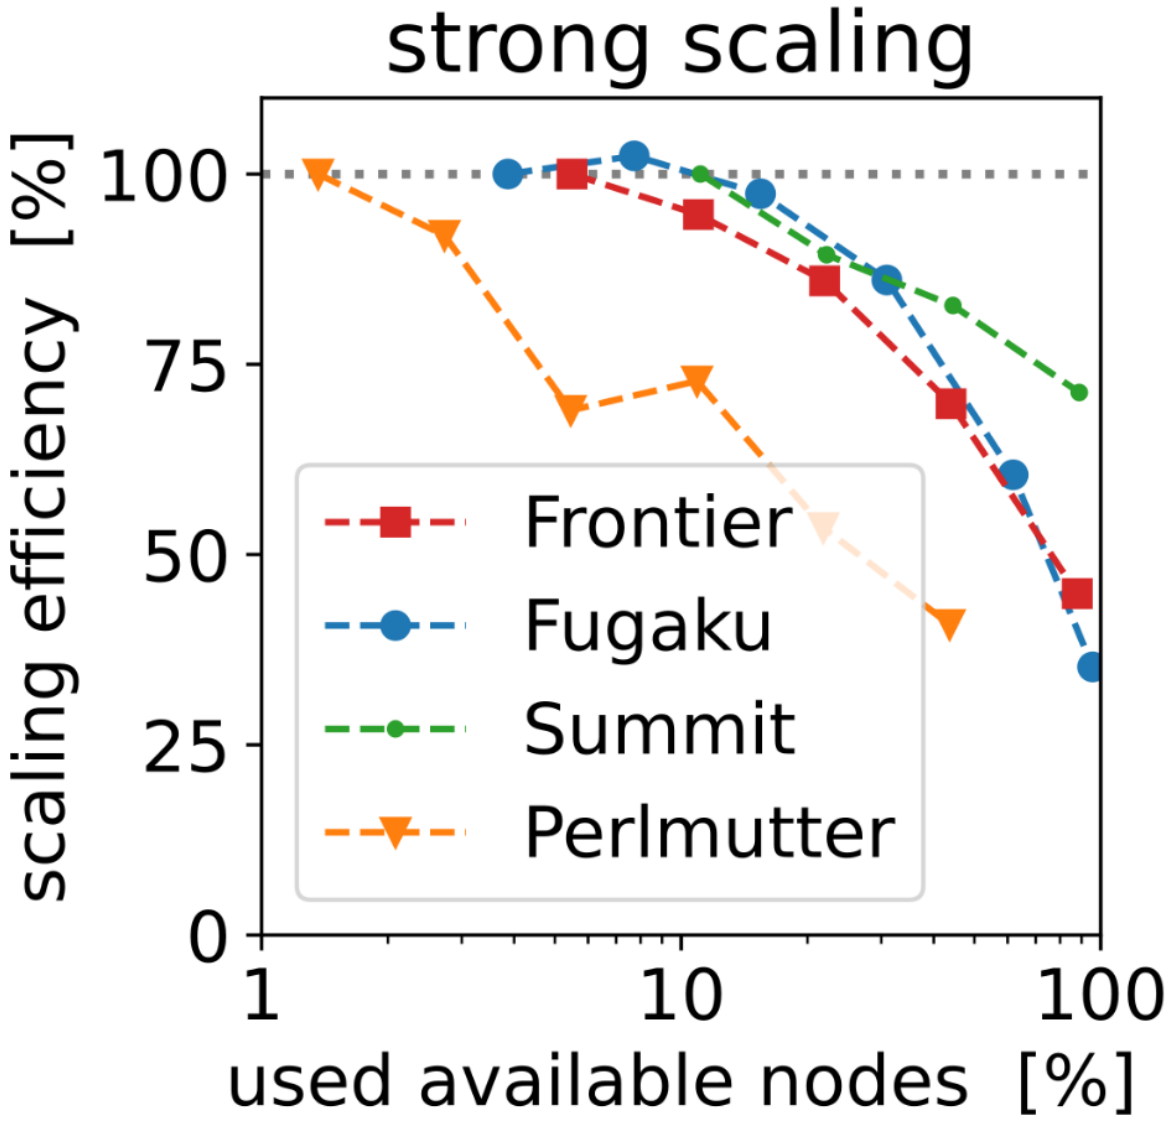

L.Fedeli et al. SC22 (2022)

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

L.Fedeli et al. SC22 (2022)

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

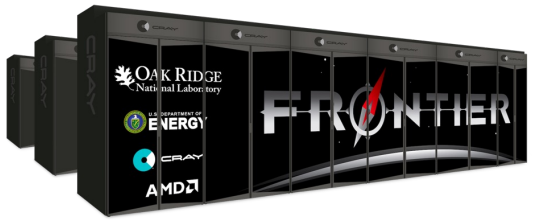

	DP PFlop/s
Perlmutter A100	3.38
Summit V100	11.79
Frontier MI250X	43.45

L.Fedeli et al. SC22 (2022)

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

	DP PFlop/s
 **Perlmutter A100**	3.38
 **Summit V100**	11.79
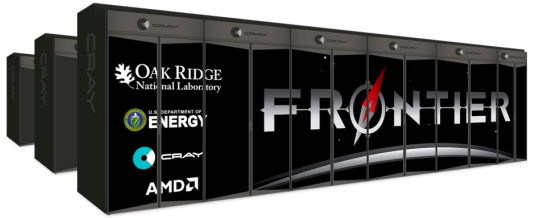 **Frontier MI250X**	43.45
 **Fugaku A64FX**	5.31

L.Fedeli et al. SC22 (2022)

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

	DP PFlop/s
 **Perlmutter A100**	3.38
 **Summit V100**	11.79
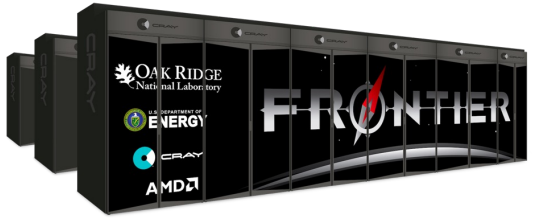 **Frontier MI250X**	43.45
 **Fugaku A64FX**	5.31

L.Fedeli et al. SC22 (2022)

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

With the help of Atos

we optimized the most expensive kernels for A64FX (single precision only)

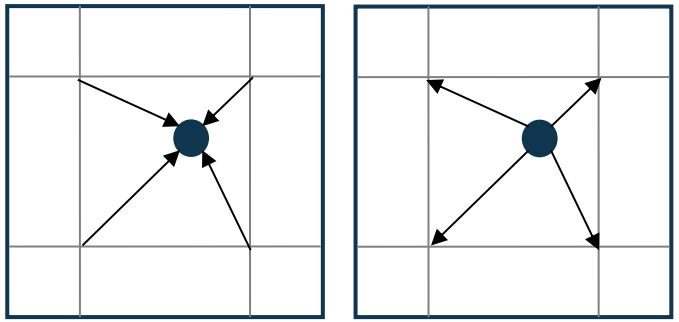

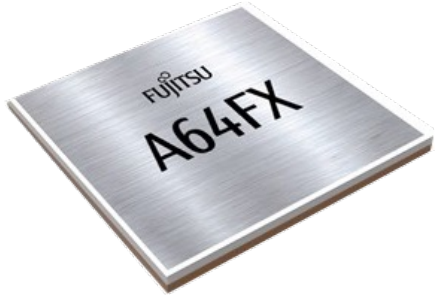

A64FX-specific optimizations

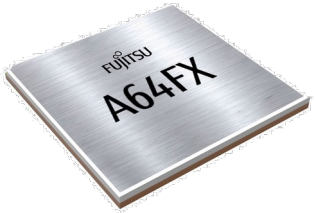

Current deposition is
a very expensive operation

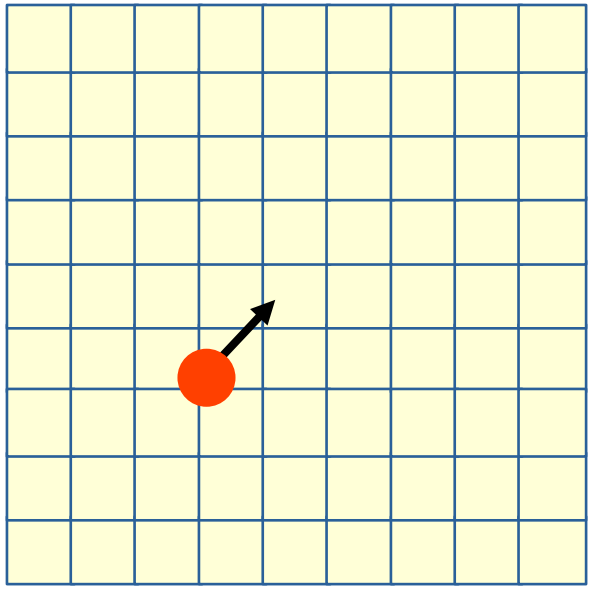

A64FX-specific optimizations

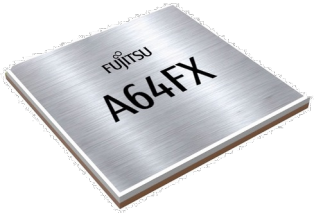

Current deposition is
a very expensive operation

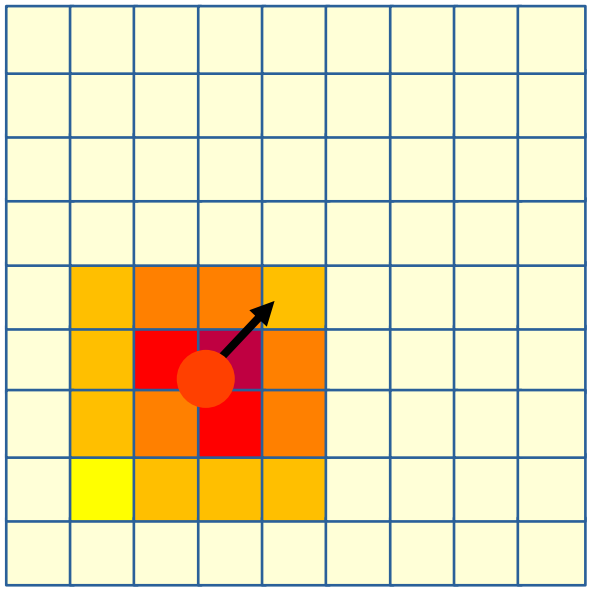

Pseudocode

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

}
```

← 64 indices in 3D, but very small loops (4x4x4)

Inefficient vectorization

A64FX-specific optimizations

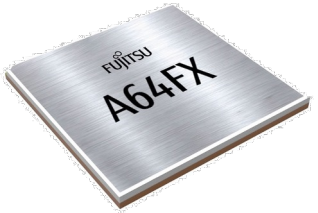

Current deposition is
a very expensive operation

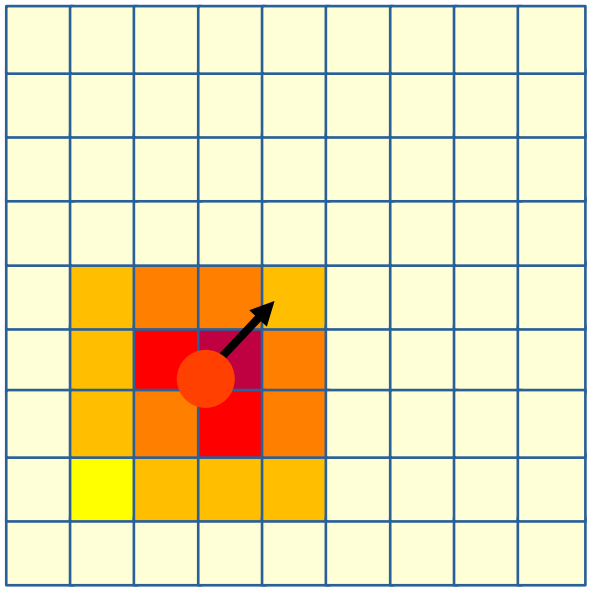

Pseudocode

