Exascale simulation for high-gradient accelerators

A. Beck, LLR, Ecole polytechnique

G. Bouchard; A. Specka

M. Grech; F. Pérez; T. Vinci

M. Lobet; F. Massimo; J. Silva Cuevas

5th EAAC

September 22nd 2021

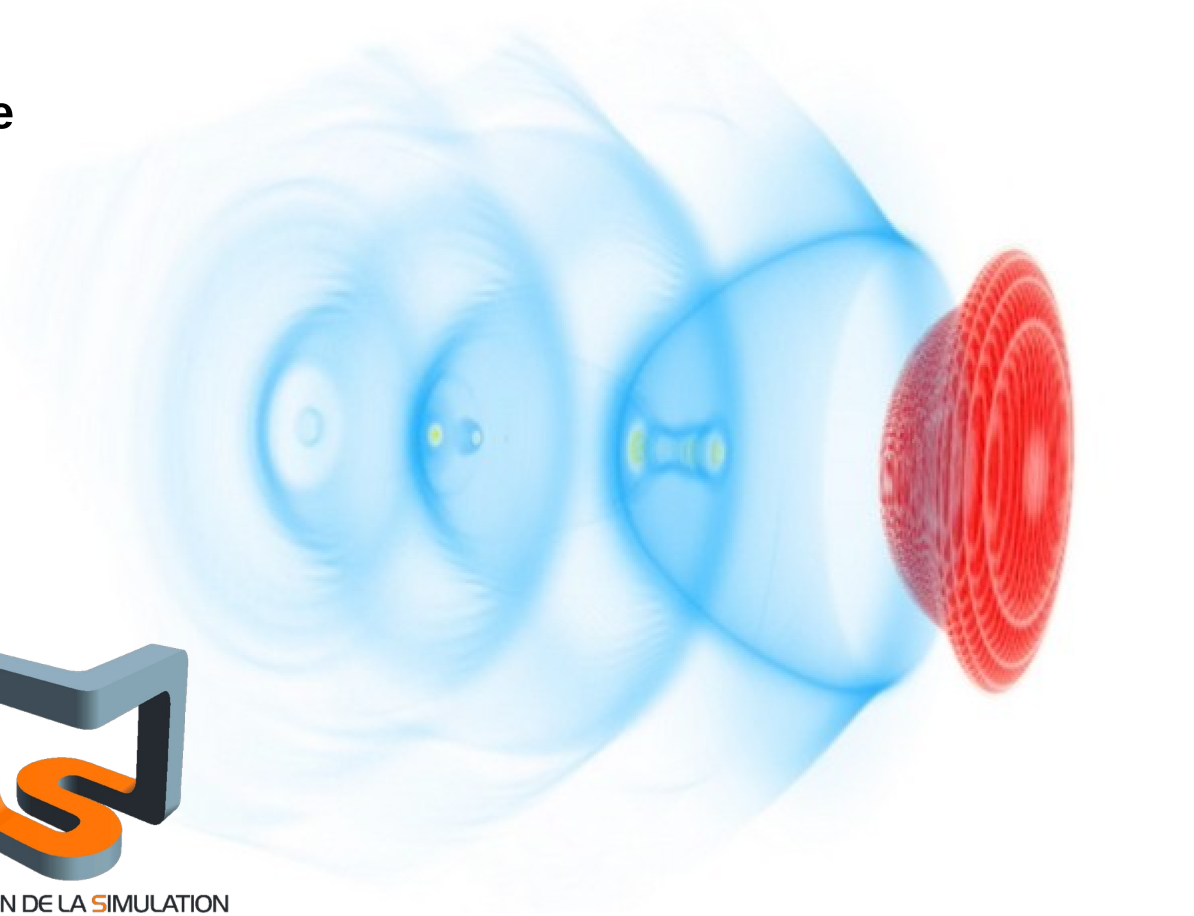

MAISON DE LA SIMULATION

What Exascale is NOT

- Not yet another generation of faster super-computer

  For almost 20 years now, processors operate at the same frequency.

- Not a larger version of our current super-computer

  All components of the super-computer (memory, network, computing) do not improve by the same factor.

Definition of Exascale

Exascale System: “Computing system capable of a theoretical peak performance of 10^{18} floating point operations per second (flop/s)”

- Based on computing only.
- Theoretical peak performance is given by the constructors and is of very little practical value …

Performance measure

- TOP 500 uses “Linpack” to measure performances and rank super-computers.
- Problem: Computers are overtrained to Linpack.

Rank	System	Peak	Linpack
1st	FUGAKU	532 Pflop/s	442 Pflop/s
2nd	SUMMIT	200 Pflop/s	148 Pflop/S

Performance measure

- HPCG provides a more realistic measure.
- Only a few percent of the computing power is used !
- Why is it so bad ?

Rank	System	Peak	Linpack	HPCG
1st	FUGAKU	532 Pflop/s	442 Pflop/s	16 Pflop/s
2nd	SUMMIT	200 Pflop/s	148 Pflop/S	2.9 Pflop/s

Exceptional service in the national interest

Sandia National Laboratories

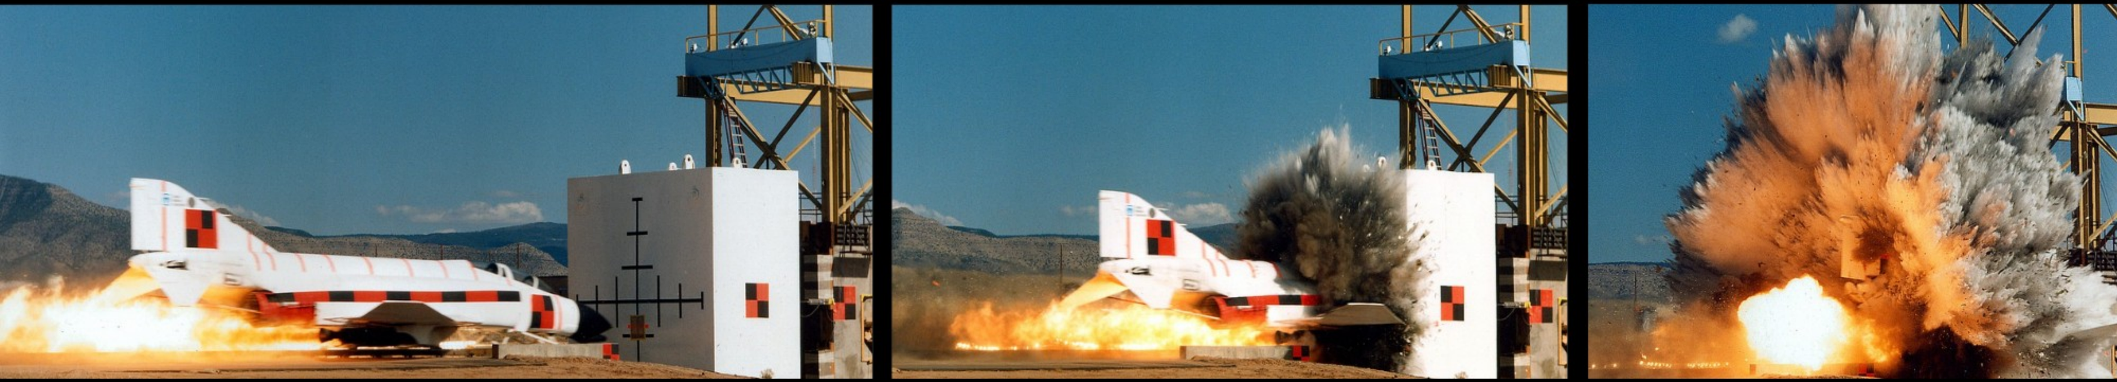

Oh, $#*@! Exascale!

The effect of emerging architectures on data analysis software

SOS 17 Panel, March 27, 2013

Kenneth Moreland, Sandia National Laboratories

SDAV

U.S. DEPARTMENT OF ENERGY

NNSA National Nuclear Security Administration

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.
SAND 2013-2003P

10 years ago in 2011...