```
for p : particle
{
  for i  x_indi
    for j  ices
          fo  : z_indices

                co  te n_ijk
      }

}
```

← 64 indices in 3D, but very small loops (4x4x4)

Inefficient vectorization

A64FX-specific optimizations

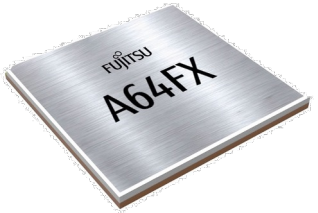

Current deposition is
a very expensive operation

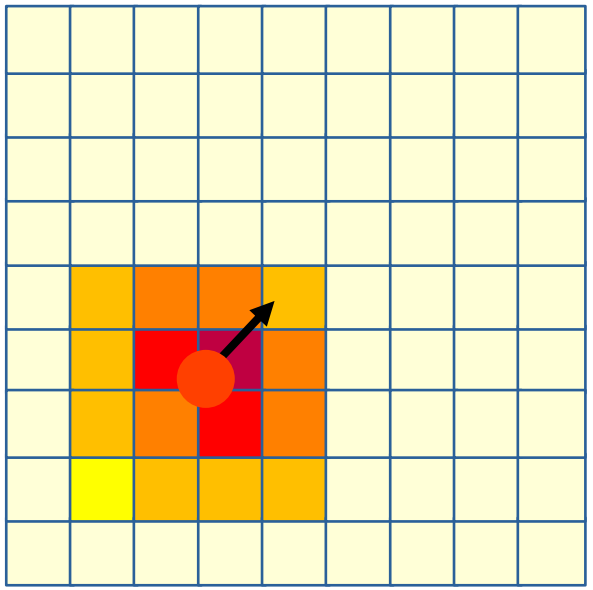

Optimized pseudocode

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

← Longer inner loop
More efficient vectorized code

However, this requires **data reorganization**

Best performances obtained using intrinsics

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

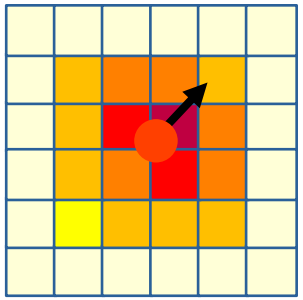

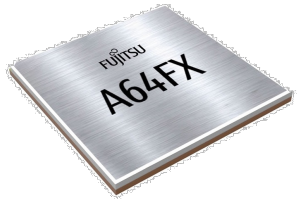

Routine	Speed up
Gather	2.63X
Deposition	4.60X

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

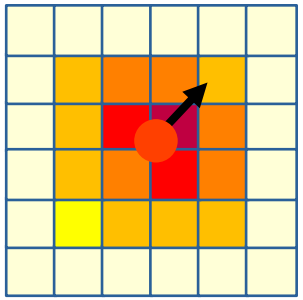

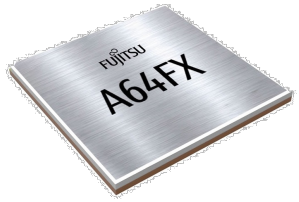

Routine	Speed up
Gather	2.63X
Deposition	4.60X

Now we want to generalize those optimizations for other CPUs architectures !

Outline

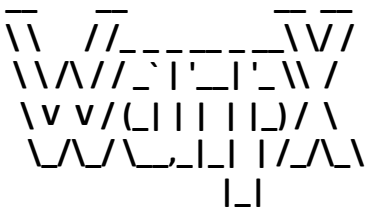

The Particle-In-Cell code WarpX

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

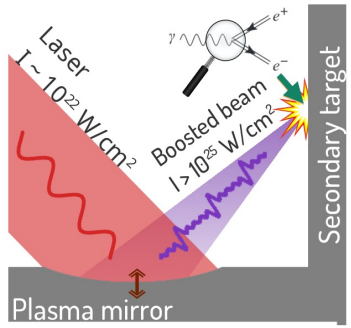

What we study with WarpX

WarpX is used for many different applications!

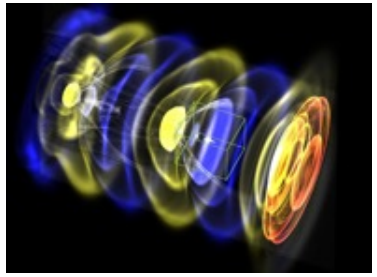

Plasma accelerators (LBNL, DESY, SLAC)

Laser-ion acceleration – advanced mechanisms (LBNL)

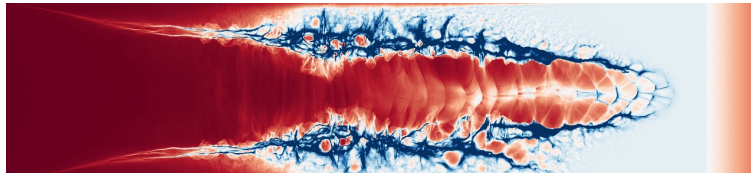

Plasma mirrors and high-field physics + QED (CEA Saclay/LBNL)

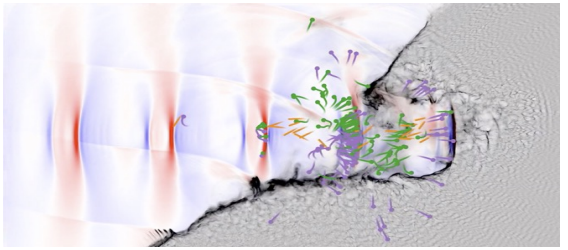

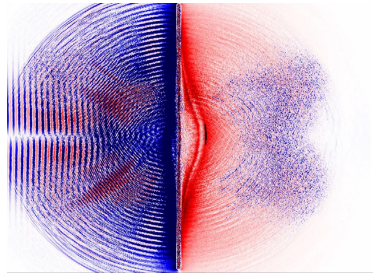

Laser-ion acceleration – laser pulse shaping (LLNL)

Fusion devices (Zap Energy, Avalanche Energy)

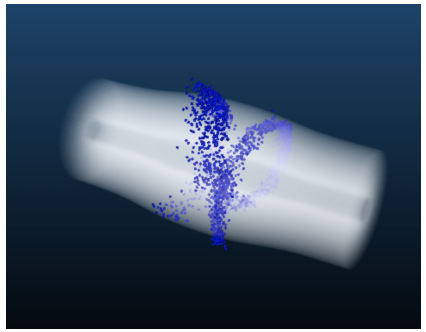

Thermionic converter (Modern Electron)

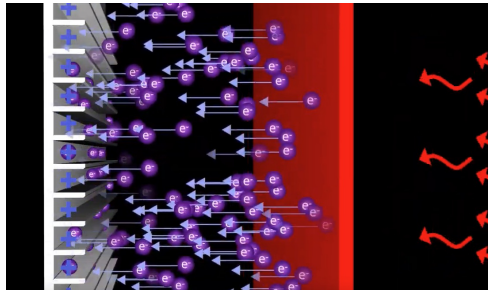

Pulsars, magnetic reconnection (LBNL)

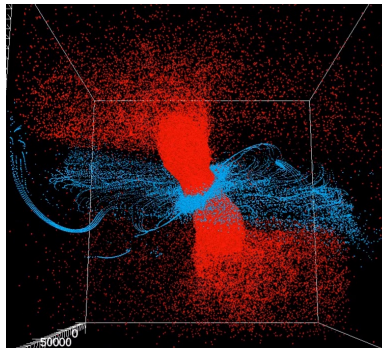

Magnetic fusion sheaths (LLNL)

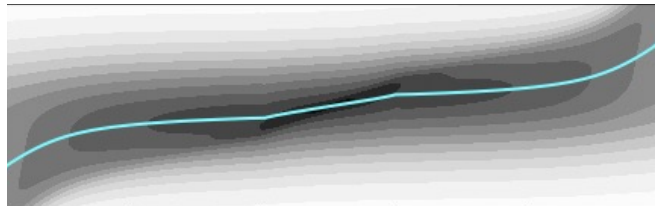

Microelectronics (LBNL) – ARTEMIS

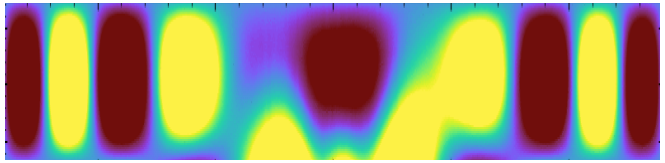

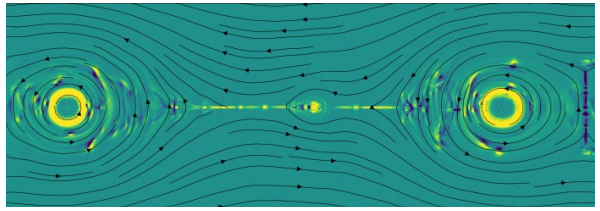

At CEA, We are mainly interest by those one !

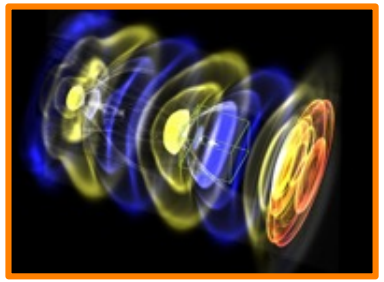

Plasma accelerators (LBNL, DESY, SLAC)

Laser-ion acceleration - advanced mechanisms (LBNL)

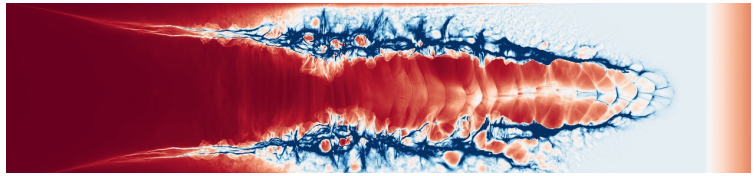

Plasma mirrors and high-field physics + QED (CEA Saclay/LBNL)

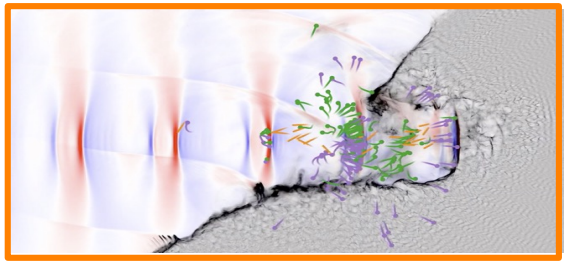

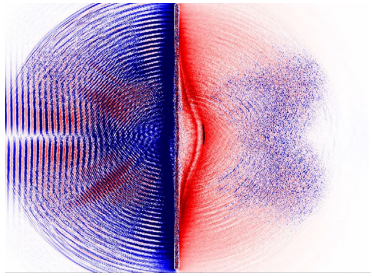

Laser-ion acceleration - laser pulse shaping (LLNL)

Fusion devices (Zap Energy, Avalanche Energy)

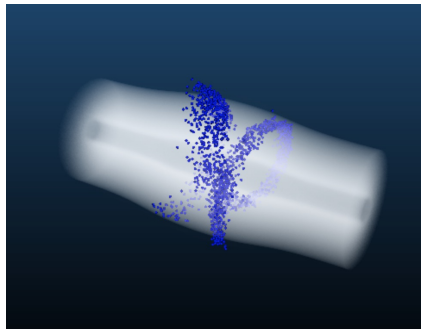

Thermionic converter (Modern Electron)

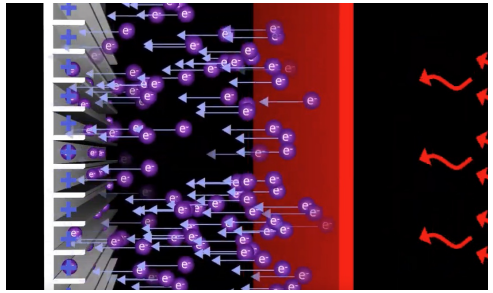

Pulsars, magnetic reconnection (LBNL)

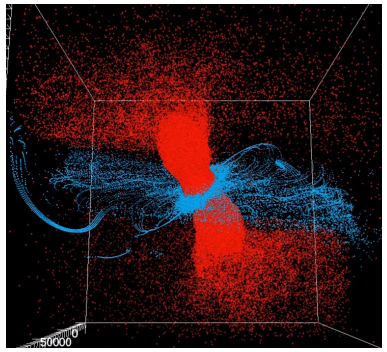

Magnetic fusion sheaths (LLNL)

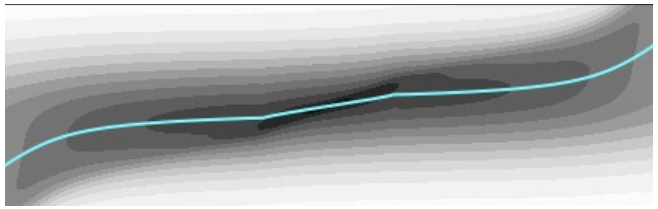

Microelectronics (LBNL) - ARTEMIS

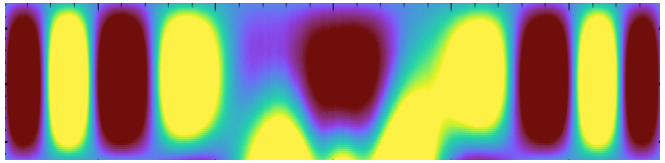

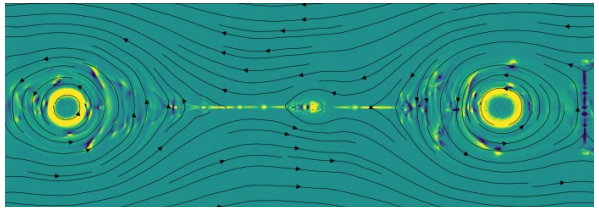

Our main research interests

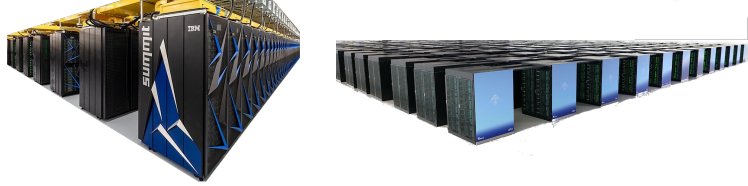

Using ultra-intense lasers to study strong-field Quantum Electrodynamics

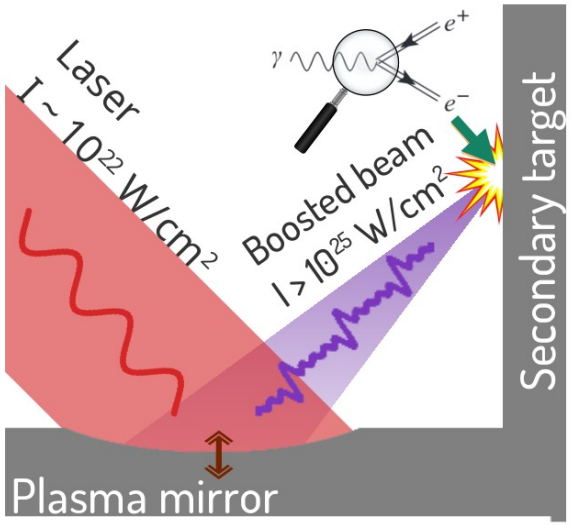

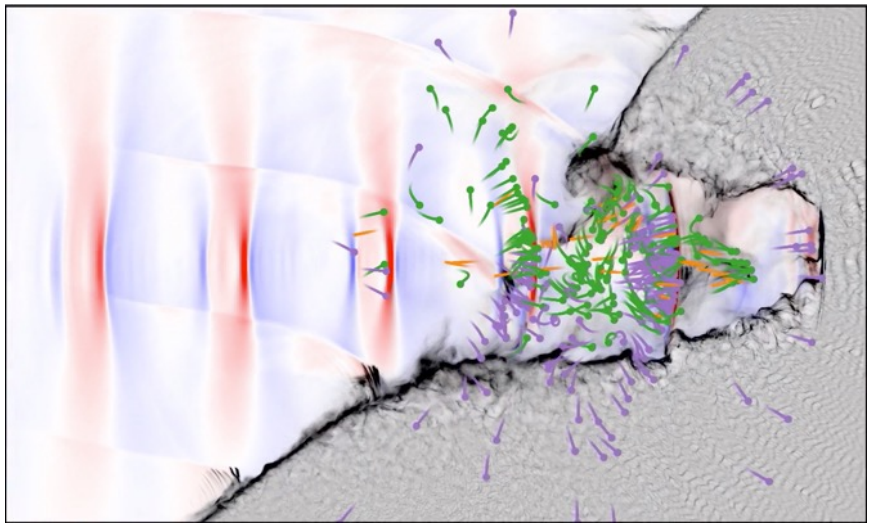

Advanced laser-driven electron sources