The so called “Slide of Doom” with expectations for exascale in 2018:

Parameter	2011	“2018”	Factor
Peak	2 Pflop/s	1 Eflop/s	500
Power	6 MW	< 20 MW	3
Concurrency	225 k	10 – 100 B	40 000 – 400 000
Memory	0.3 PB	32 - 64 PB	100 - 200
Network	1.5 GB/s	100 - 1000 GB/s	66 - 660

… and today

Parameter	2011	“2018”	2021
Peak	2 Pflop/s	1 Eflop/s	0.5 Eflop/s
Power	6 MW	< 20 MW	30 MW
Concurrency	225 k	10 – 100 B	7.6 M
Memory	0.3 PB	32 - 64 PB	5 PB
Network	1.5 GB/s	100 - 1000 GB/s	40 GB/s

Power wall

- 30 MW => 30 M€ per year at 0.10€ / kWh
- Power => heat : Expensive cooling systems are necessary
- In the context of global warming, massively power consuming devices should be avoided and probably won't be supported

Transistor Power $P = C V^2 F$ (Conductive load, Voltage, Frequency)

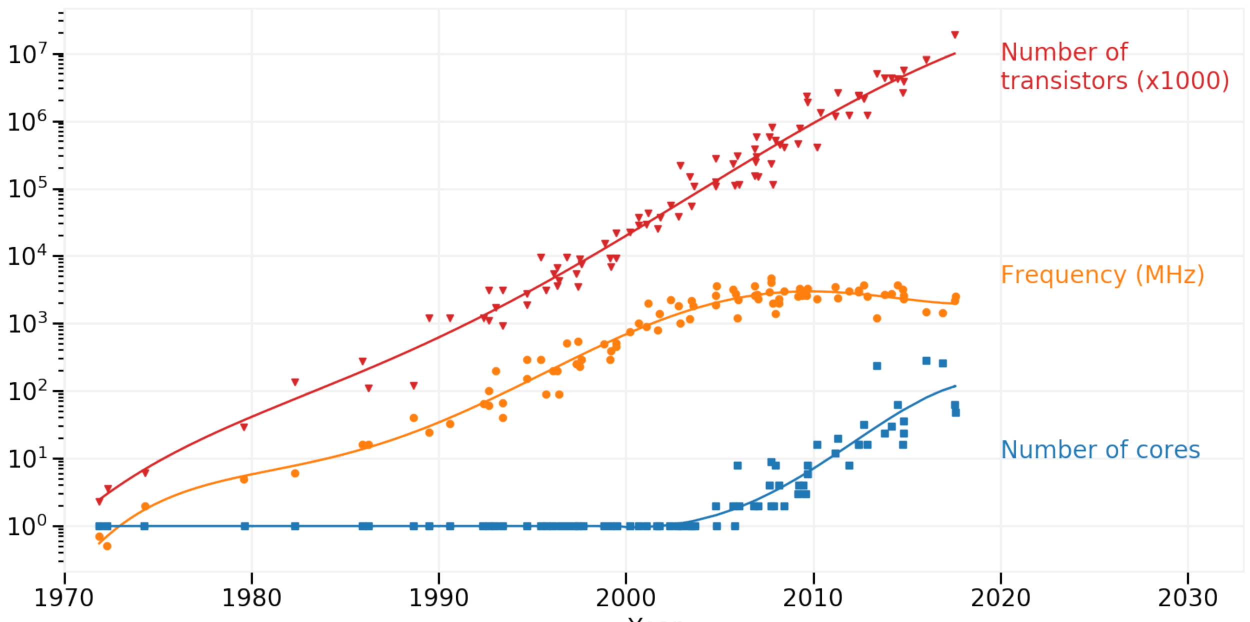

- Limit Frequency
- More cores
- **More compute** => SIMD, GPU, …
- **Less of the rest** => Logic, data movement,...

Memory wall

Parameter	2011	“2018”	Factor	2021
Peak	2 Pflop/s	1 Eflop/s	500	0.5 Eflop/s
Concurrency	225 k	10 – 100 B	40 000 – 400 000	7.6 M
Memory	0.3 PB	32 - 64 PB	100 - 200	5 PB

Computing and concurrency grows much more than memory

- Less memory per thread.
- Managing the memory becomes the dominant problem.
- End of weak scaling. Focus on strong scaling.
- Very fine grain parallelism is mandatory.

Exascale fundings

USA: Massive fundings since 2016.

- Exascale Computing initiative => exascale supercomputer expected **within a year**. $475 million for year 2021.
- Exascale Computing Project => applications and software stack. $250 million for year 2021. Plasma wakefield accelerator is listed as an objective and it supports the development of the American code WarpX.

Europe: Just starting to invest.

- EuroHPC: 8 Billion € for 2021-2030. It is not clear how this will be shared between system, software stack and applications.
- 2 Exascale systems announced for 2023/2024. Scientific fields and pilot applications relevant to exascale are being identified just now. **It is important that plasma acceleration is recognized as eligible for European exascale as it has been in the US.**

Smilei)

Open-source
PIC code

M. Grech
F. Perez
T. Vinci

M. Lobet
F. Massimo
J. Silva Cuevas

MAISON DE LA SIMULATION

A. Beck
G. Bouchard

Collaborative, user-friendly

GitHub • Python interface

Educational resources

Online documentation • Tutorials

High-performance

MPI-OpenMP • Load balancing • vectorization

Physics

Ionisation • Collisions • Strong-field QED

Advanced solvers

Spectral solvers • Multi-geometries • Laser envelope

Derouillat et al., CPC** 222 **(2018)

Sharing knowledge

Next SMILEI Workshop March 2022
Registration starting soon at
https://indico.math.cnrs.fr/event/6911/