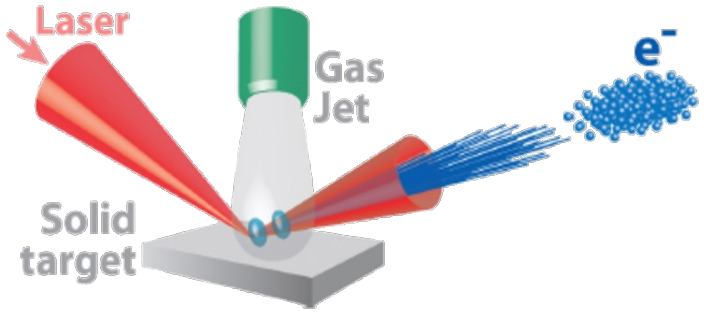

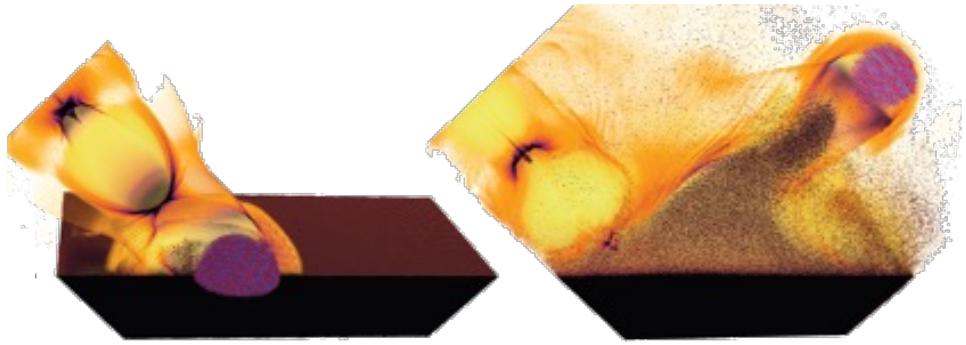

Our main research interests

Using ultra-intense lasers to study strong-field Quantum Electrodynamics

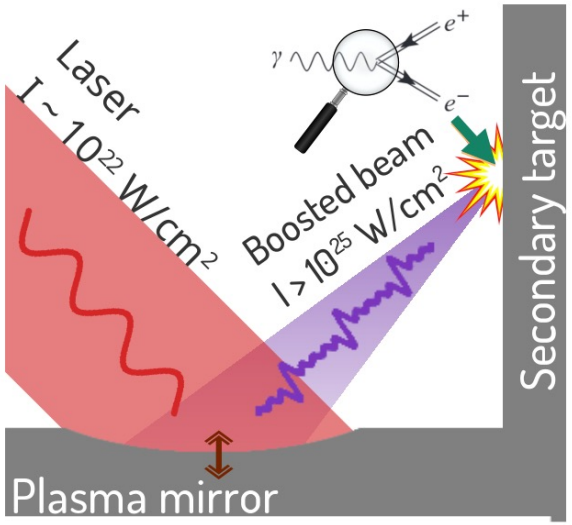

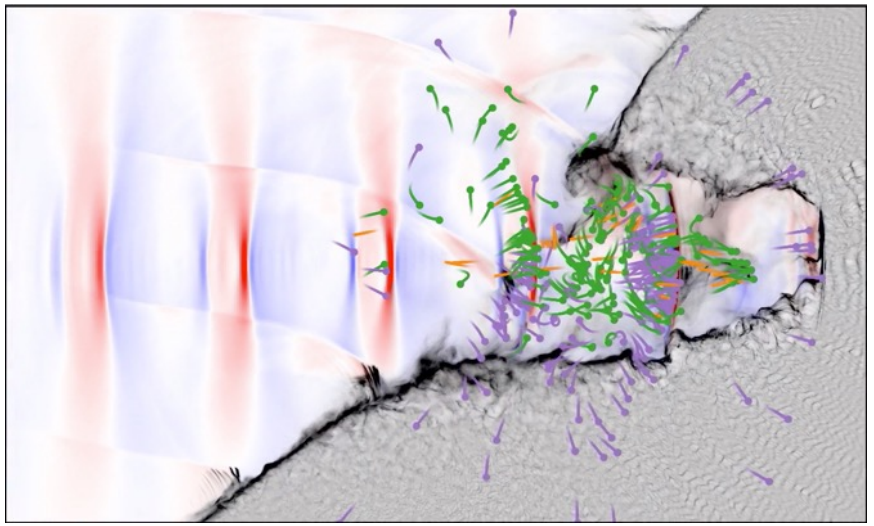

Advanced laser-driven electron sources

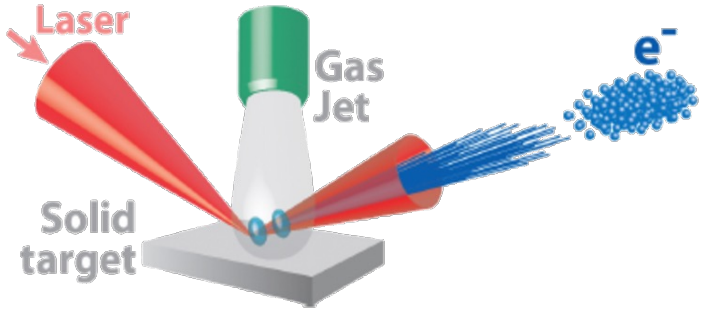

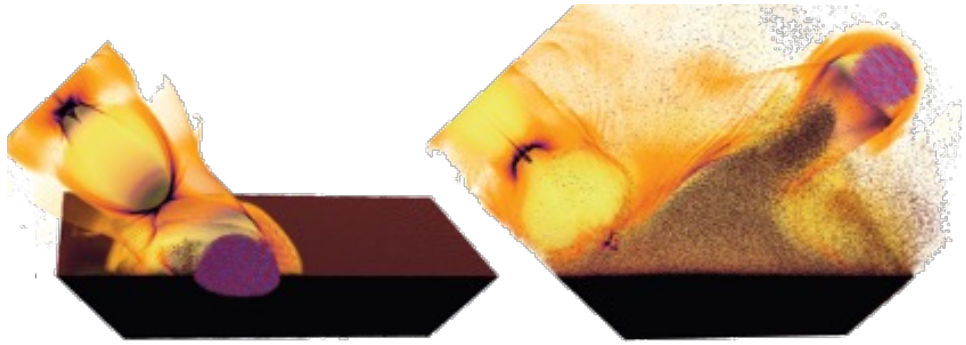

Strong-field QED is relevant for extreme astrophysical scenarios

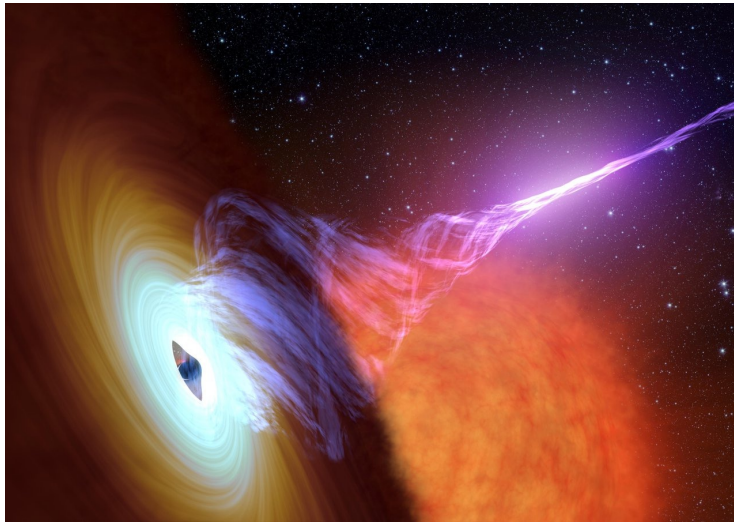

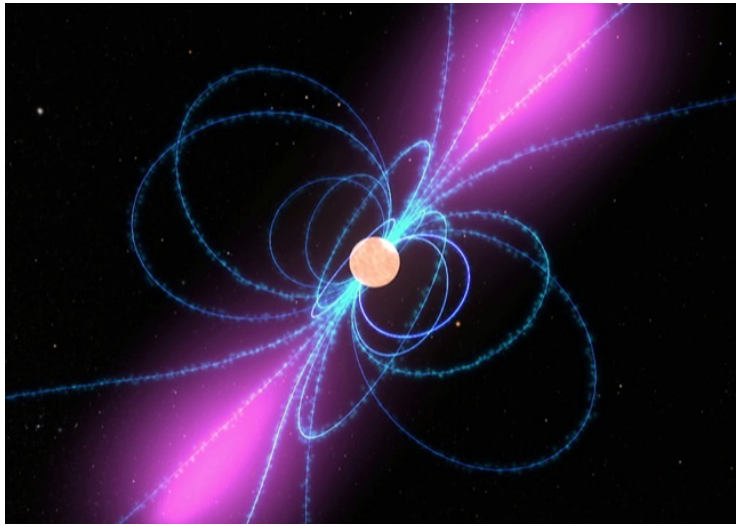

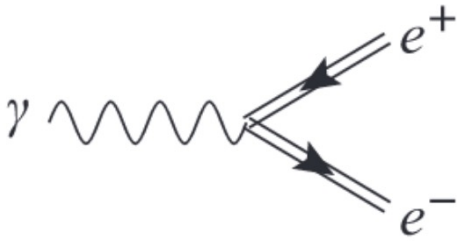

Nonlinear Breit-Wheeler pair production

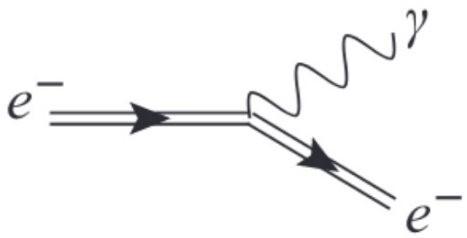

Nonlinear Compton scattering

Very difficult to study on Earth!

When we say “strong-field” we really mean that!

$$E_{Schwinger} = \frac{m^2 c^3}{e\hbar} \approx 1.3 \times 10^{18}\ V/m$$

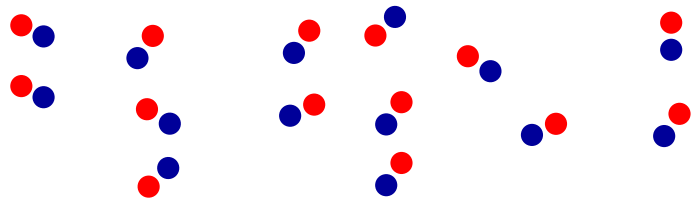

Vacuum breakdown!

Virtual e+ e- pairs become real!

When we say “strong-field” we really mean that!

$$E_{Schwinger} = \frac{m^2 c^3}{e\hbar} \approx 1.3 \times 10^{18}\ V/m$$

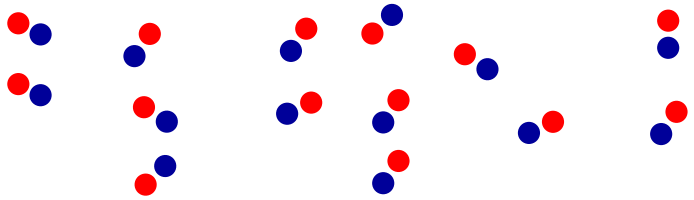

Vacuum breakdown!

Virtual e+ e- pairs become real!

We are **very** far from this “Schwinger field”!

Even the more intense laser is not nearly enough to reach the Schwinger field!

$$E_{Schwinger} = \frac{m^2 c^3}{e\hbar} \approx 1.3 \times 10^{18}\ V/m$$

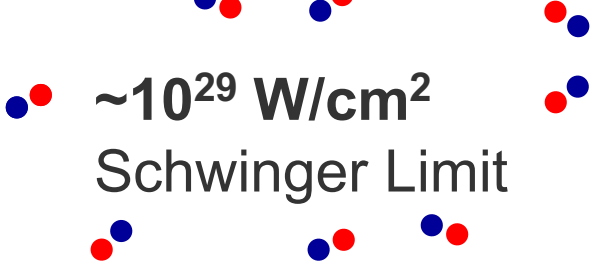

~10^{29} W/cm^2
Schwinger Limit

Even the more intense laser is not nearly enough to reach the Schwinger field!

$$E_{Schwinger} = \frac{m^2 c^3}{e\hbar} \approx 1.3 \times 10^{18}\ V/m$$

Six orders of magnitude gap!

~10^{29} W/cm^2
Schwinger Limit

Current record:
I ~ 10^{23} W/cm^2

Jin Woo Yoon & al. Optica 8, 5, 2019

Even the more intense laser is not nearly enough to reach the Schwinger field!

$$E_{Schwinger} = \frac{m^2 c^3}{e\hbar} \approx 1.3 \times 10^{18}\ V/m$$

Six orders of magnitude gap!

~10^{29} W/cm^2
Schwinger Limit

Current record:
I ~ 10^{23} W/cm^2

Jin Woo Yoon & al. Optica 8, 5, 2019

However !

Approaching the Schwinger field is easier in the reference frame of relativistic particles

$$\gamma = \frac{1}{\sqrt{1 - v^2/c^2}}$$

Approaching the Schwinger field is easier in the reference frame of relativistic particles

$$\gamma = \frac{1}{\sqrt{1 - v^2/c^2}}$$

$$v \to c \quad \Longrightarrow \quad E' \approx 2\gamma E$$

$$\chi = E'/E_{Schwinger}$$

$$\chi = \frac{|p_\mu F^{\mu\nu}|}{m_e E_{Schwinger}}$$

With particles:
I > 10^{23} W/cm^2 + 100 MeV e$^-$
just to reach $\chi \sim 1$

Approaching the Schwinger field is easier in the reference frame of relativistic particles

$$\gamma = \frac{1}{\sqrt{1 - v^2/c^2}}$$

$$v \to c \quad \Longrightarrow \quad E' \approx 2\gamma E$$

$$\chi = E'/E_{Schwinger}$$

$$\chi = \frac{|p_\mu F^{\mu\nu}|}{m_e E_{Schwinger}}$$

With particles:
I > 10^{23} W/cm^2 + 100 MeV e$^-$
just to reach $\chi \sim 1$

Experiments have probed
at most $\chi \sim 1$

The QED component of PICSAR provides **WarpX** with a Monte Carlo module to simulate QED processes

The QED component of PICSAR provides **WarpX** with a Monte Carlo module to simulate QED processes

P-I-C-S-A-R
Particle-In-Cell Scalable Application Resource