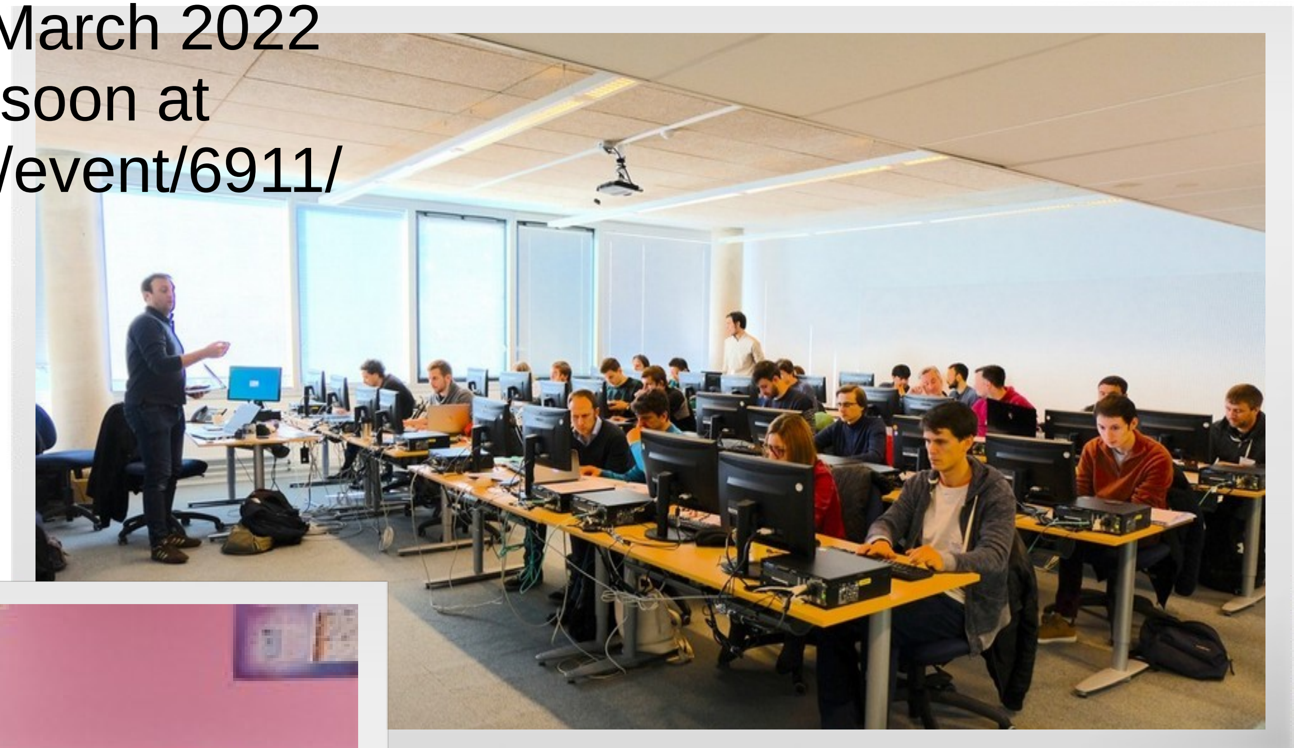

Practicals for master students

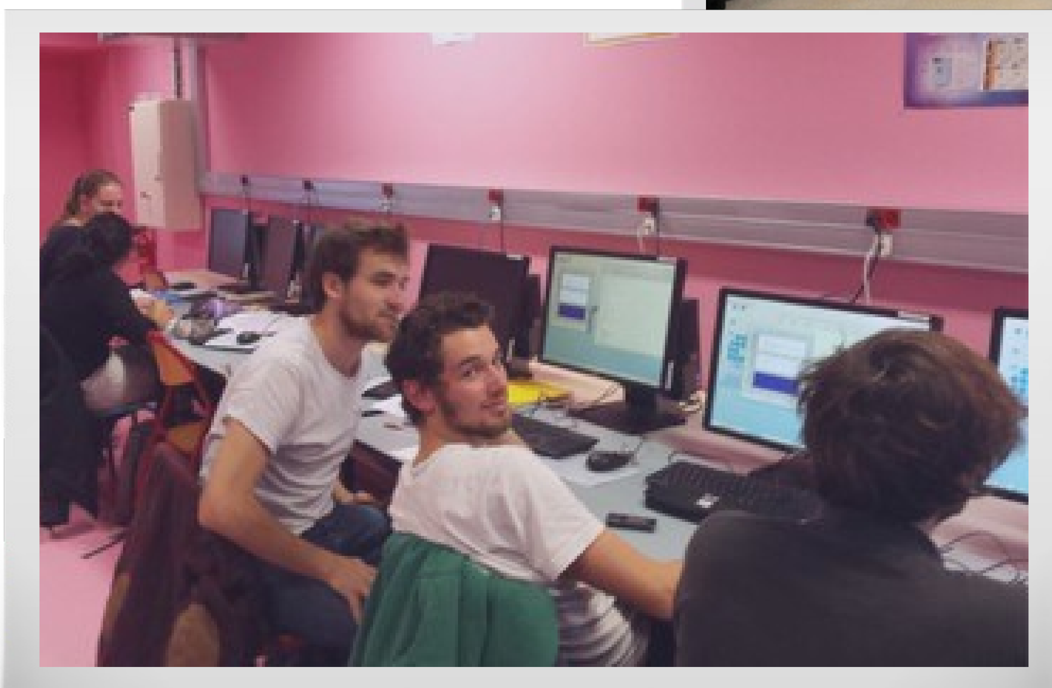

User training & workshop

Tutorials available at
smileipic.github.io/tutorials

International Users Community

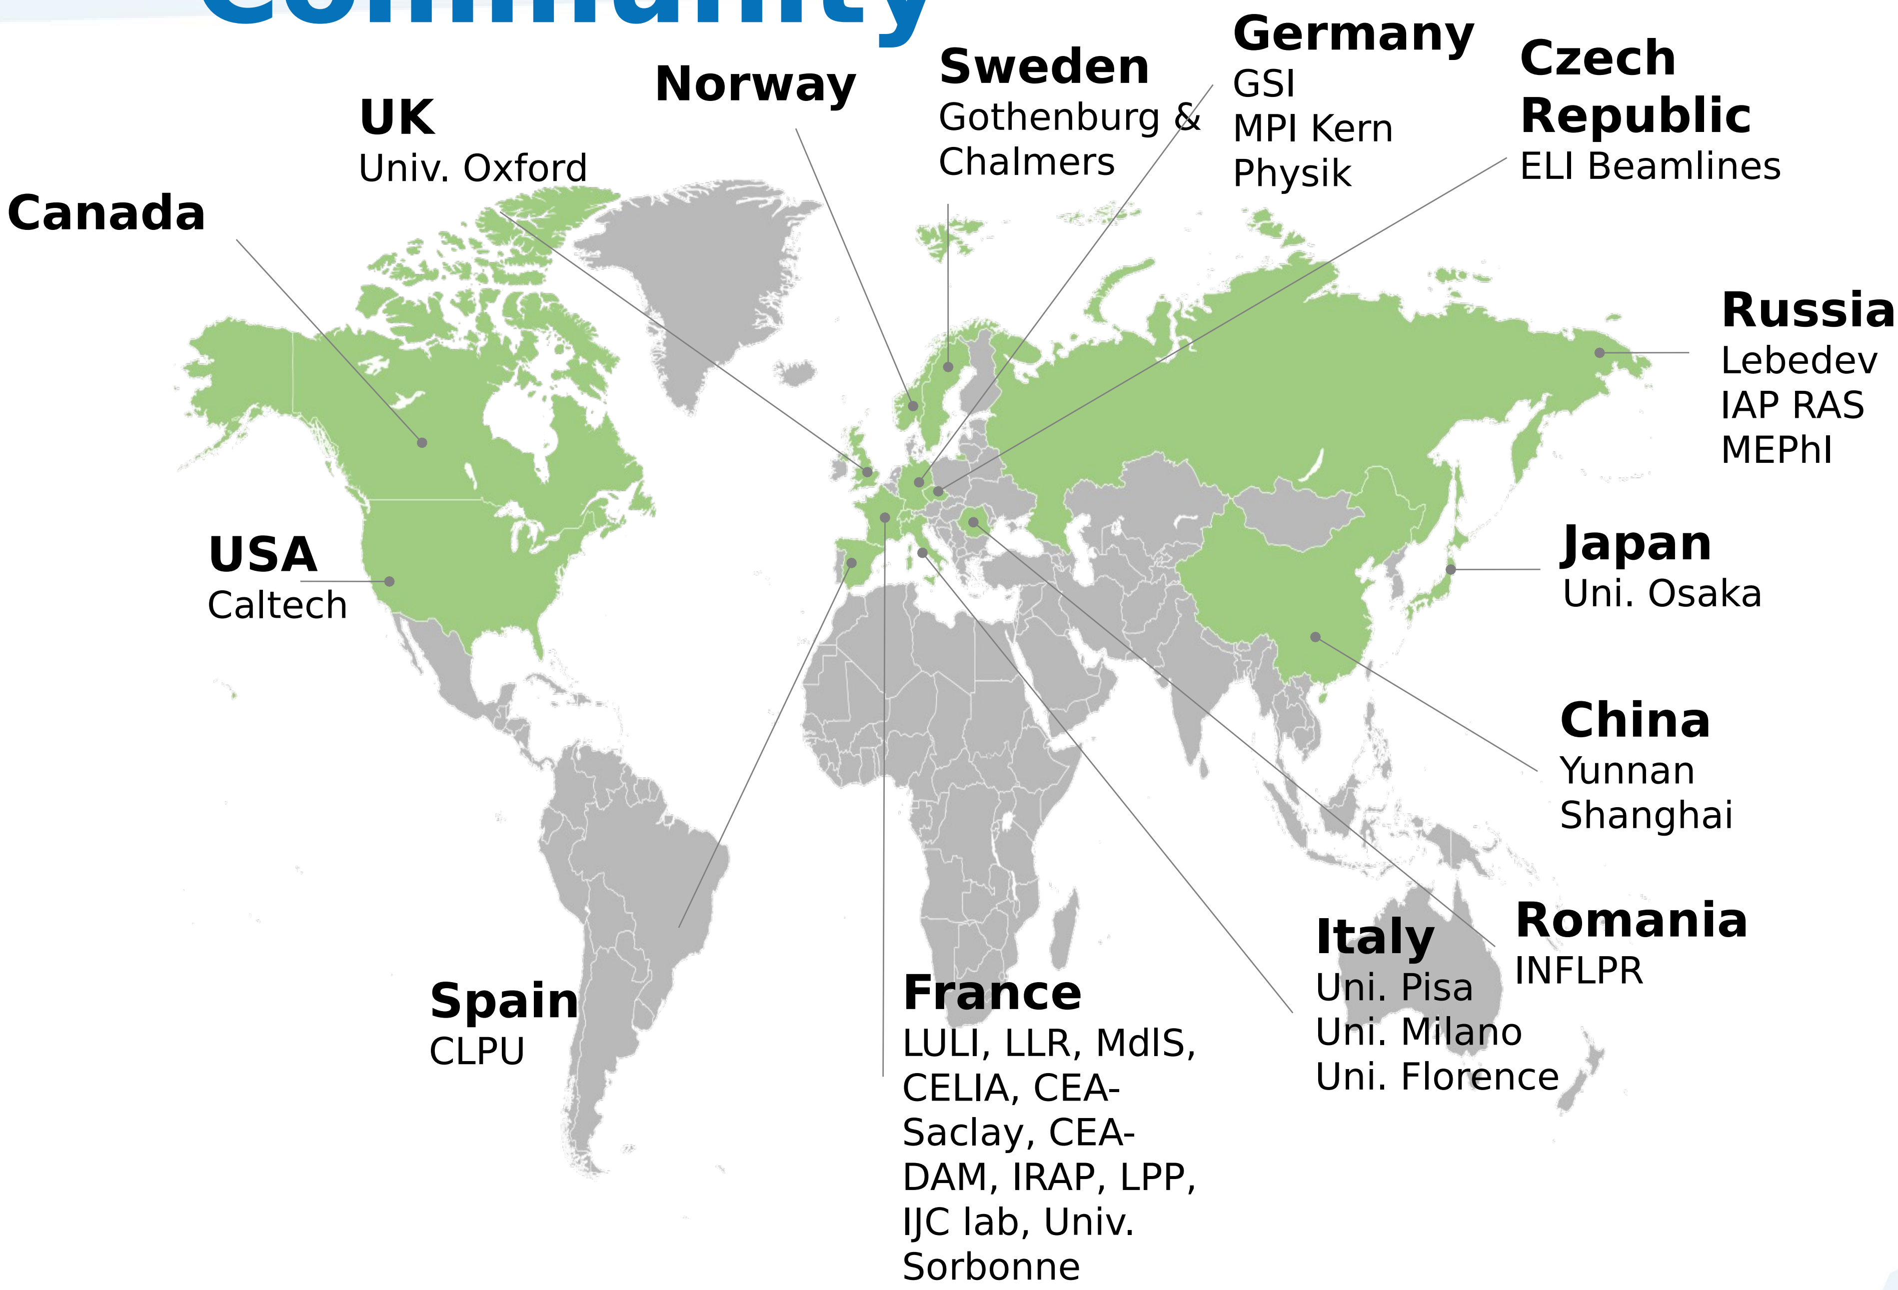

Joint effort between communities

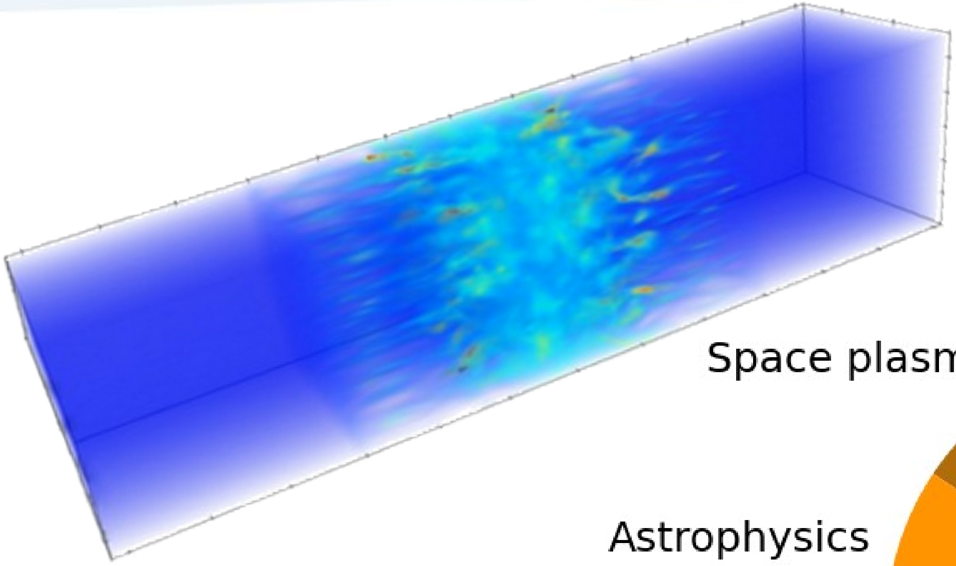

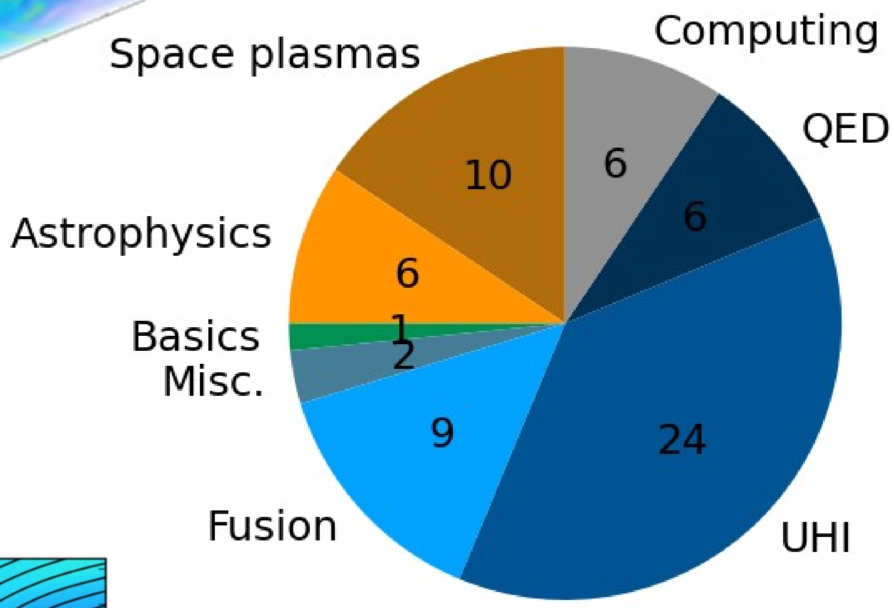

Publications distribution