Includes **the most relevant QED** processes for PIC codes:
synchrotron emission, nonlinear Breit-Wheeler pair production

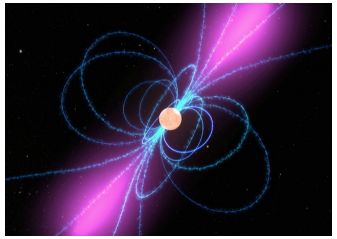

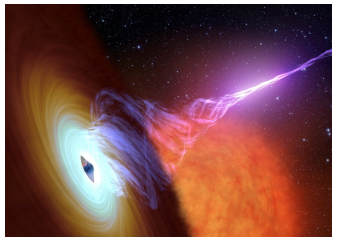

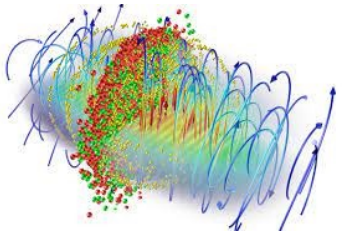

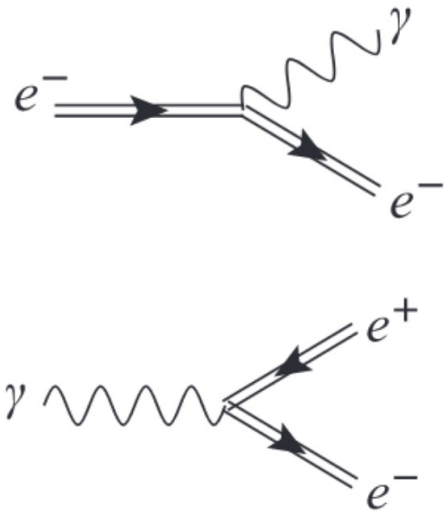

The QED component of PICSAR provides WarpX with a Monte Carlo module to simulate QED processes

P-I-C-S-A-R
Particle-In-Cell Scalable Application Resource

Includes **the most relevant QED** processes for PIC codes:
synchrotron emission, nonlinear Breit-Wheeler pair production

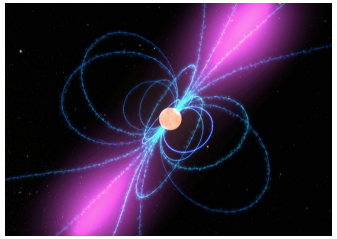

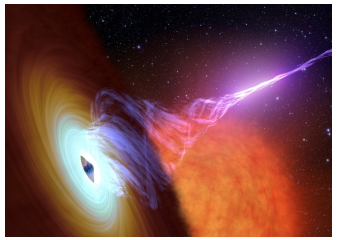

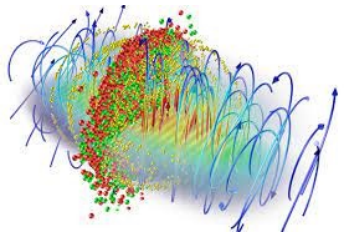

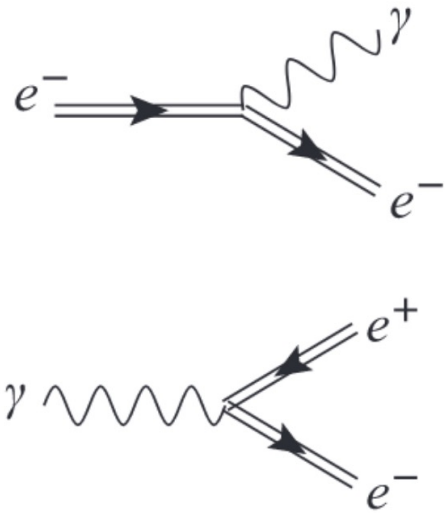

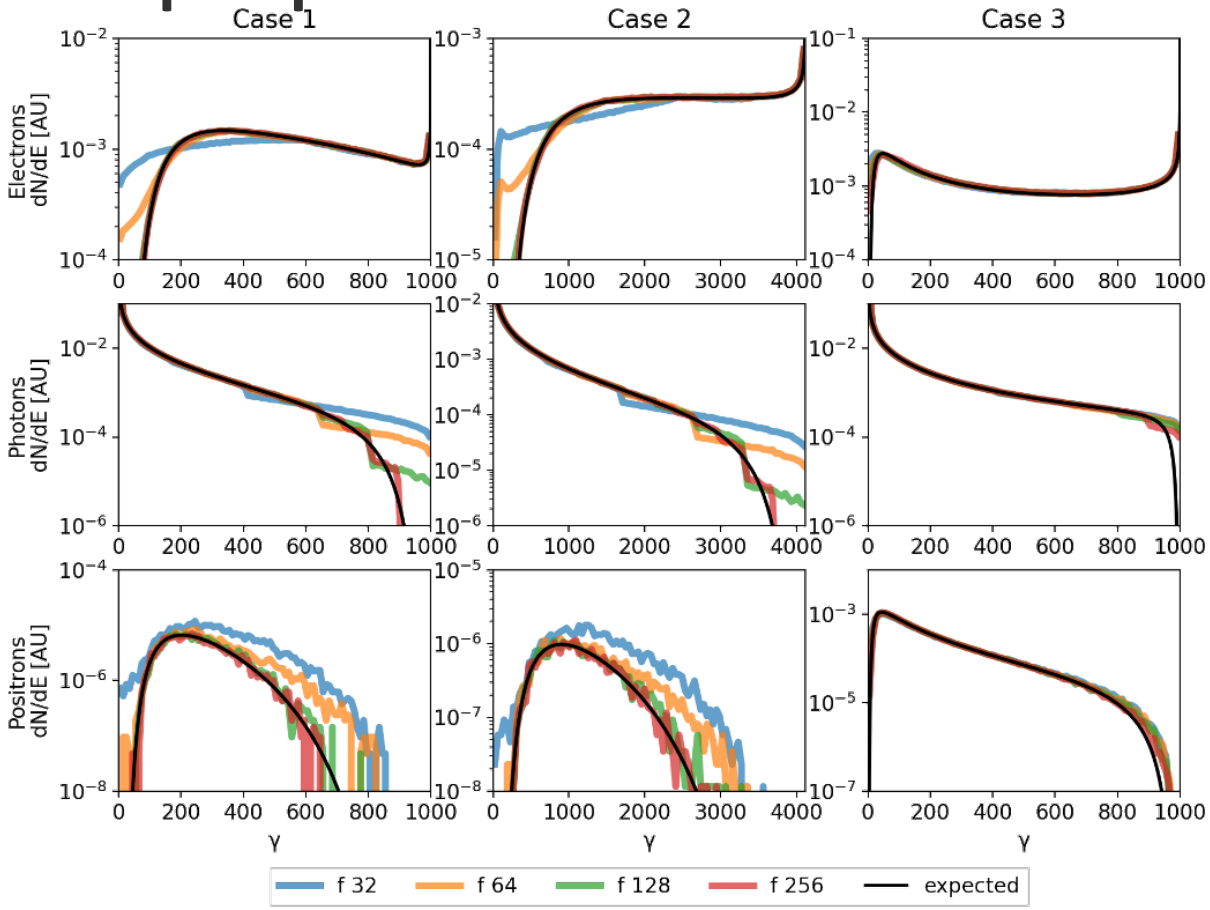

Has been fully benchmarked in one of our recent papers:

L. Fedeli et al. New J. Phys. 24 025009, 2022

The QED component of PICSAR provides **WarpX** with a Monte Carlo module to simulate QED processes

P-I-C-S-A-R
Particle-In-Cell Scalable Application Resource

Includes **the most relevant QED** processes for PIC codes:
synchrotron emission, nonlinear Breit-Wheeler pair production

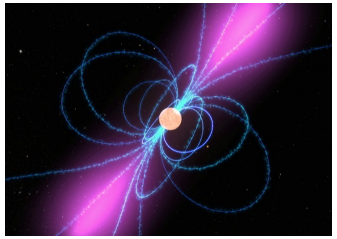

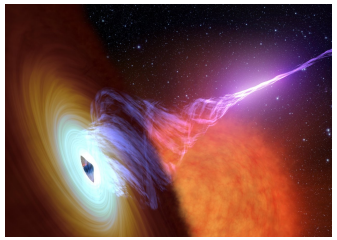

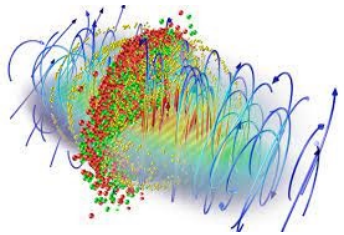

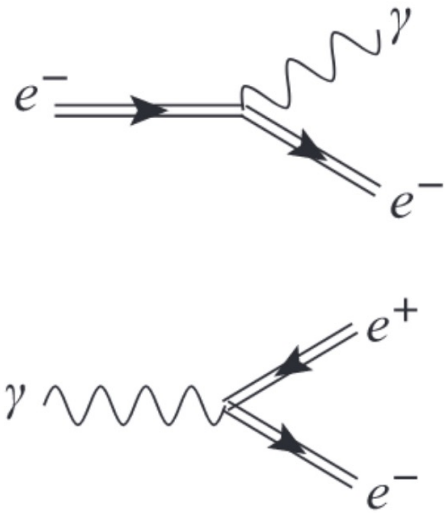

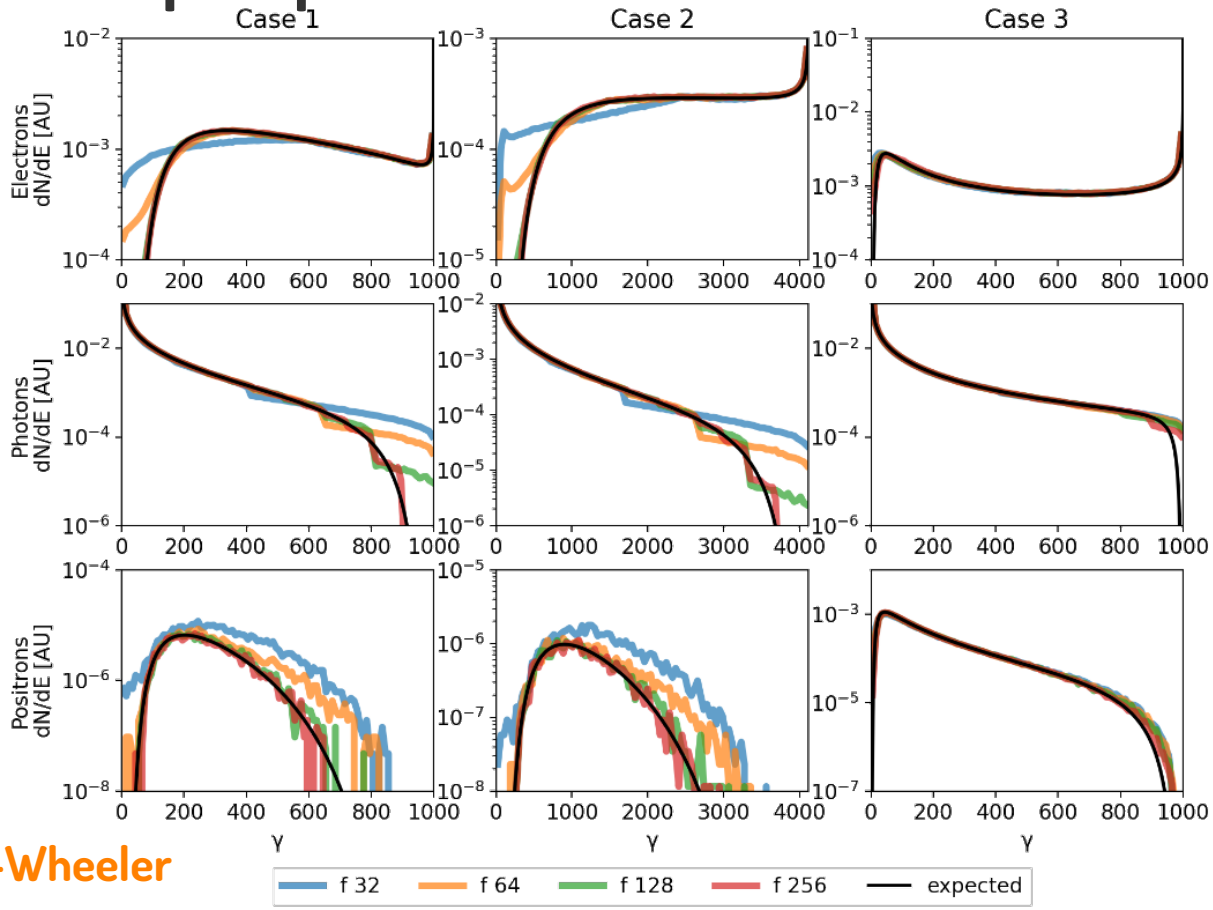

Has been fully benchmarked in one of our recent papers:

L. Fedeli et al. New J. Phys. 24 025009, 2022

We are currently implementing other processes as linear Breit-Wheeler

Our Scheme to boost the intensity of existing ultra-intense lasers

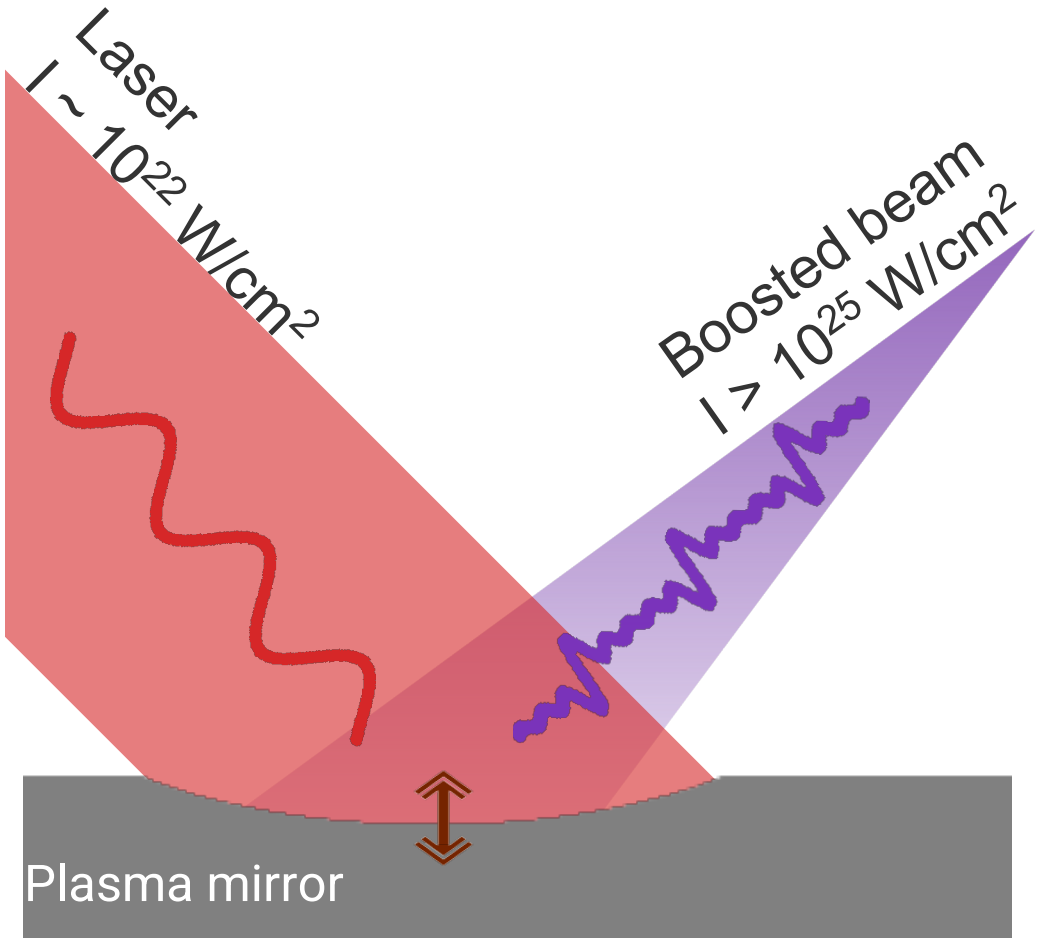

We can boost a laser beam using a curved plasma mirror

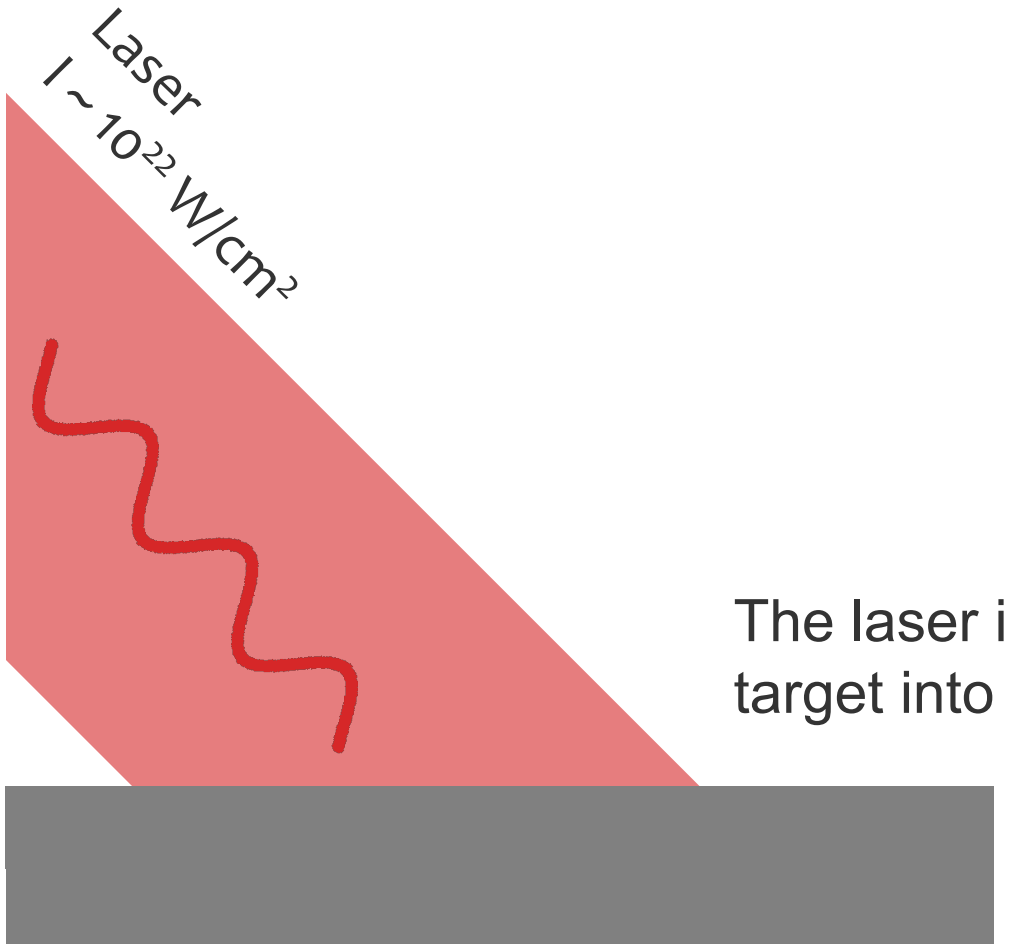

The laser instantly turns the target into a plasma

H.Vincenti. Phys. Rev. Lett. 123, 105001, 2019

We can boost a laser beam using a curved plasma mirror

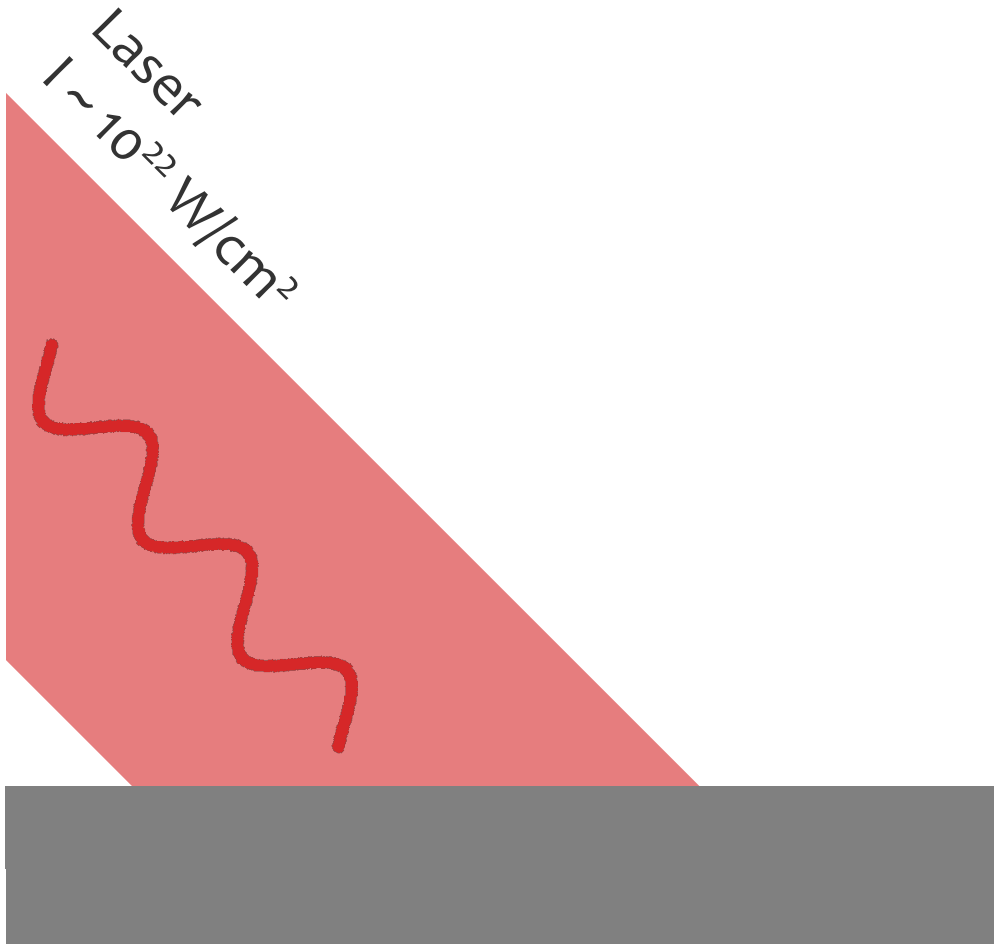

H.Vincenti. Phys. Rev. Lett. 123, 105001, 2019

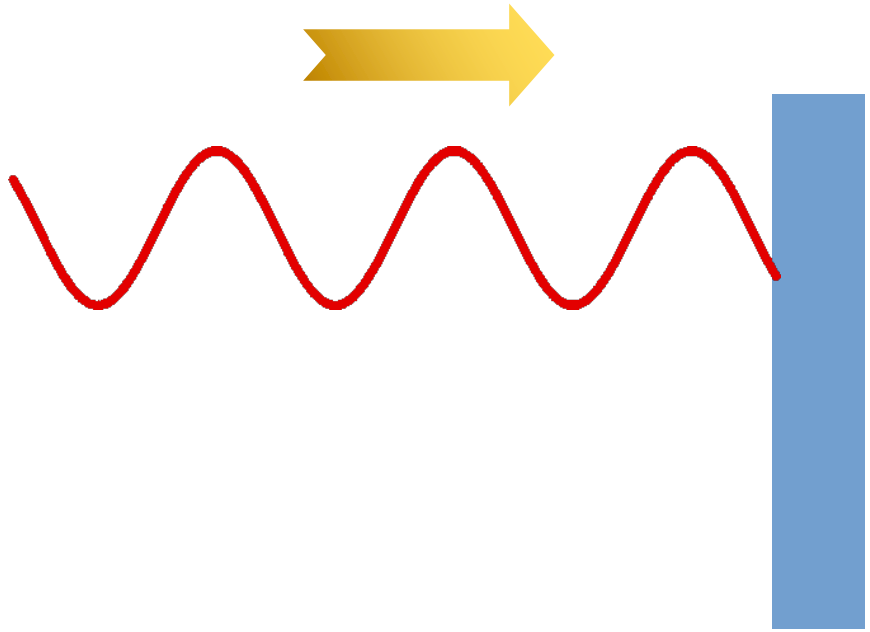

The target behaves like a mirror

We can boost a laser beam using a curved plasma mirror

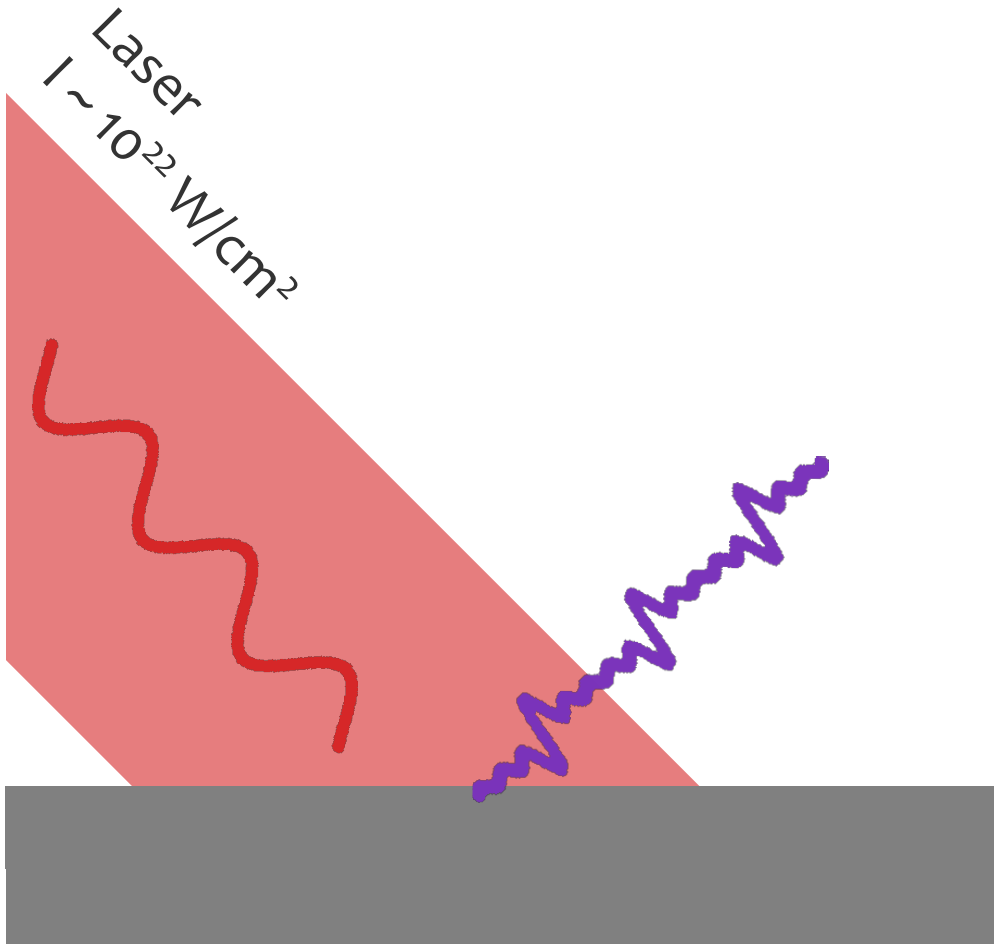

H.Vincenti. Phys. Rev. Lett. 123, 105001, 2019

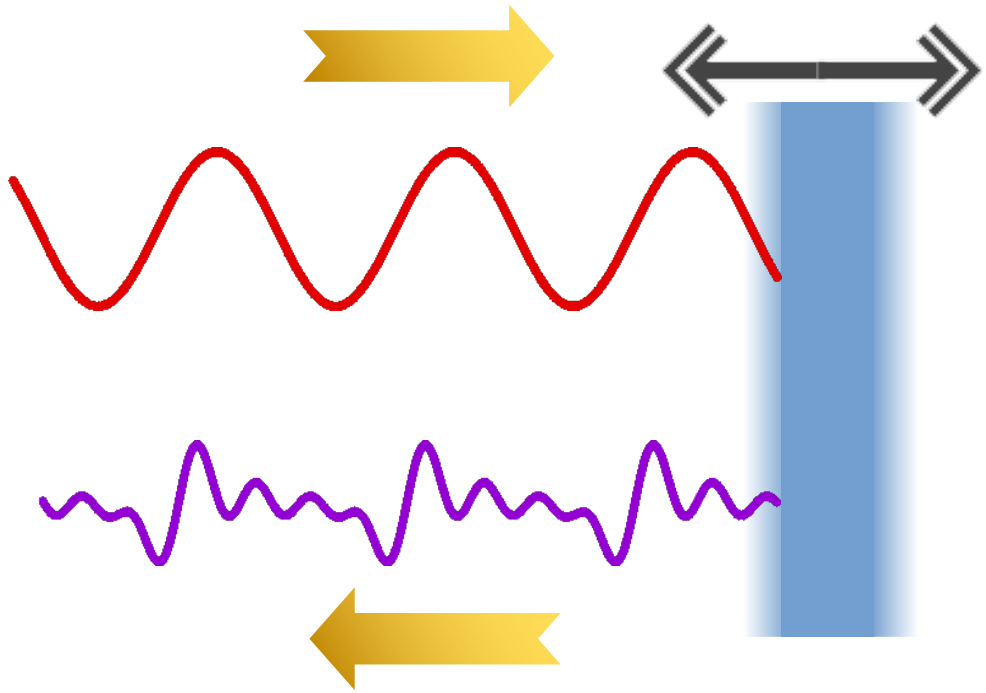

Actually a mirror oscillating at relativistic velocities (very well known in experiments)

We can boost a laser beam using a curved plasma mirror

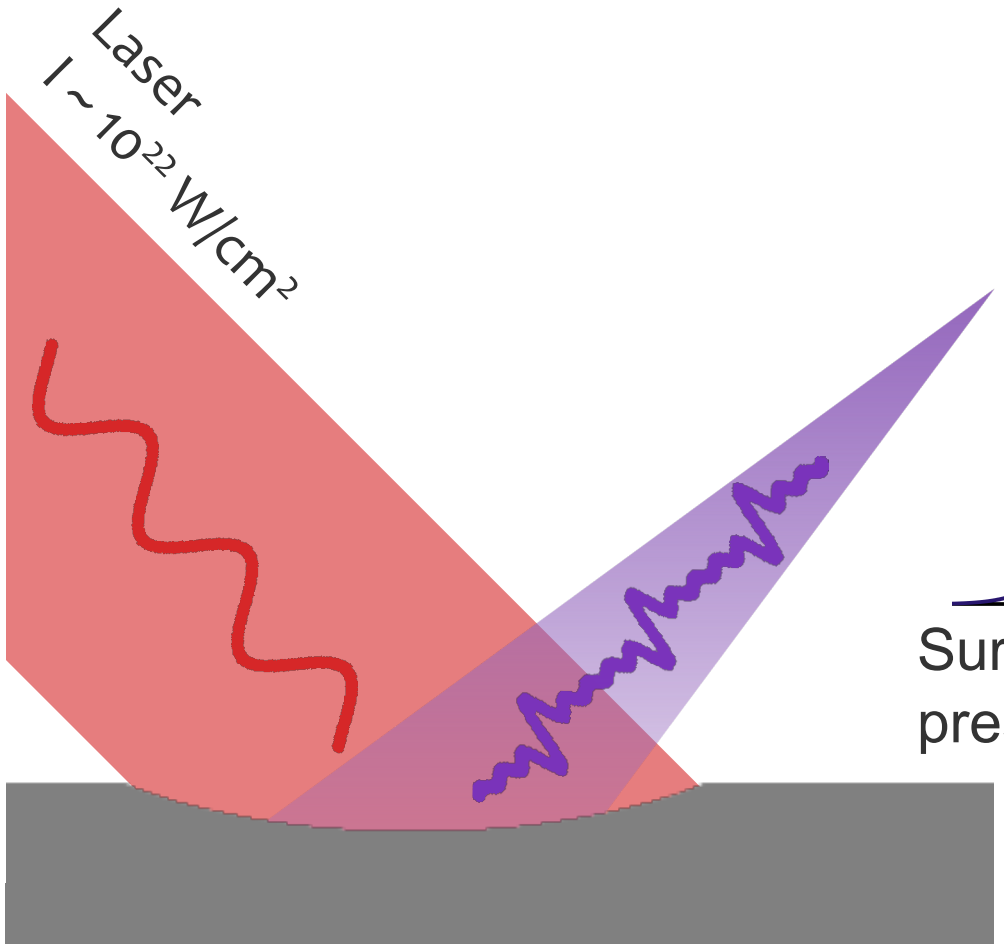

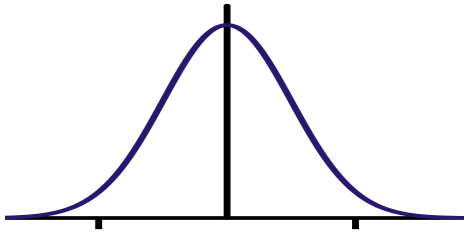

Surface curved by radiation pressure

Laser intensity can be boosted by 3 orders of magnitude !

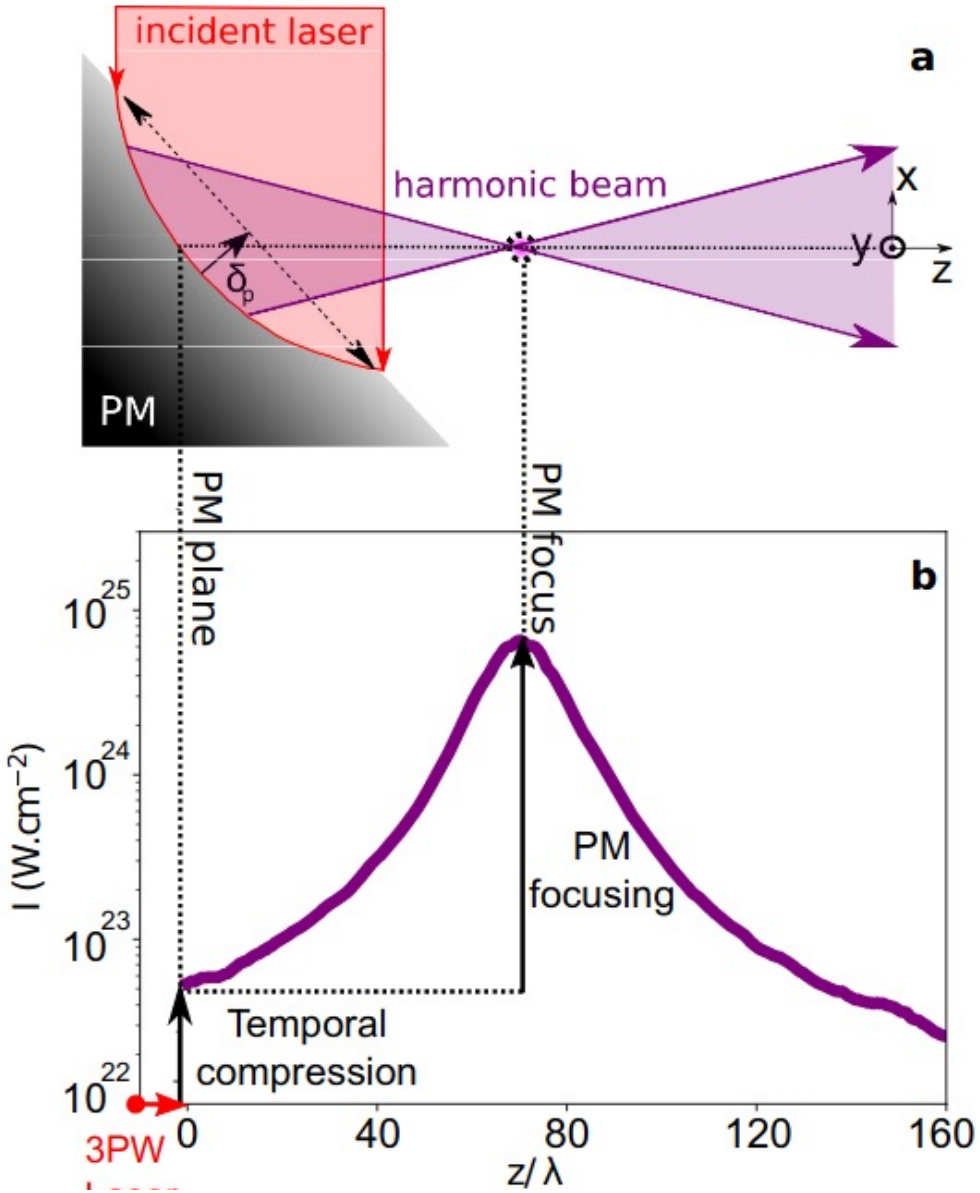

Laser intensity can be boosted by 3 orders of magnitude !

(and we are working to do **significantly better**)

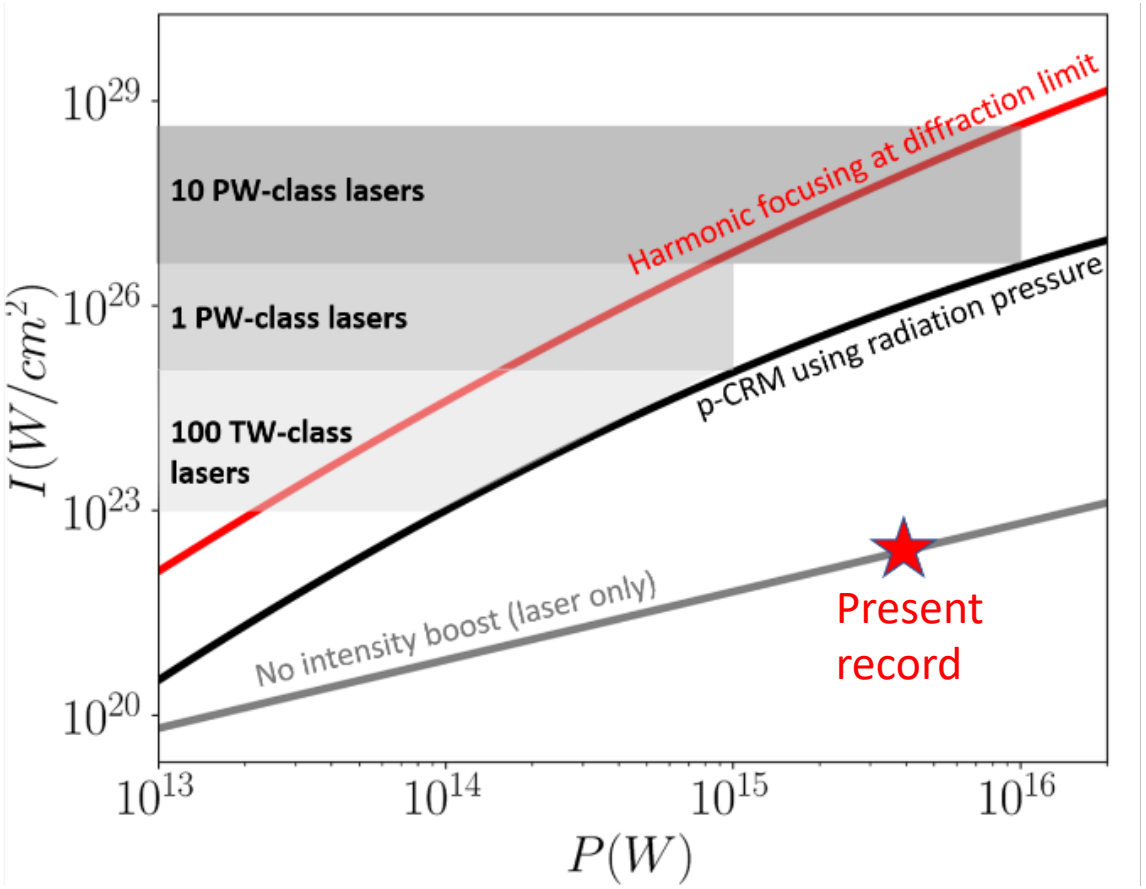

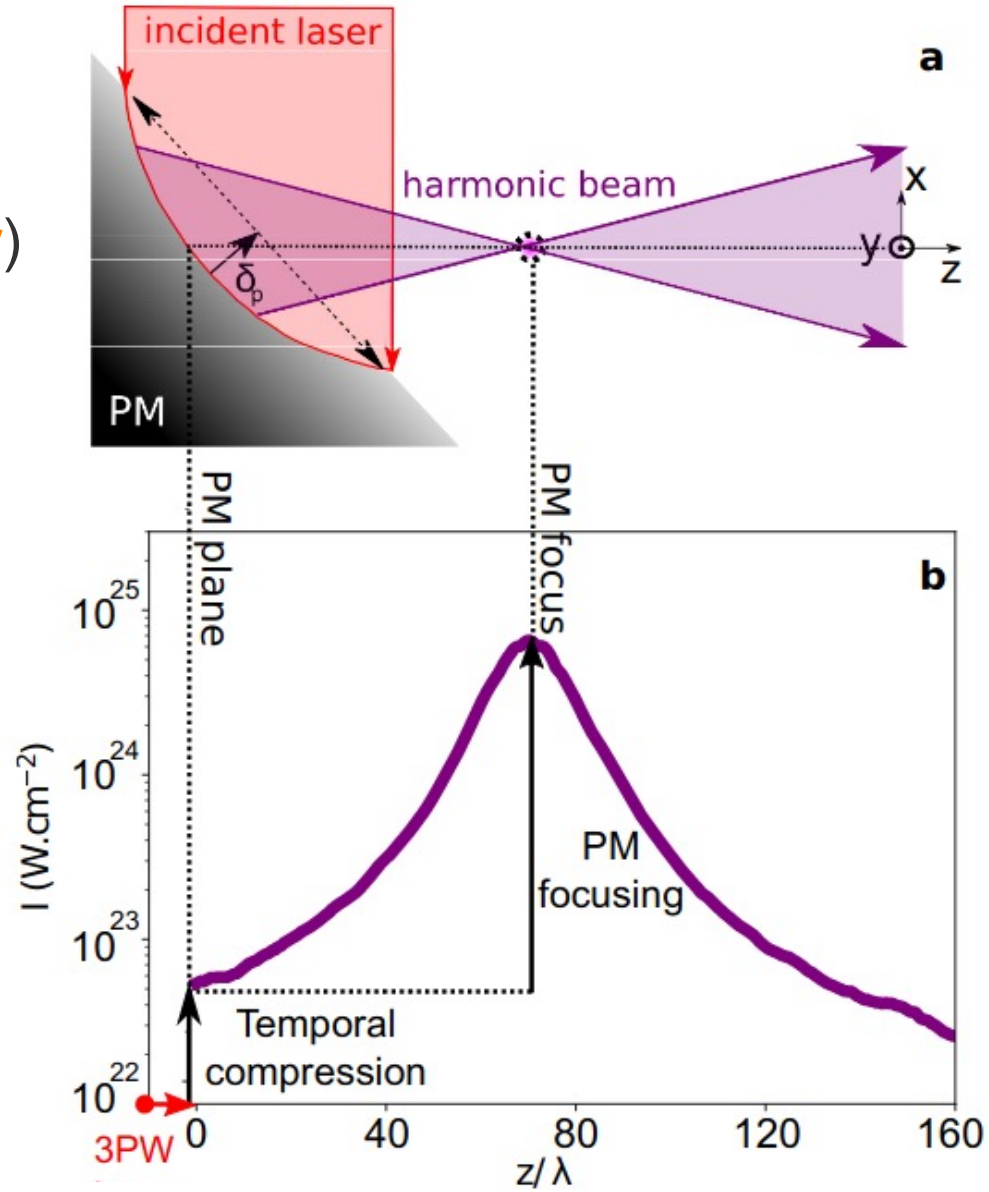

F.Quéré & H.Vincenti HPLSE, 2021

Now we want to design a QED experiment on Apollon facility

Now we want to design a QED experiment on Apollon facility

Lead by Adrien Leblanc's team, UPX

Now we want to design a QED experiment on Apollon facility

loa

Lead by Adrien Leblanc's team, UPX

Laser-Plasma Accelerator

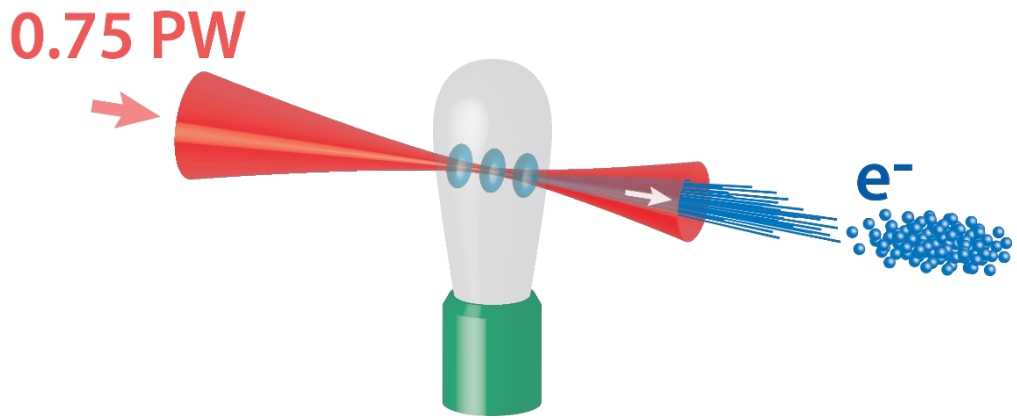

Now we want to design a QED experiment on Apollon facility

Lead by Adrien Leblanc's team, UPX

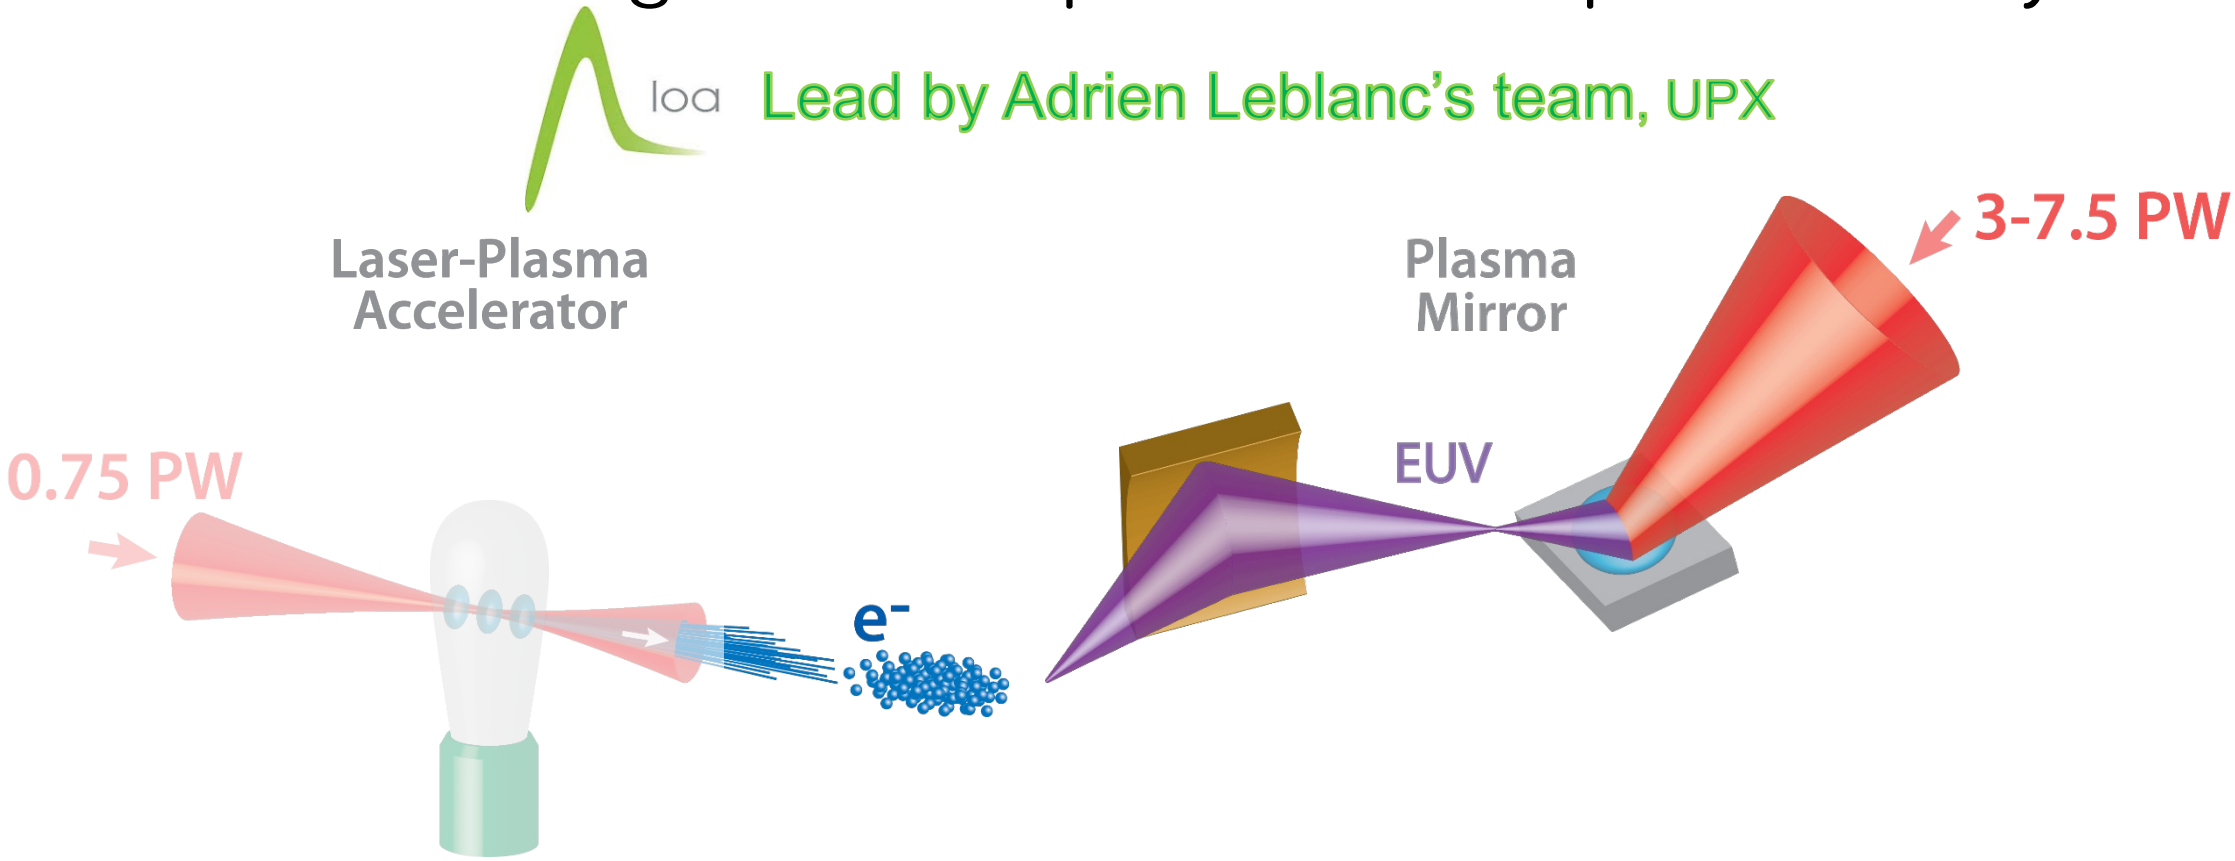

Now we want to design a QED experiment on Apollon facility

Lead by Adrien Leblanc's team, UPX

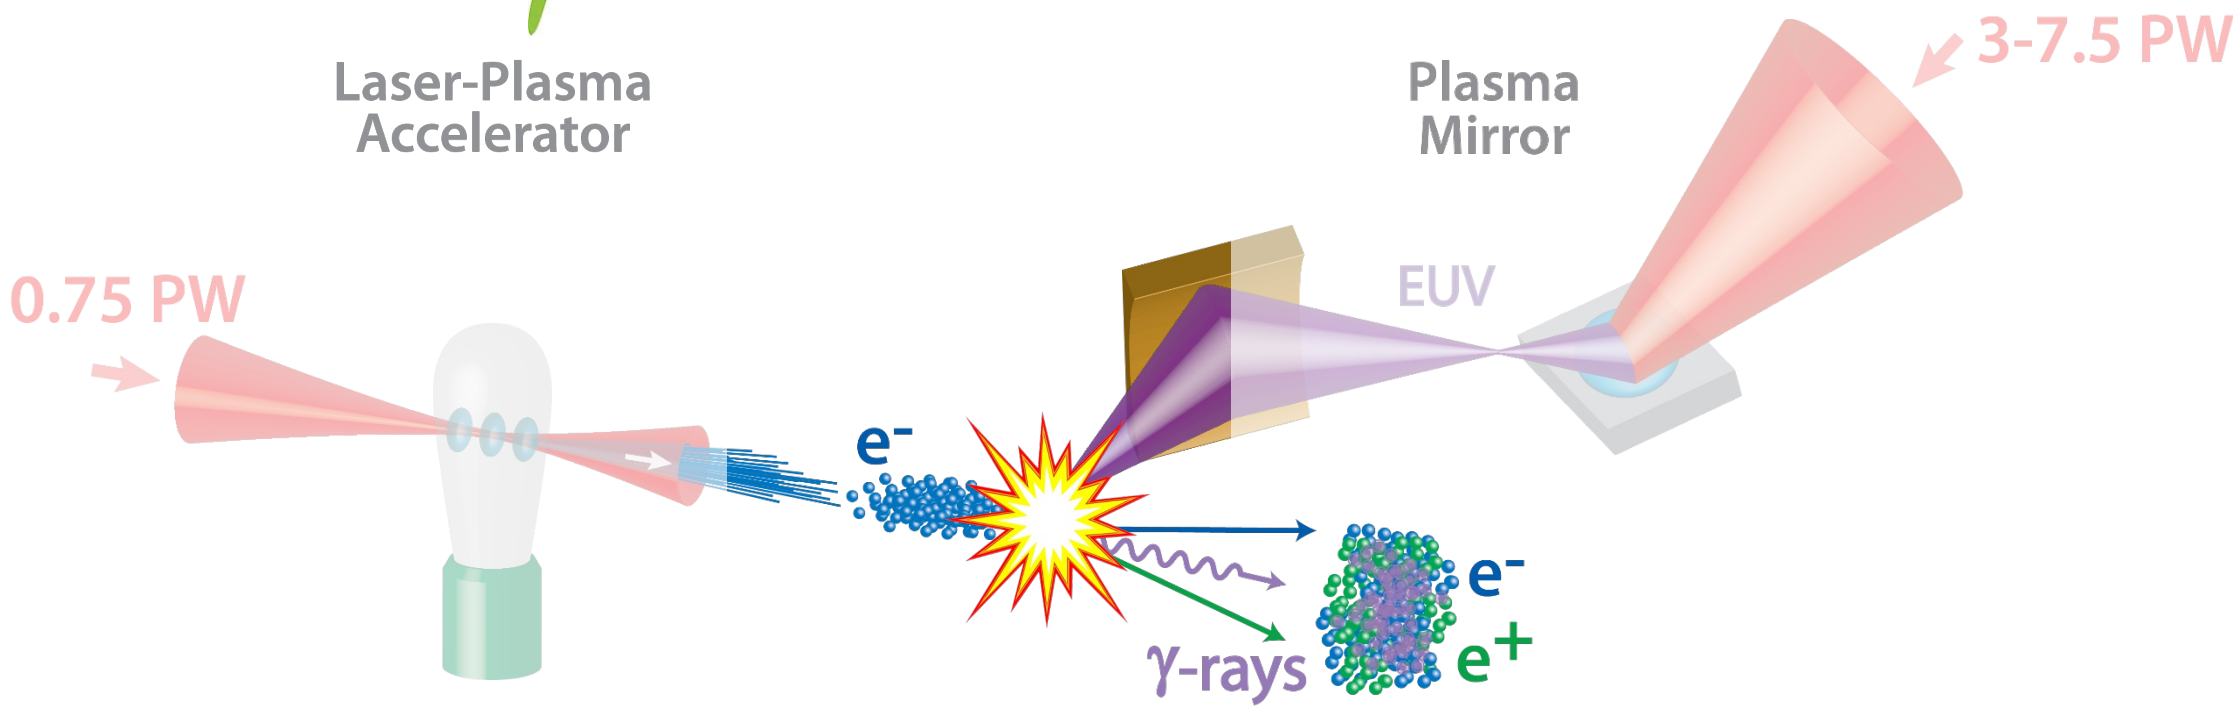

Now we want to design a QED experiment on Apollon facility

Lead by Adrien Leblanc's team, UPX

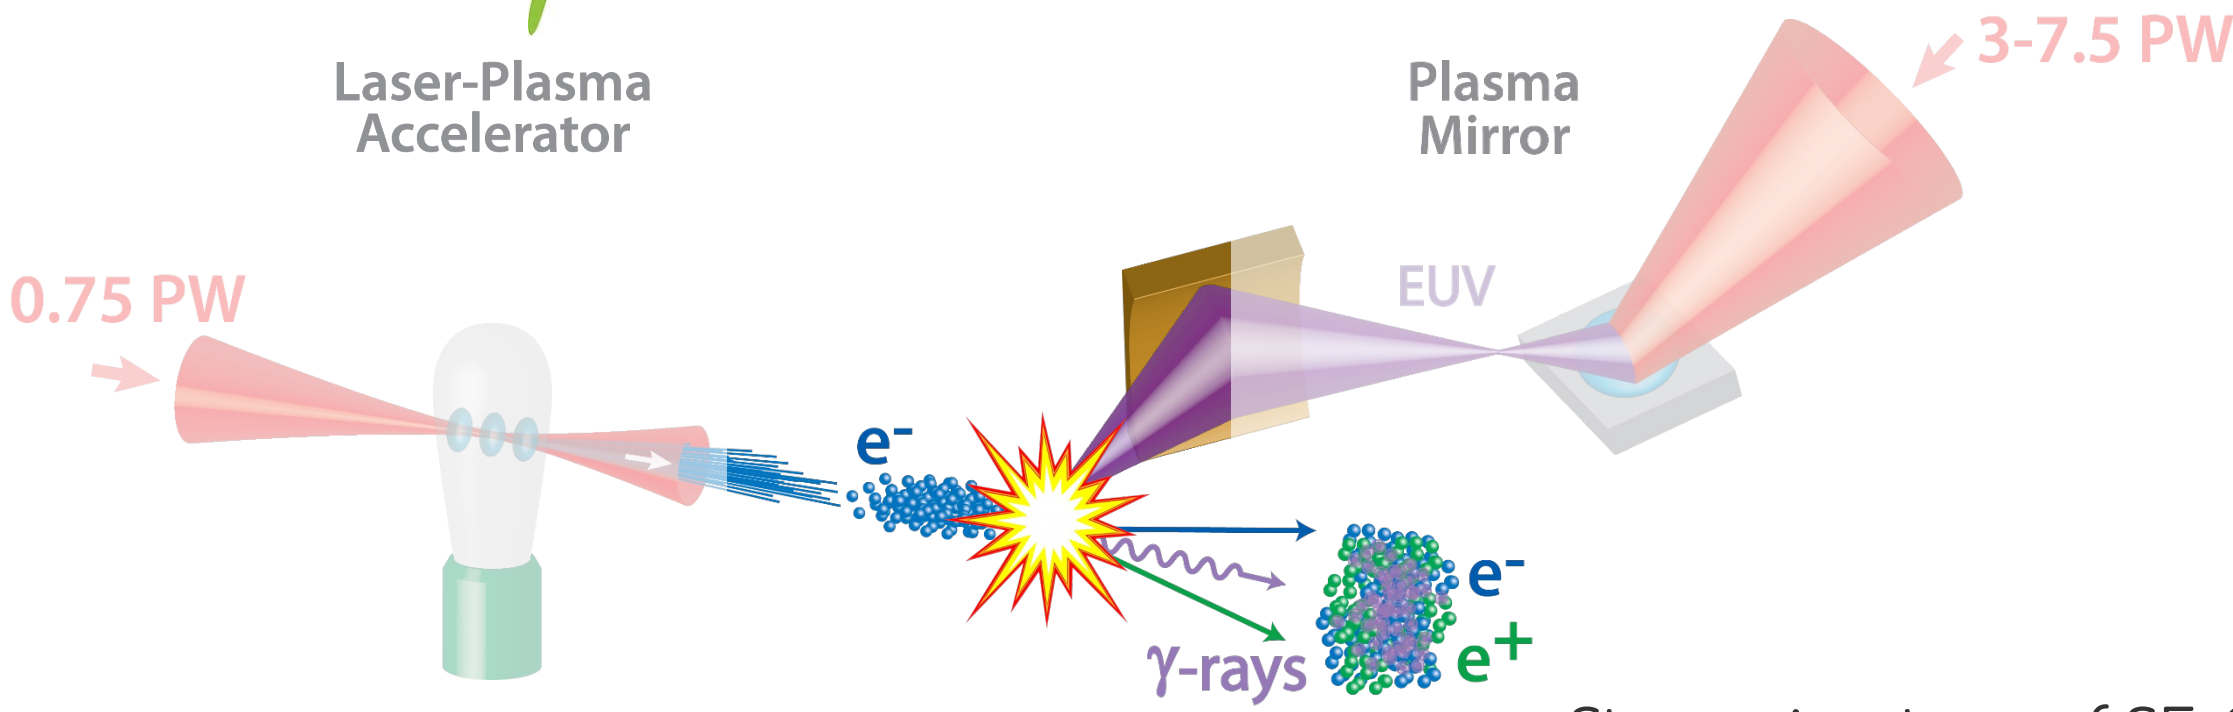

Strong signatures of SF-QED !

Our main research interests

Using ultra-intense lasers to study strong-field Quantum Electrodynamics

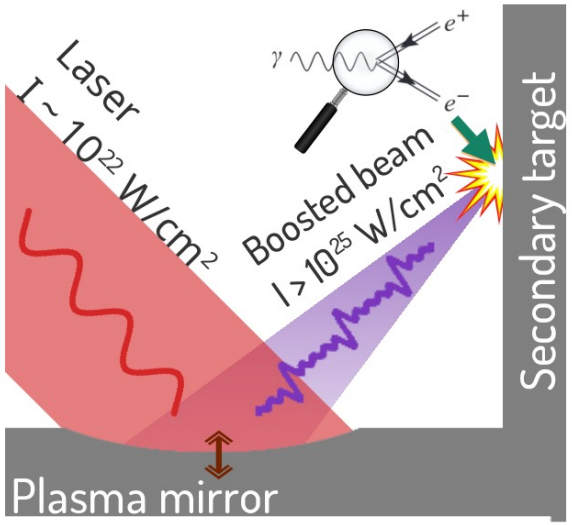

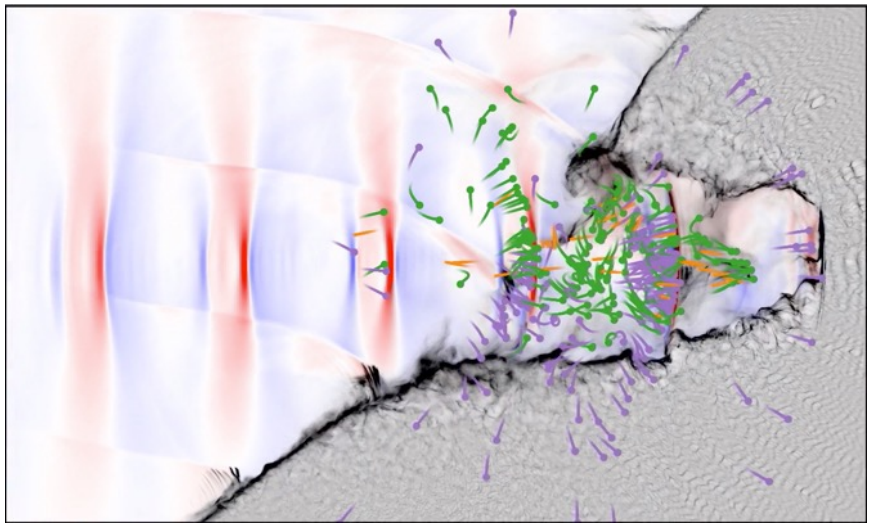

Advanced laser-driven electron sources

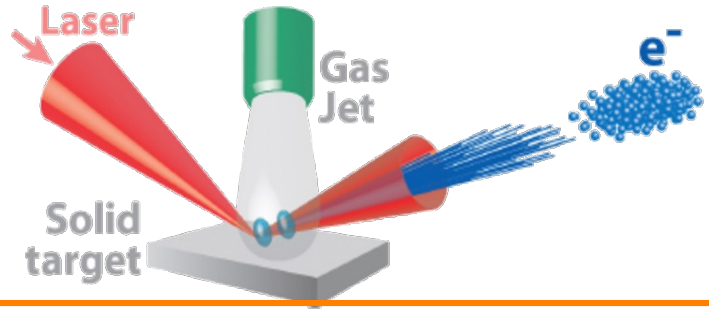

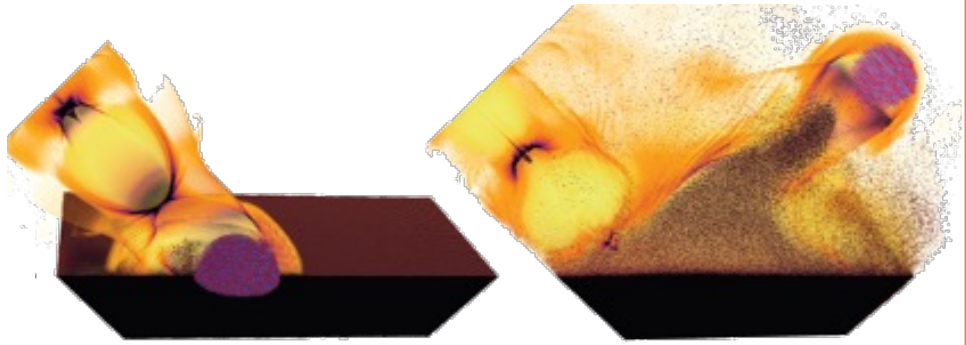

Where the idea of the Hybrid Target comes from ?

Where the idea of the Hybrid Target comes from ?

Electron acceleration in gas

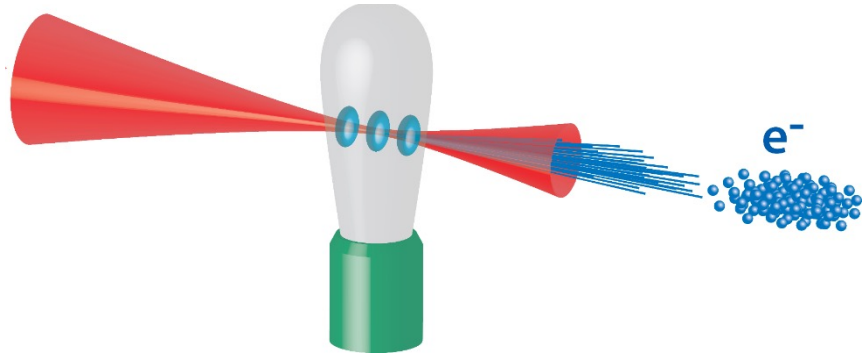

- Low charge : 10s to 100 pC

But

- ✓ High energy : 100s Mev to GeV
- ✓ Low divergence

Where the idea of the Hybrid Target comes from ?

Electron acceleration in gas

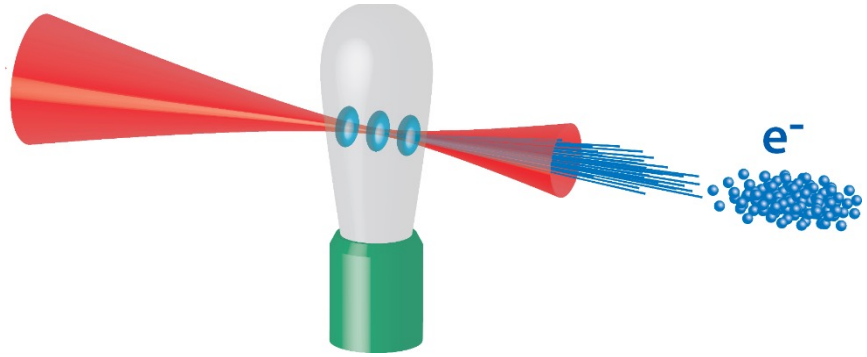

- Low charge : 10s to 100 pC

But

- ✓ High energy : 100s Mev to GeV
- ✓ Low divergence

Electron acceleration with a plasma mirror

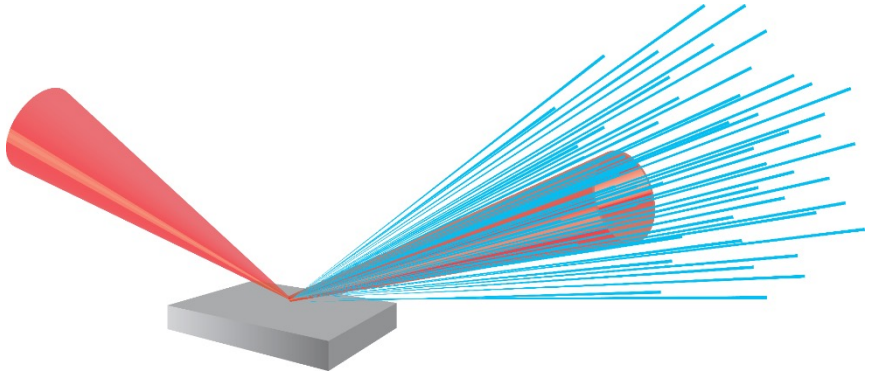

- ✓ High charge

But

- Low energy : 10 MeV
- High divergence

Where the idea of the Hybrid Target comes from ?

A two-step process :

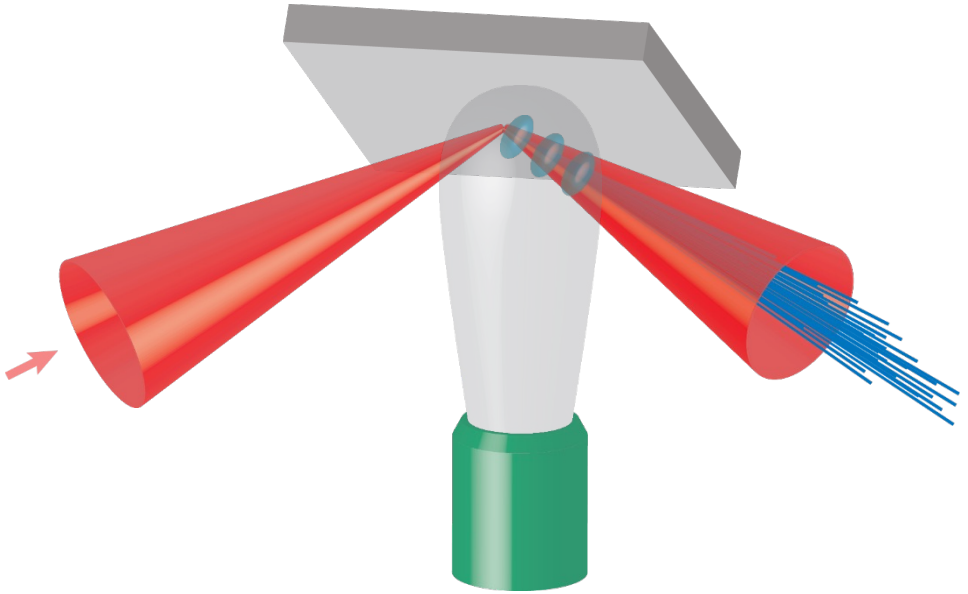

Where the idea of the Hybrid Target comes from ?

A two-step process :

1) Injection from the solid target

2) Acceleration in the gas

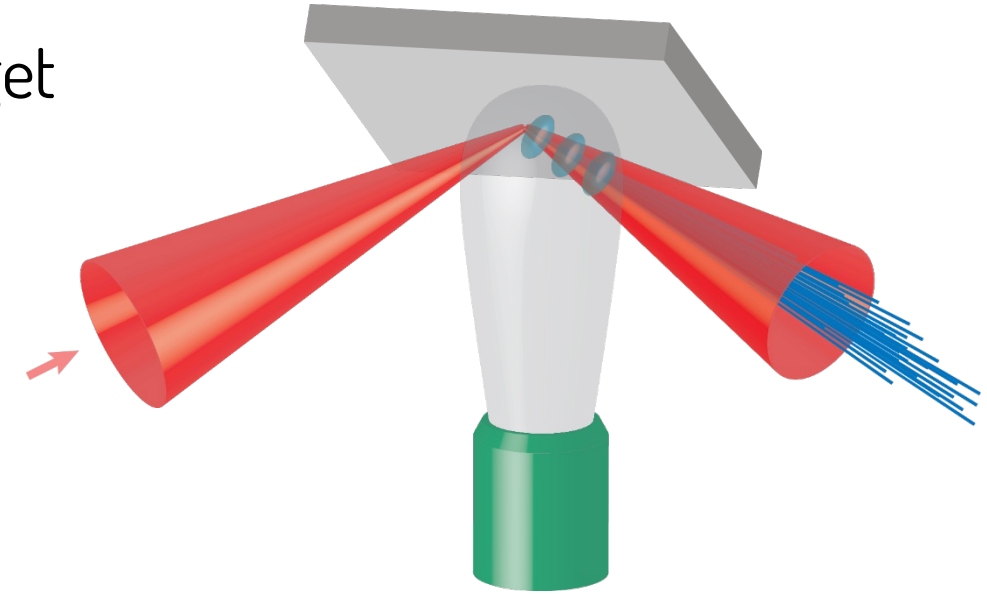

Where the idea of the Hybrid Target comes from ?

A two-step process :

1) Injection from the solid target

2) Acceleration in the gas

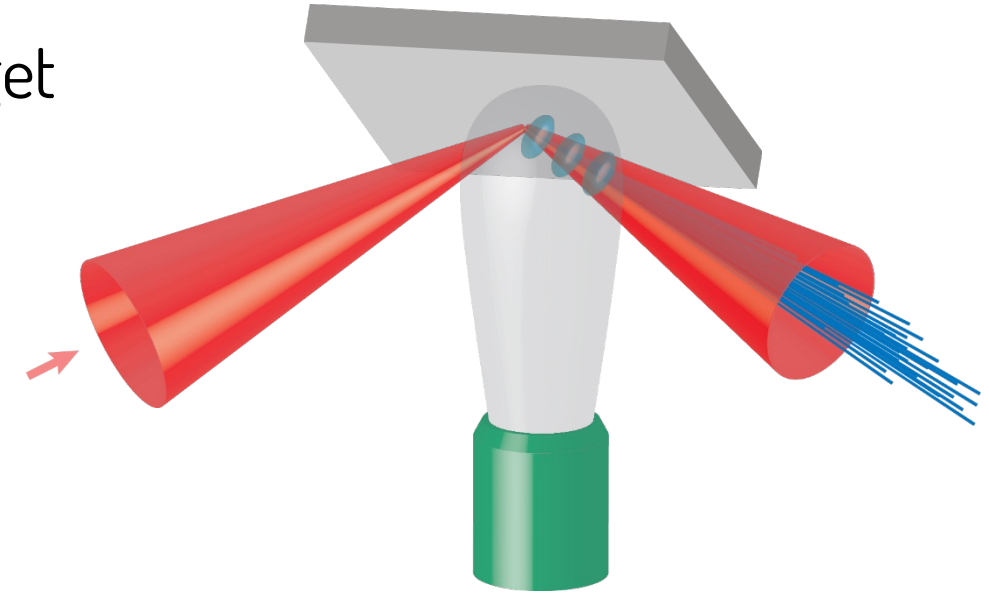

It should provide :

- ✓ A high charge from the high density of the solid target
- ✓ A high quality since the injection is localized at solid surface

What does it gives in simulation ?

A movie from our 3D simulations

A movie from our 3D simulations

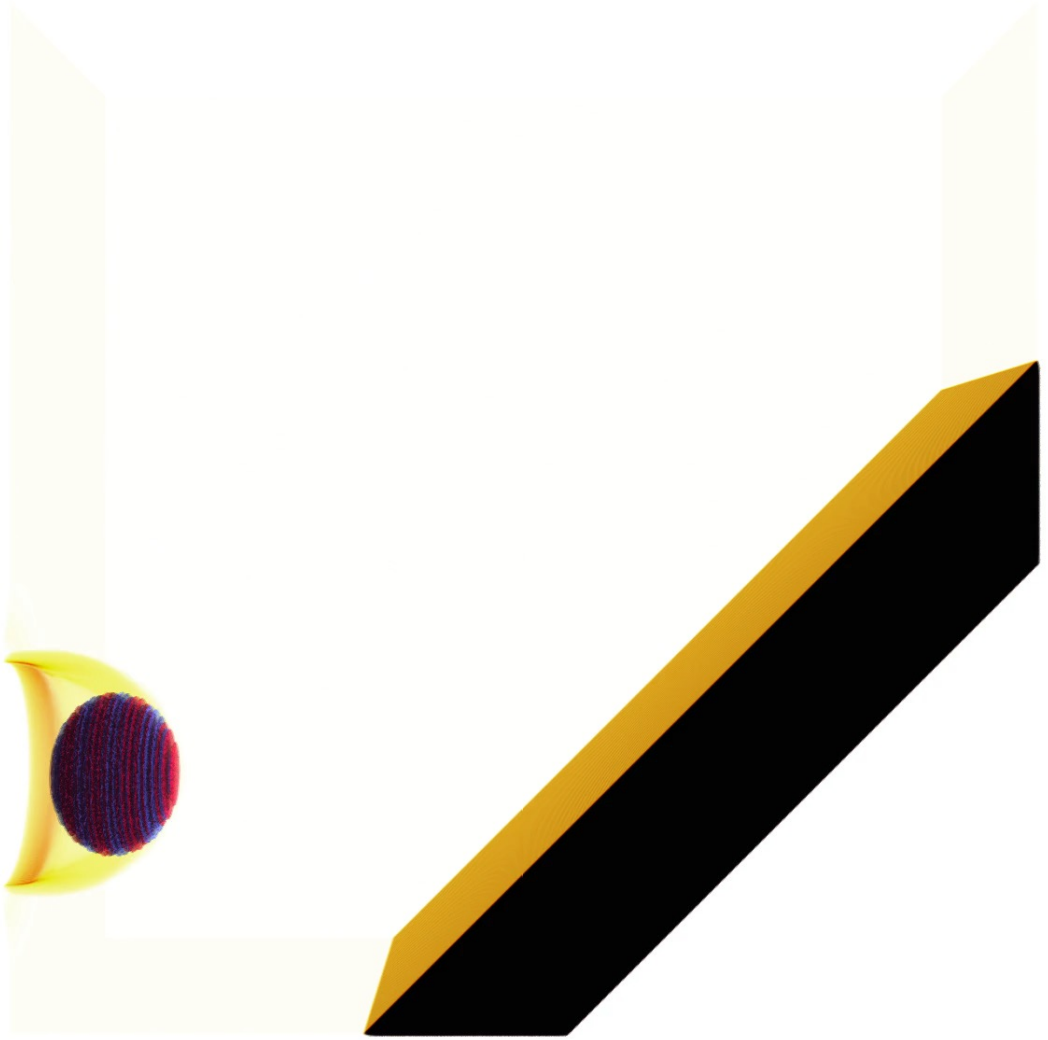

A movie from our 3D simulations

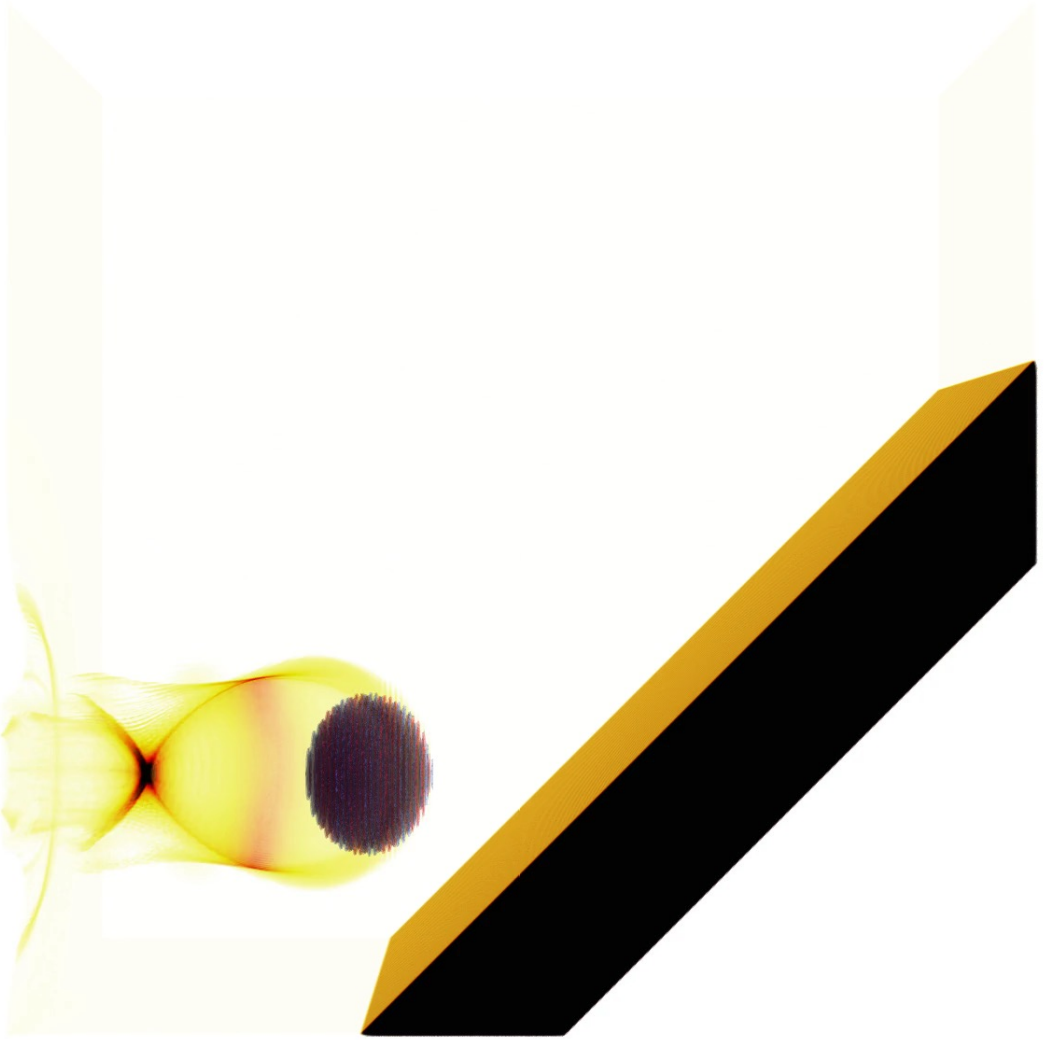

A movie from our 3D simulations

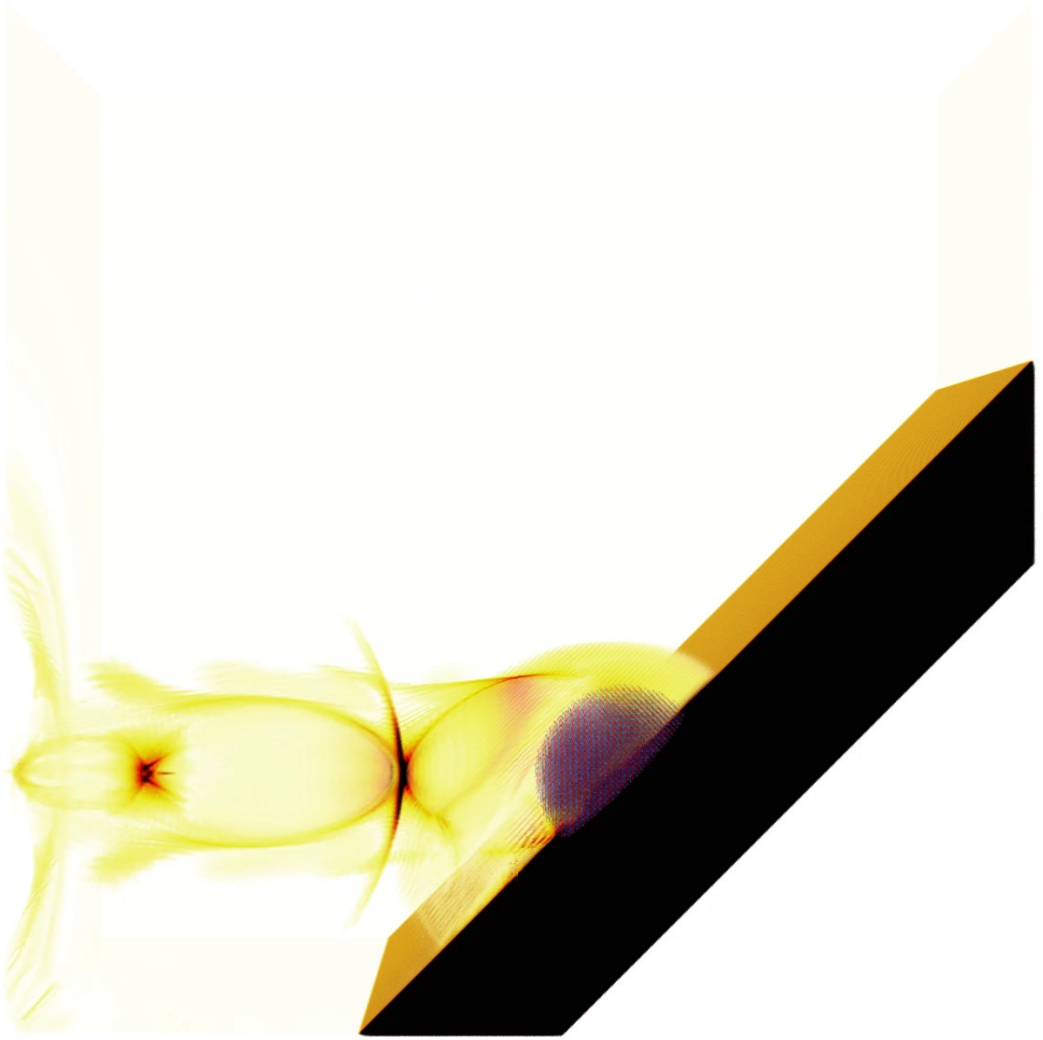

A movie from our 3D simulations

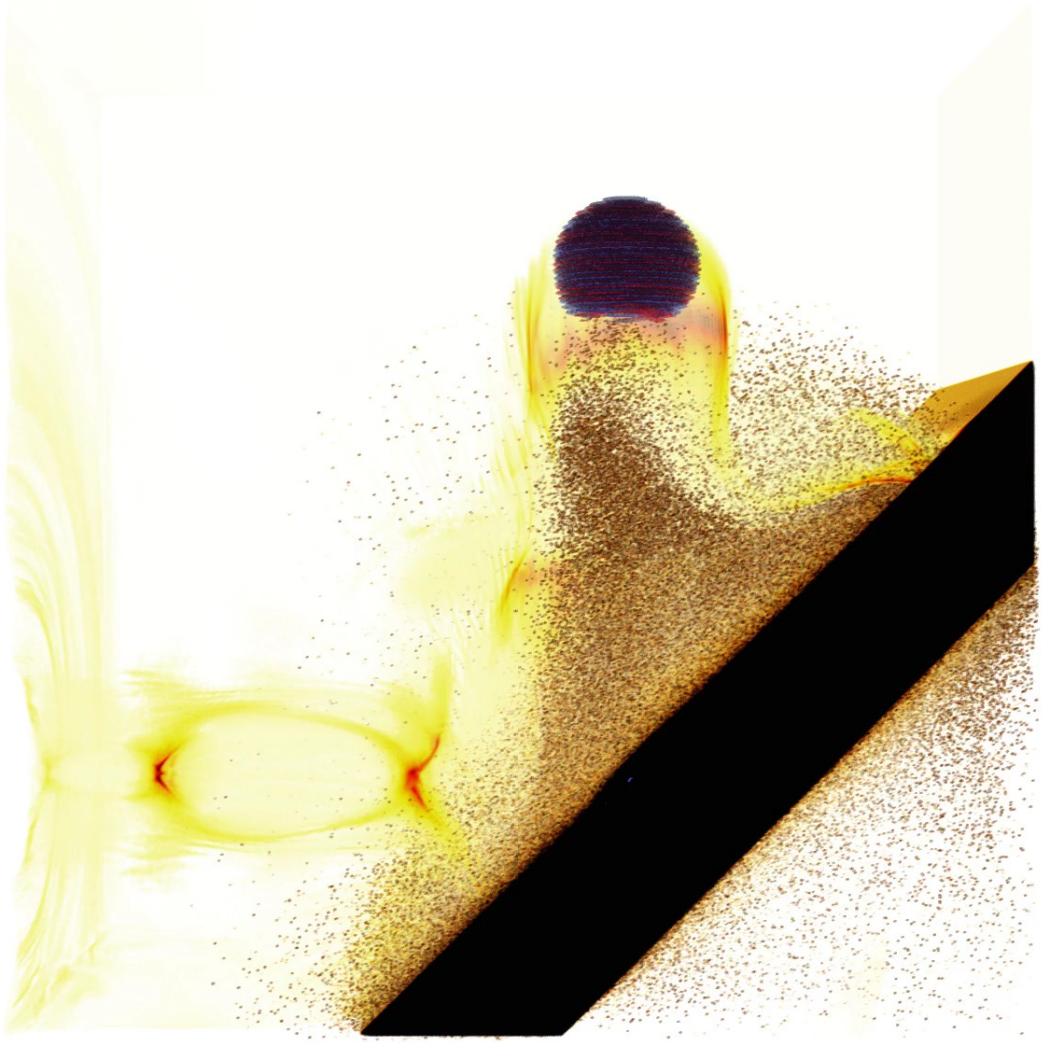

A movie from our 3D simulations

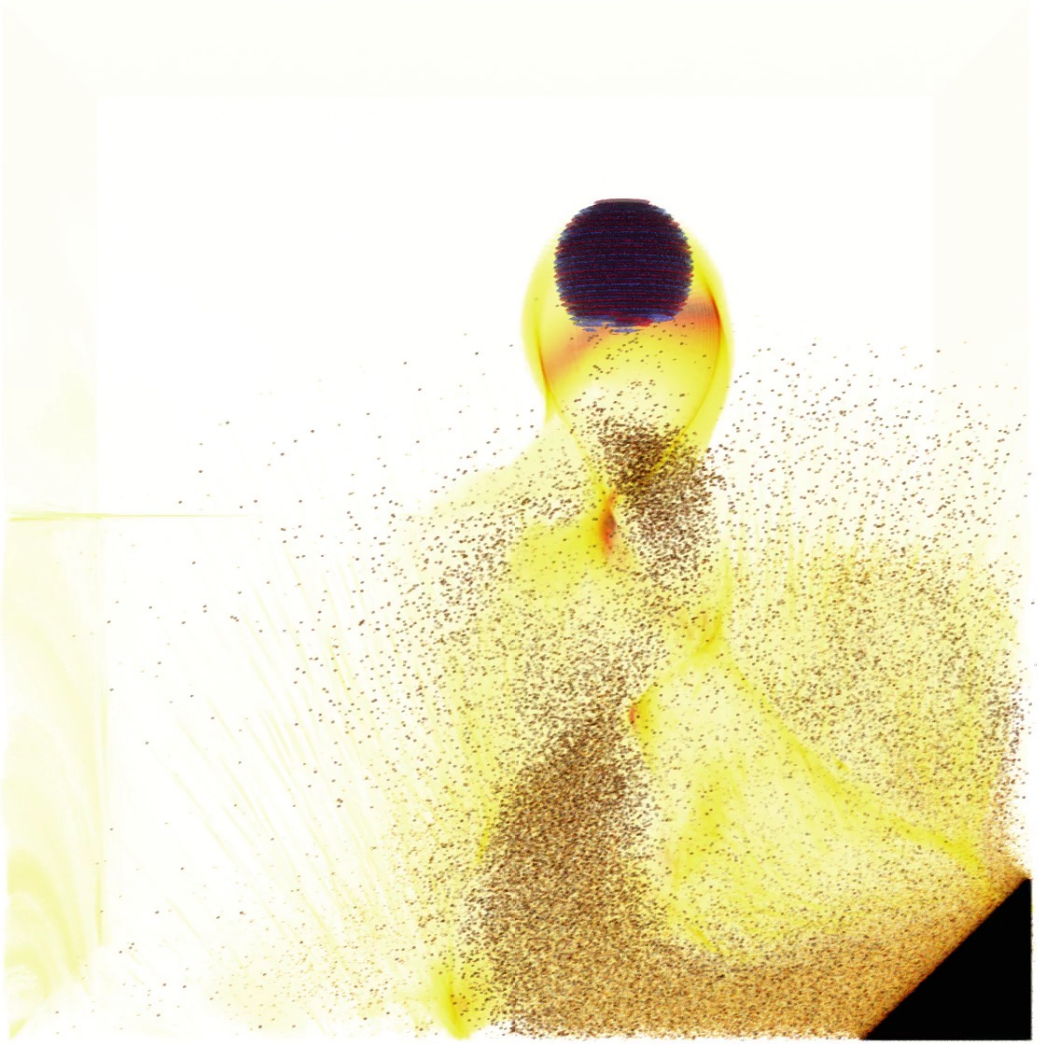

A movie from our 3D simulations

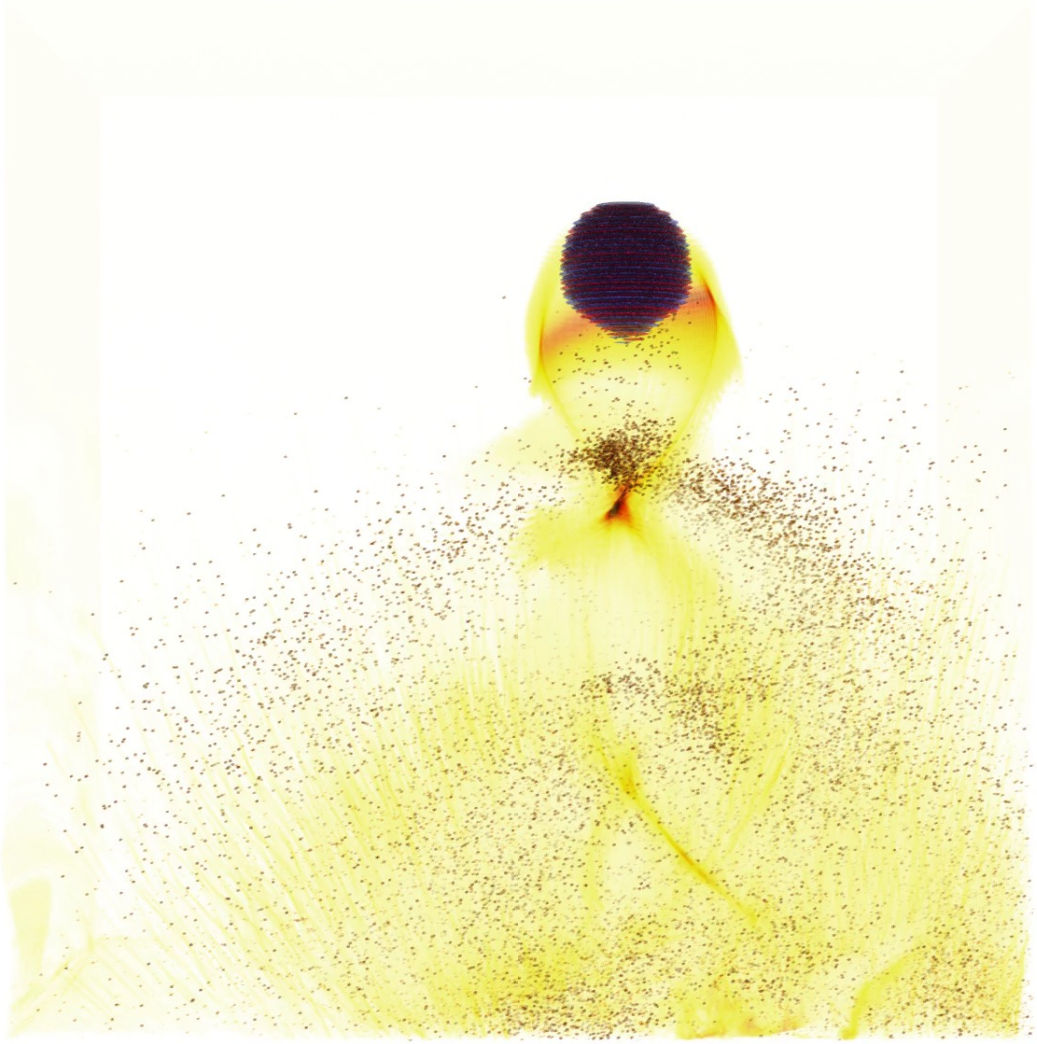

A movie from our 3D simulations

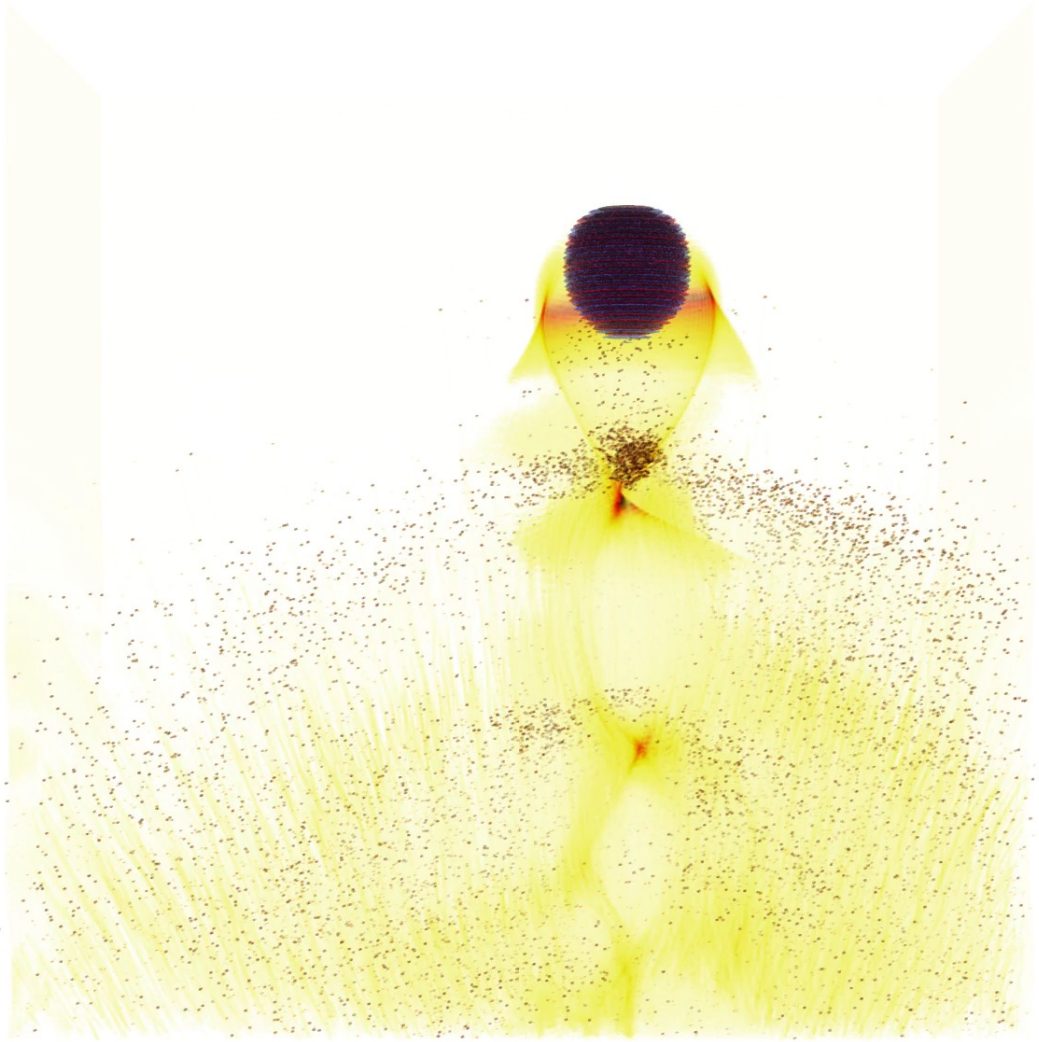

A movie from our 3D simulations

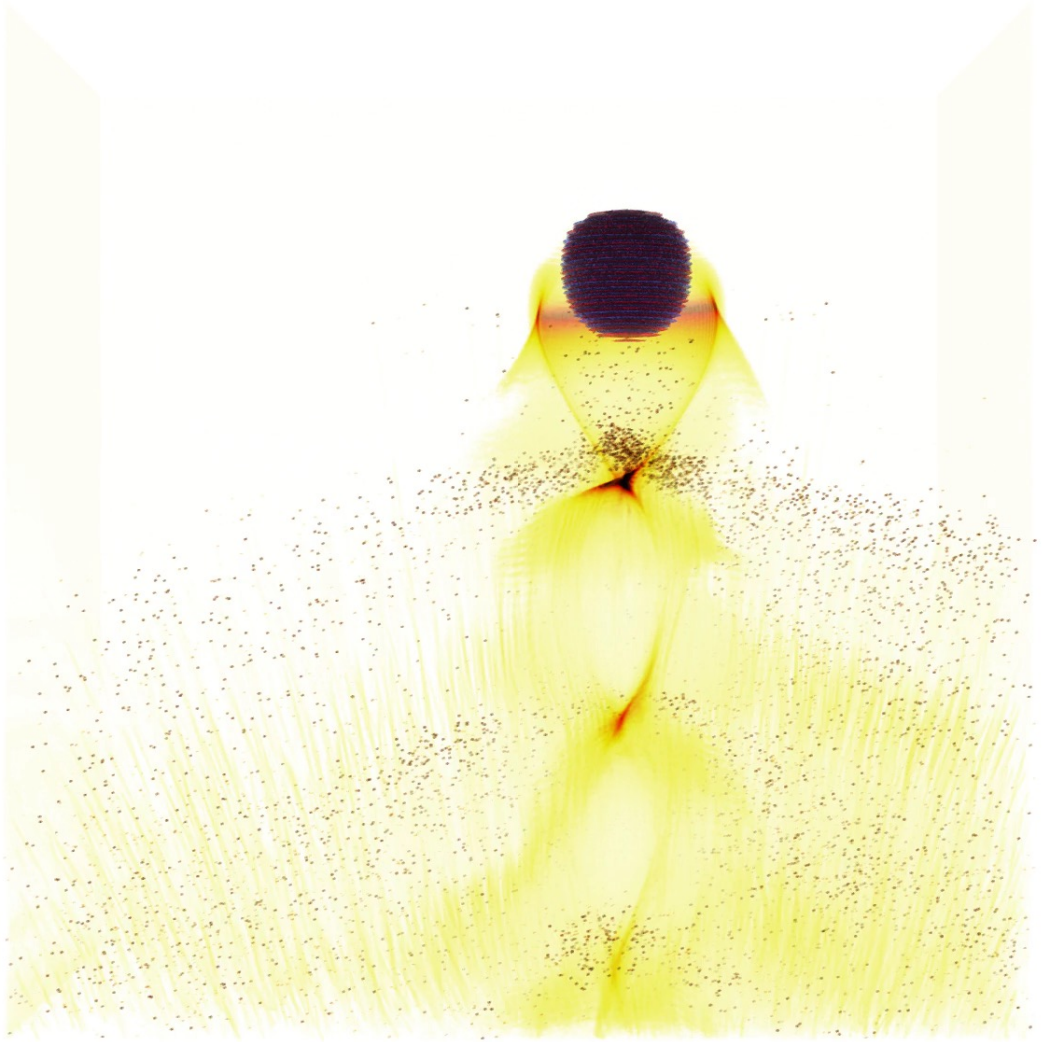

A movie from our 3D simulations

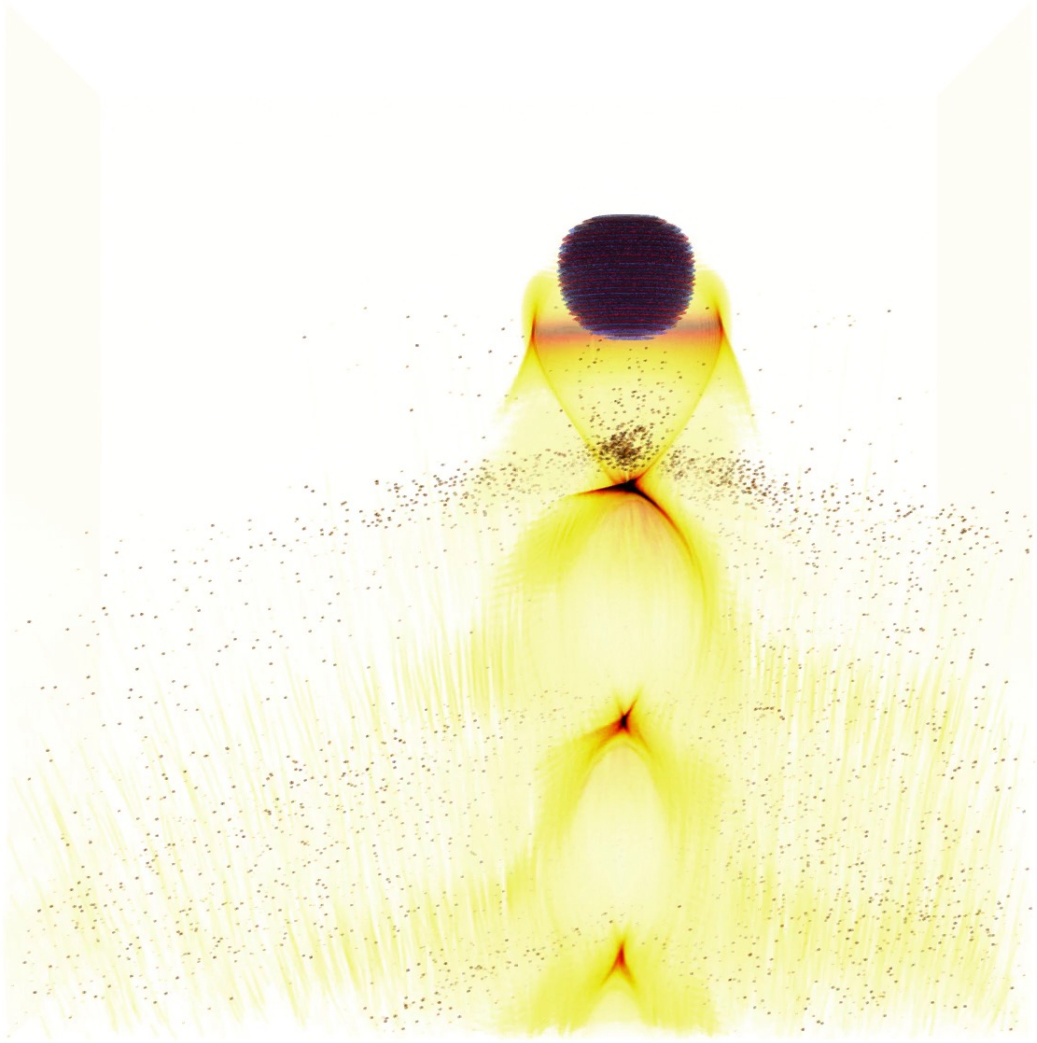

A movie from our 3D simulations

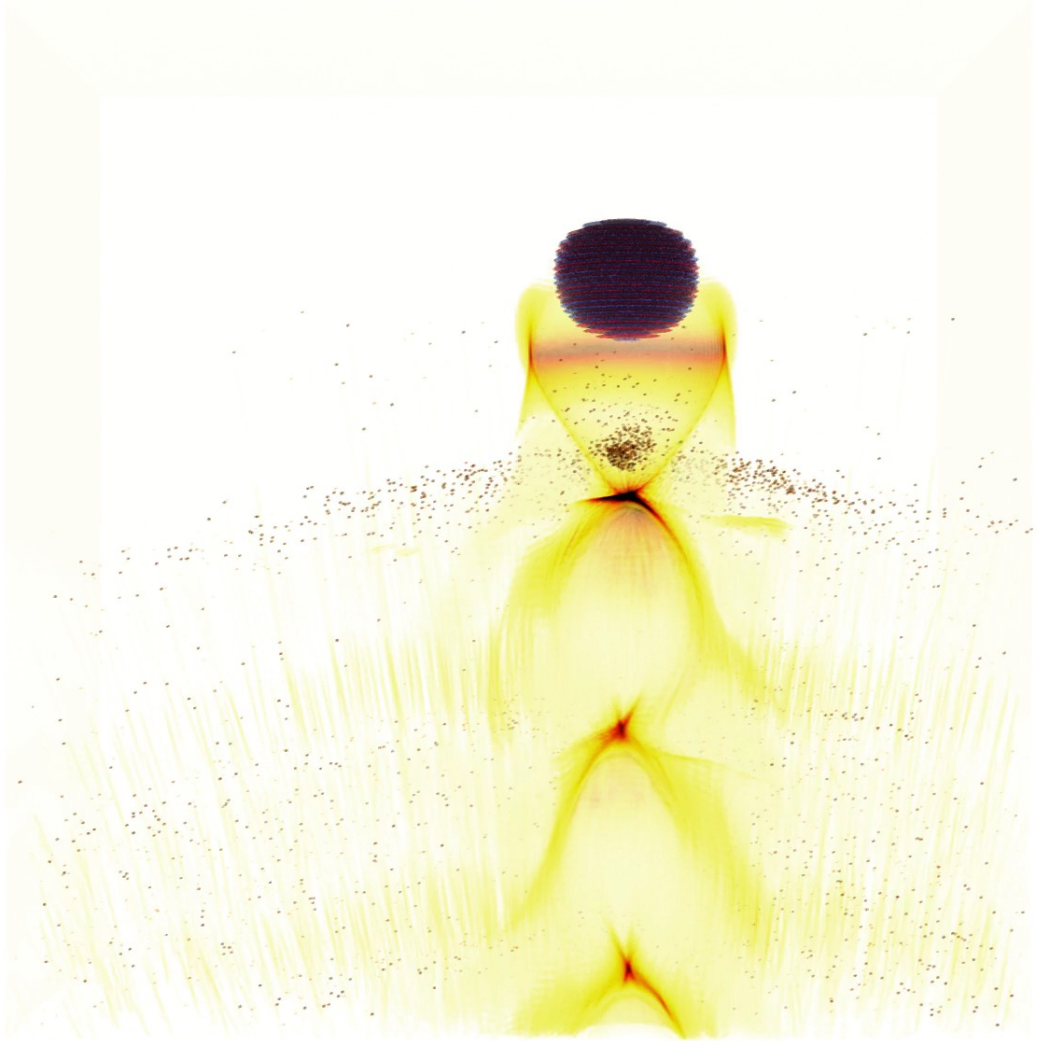

A movie from our 3D simulations

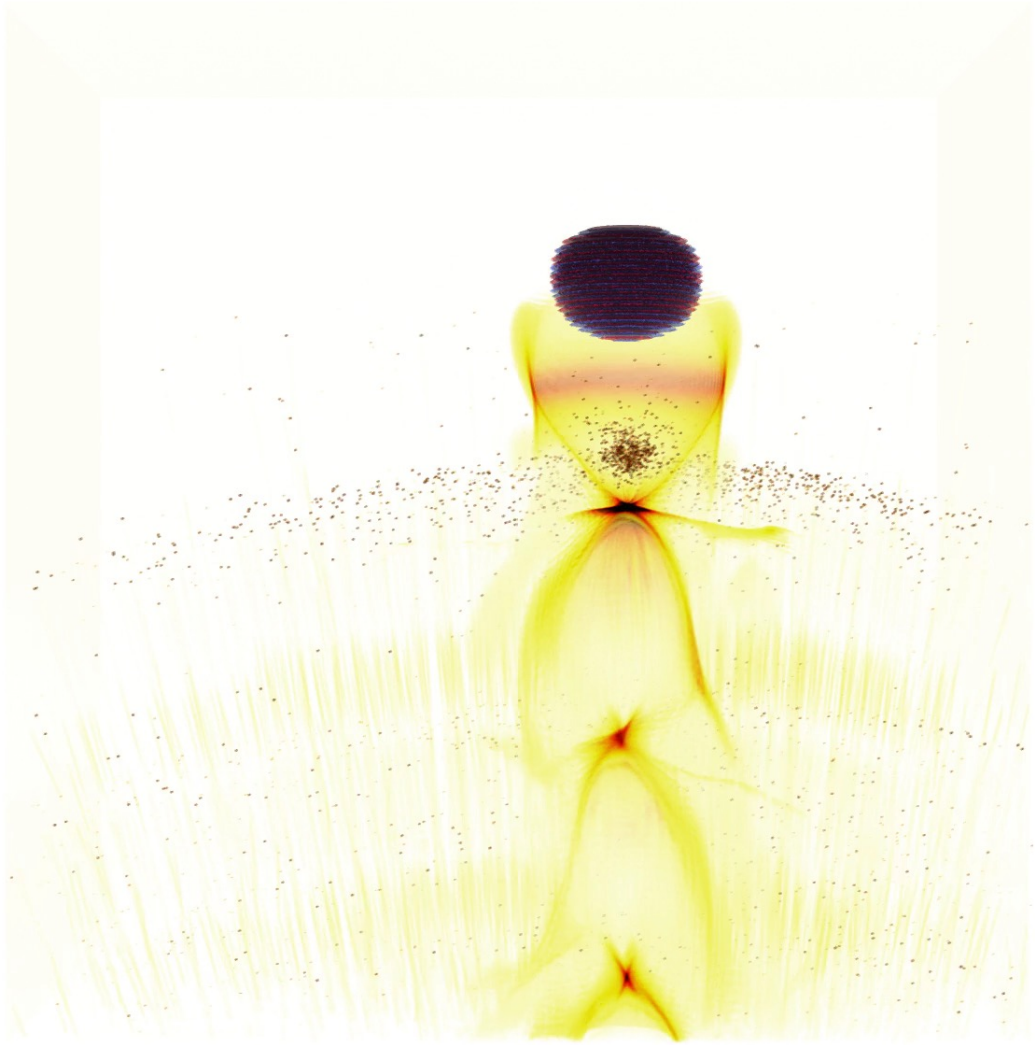

A movie from our 3D simulations

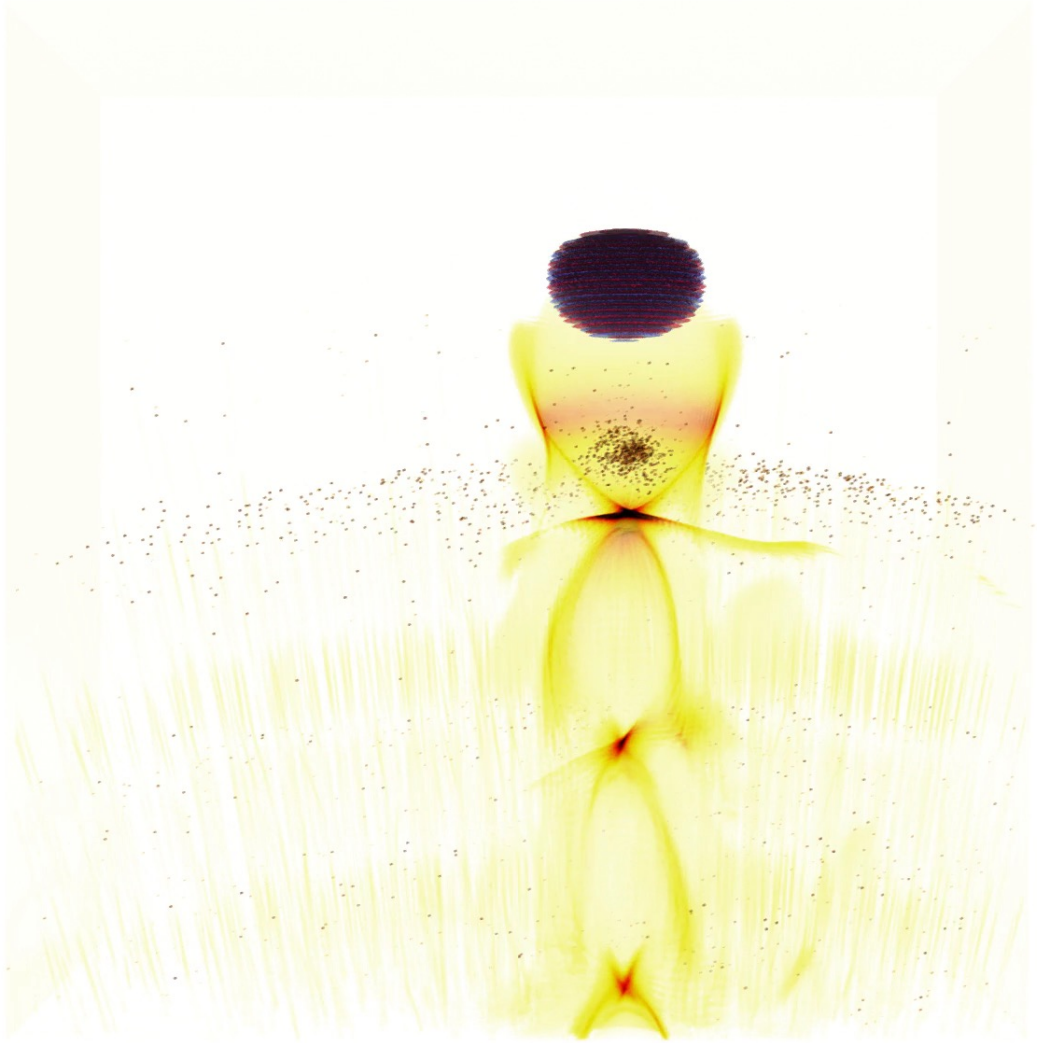

A movie from our 3D simulations

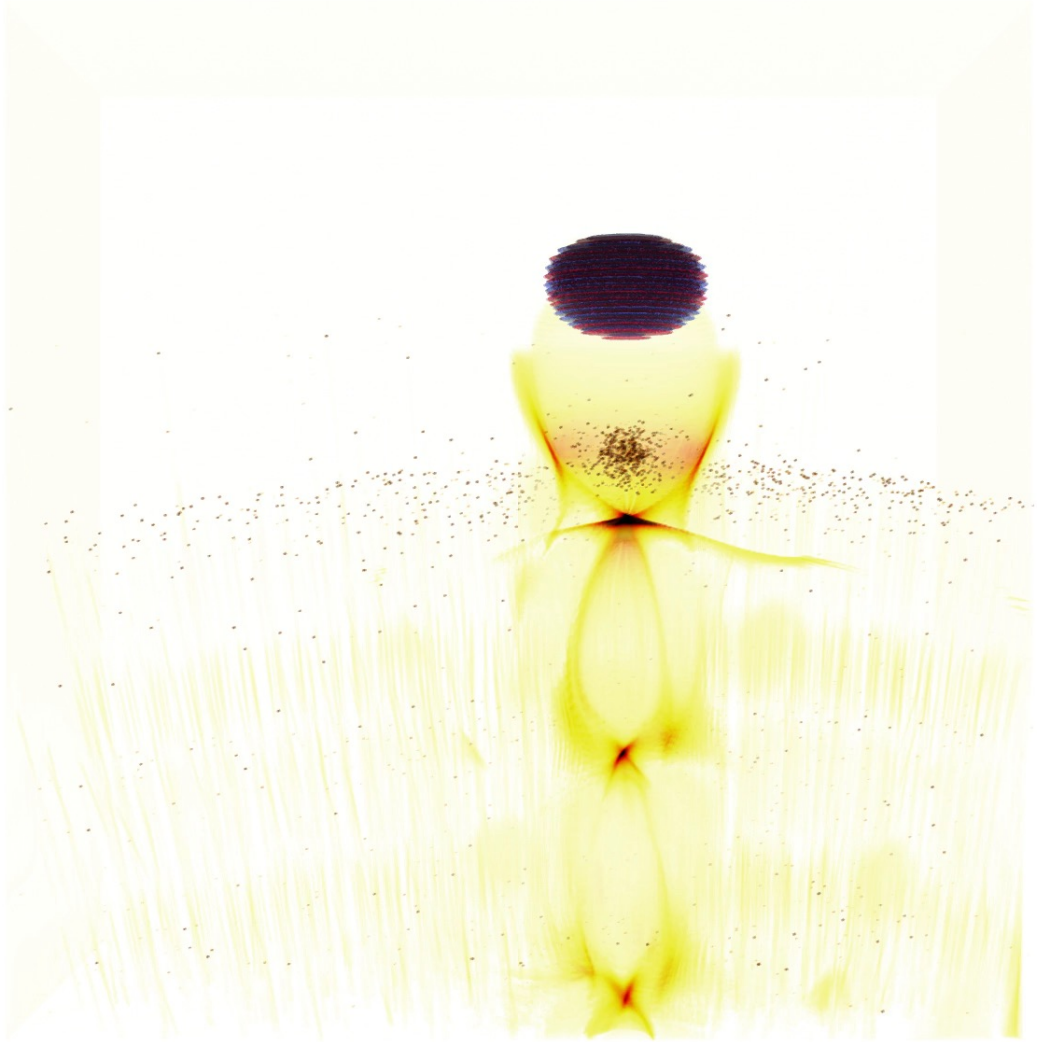

A movie from our 3D simulations

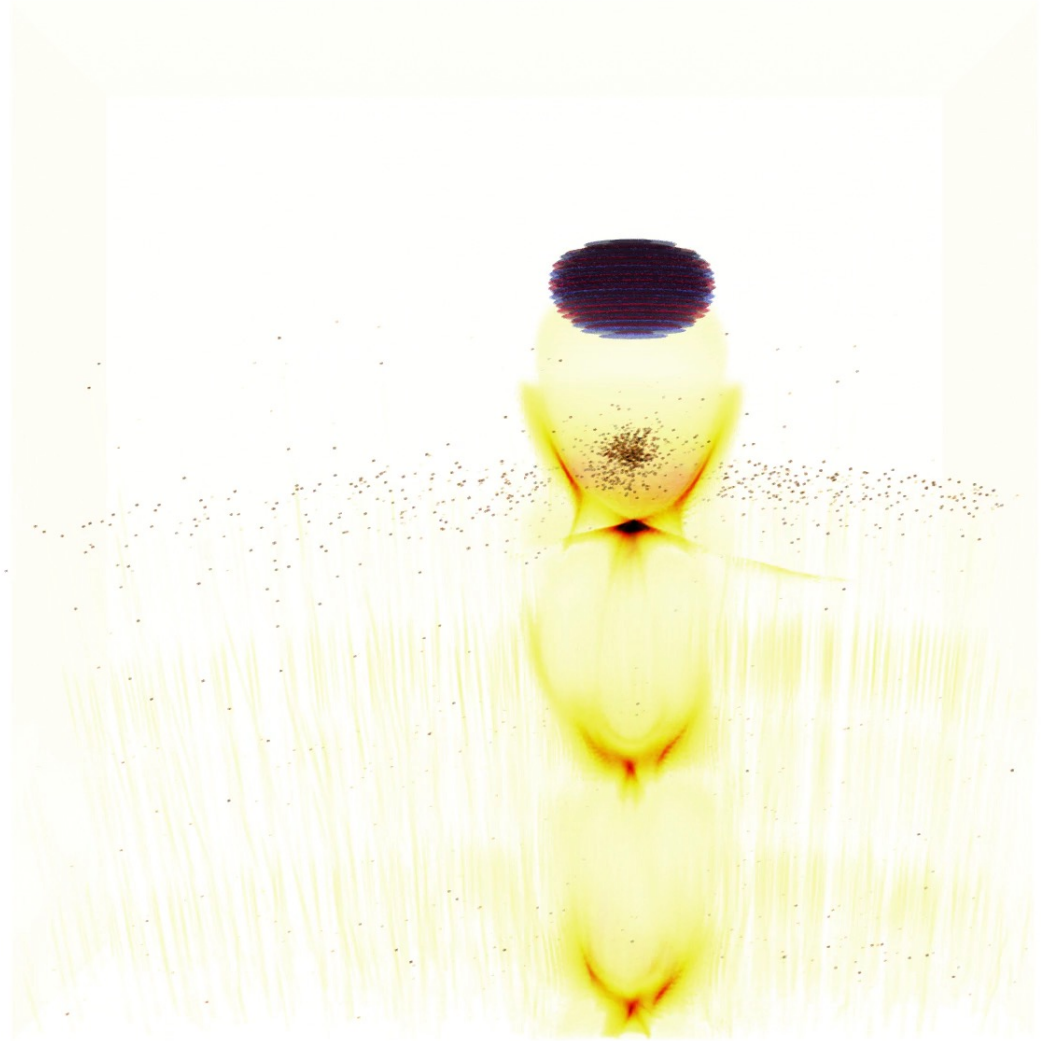

A movie from our 3D simulations

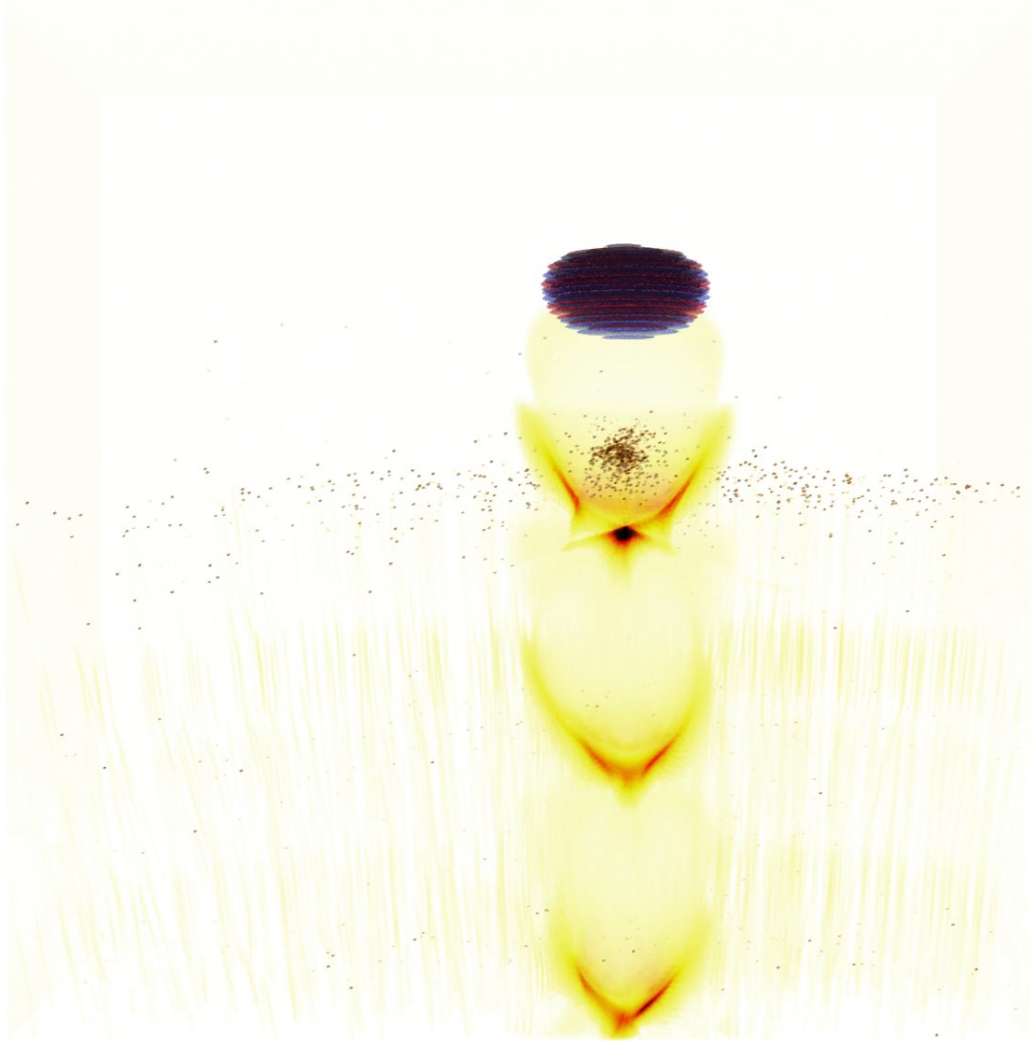

A movie from our 3D simulations

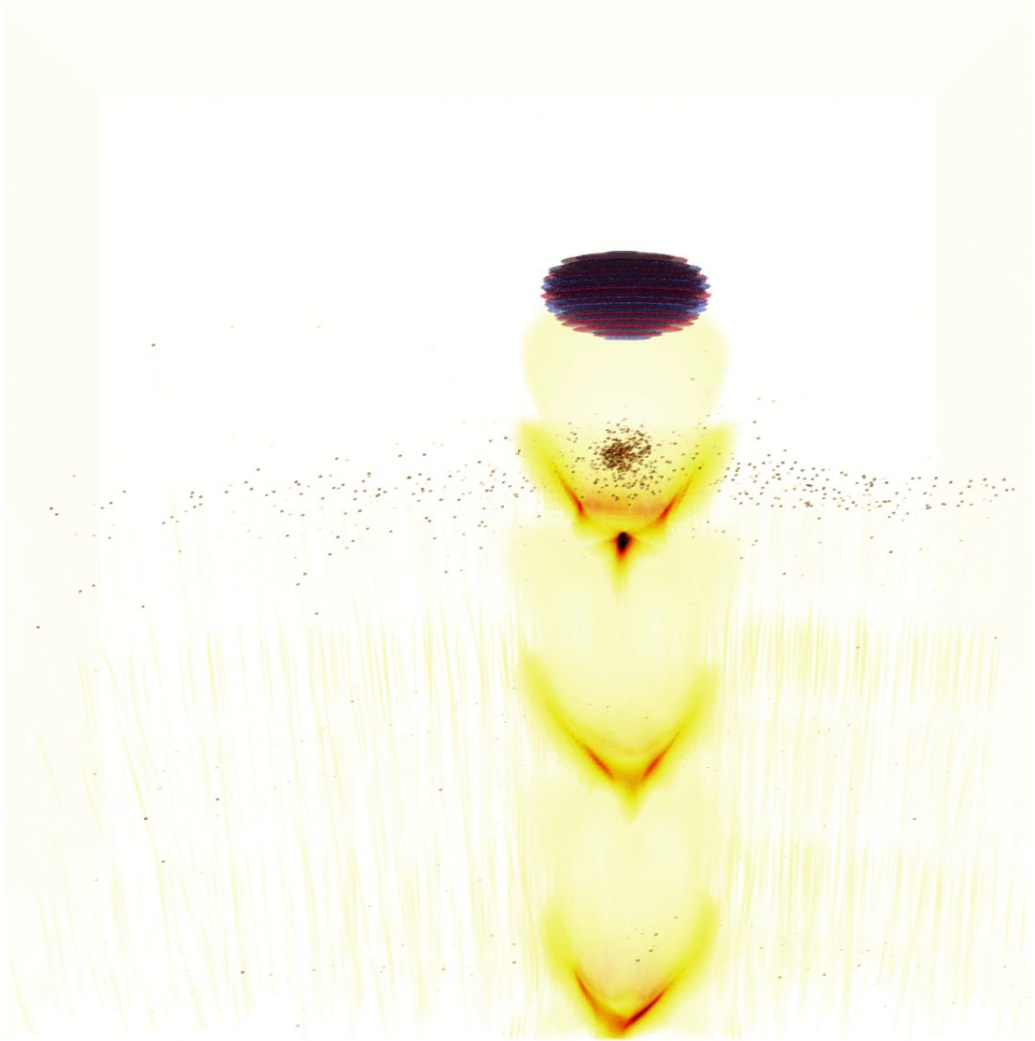

Why those simulations are so expensive ?

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

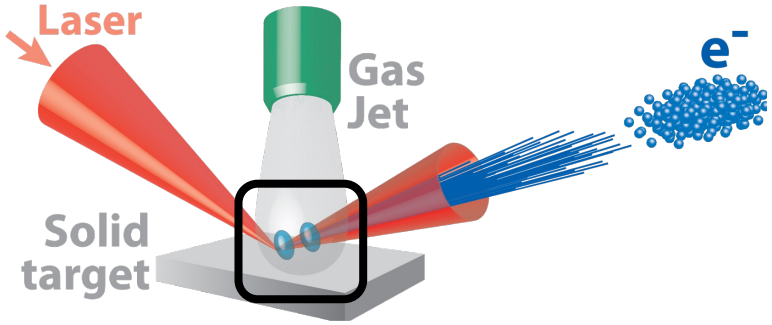

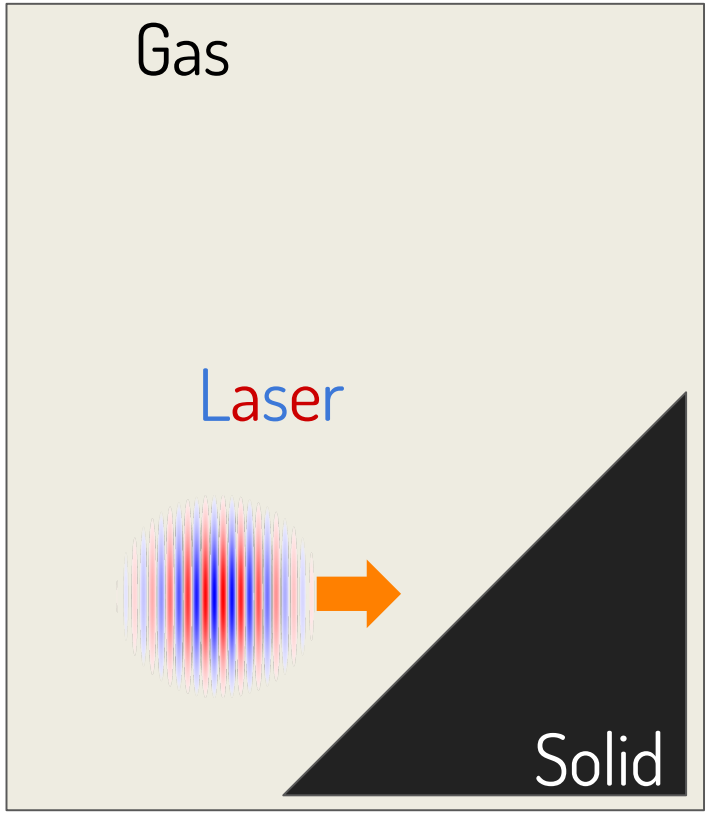

The laser pushes electrons away and generates a positively charged "bubble"

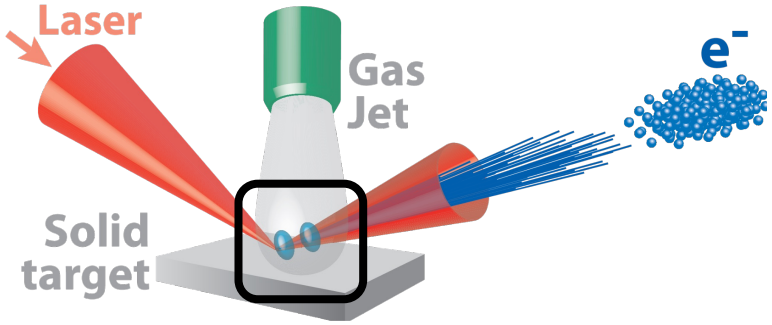

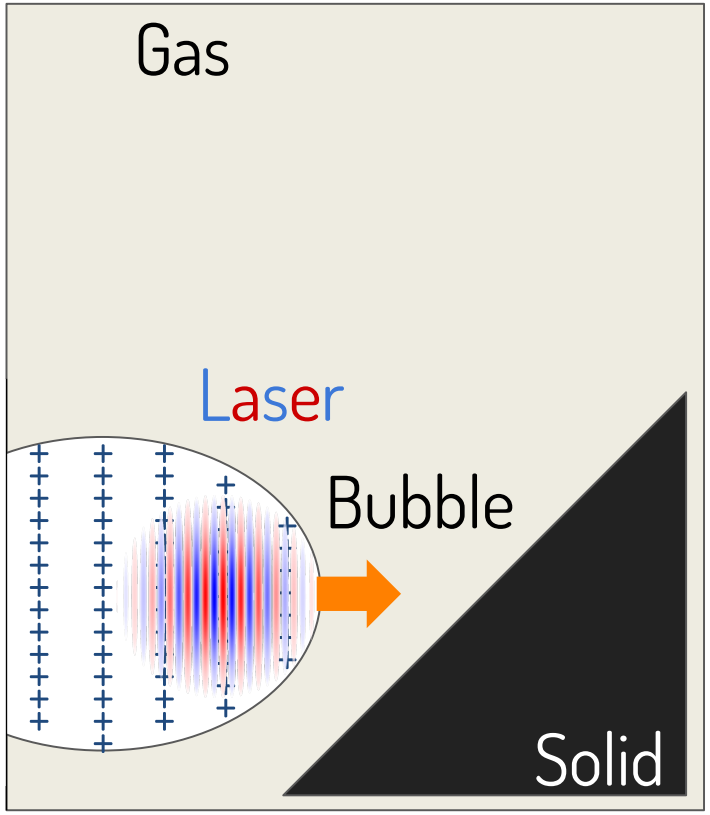

The laser is reflected by the high-density plasma and the bubble traps some of its electrons

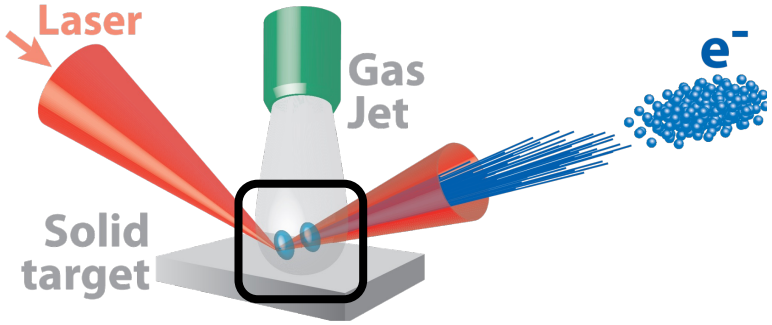

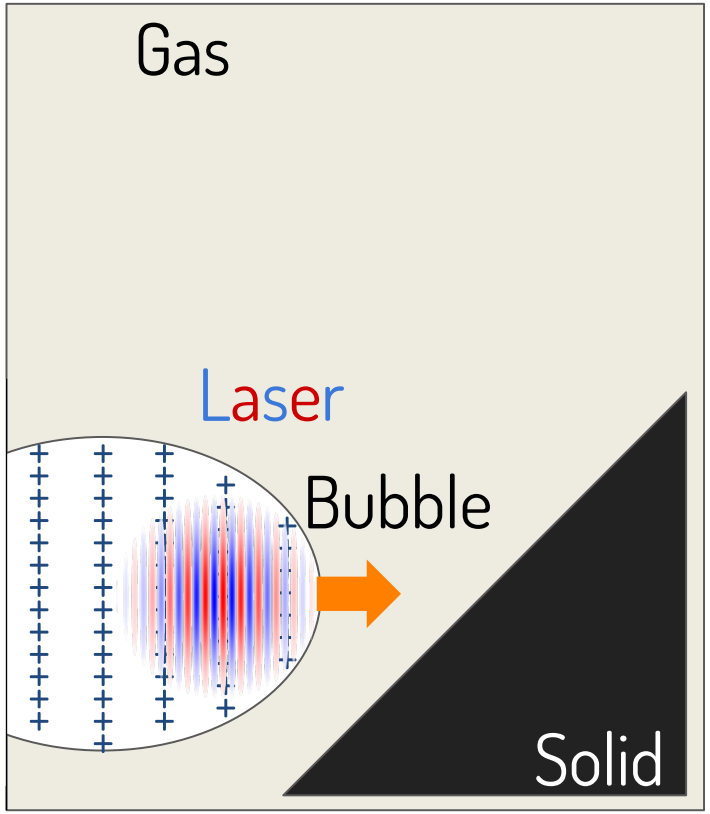

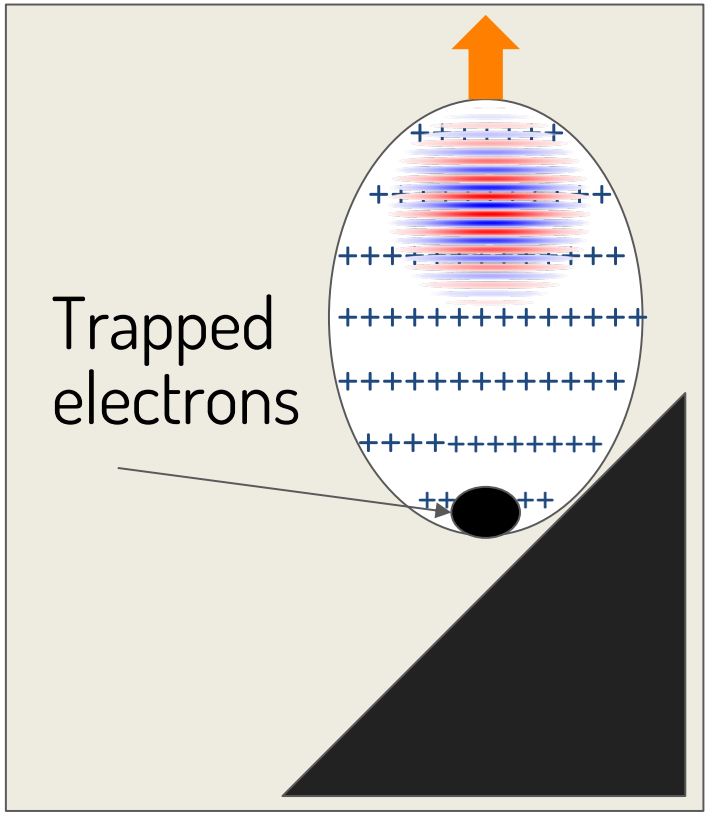

The bubble accelerates electrons over few millimeters

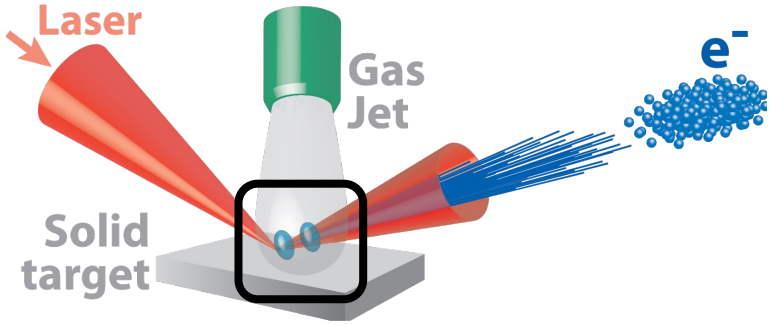

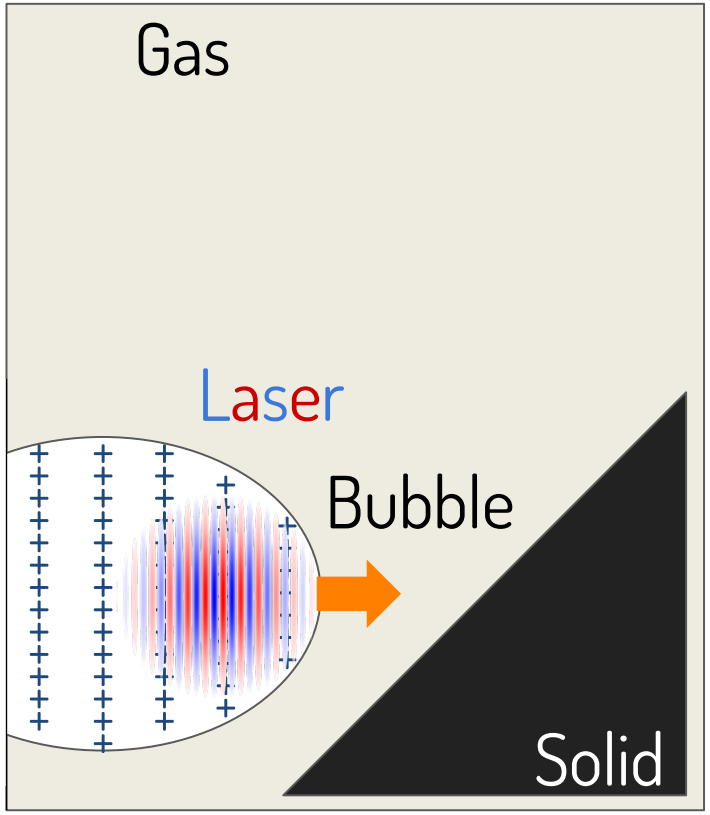

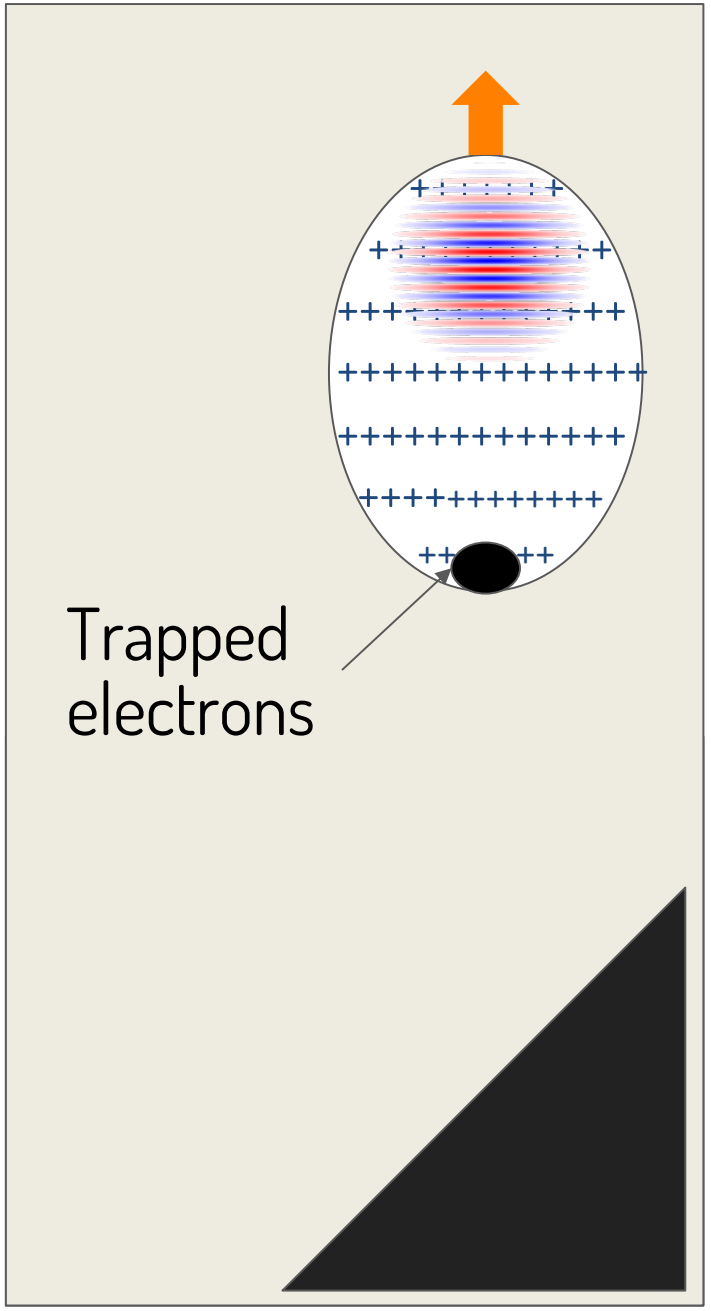

We can have smaller simulation boxes with a "moving window"

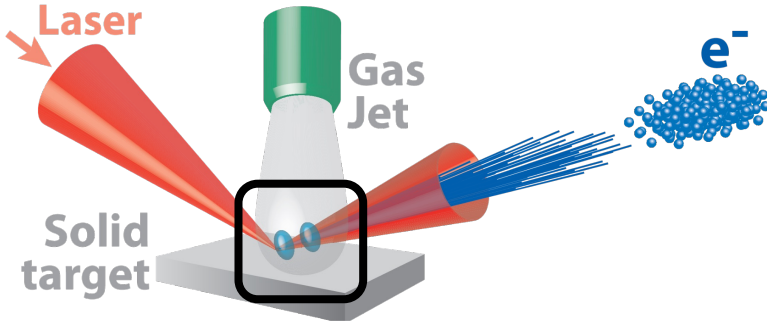

We only need to simulate

~ 100x100x100 μm^3 →

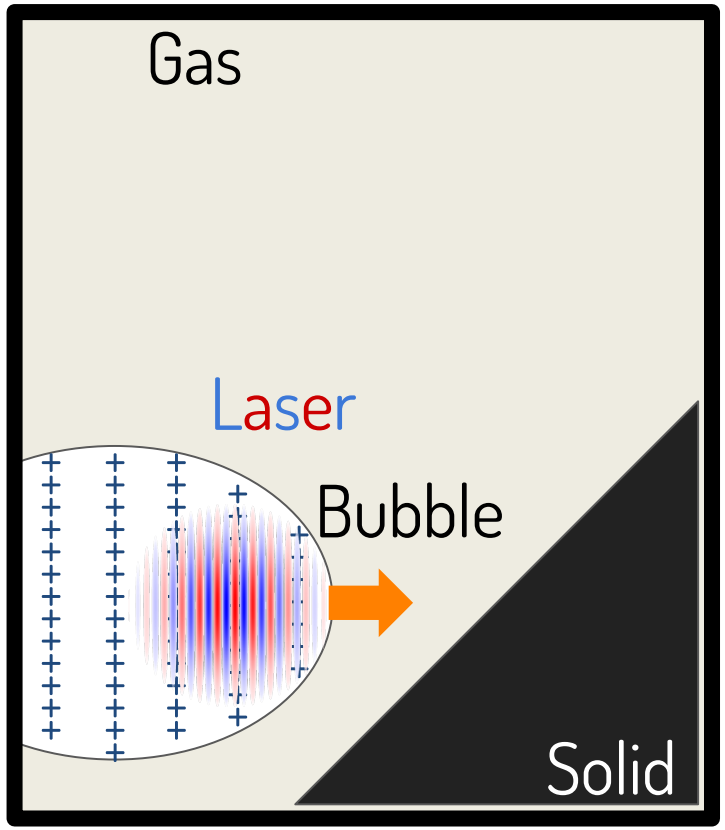

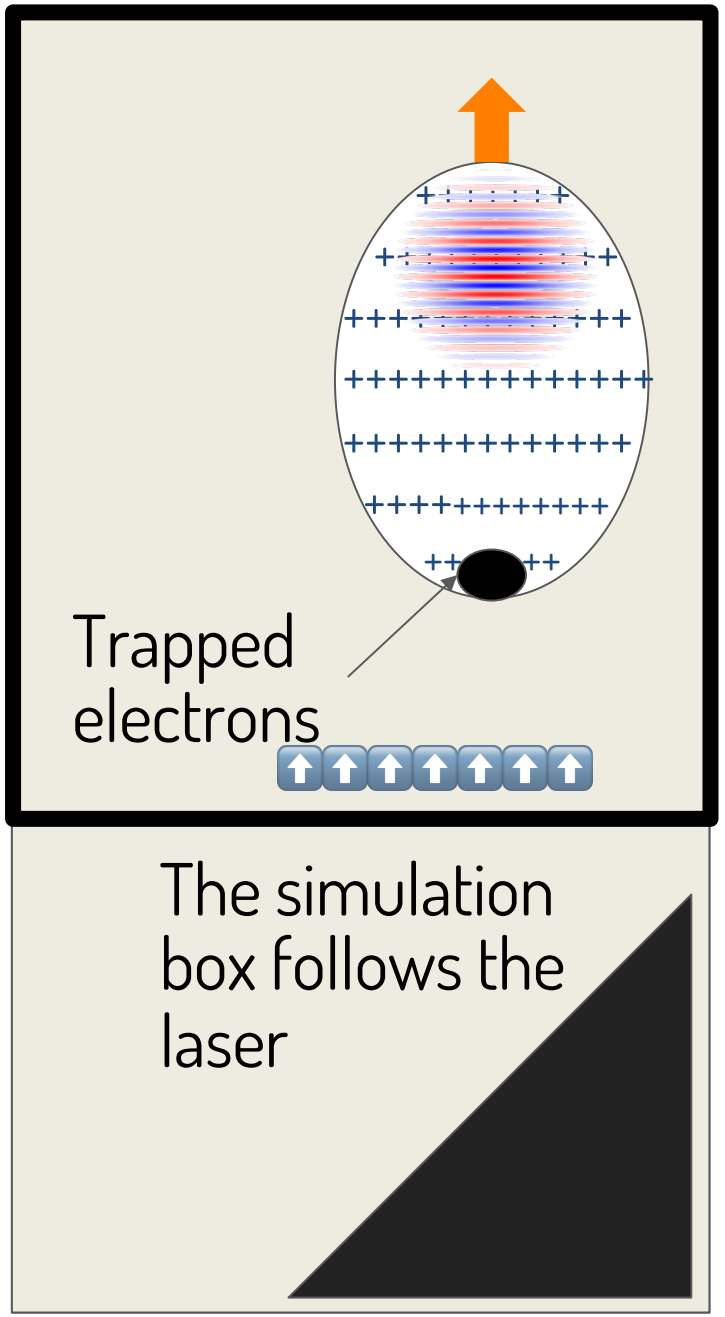

The main challenge concerns laser-solid interaction

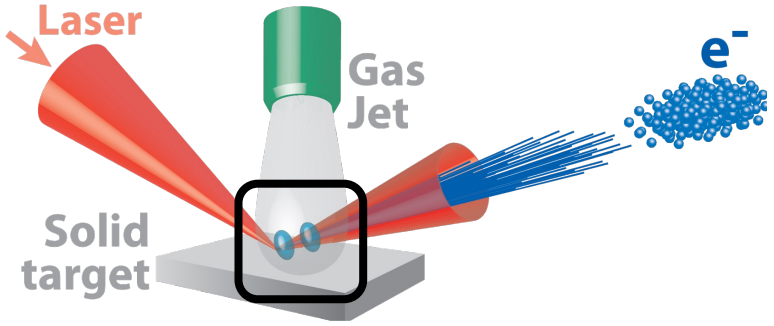

We need a resolution of **few 10s nanometers** for laser-solid interaction →

We only need to simulate

~ 100x100x100 µm³ →

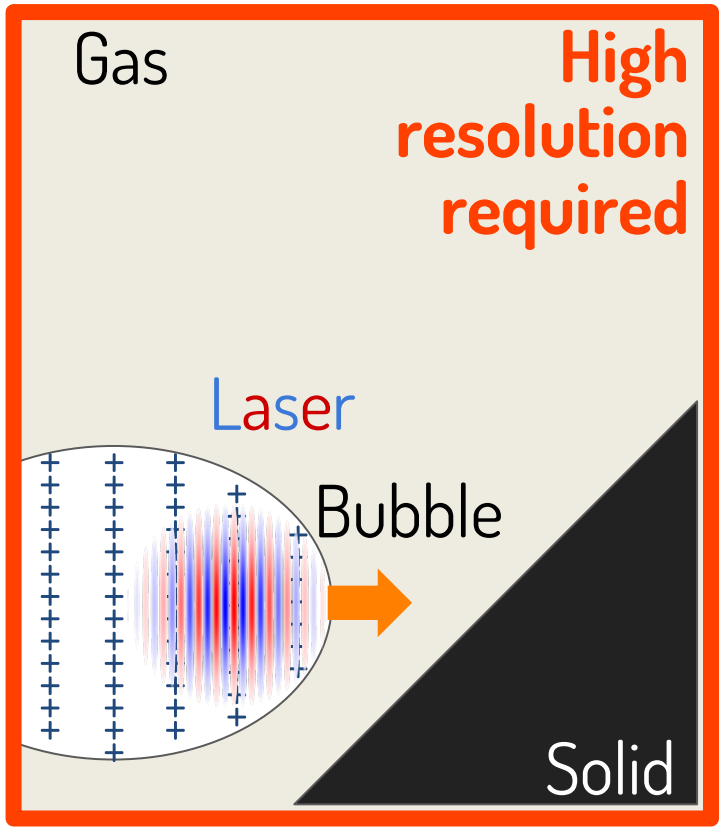

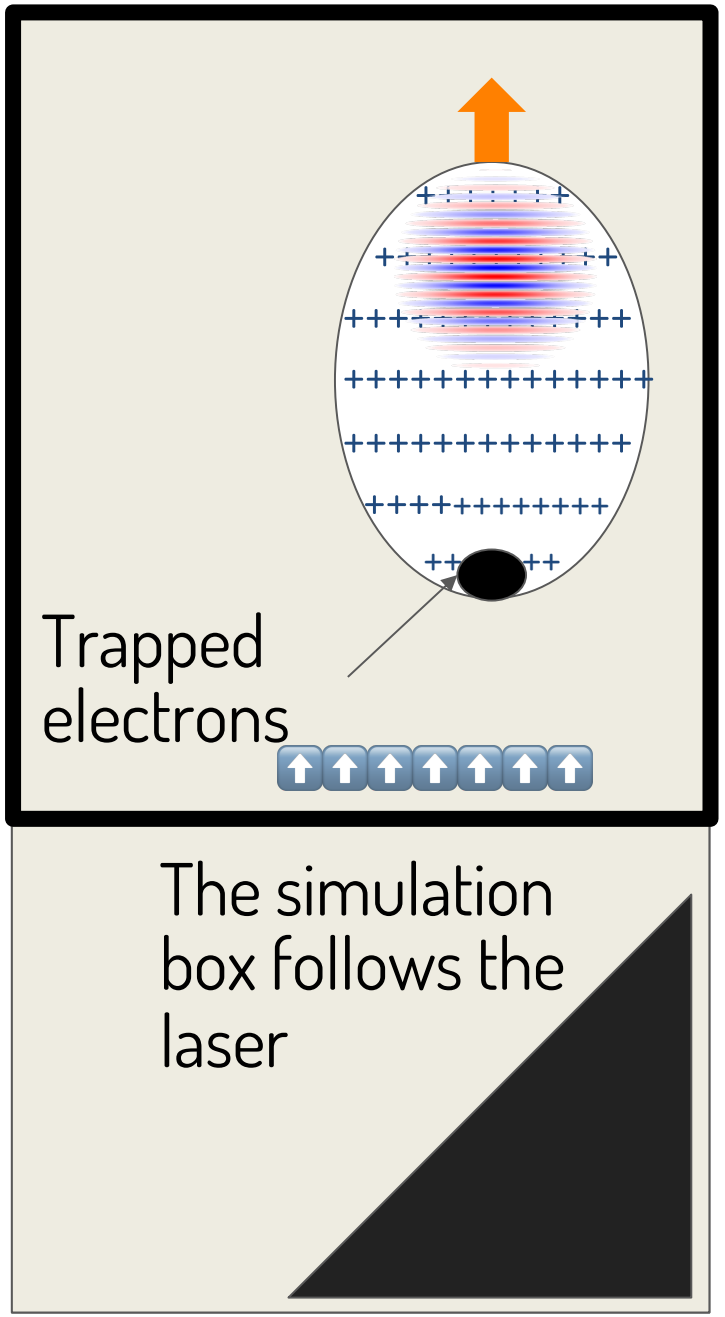

The main challenge concerns laser-solid interaction

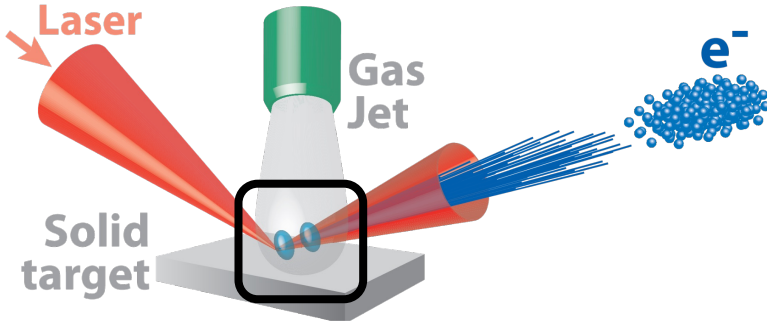

We need a resolution of **few 10s nanometers** for laser-solid interaction →

but

Hefty price to pay:

dt ~ dx and **size ~ (1/dx)³**

We only need to simulate

~ 100x100x100 µm³ →

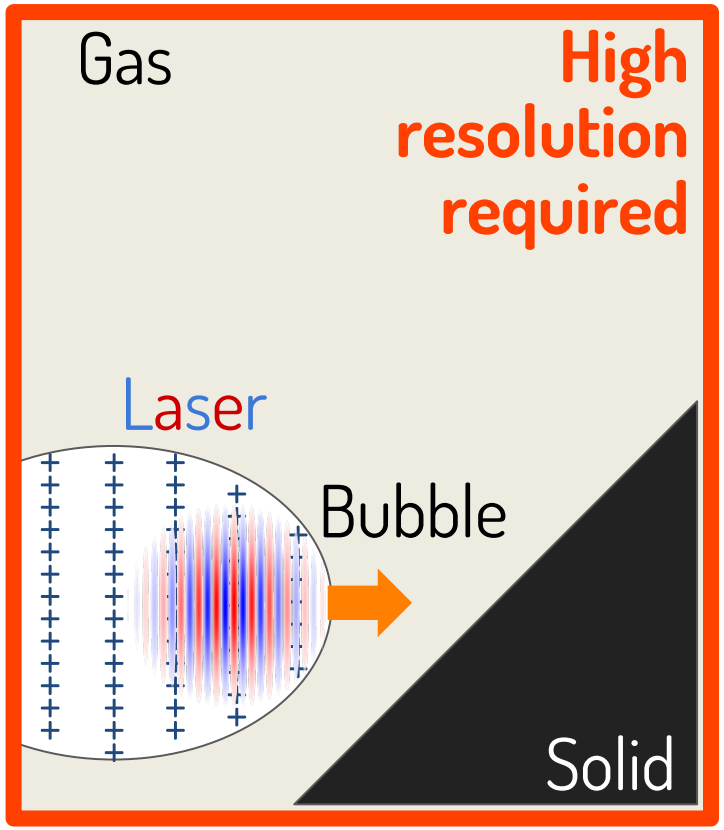

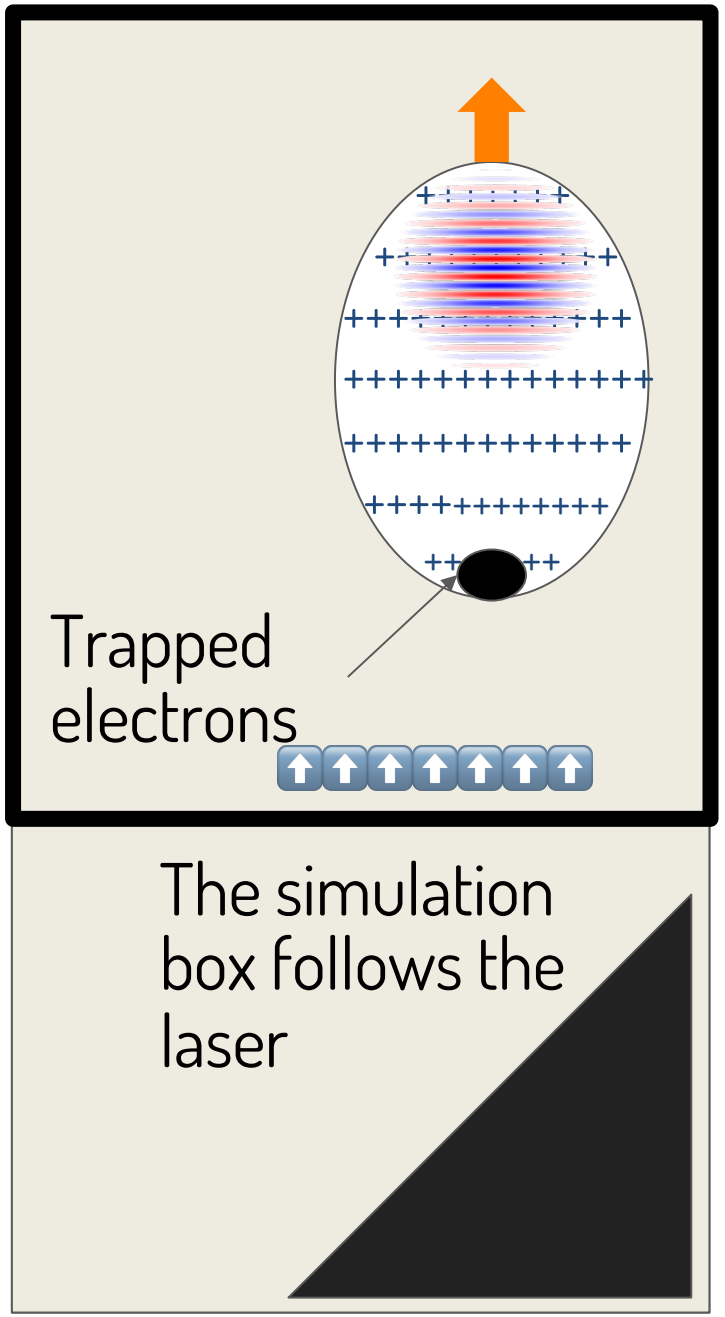

The main challenge concerns laser-solid interaction

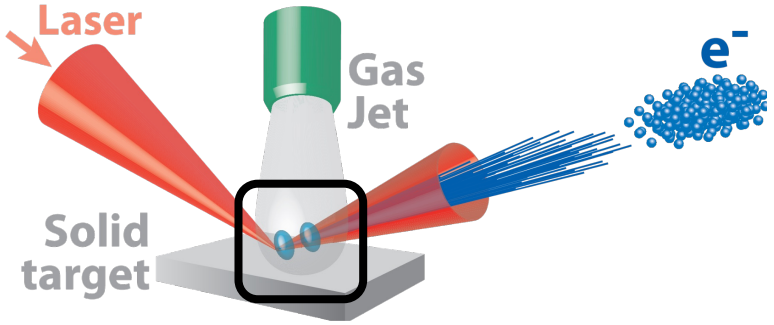

We only need to simulate

~ 100x100x100 µm³ →

Enabled by **very good weak scaling** →

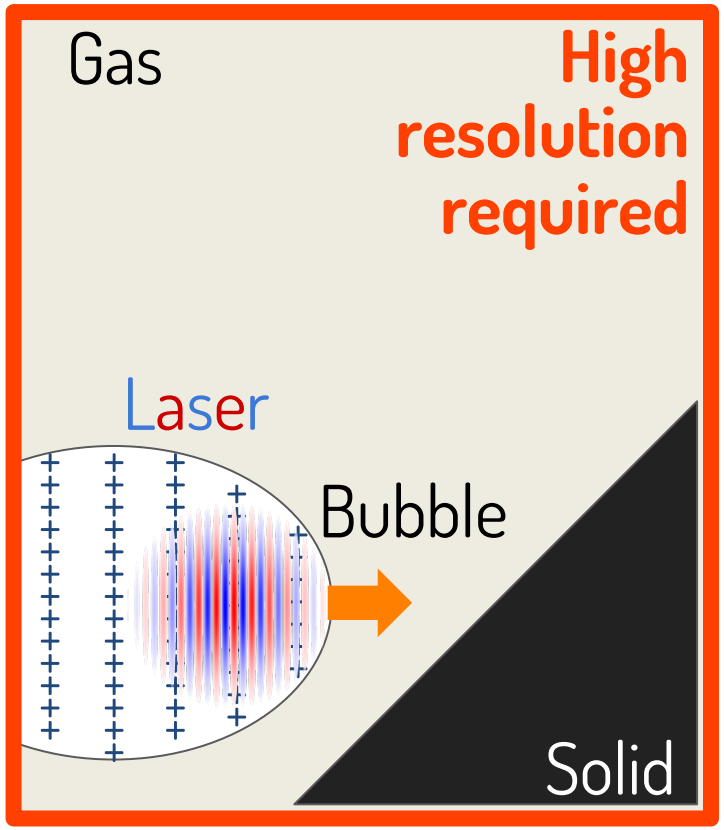

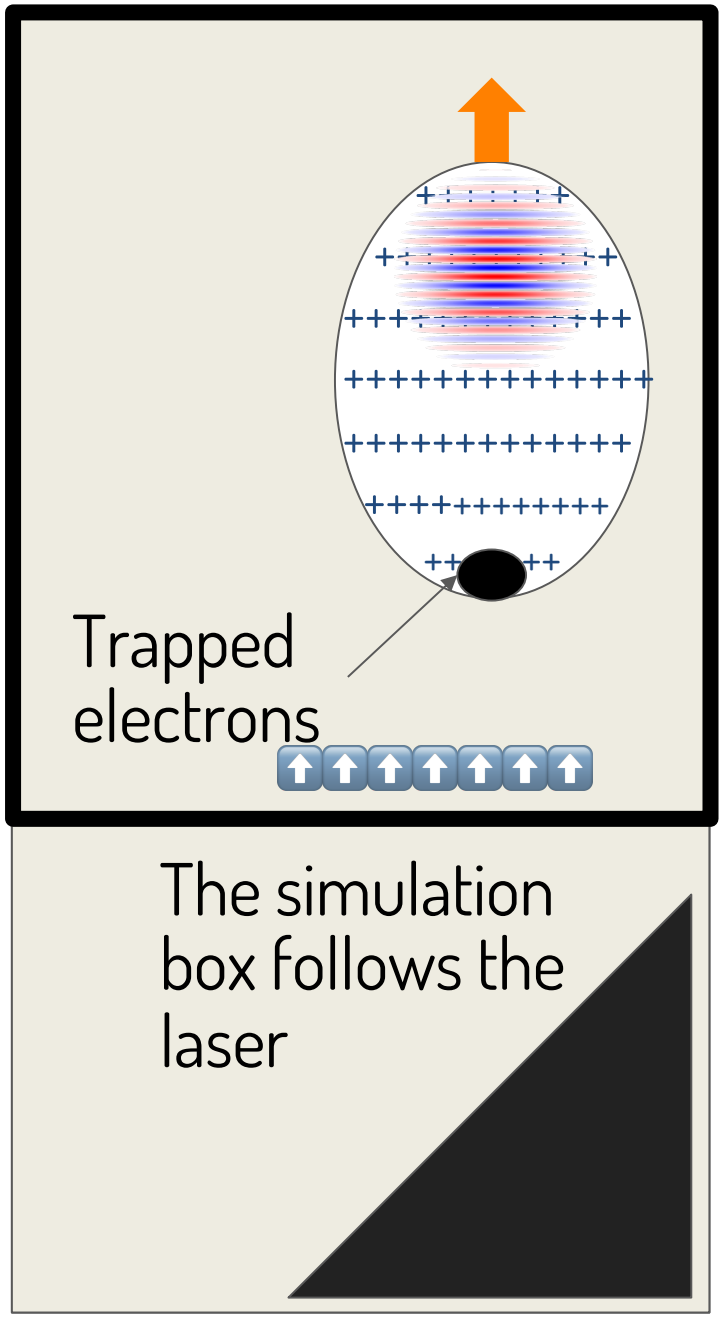

The main challenge concerns laser-solid interaction

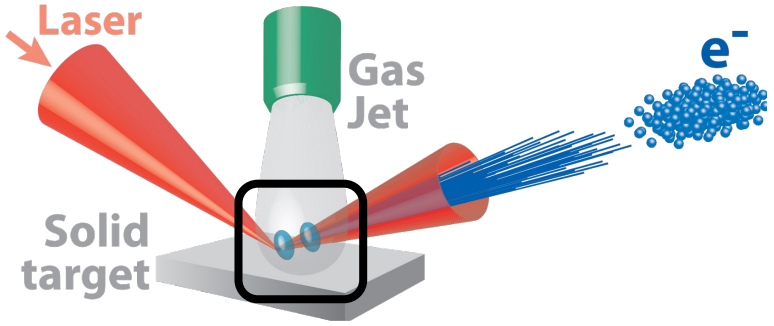

We only need to simulate

~ 100x100x100 µm³ →

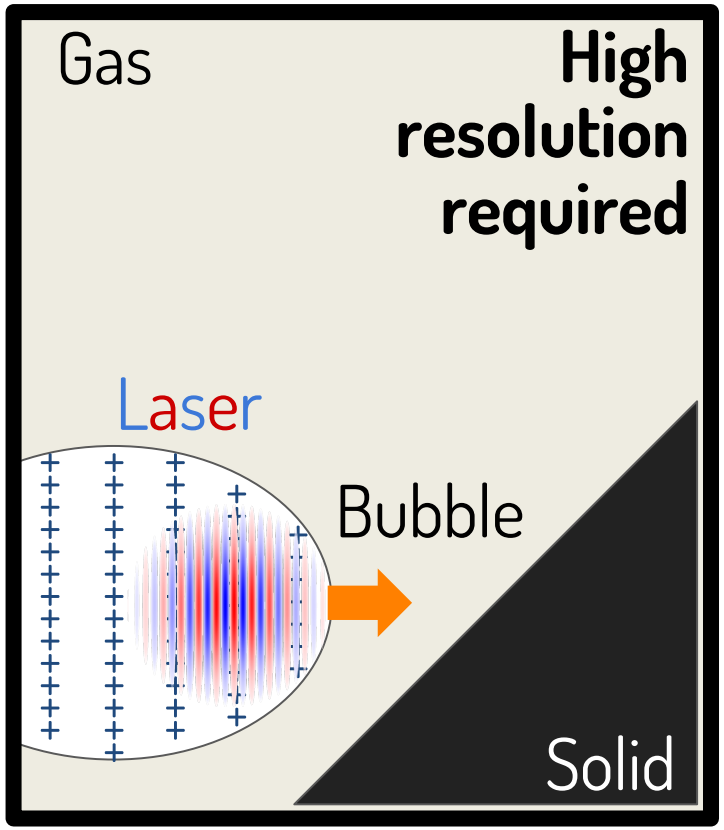

Enabled by **very good weak scaling** →

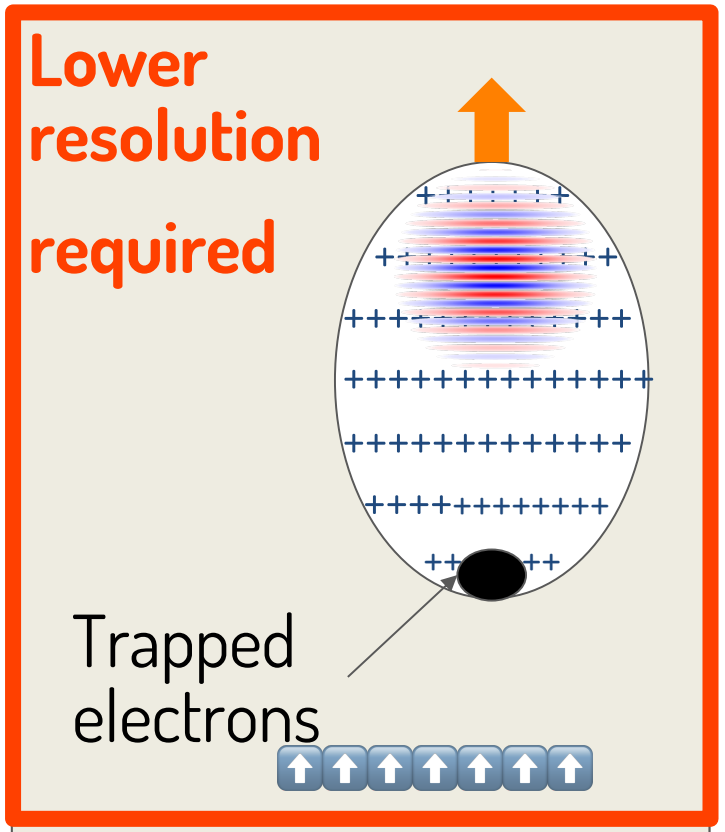

The simulation box follows the laser

The main challenge concerns laser-solid interaction

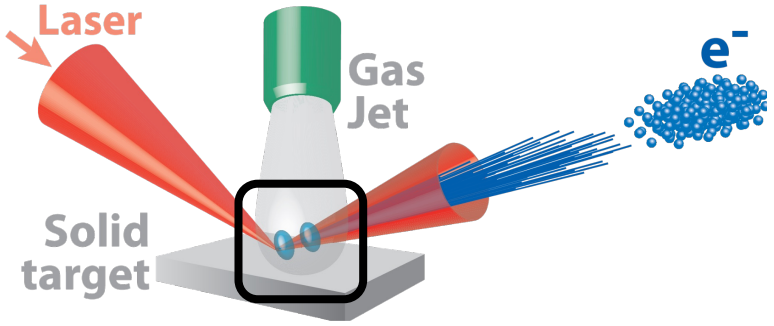

We need to simulate it **very good strong scaling** →

Enabled by **very good weak scaling** →

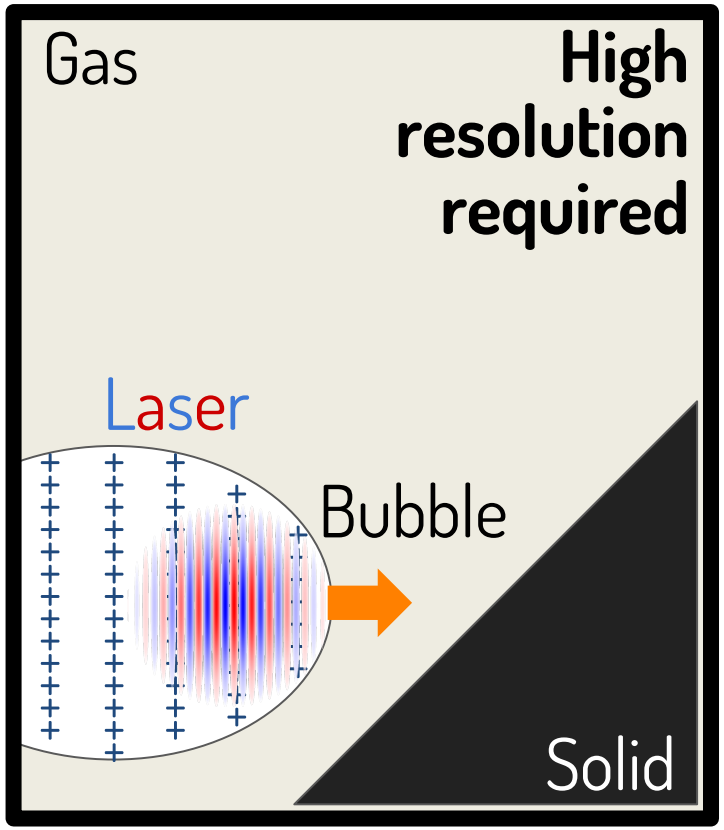

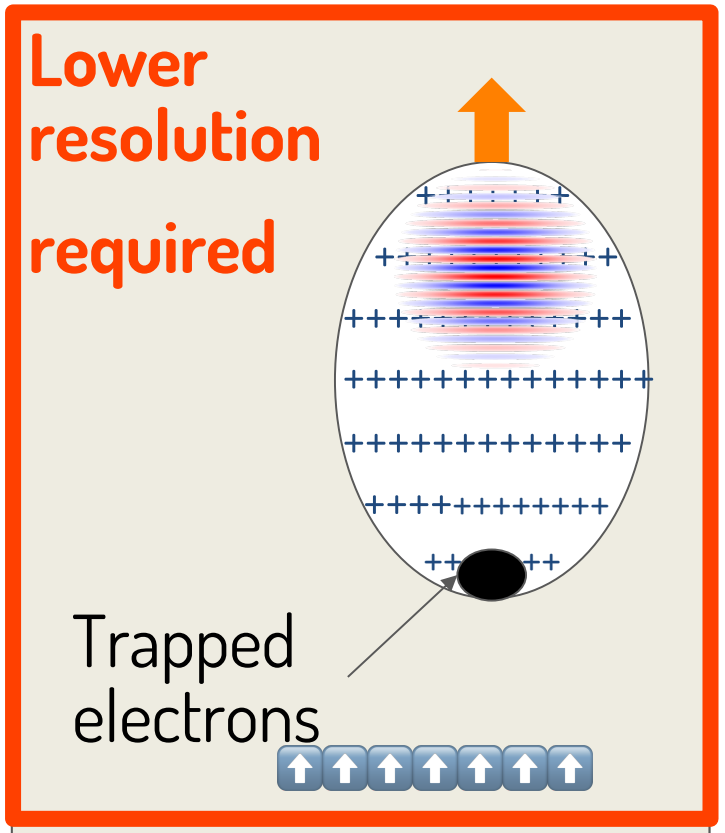

The simulation box follows the laser

The main challenge concerns laser-solid interaction

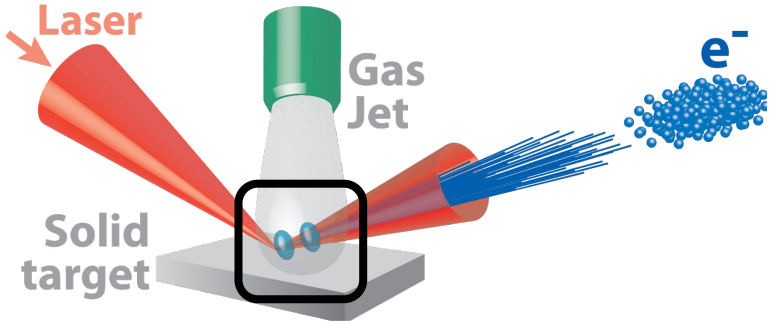

We need to simulate it **on mm/cm→**

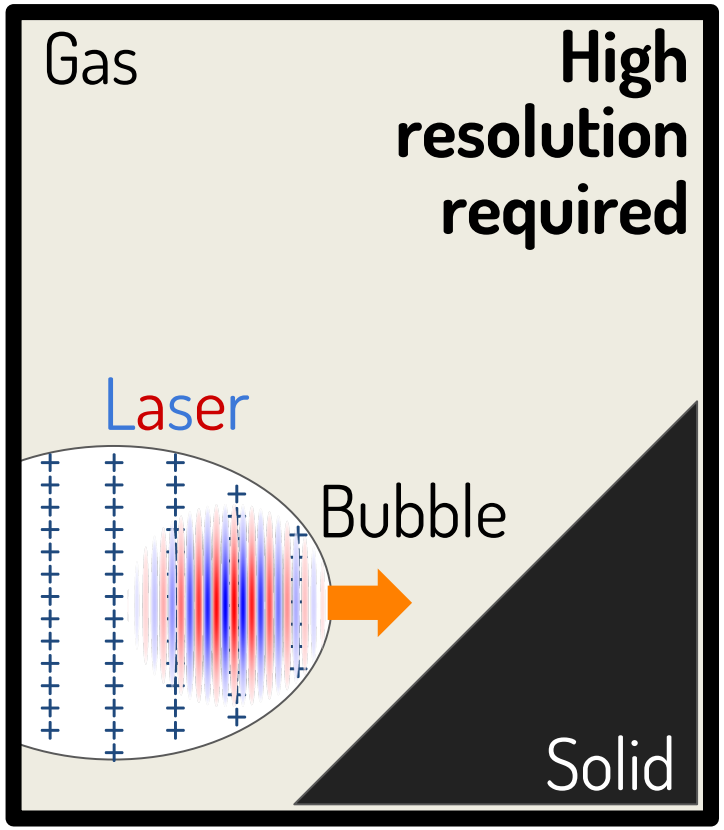

Enabled by **very good weak scaling →**

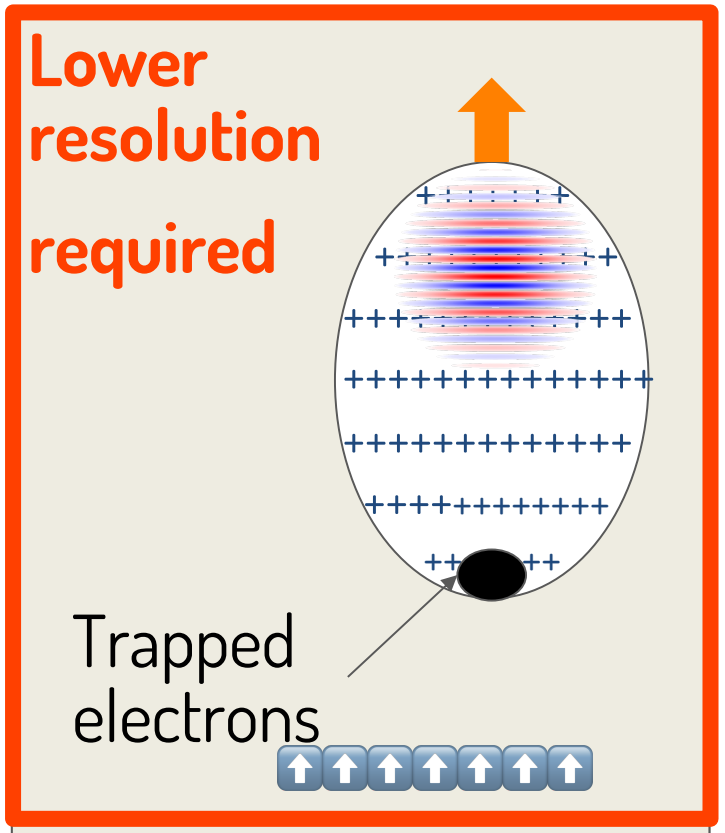

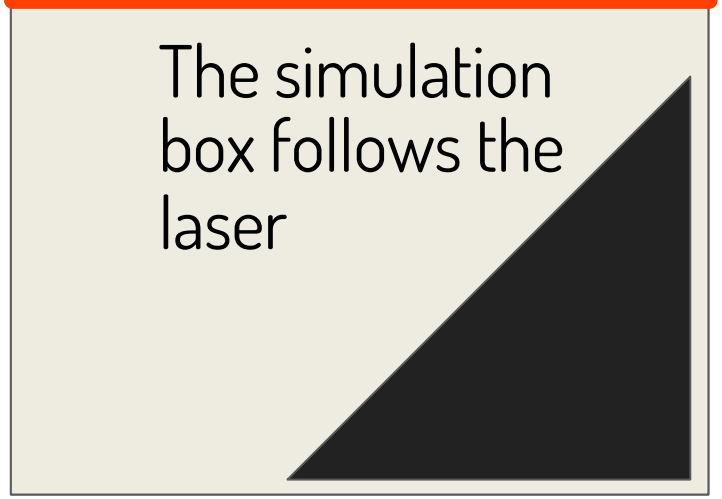

How do we switch resolution in the middle of the simulation?

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

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

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

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…)

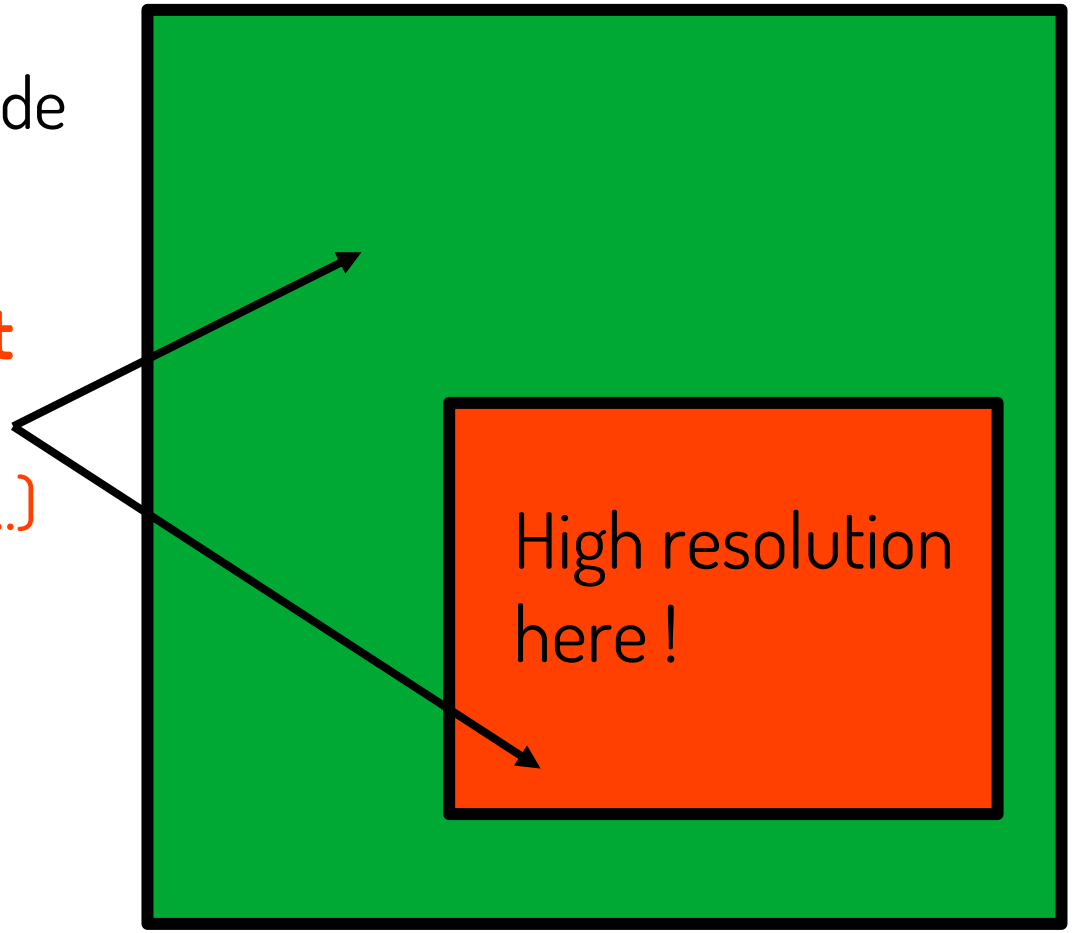

Mesh refinement, one of the most advanced WarpX features, comes to help

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…)

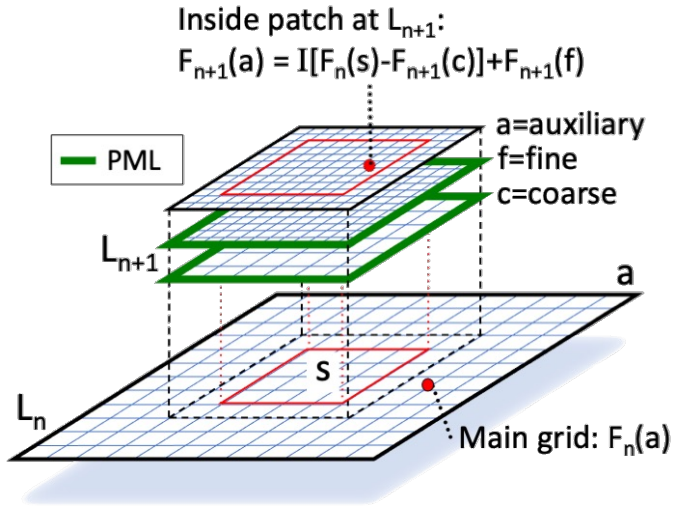

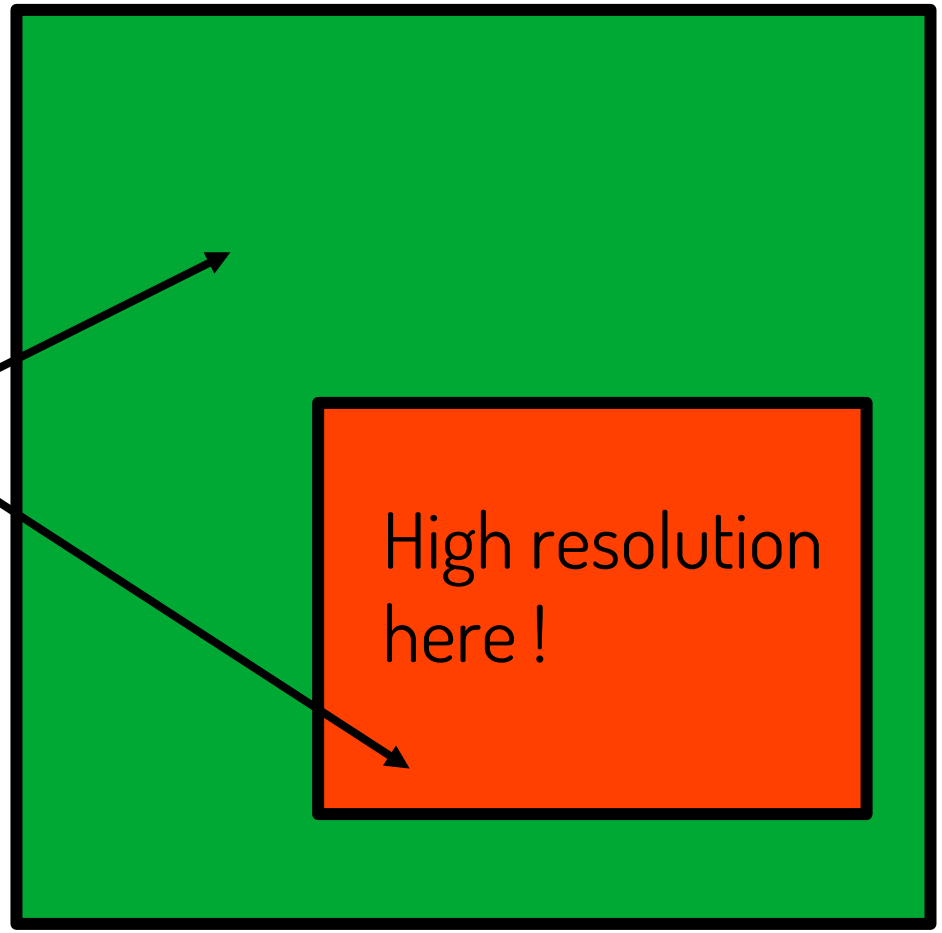

J.-L. Vay et al, Phys. Plasmas 11, 2928 (2004)
R. Lehe et al, Phys. Rev. E 106, 045306 (2022)

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

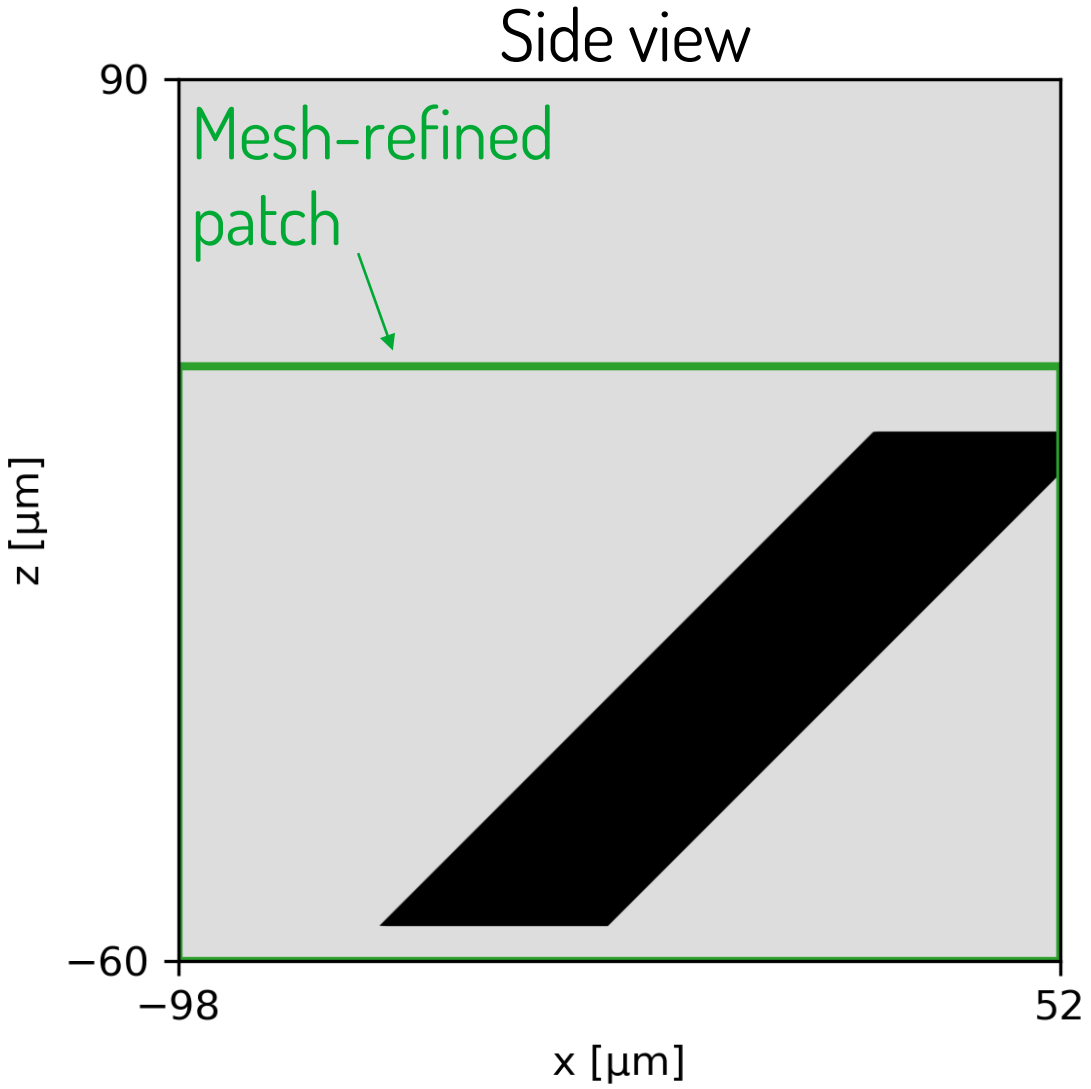

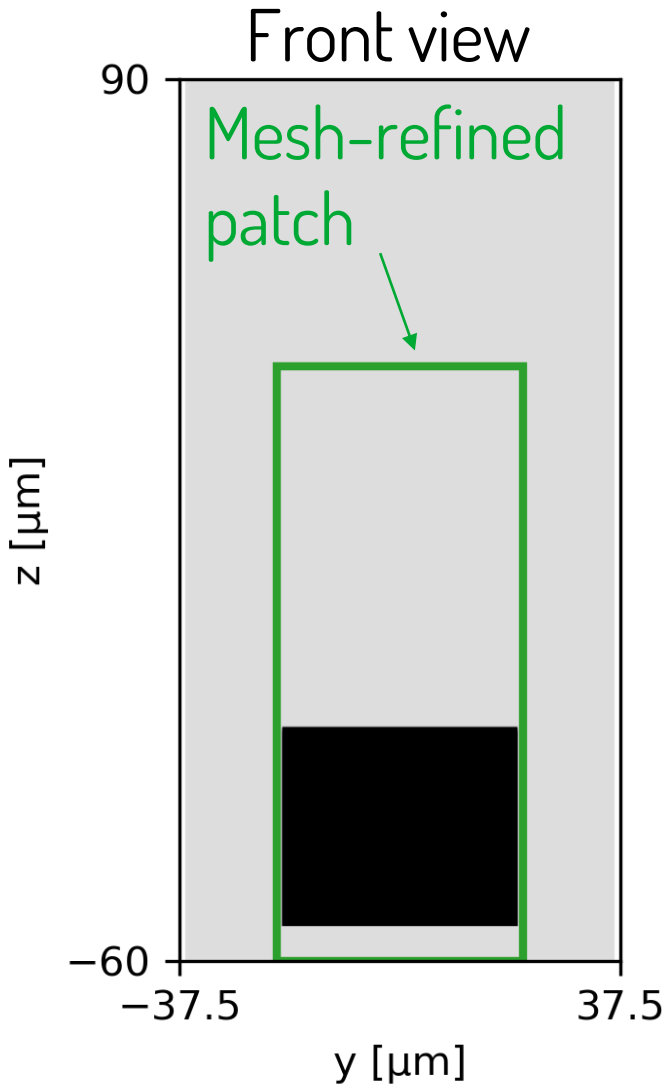

← 3D simulation on 4096 Summit nodes

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

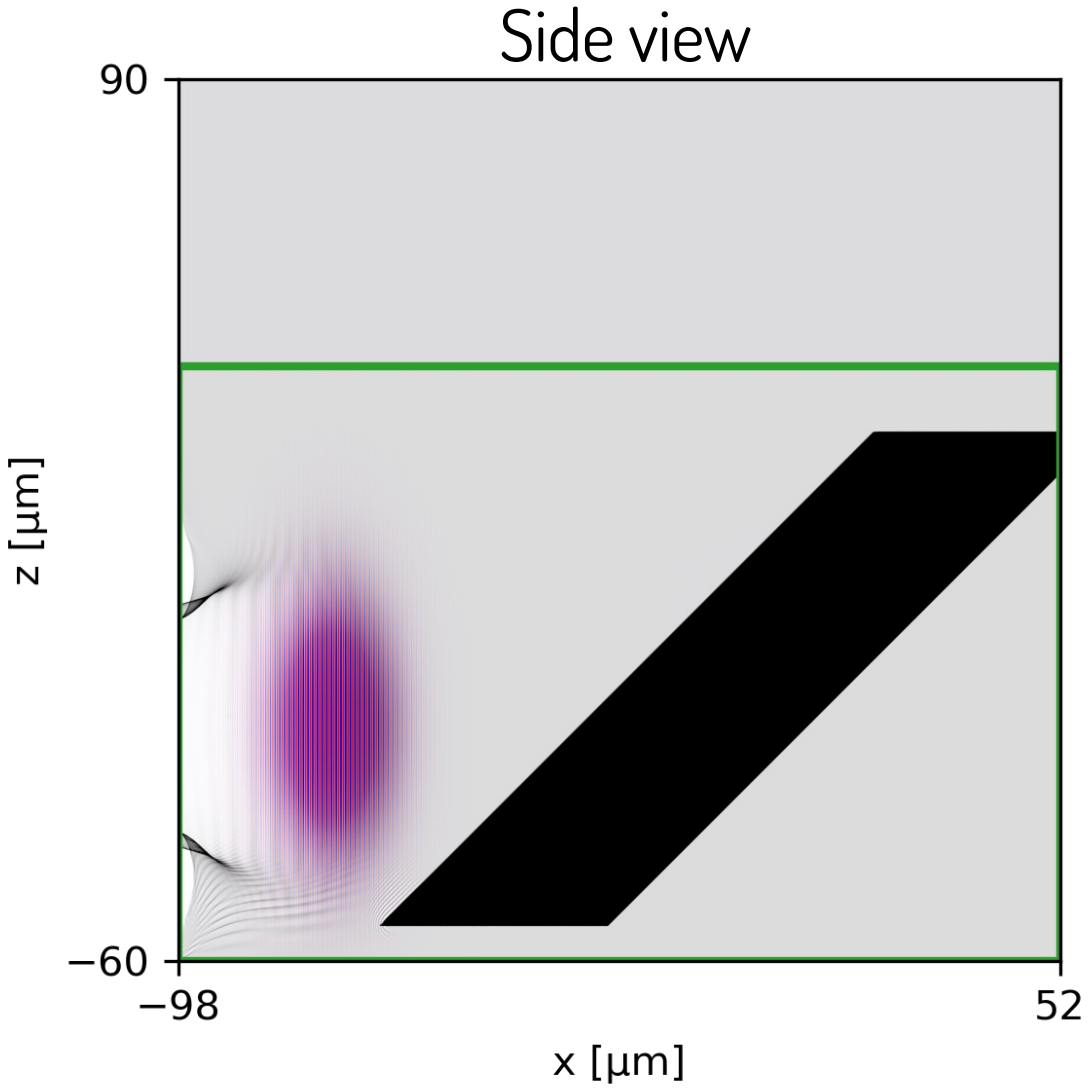

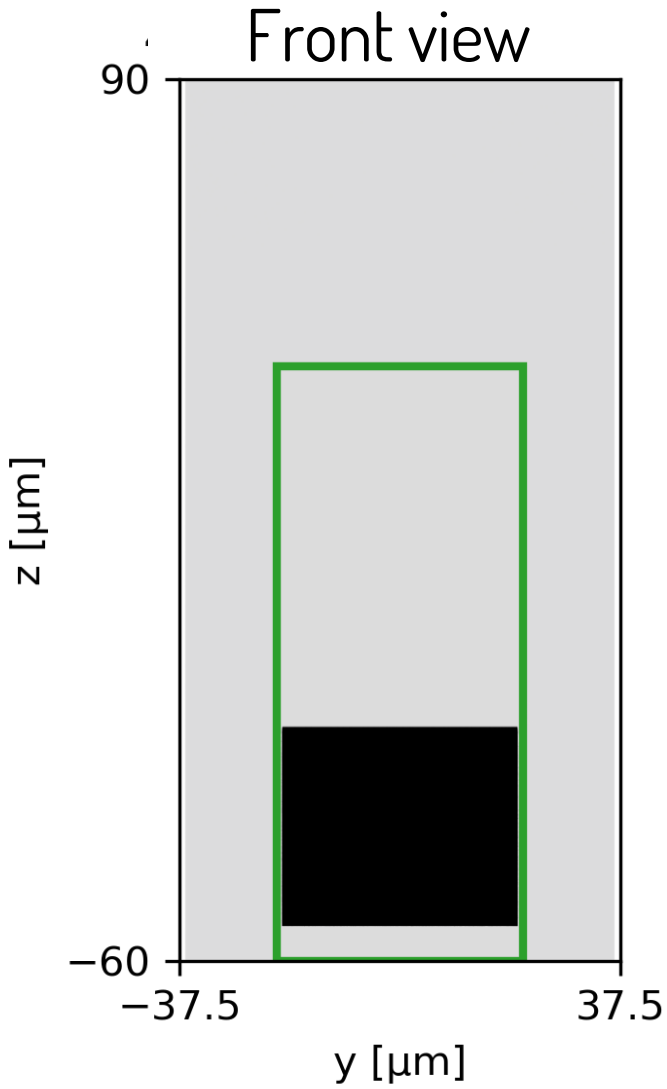

← 3D simulation on 4096 Summit nodes

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

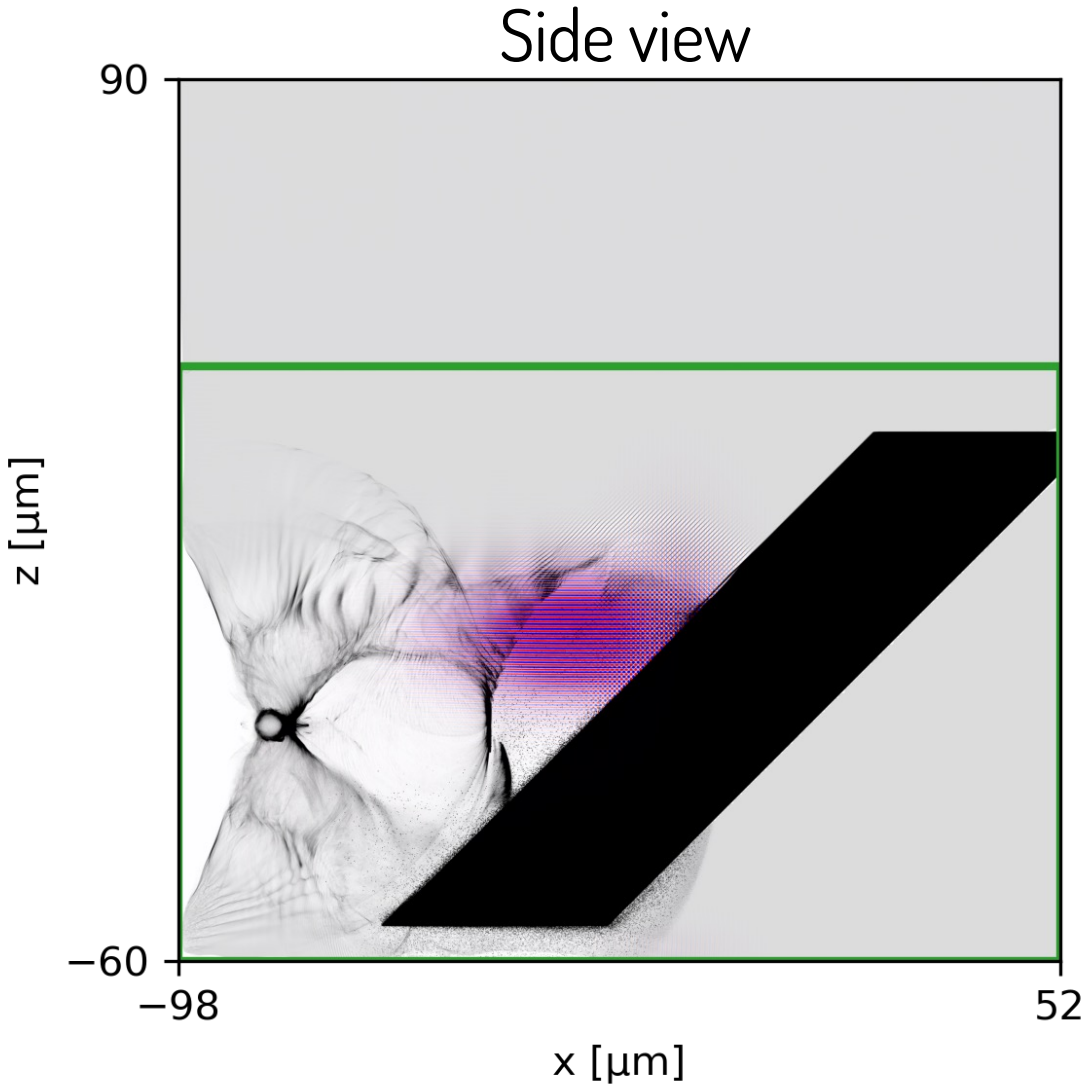

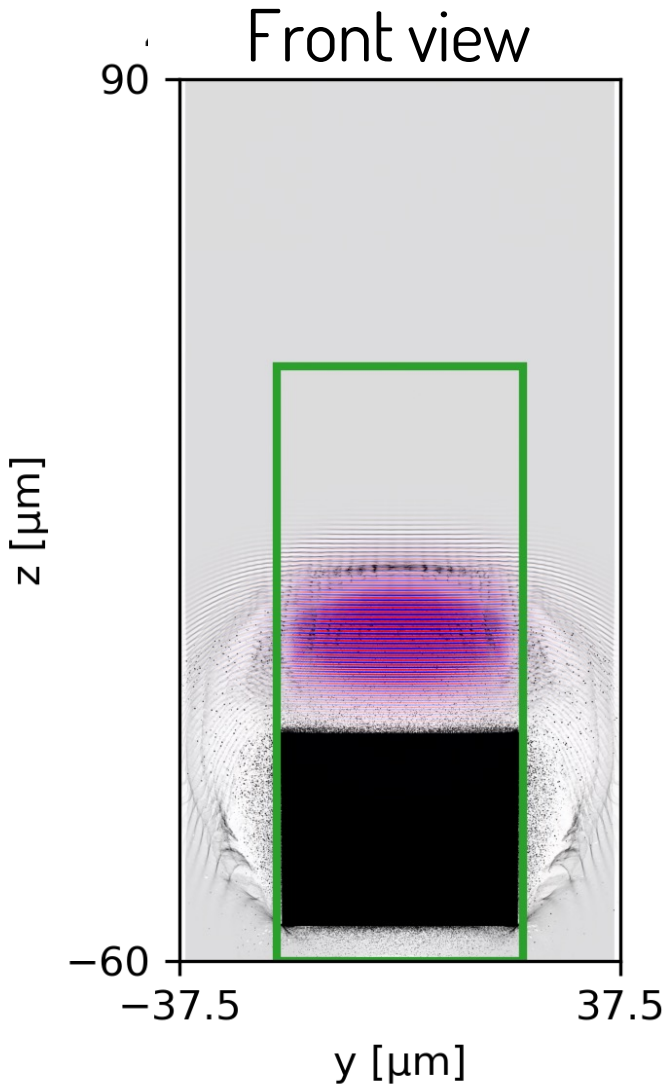

← 3D simulation on 4096 Summit nodes

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

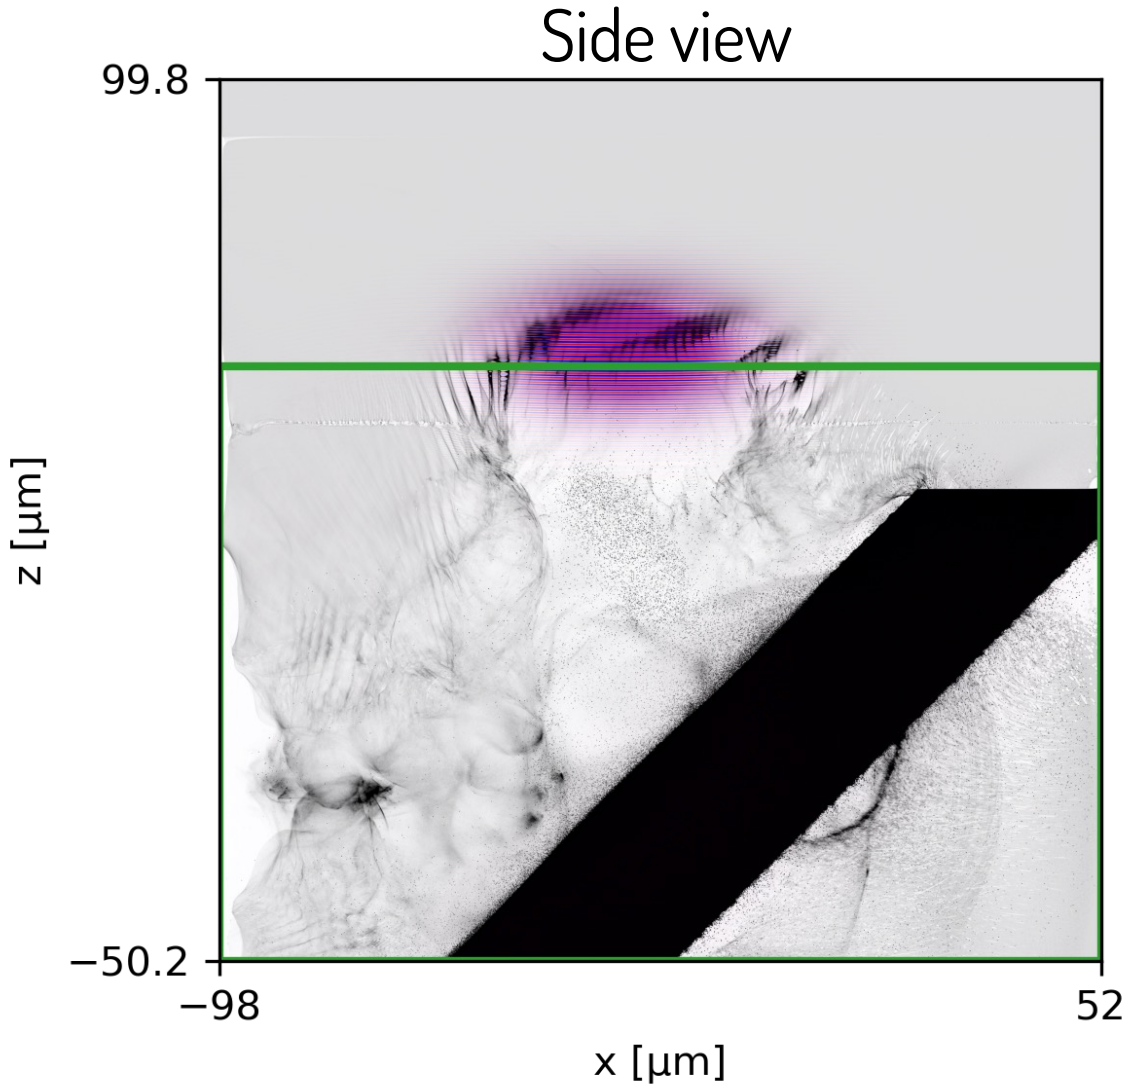

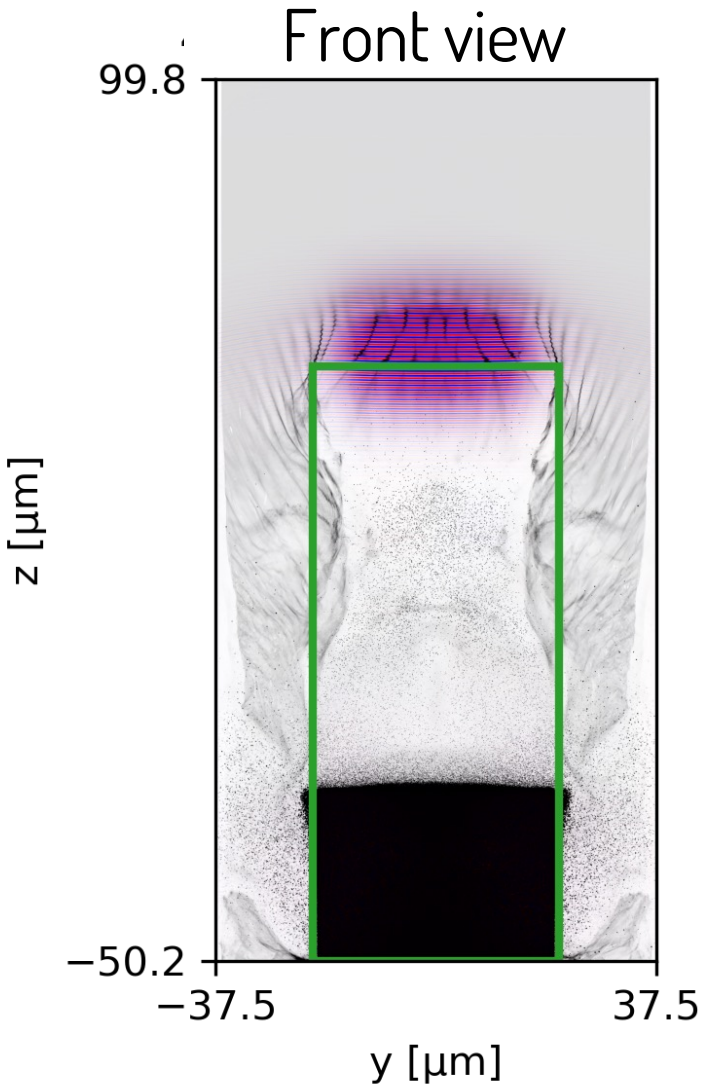

← 3D simulation on 4096 Summit nodes

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

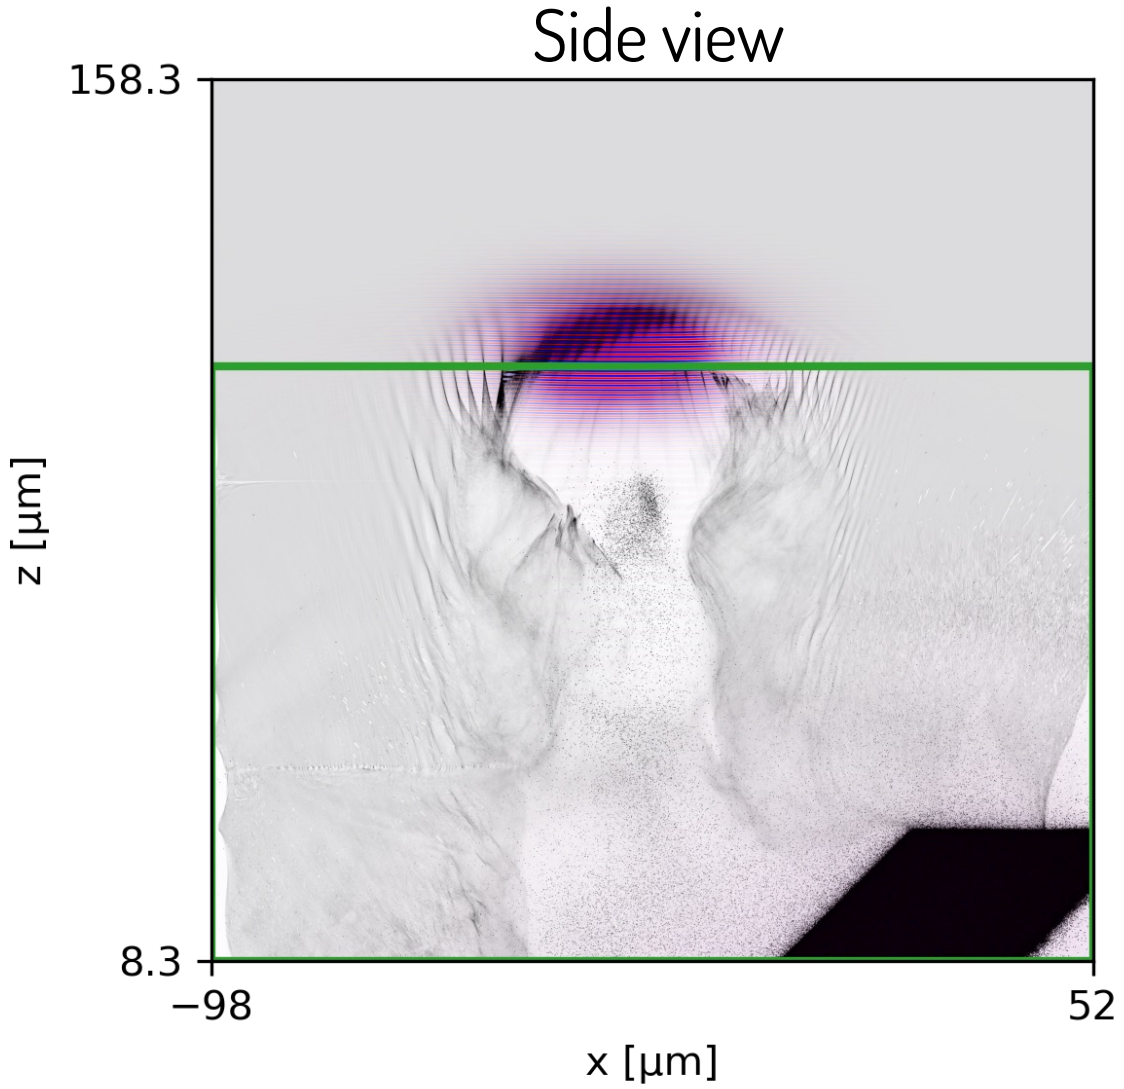

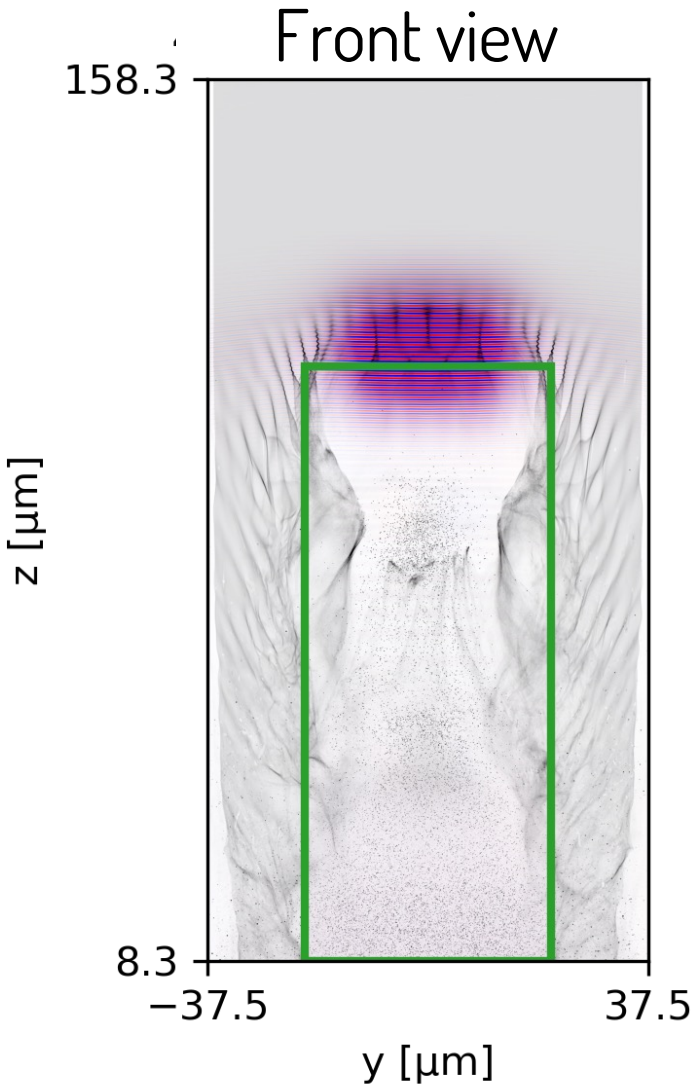

← 3D simulation on 4096 Summit nodes

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

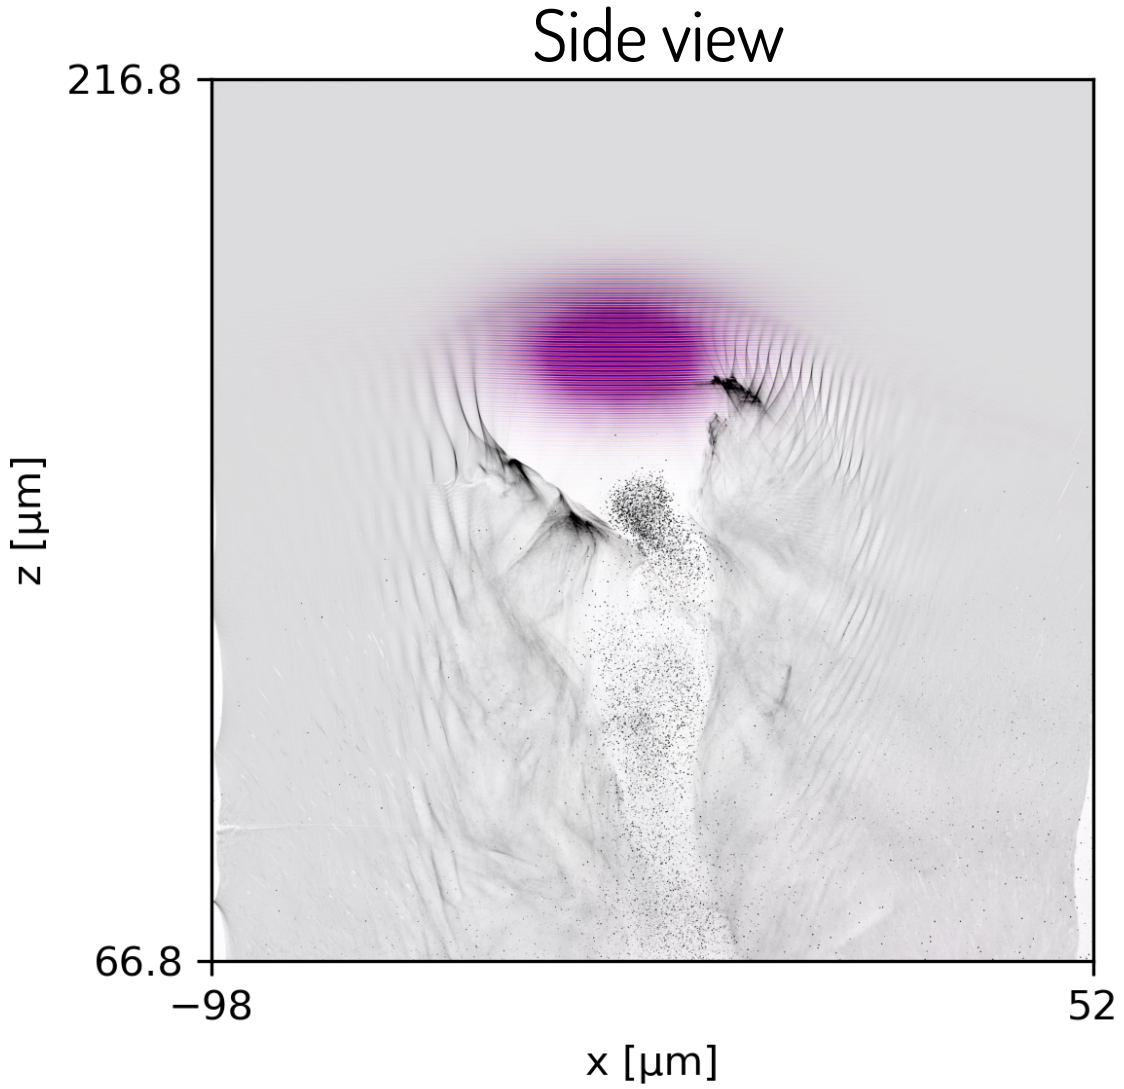

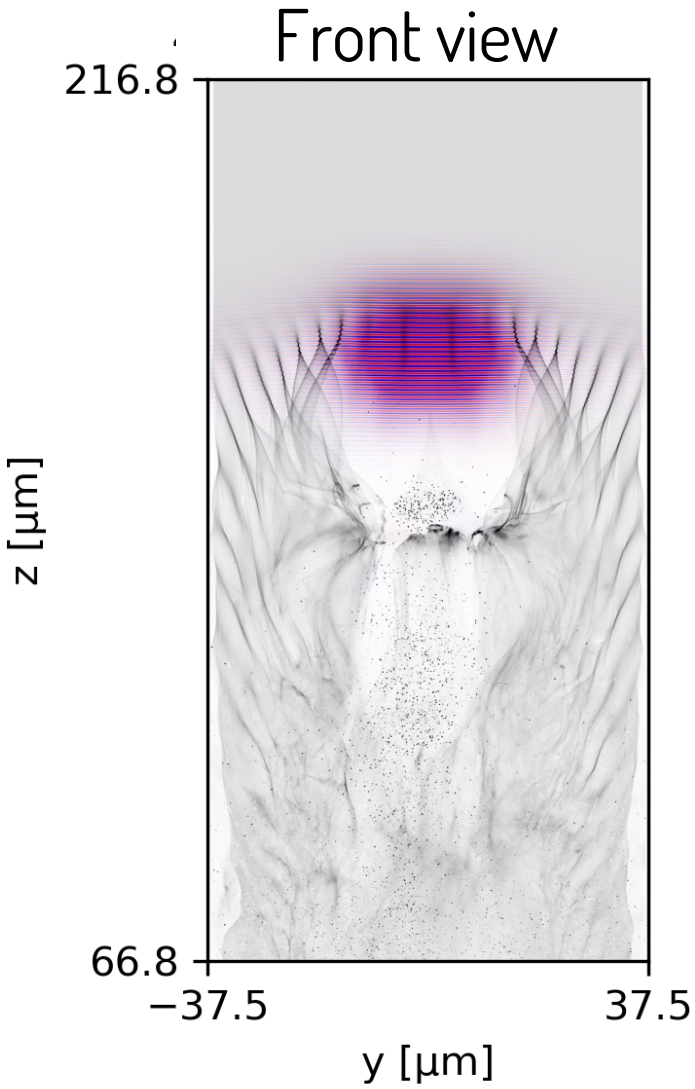

← 3D simulation on 4096 Summit nodes

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

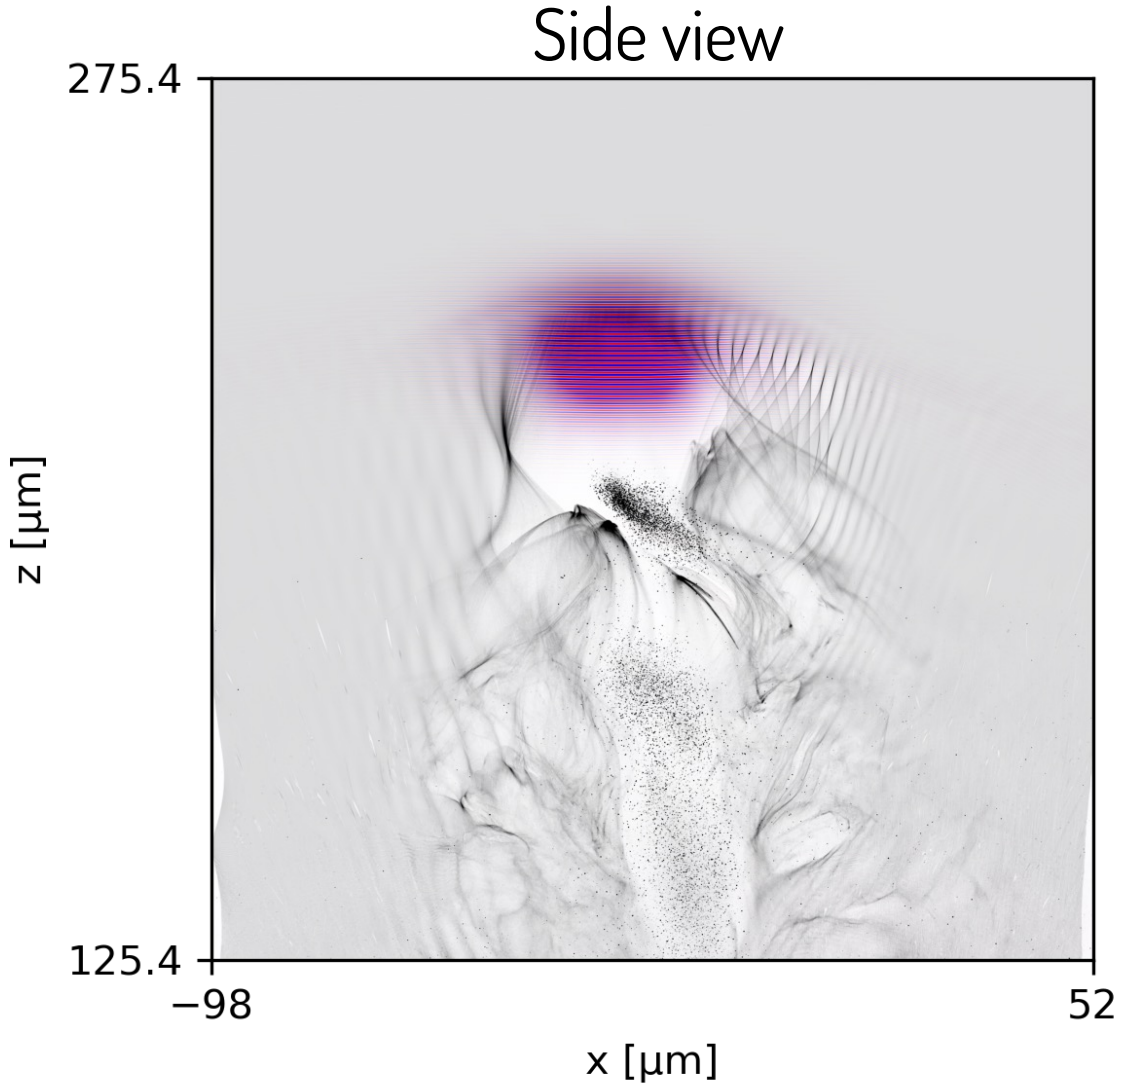

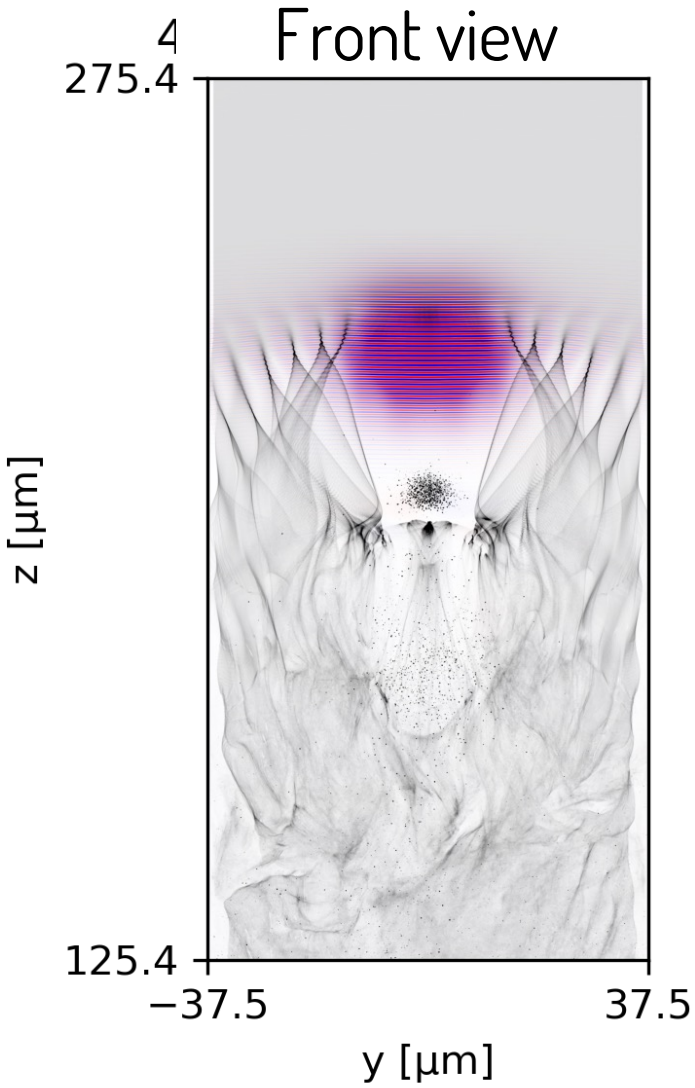

← 3D simulation on 4096 Summit nodes

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

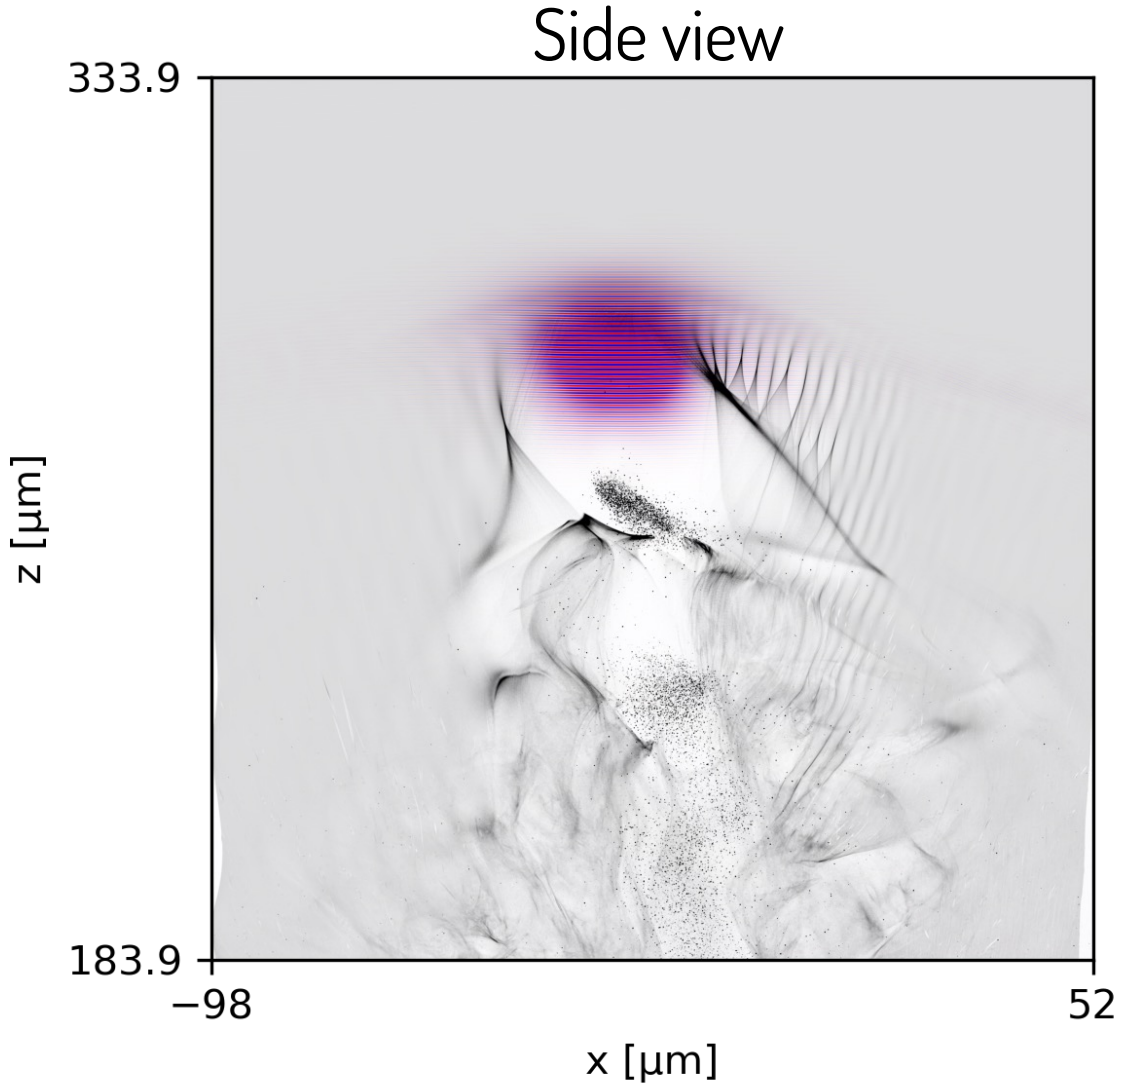

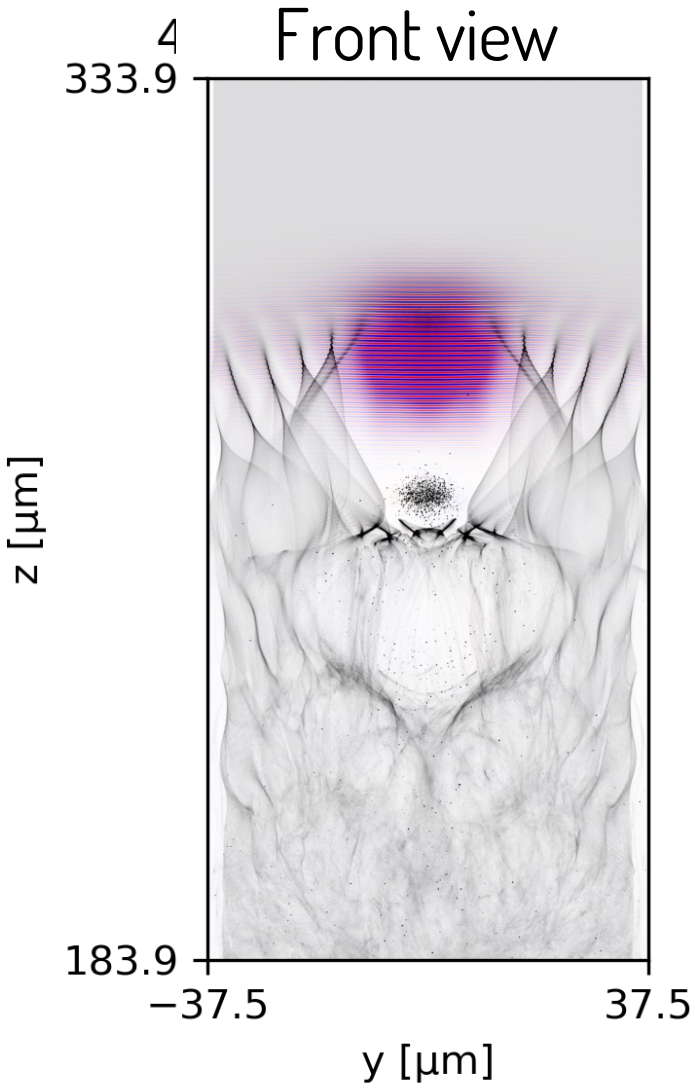

← 3D simulation on 4096 Summit nodes

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

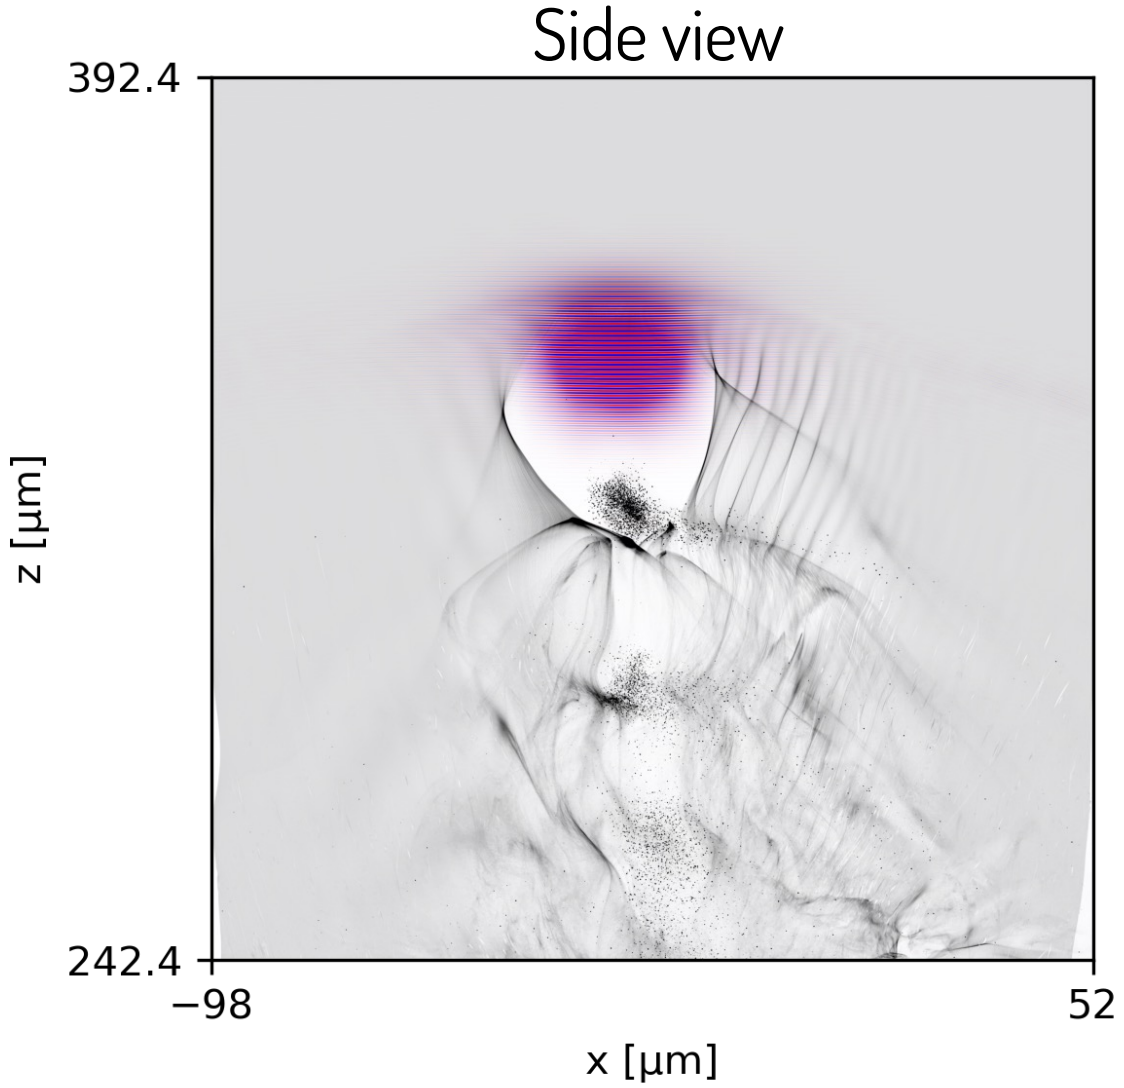

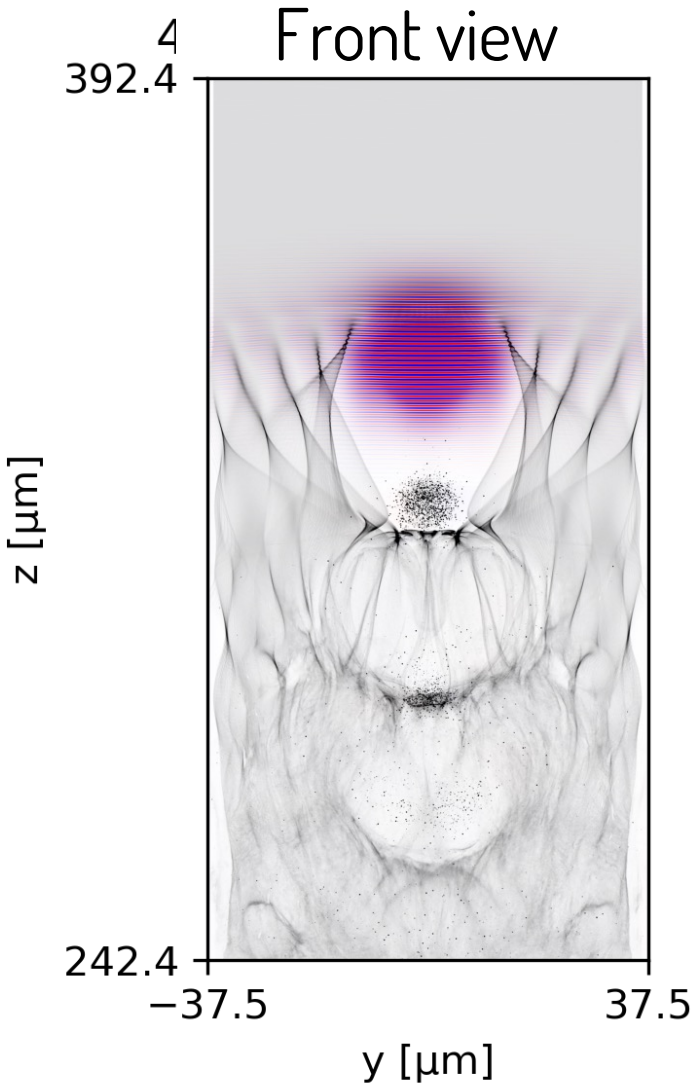

← 3D simulation on 4096 Summit nodes

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

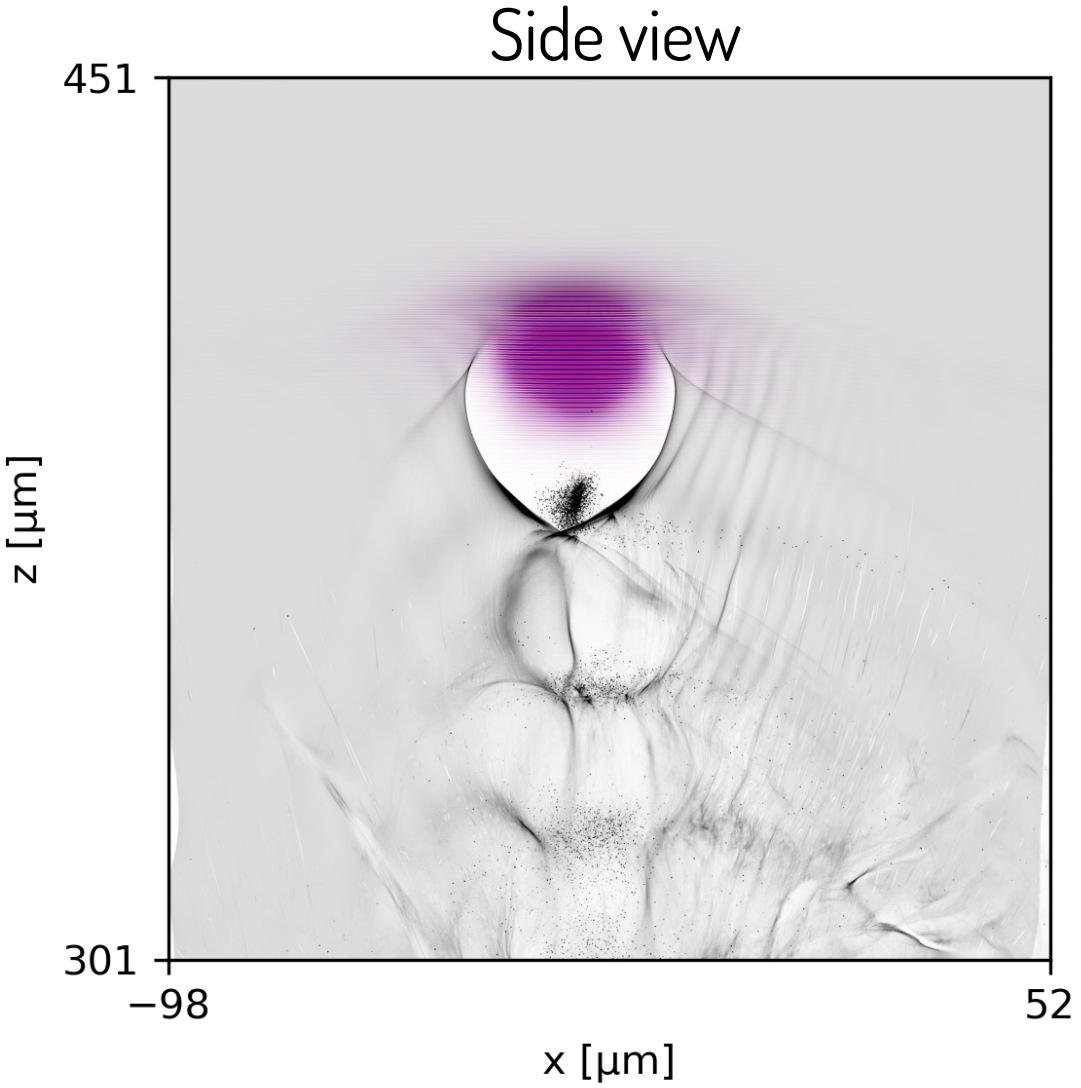

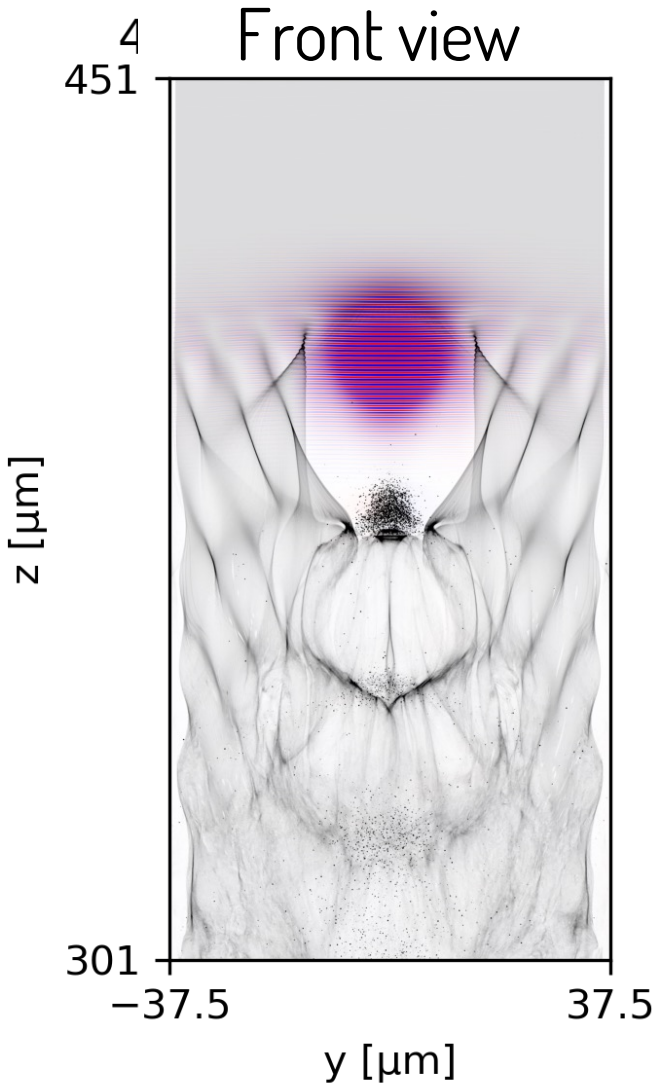

← 3D simulation on 4096 Summit nodes

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

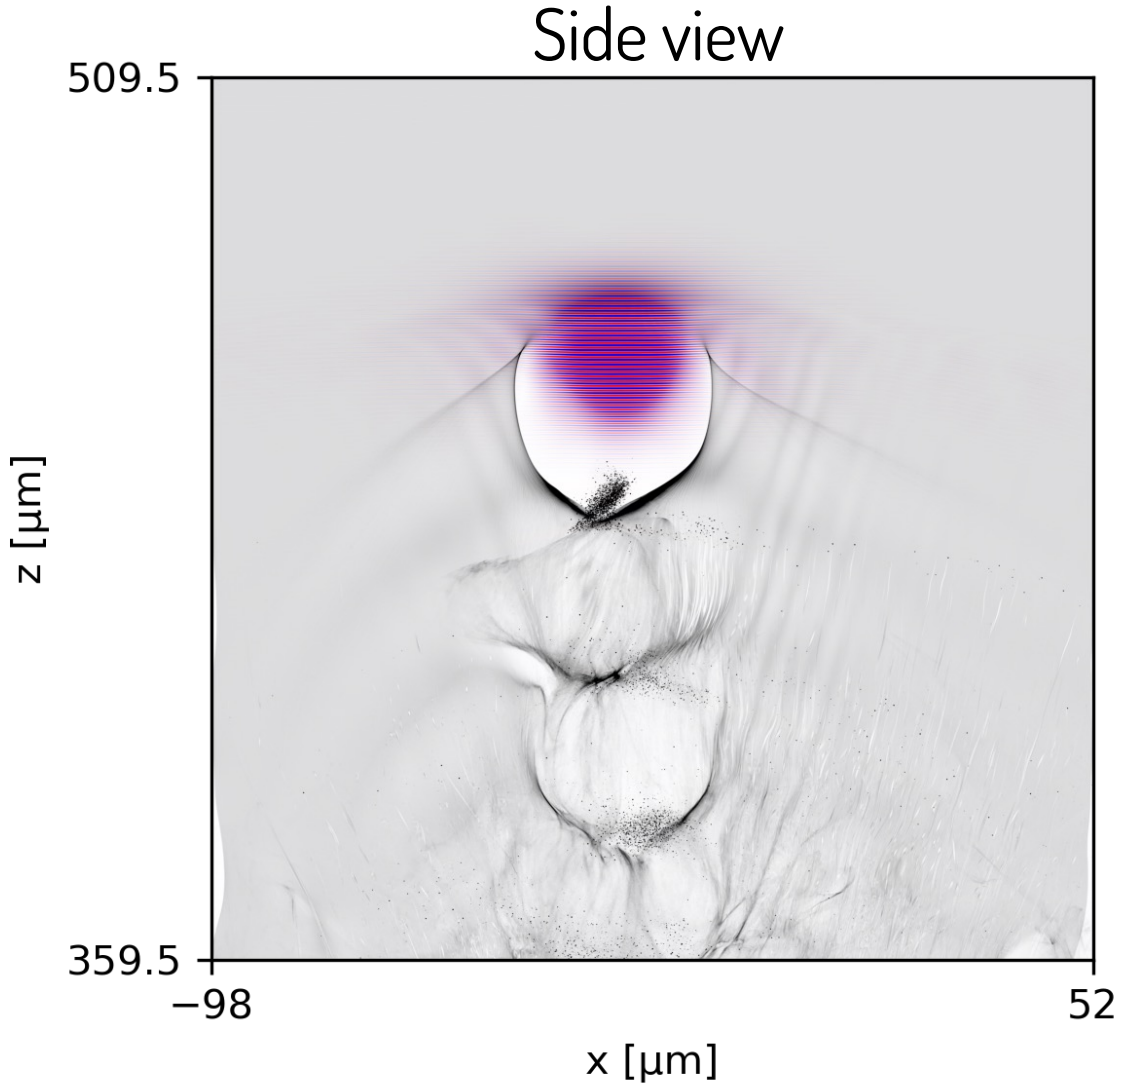

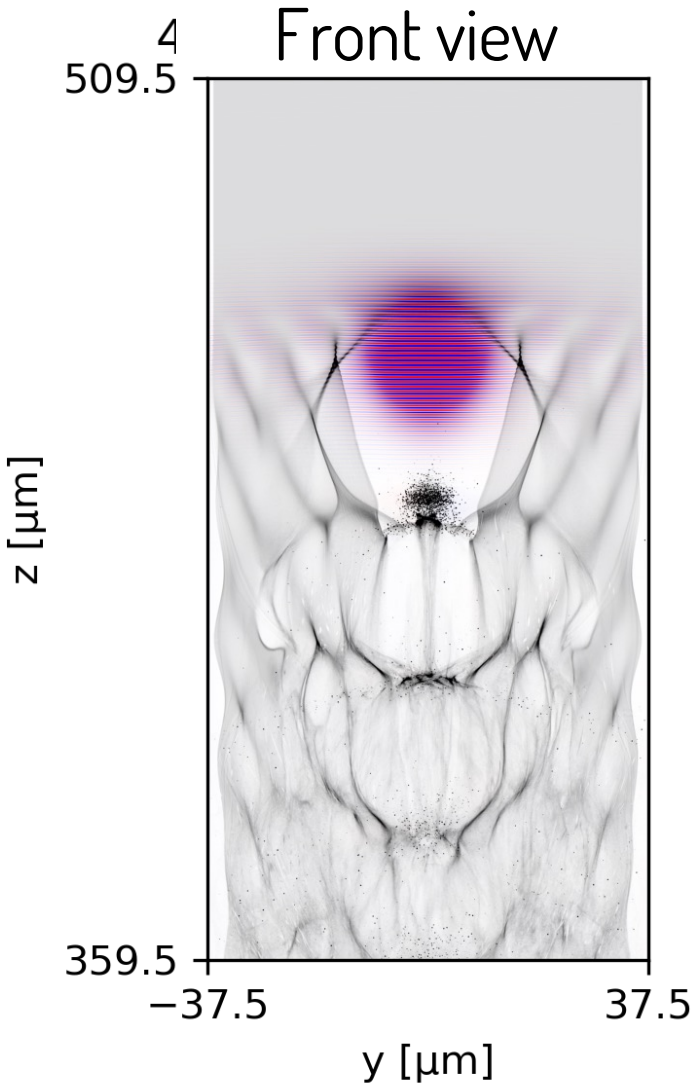

← 3D simulation on 4096 Summit nodes

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

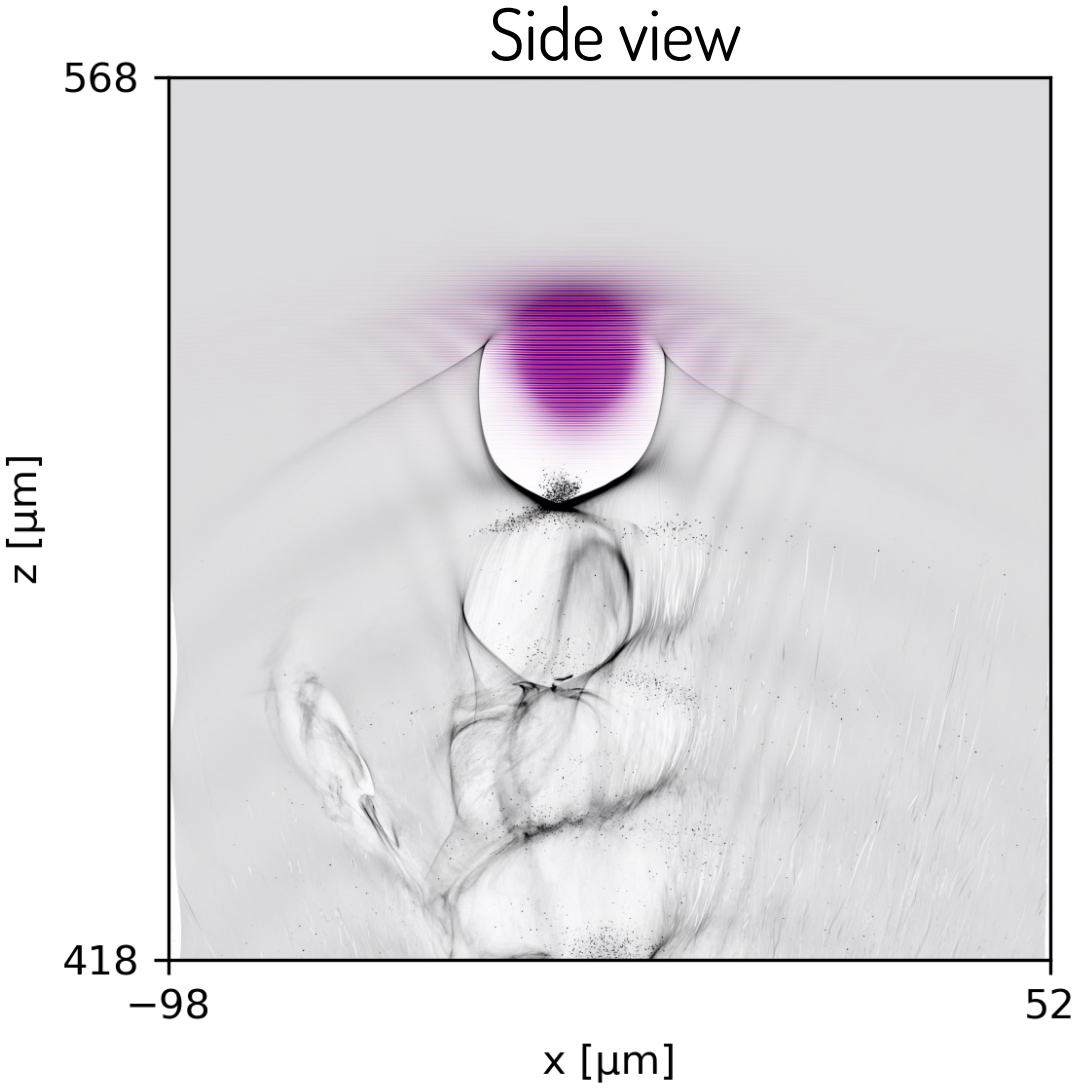

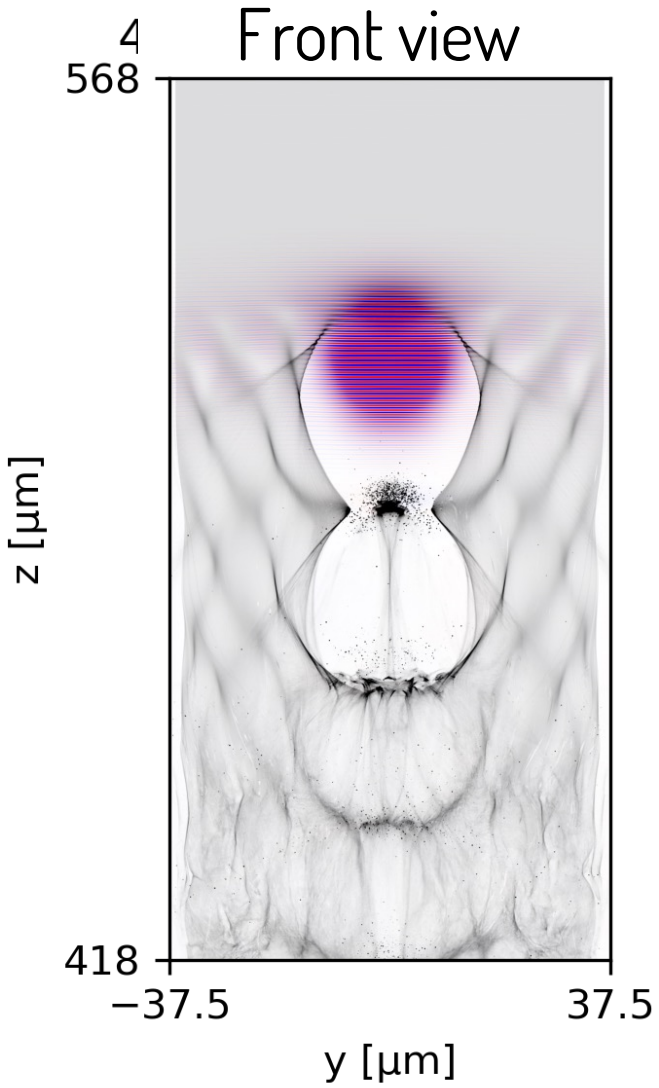

← 3D simulation on 4096 Summit nodes

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

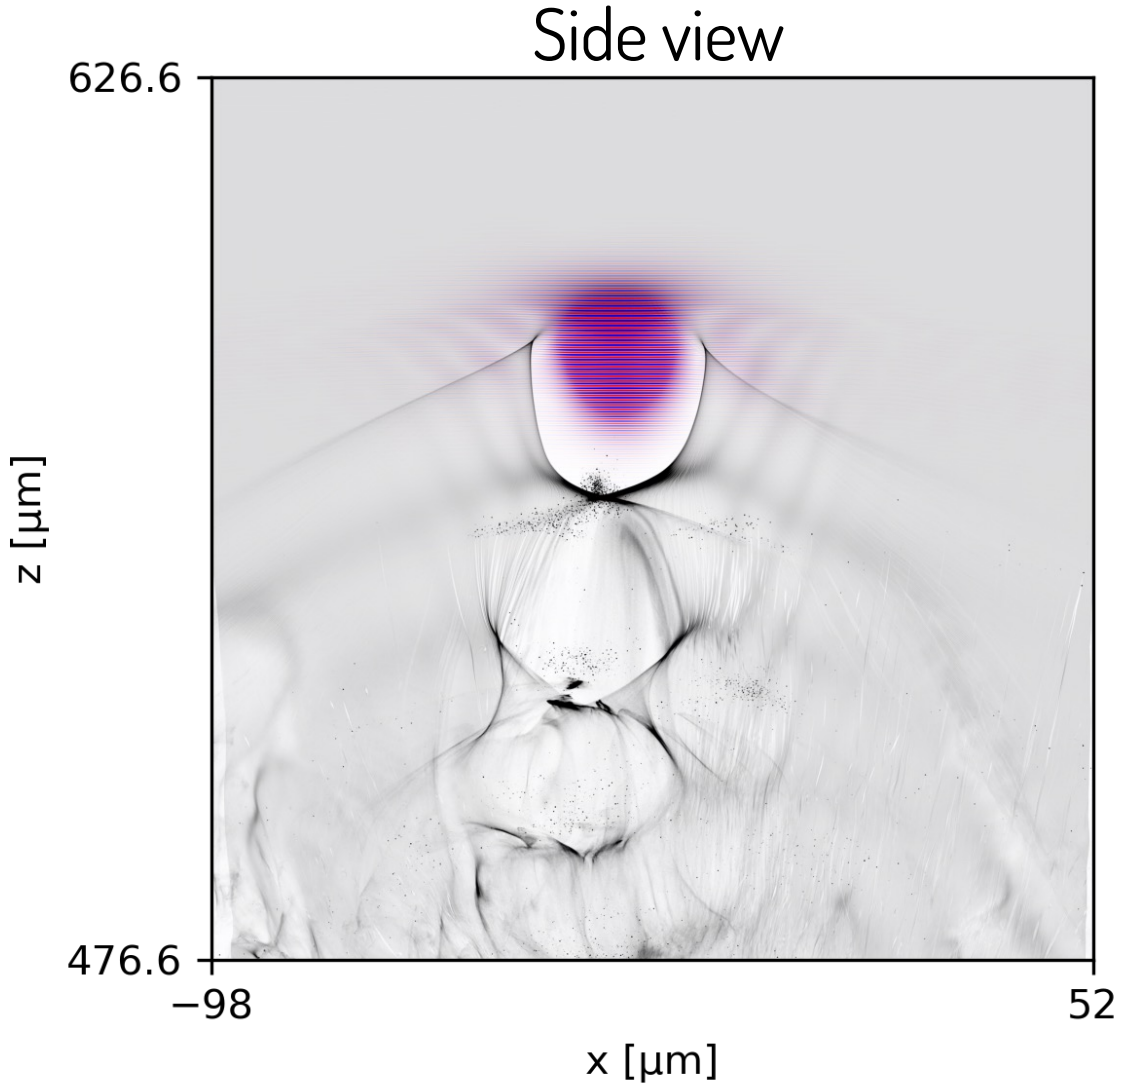

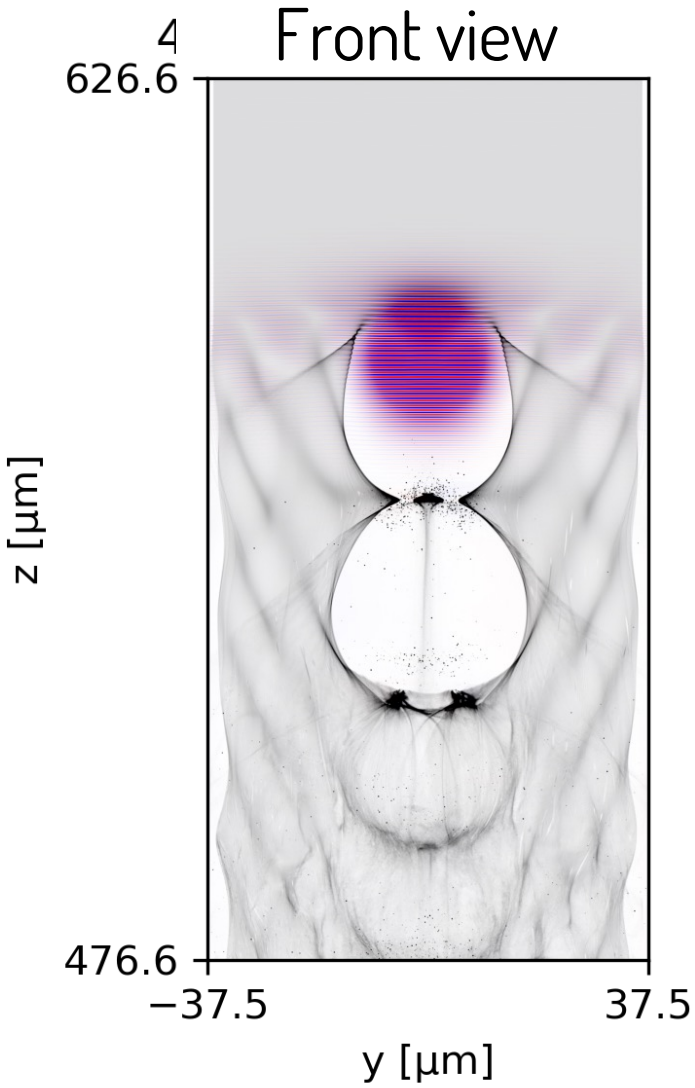

← 3D simulation on 4096 Summit nodes

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

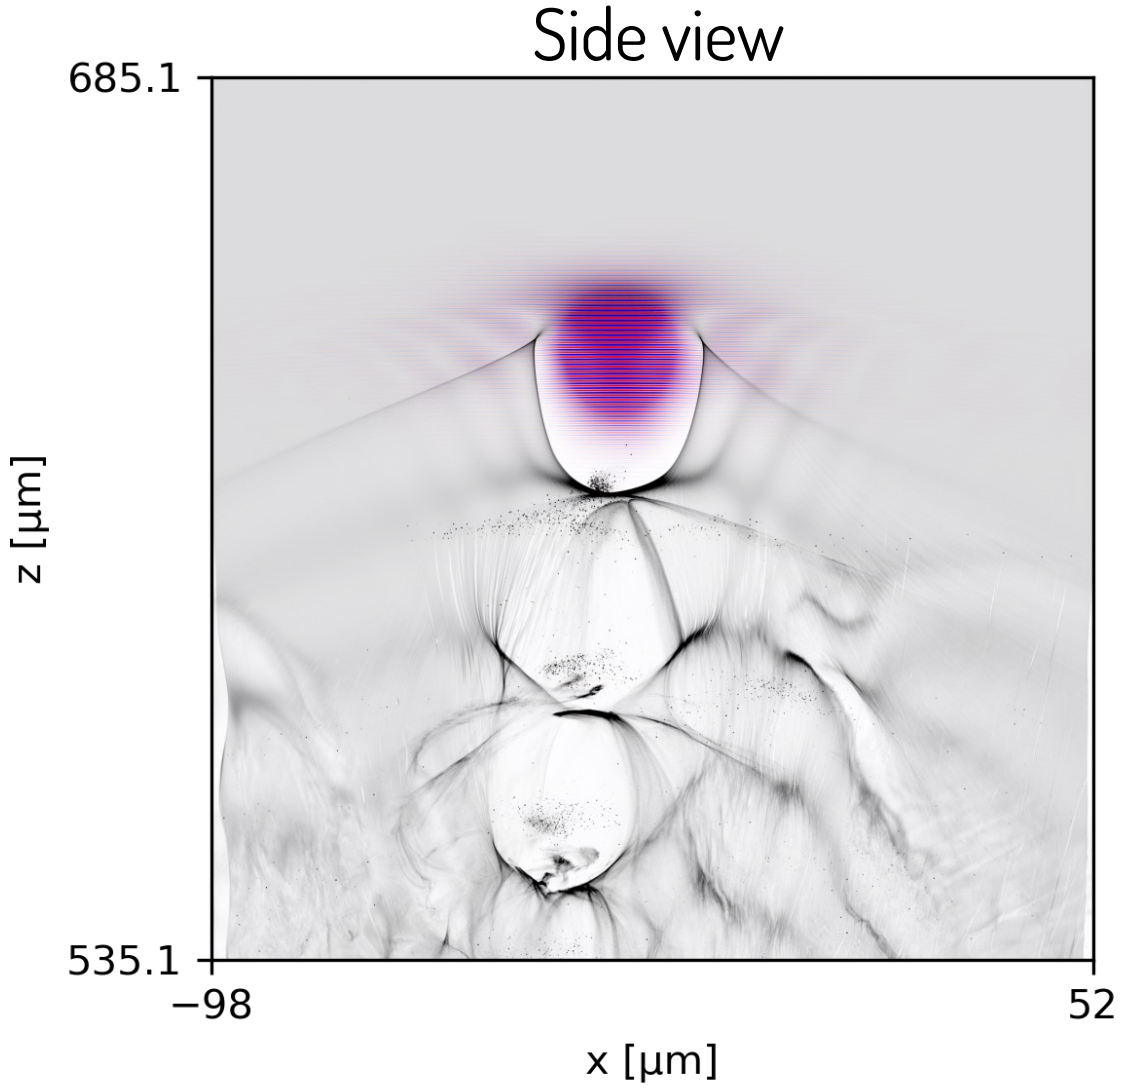

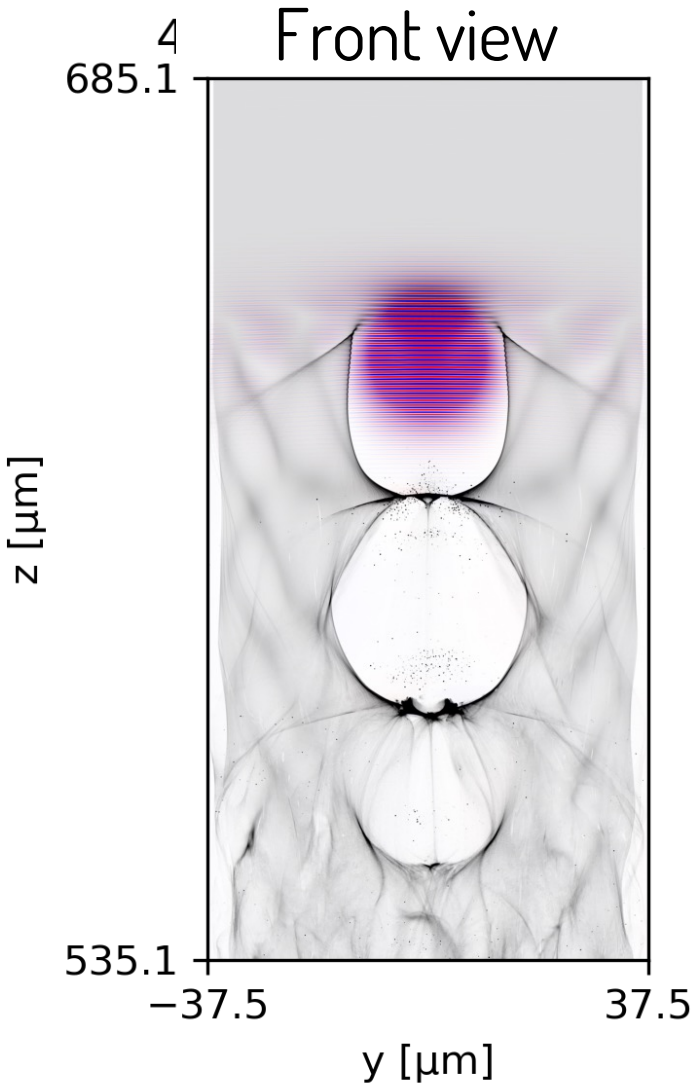

← 3D simulation on 4096 Summit nodes

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

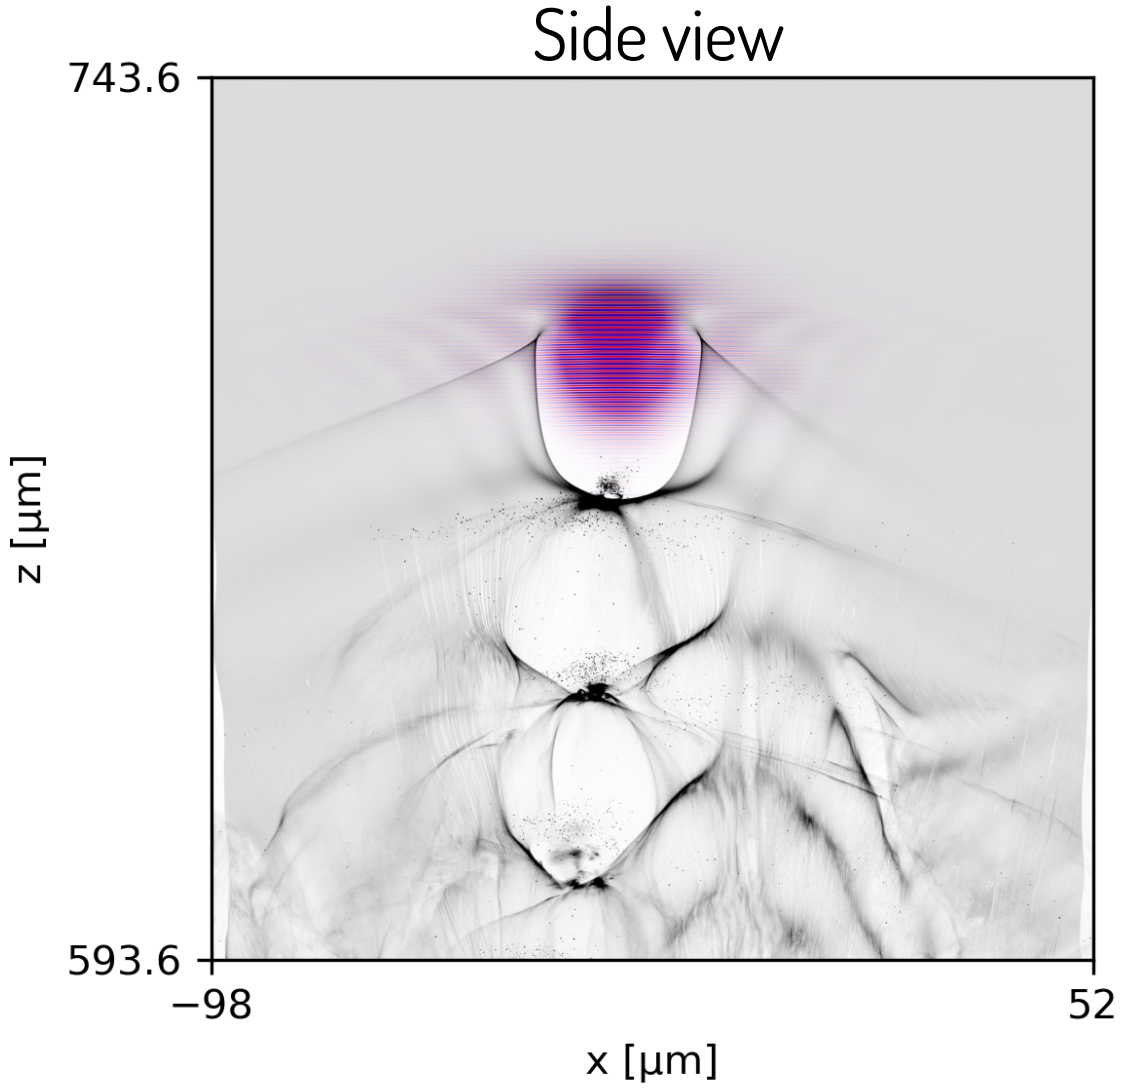

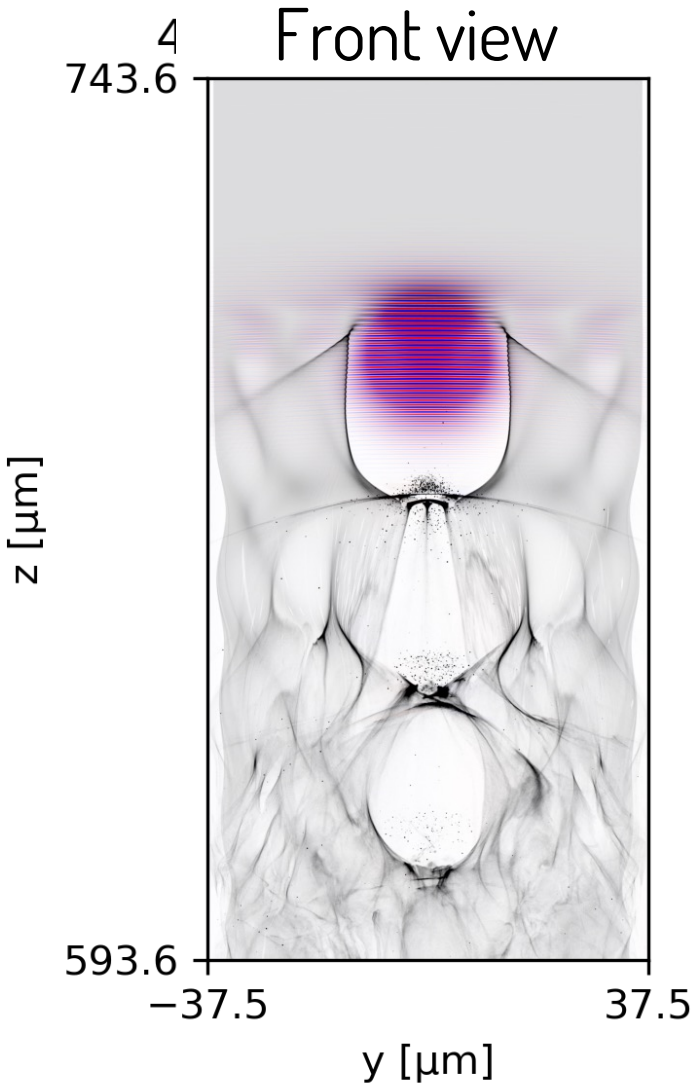

← 3D simulation on 4096 Summit nodes

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

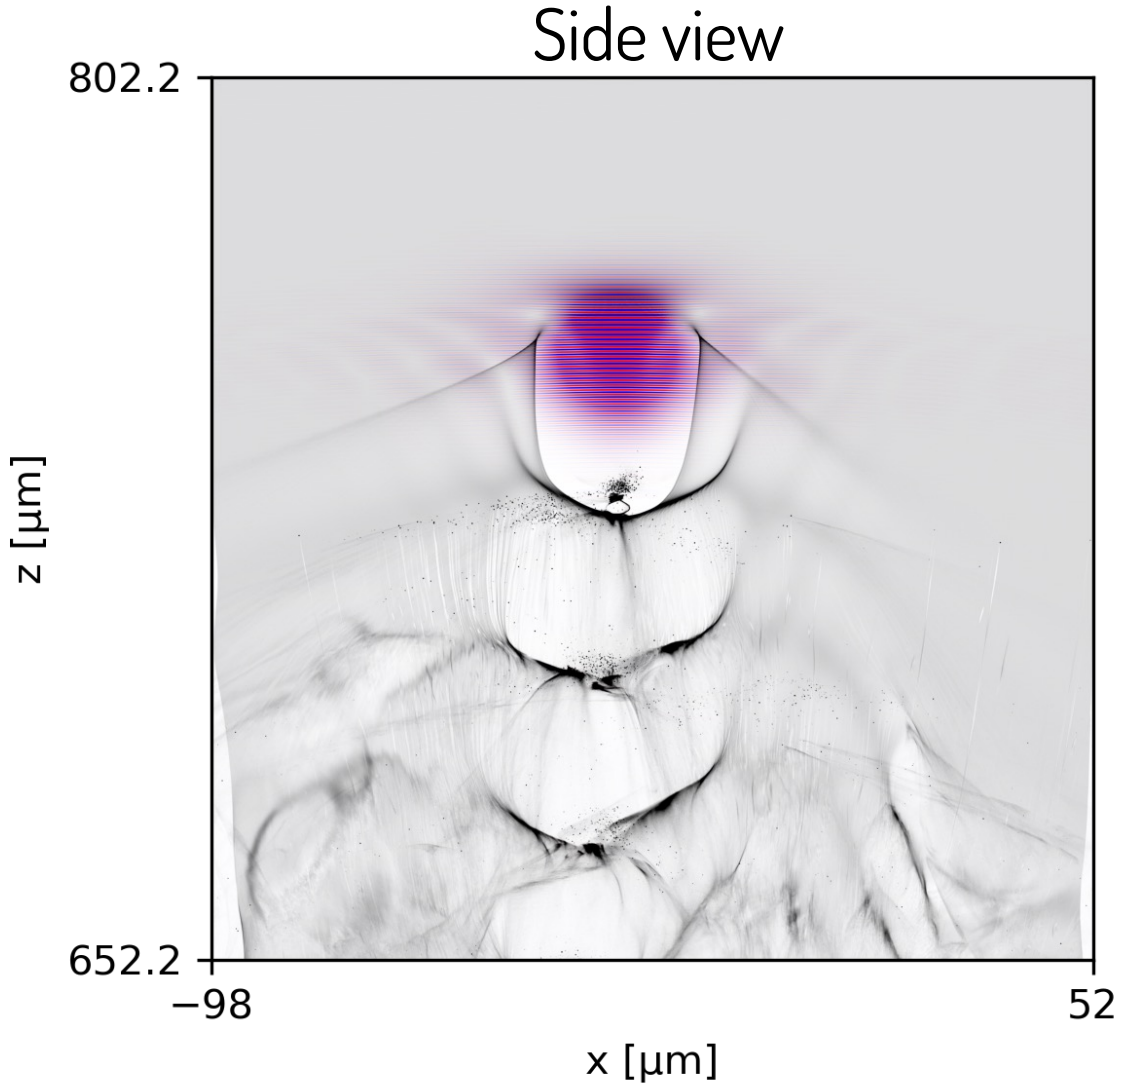

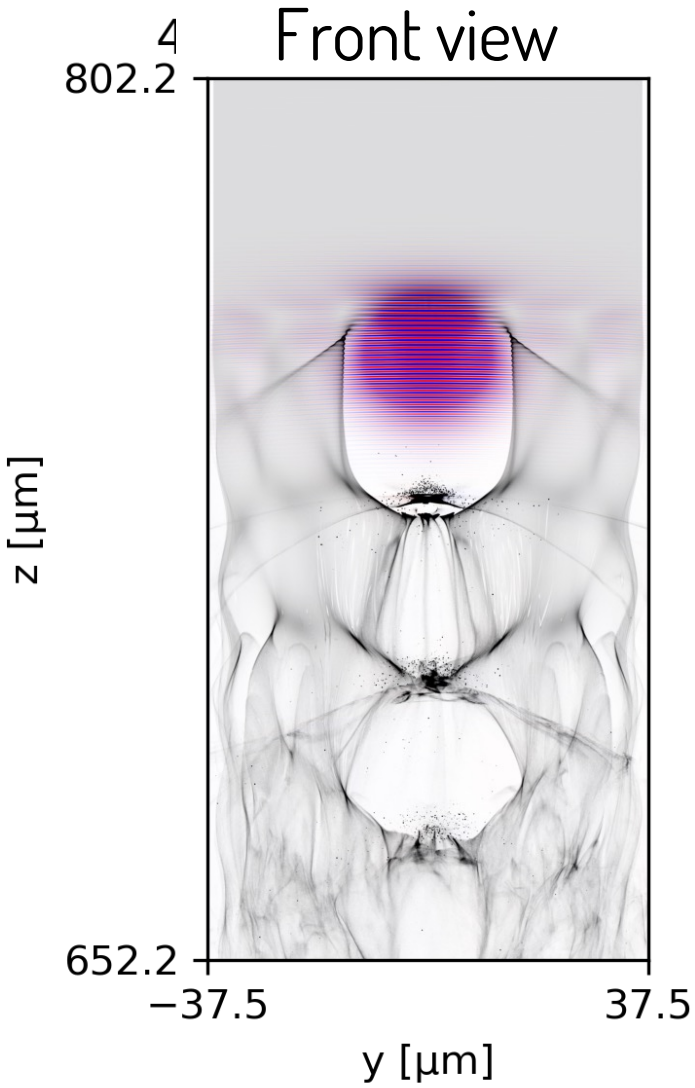

← 3D simulation on 4096 Summit nodes

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

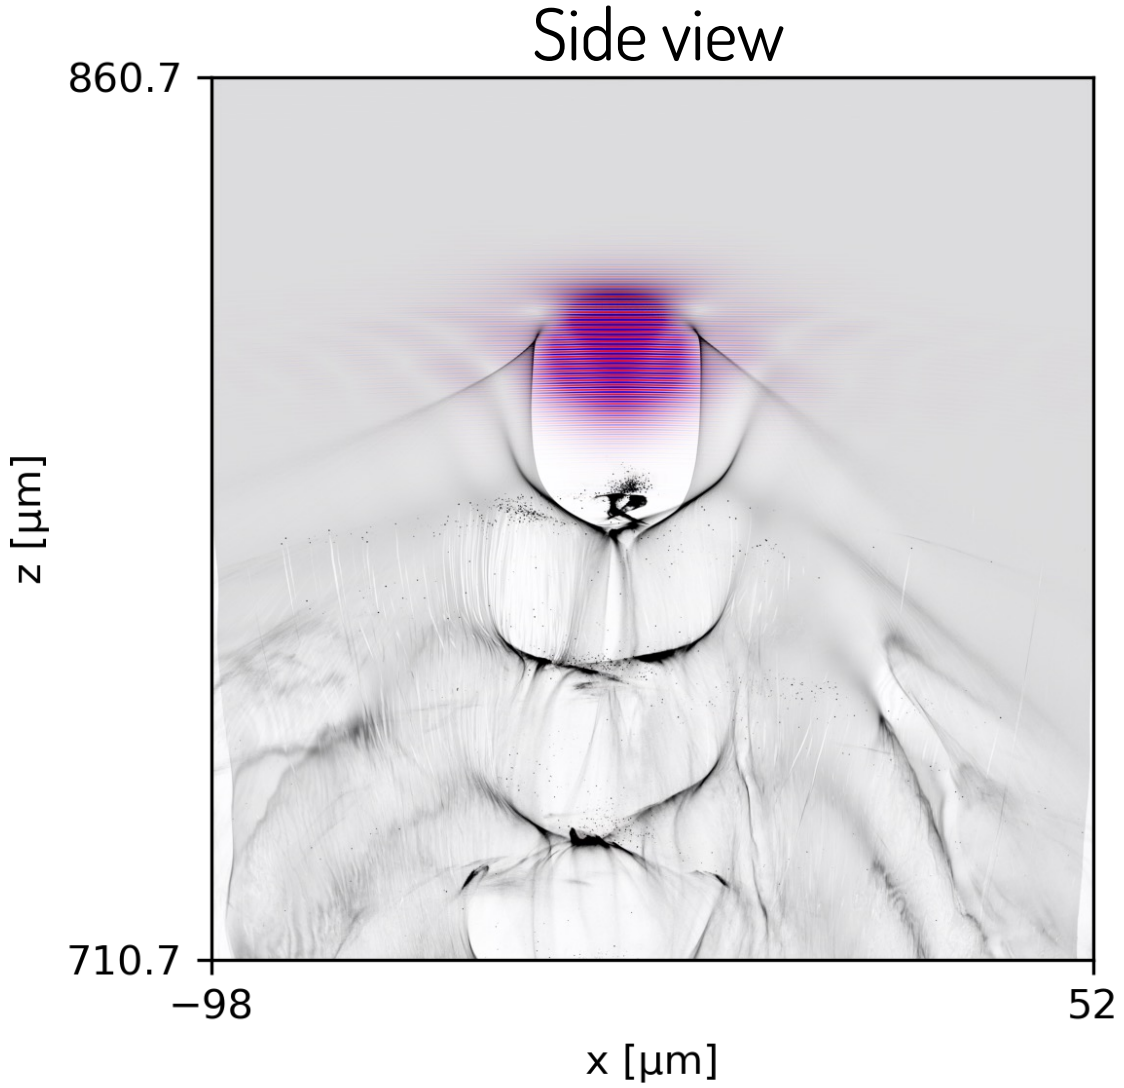

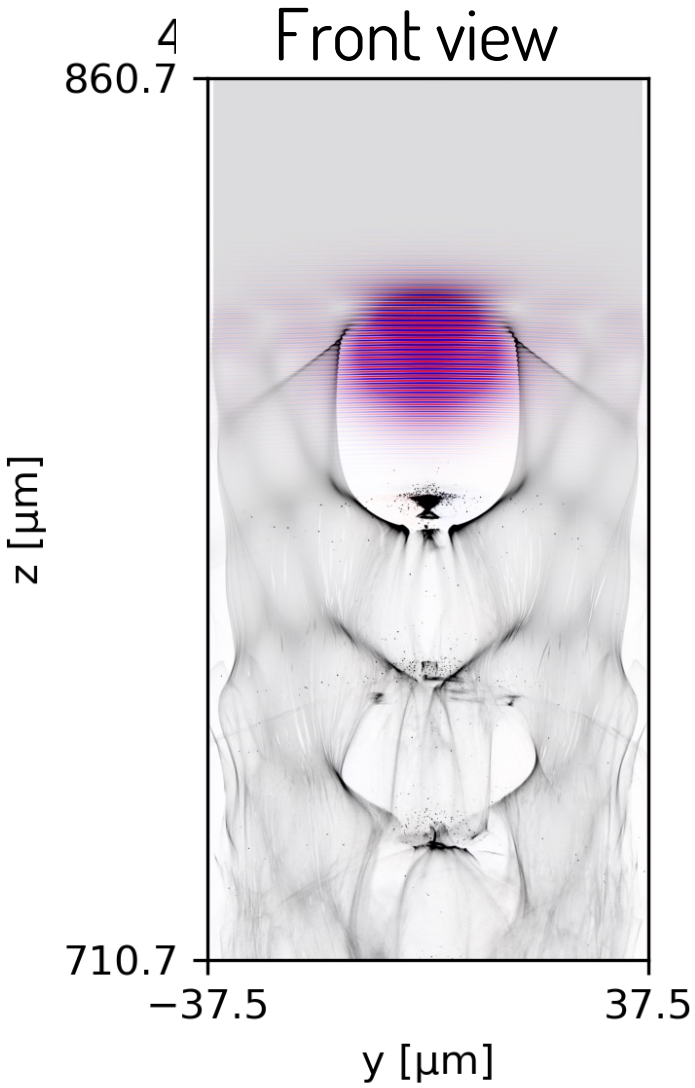

← 3D simulation on 4096 Summit nodes

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

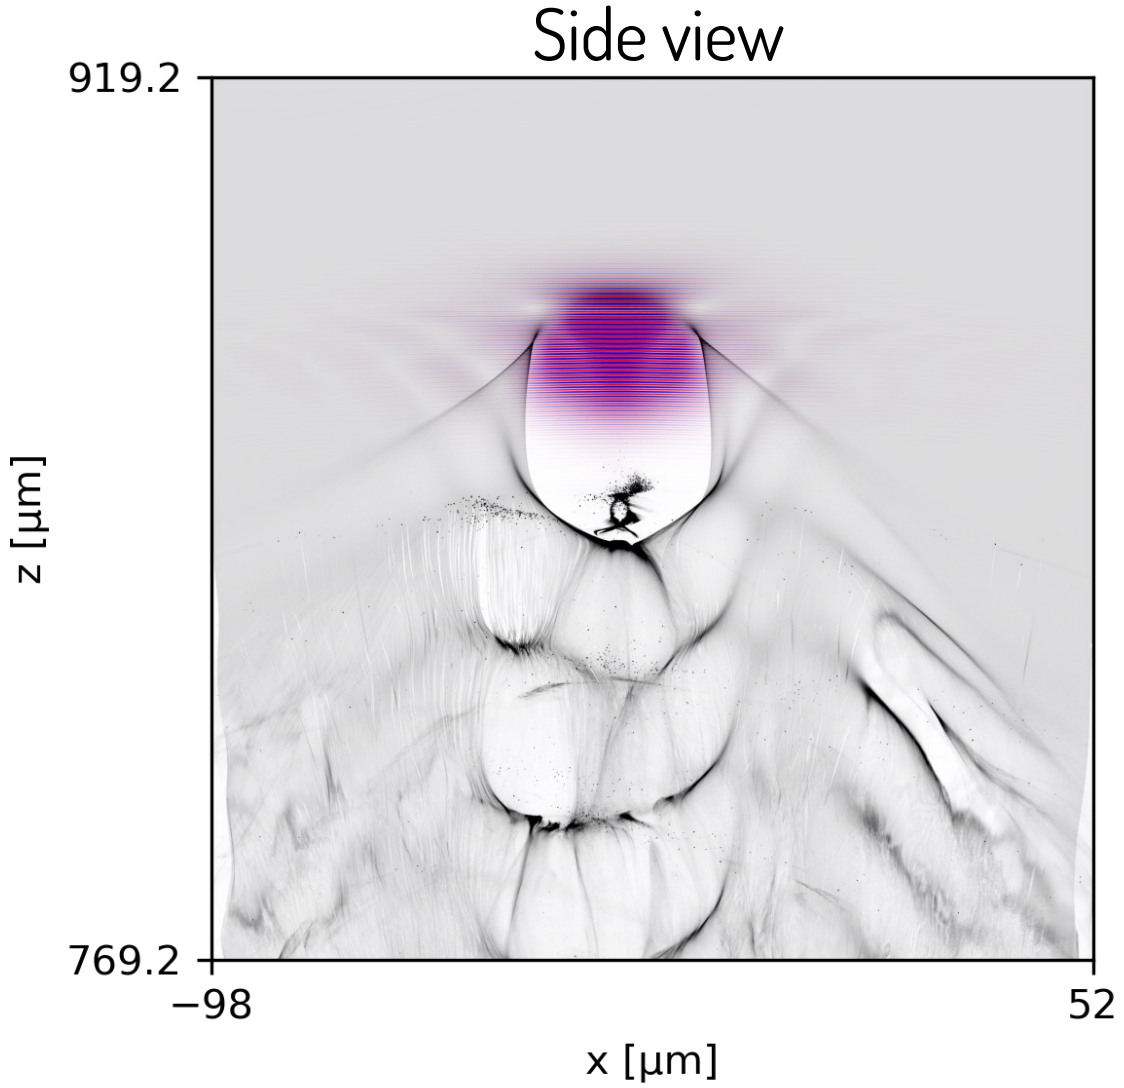

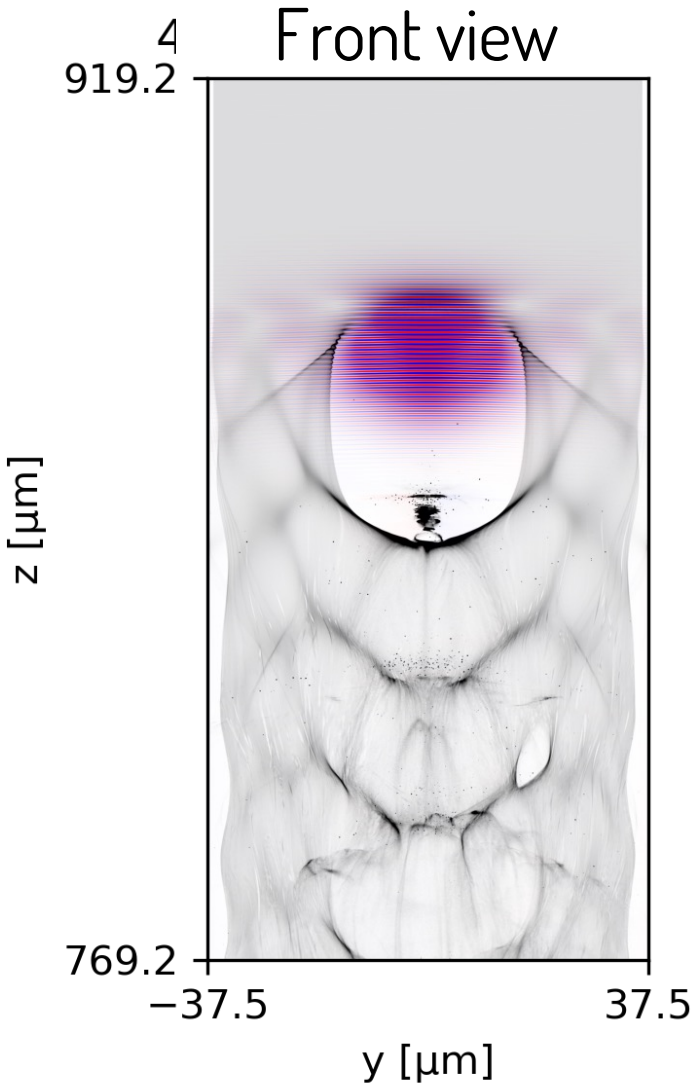

← 3D simulation on 4096 Summit nodes

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

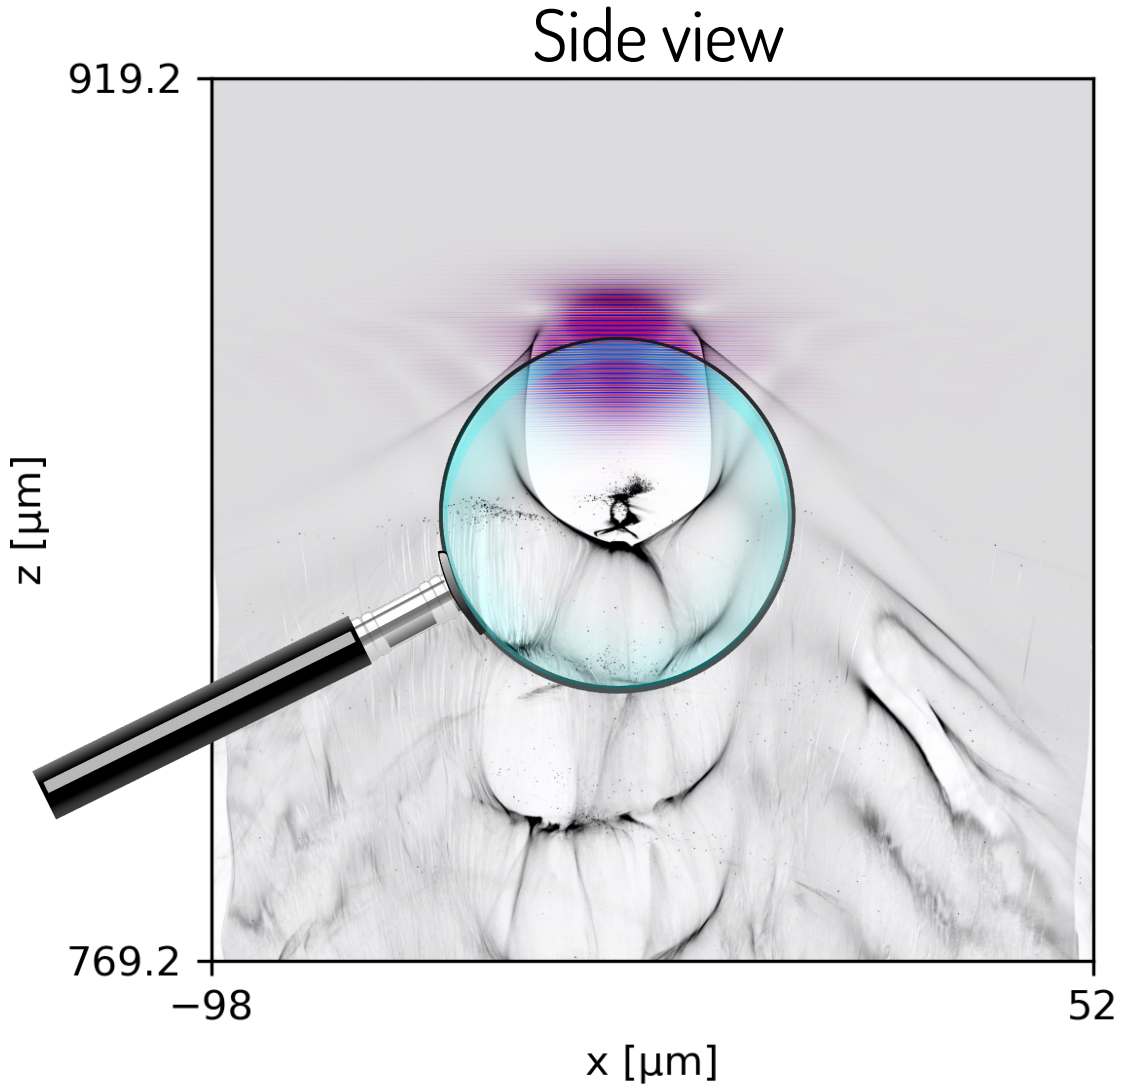

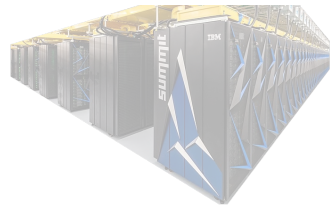

←We are mainly concerned with the properties of these electrons

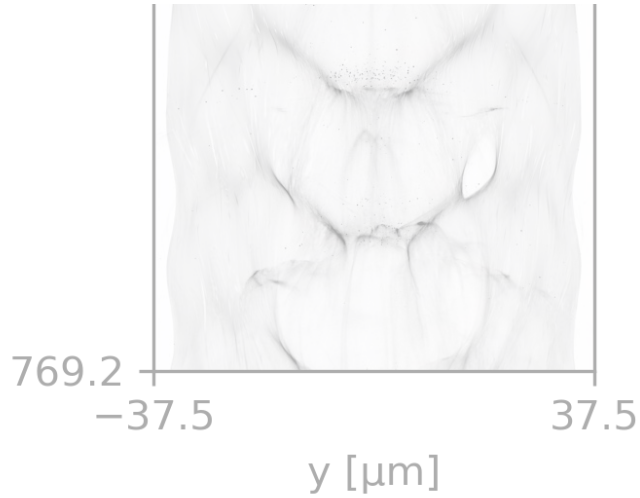

Our simulations with a PW-class laser show that we can accelerate a substantial amount of charge with high quality

After ~ 1mm
(acceleration still in progress)

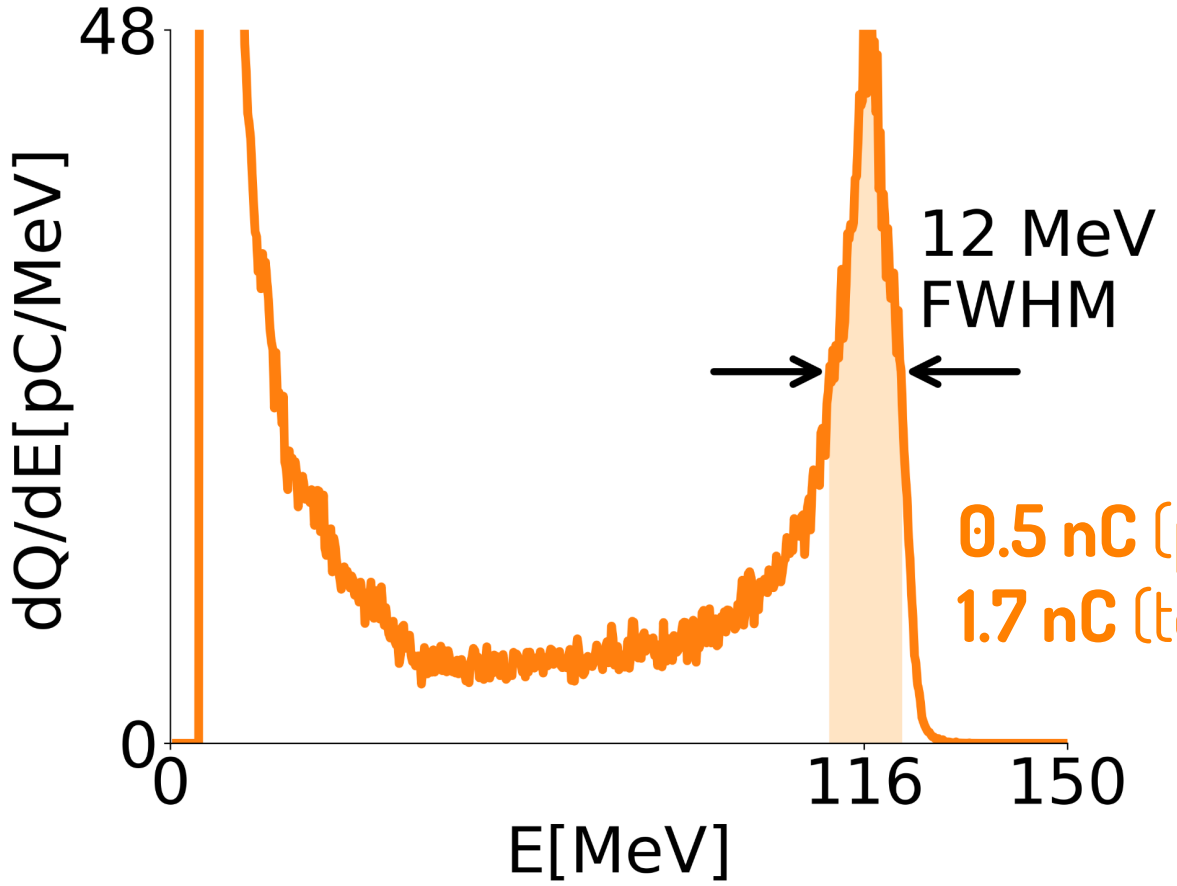

0.5 nC (peak) for a
1.7 nC (total) **PW laser**

Production runs on
Frontier, Fugaku and Summit

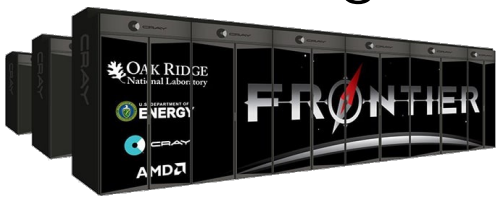

Exascale simulations informed the design of the first experimental validation of our concept (at LOA)

In collaboration with Adrien Leblanc, UPX

Laser parameters

E = 400 mJ

waist = 17 µm

P_{peak} = 10 TW

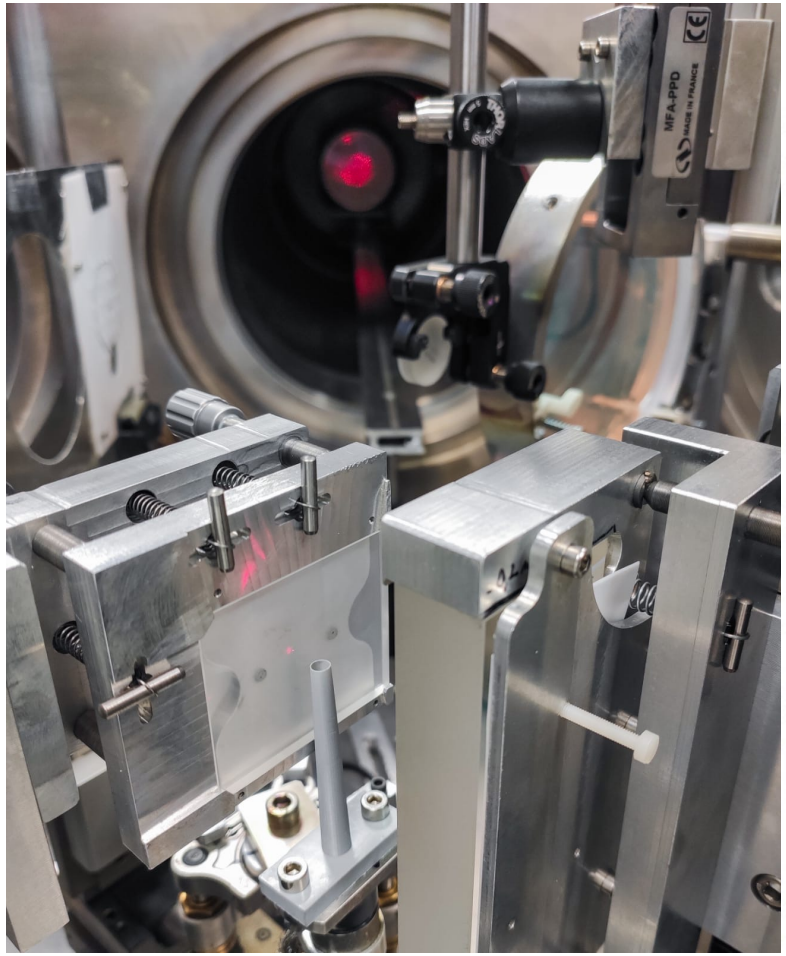

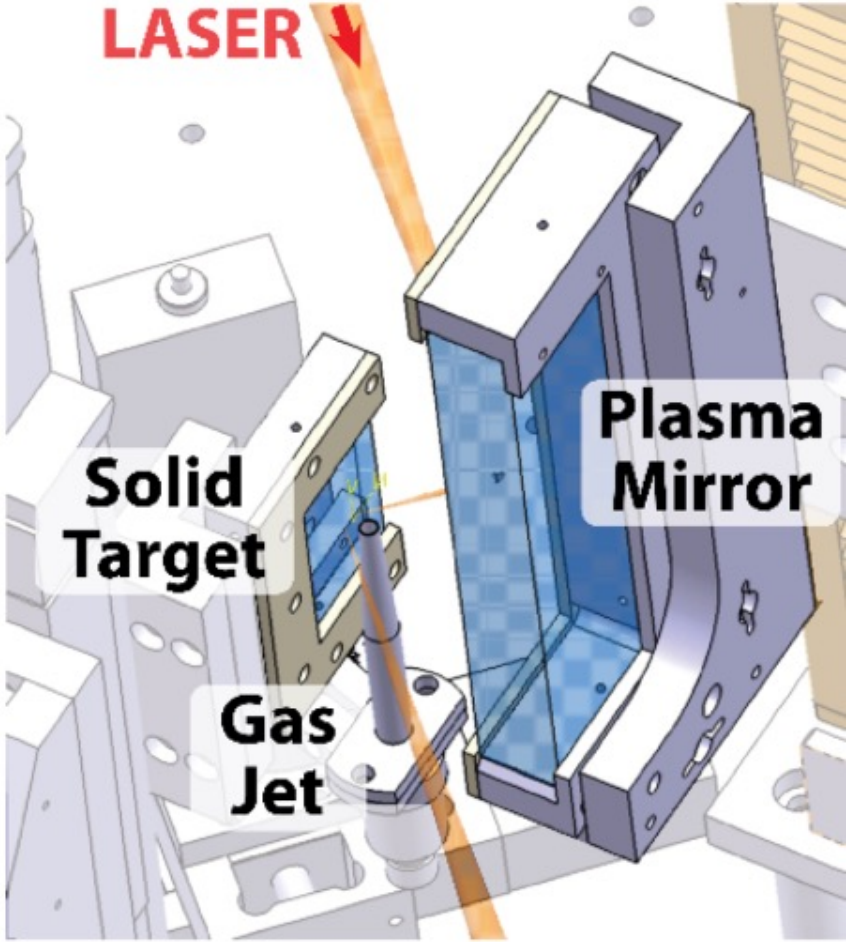

Exascale simulations informed the design of the first experimental validation of our concept (at LOA)

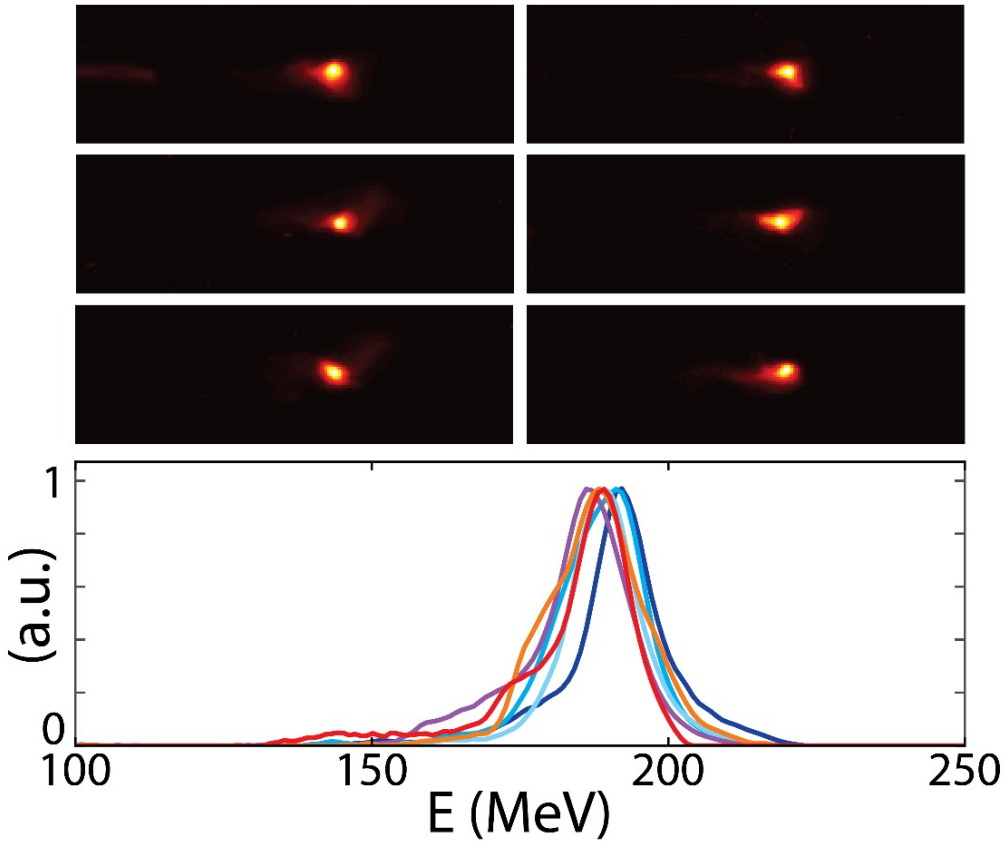

Results

Q= 17 pC

dE/E_{peak} = 8%

Divergence = 6 mrad

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

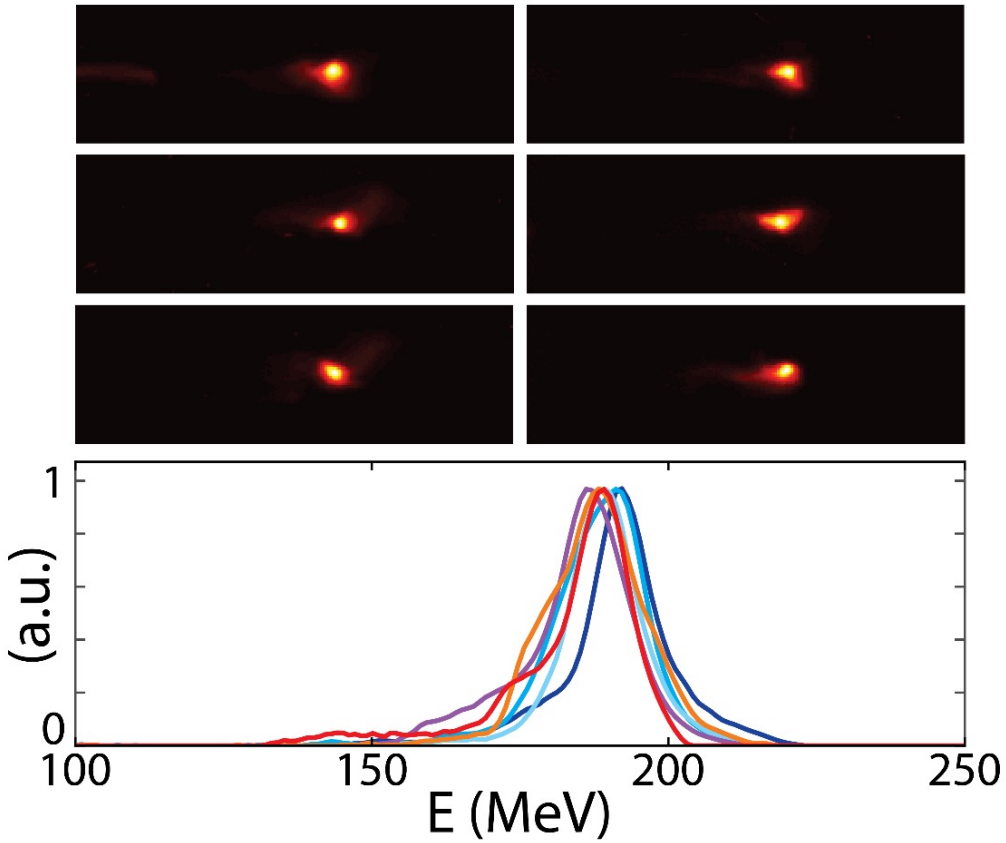

Results

Q= 17 pC

dE/E_{peak} = 8%

Divergence = 6 mrad

And…

Stability shot by shot !

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

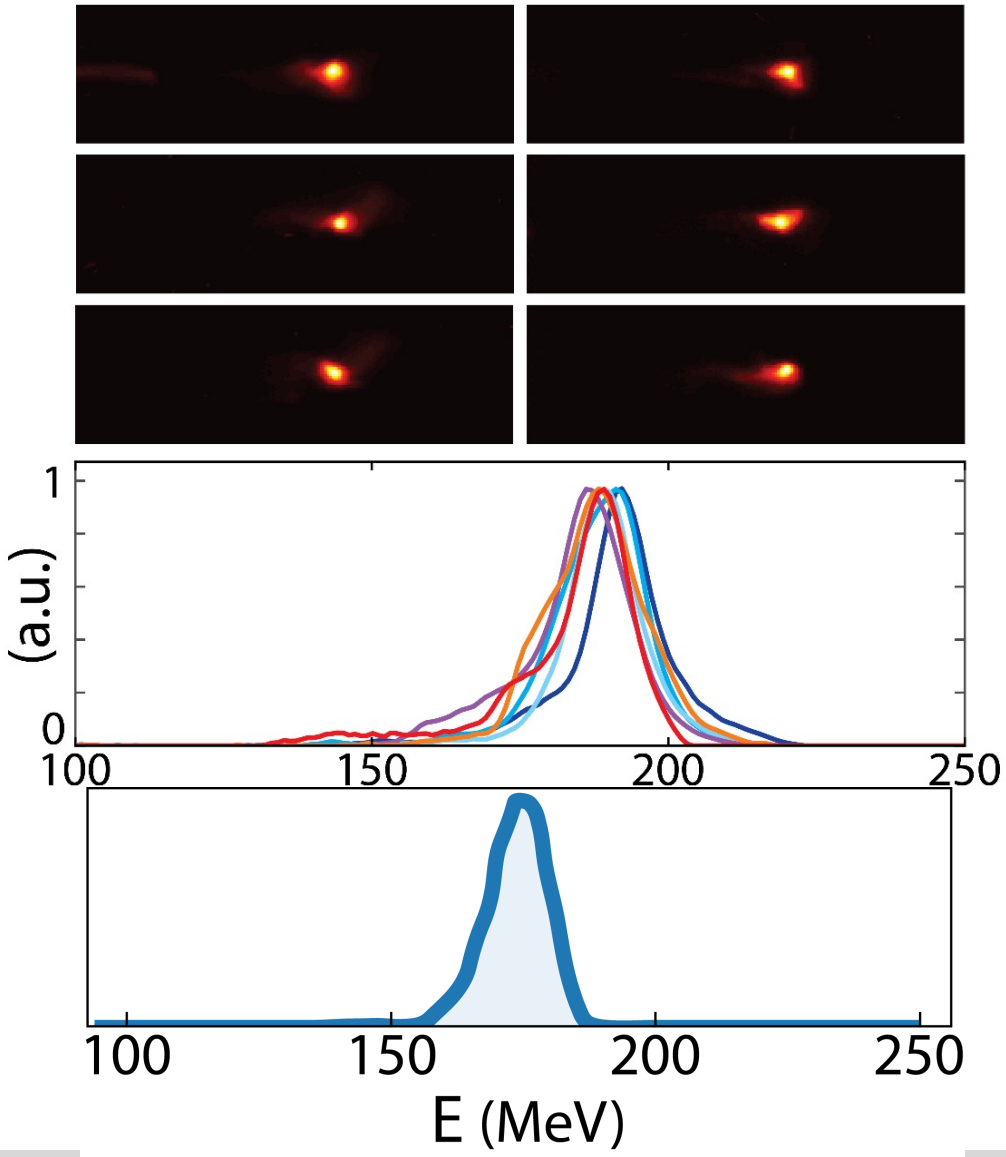

Results

Q= 17 pC

dE/E_{peak} = 8%

Divergence = 6 mrad

And…

Stability shot by shot !

Validated with simulations

Q= 26 pC

dE/E_{peak} = 9%

Divergence = 10 mrad

A promising system!

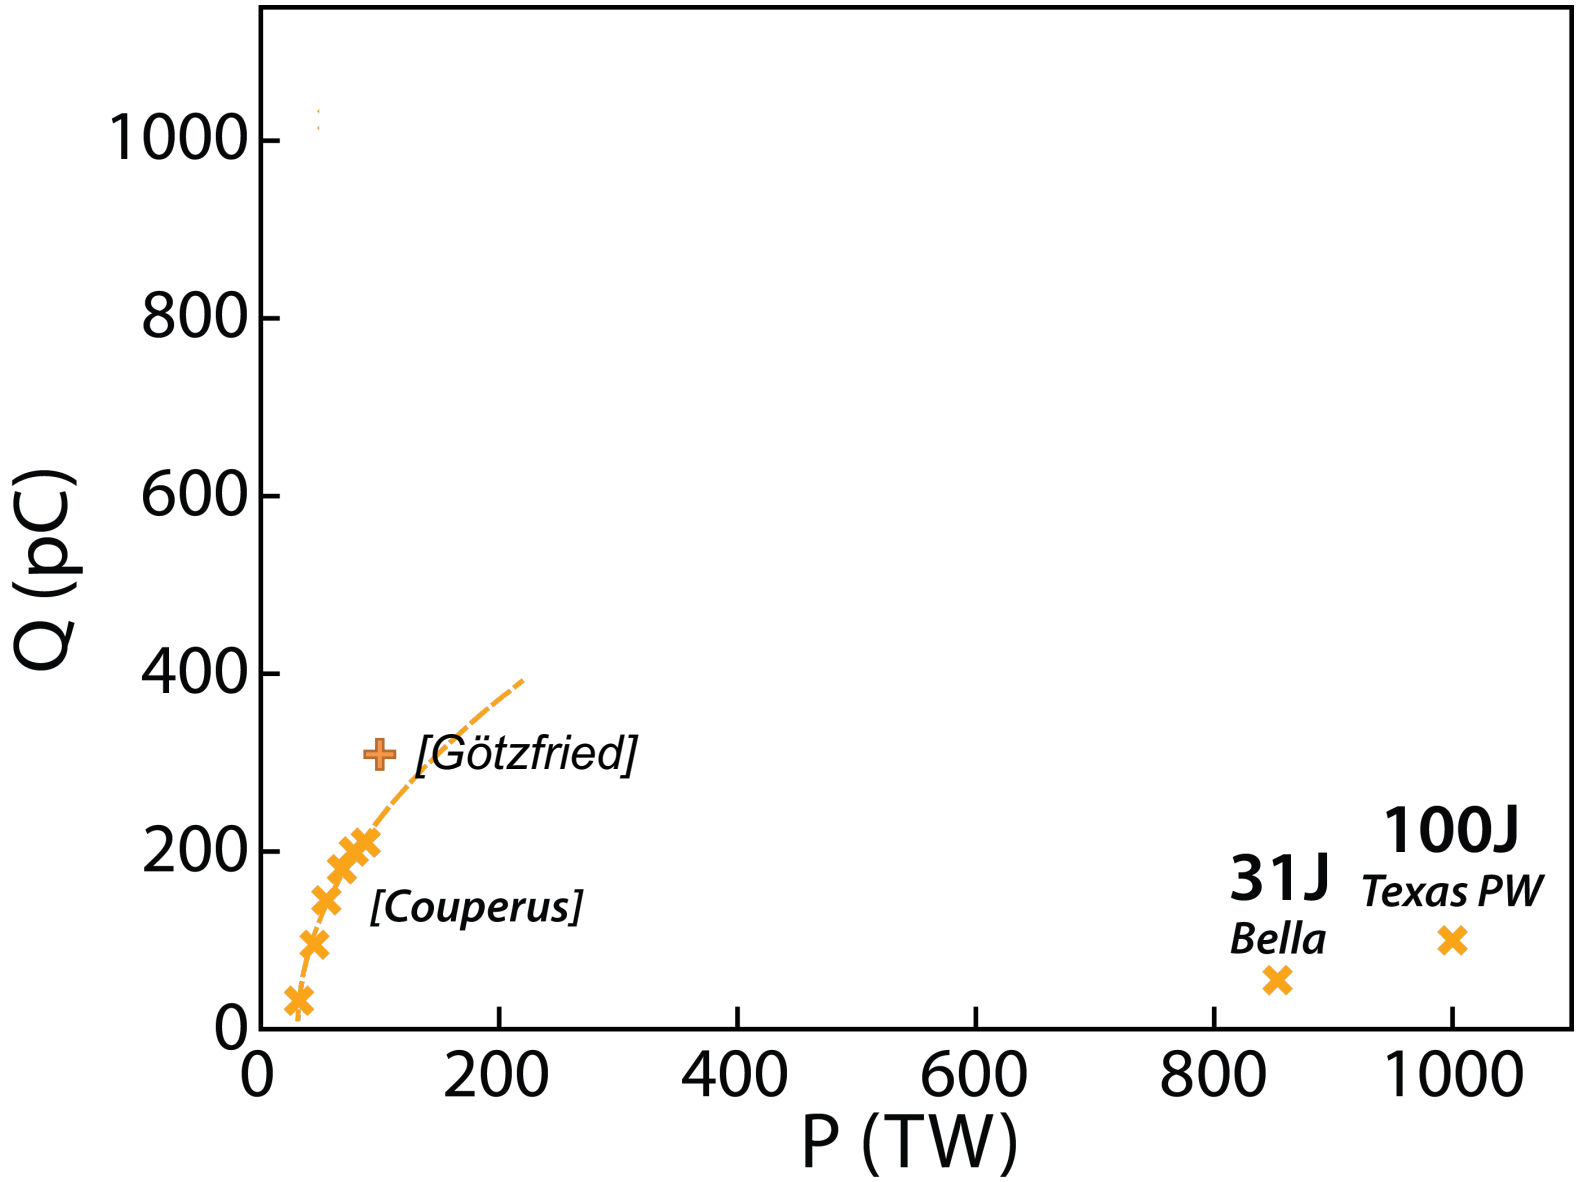

A promising system!

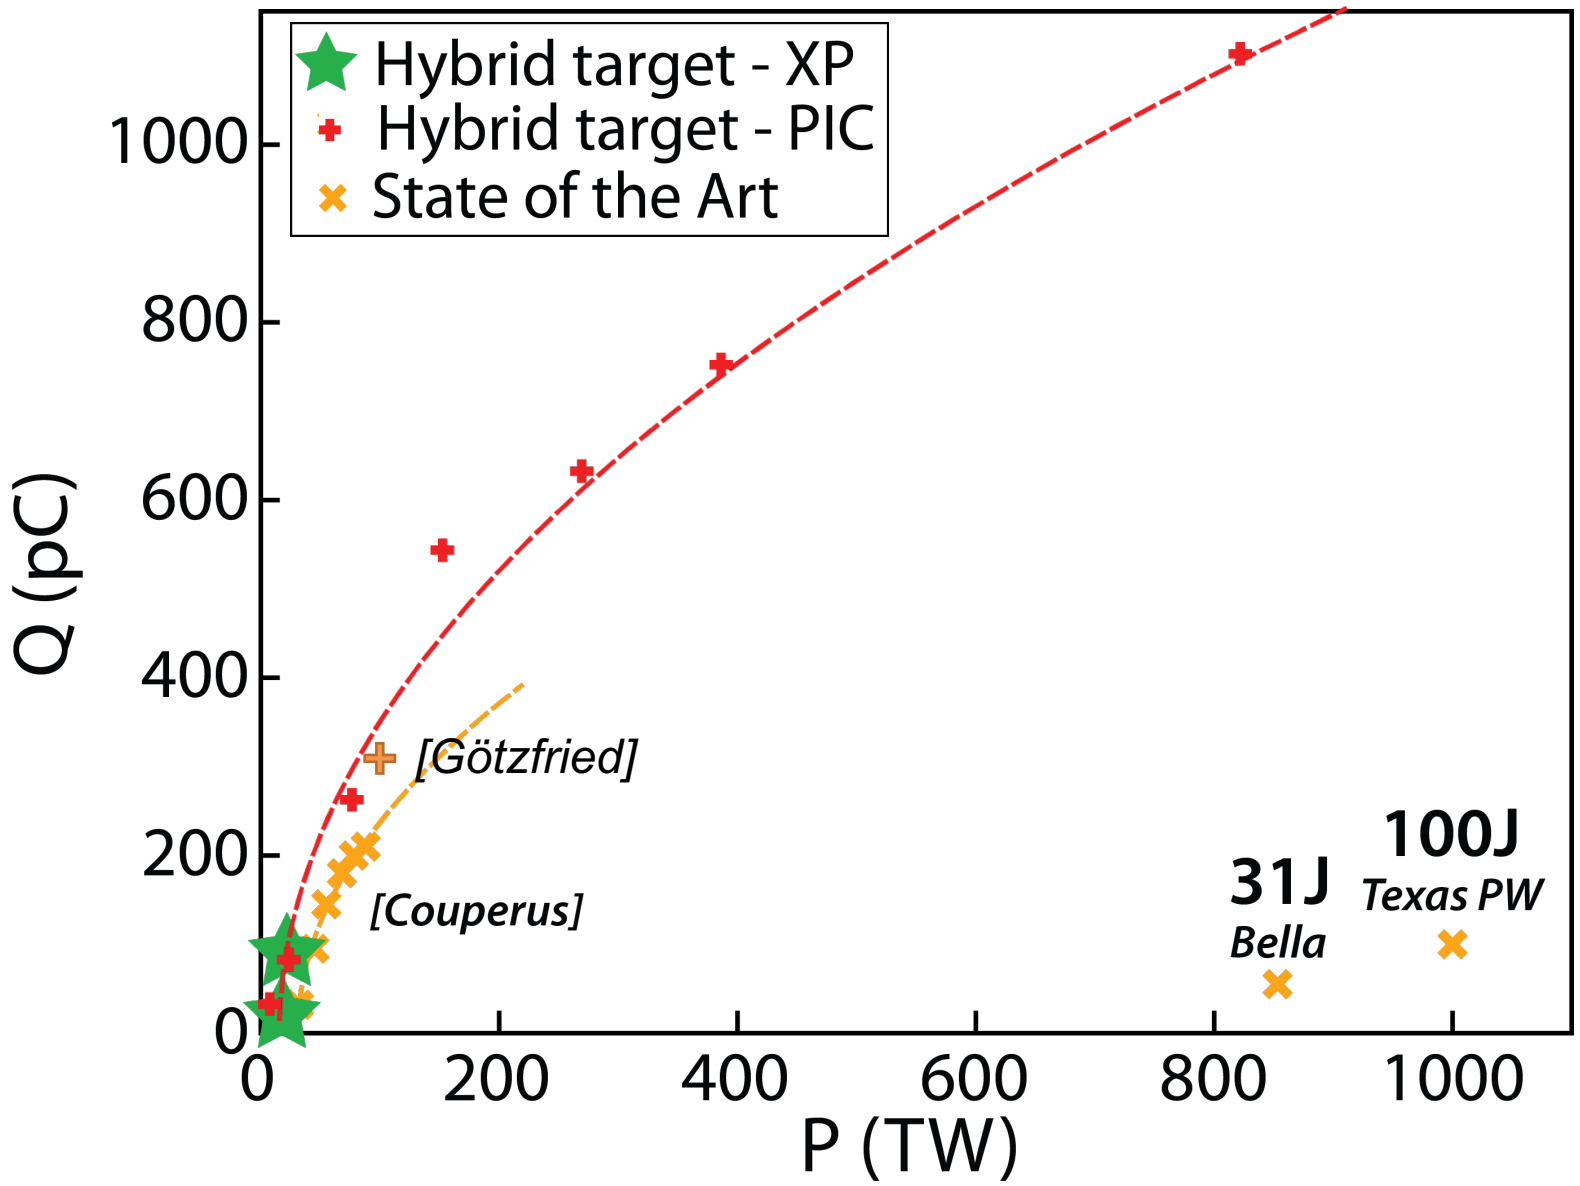

A promising system!

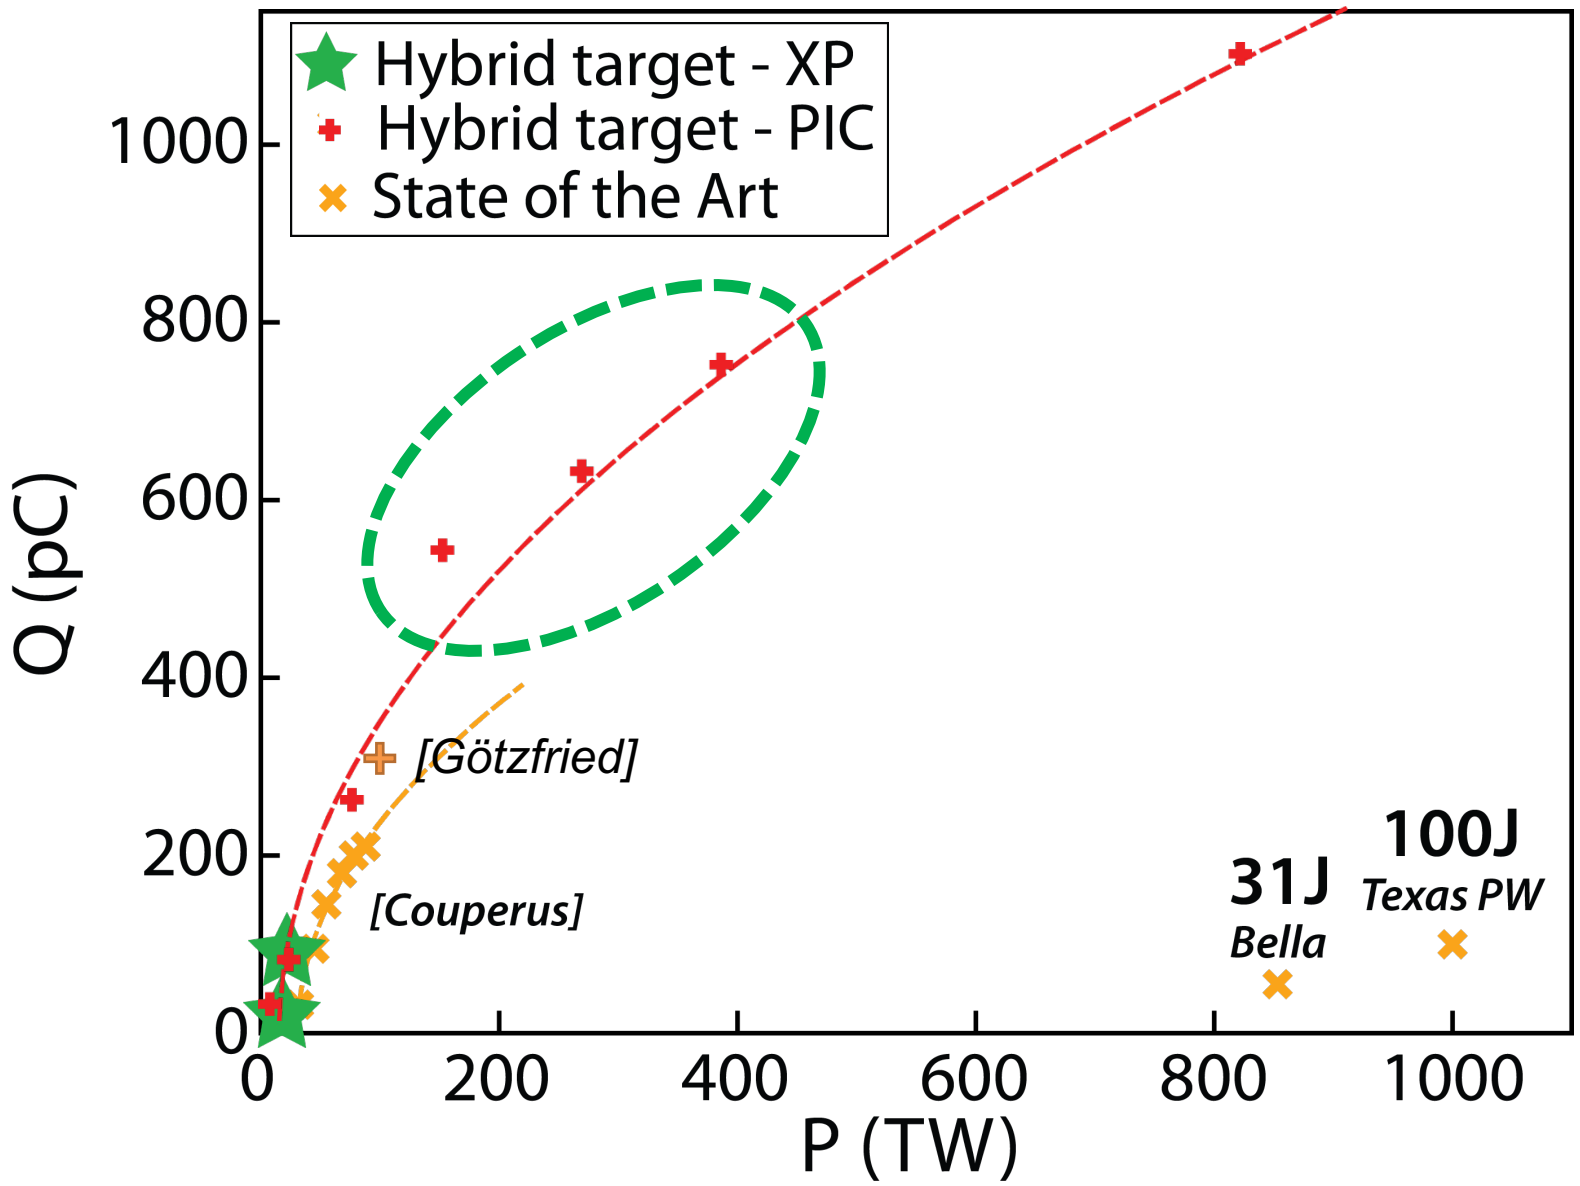

Next experimental campaign at apollon !

A promising system!

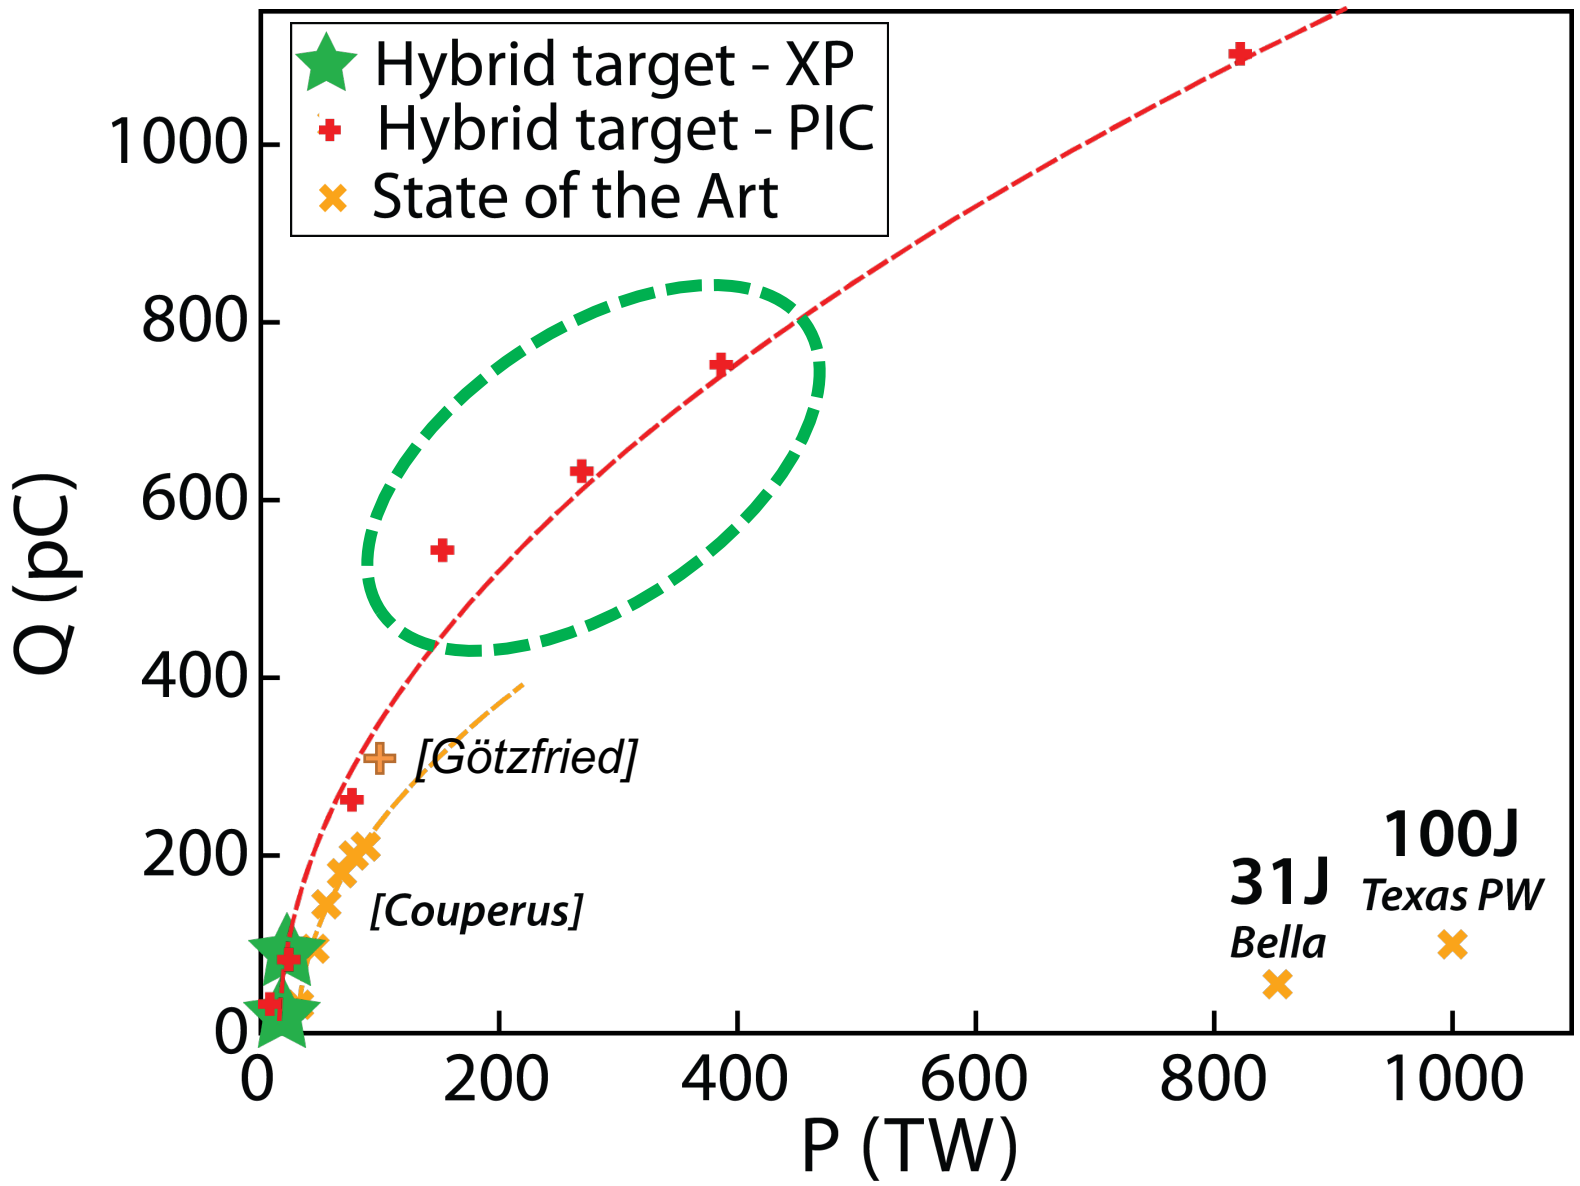

Next experimental campaign at apollon !

Preliminary results :

Q=600pC

E = 600 MeV

dE/E_{peak} = 4%

Conclusions and perspectives

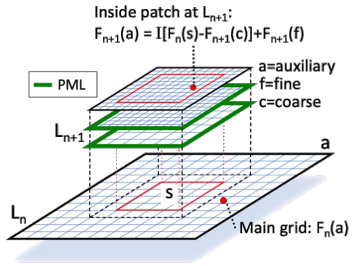

- WarpX is a state-of-the-art open-source Particle-In-Cell code implementing sophisticated numerical algorithms

Conclusions and perspectives

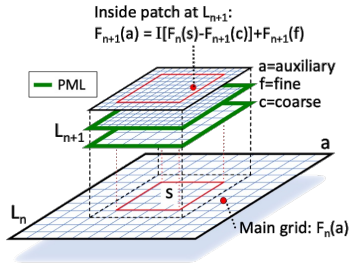

- WarpX is a state-of-the-art open-source Particle-In-Cell code implementing sophisticated numerical algorithms

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

Conclusions and perspectives

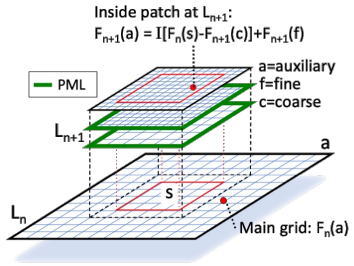

- WarpX is a state-of-the-art open-source Particle-In-Cell code implementing sophisticated numerical algorithms

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

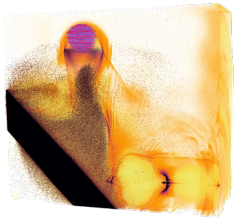

- WarpX help us to study and design new experiment as novel electron acceleration strategies

The WarpX code:
Particle-In-Cell simulations at the exascale

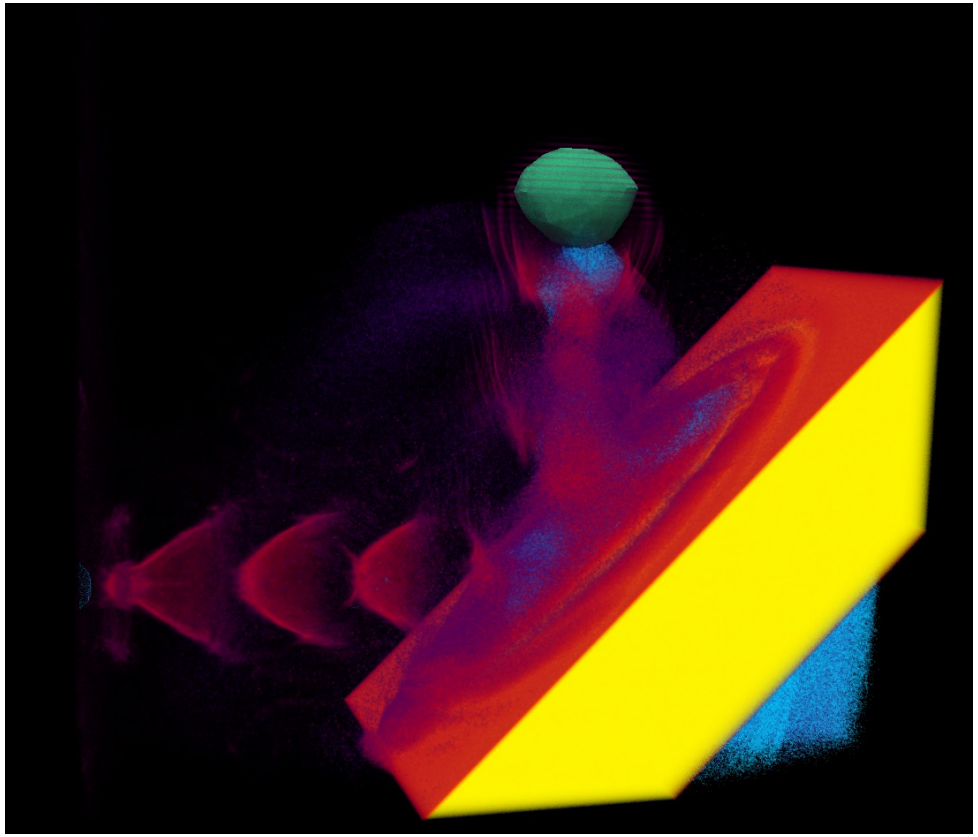

Luca Fedeli, Axel Huebl, France Boillod-Cerneux, Thomas Clark, Kevin Gott, Conrad Hillairet, Stephan Jaure, Adrien Leblanc, Rémi Lehe, Andrew Myers, Christelle Piechurski, Mitsuhisa Sato, Neil Zaim, Weiqun Zhang, Jean-Luc Vay, Henri Vincenti