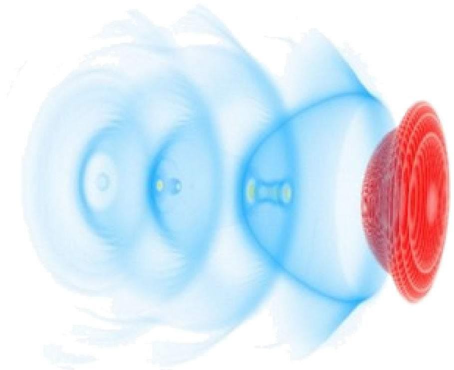

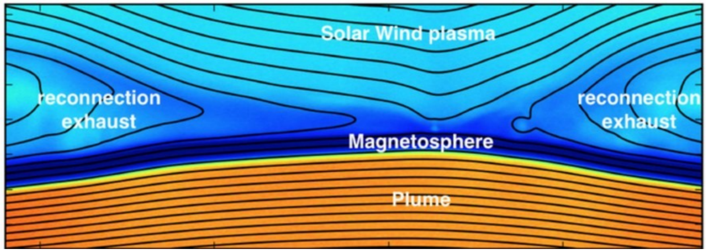

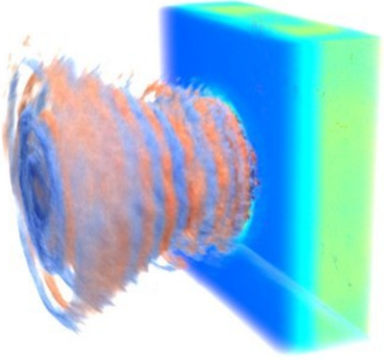

Features relevant for plasma acceleration

Physics	Numerics	HPC
• Ionization • QED	• Dispersion free solver • Envelope • Moving window • Azimuthal Decomposition • PML boundaries • Dynamic Load Balancing • Arbitrary laser/plasma profiles • Particles injection	• MPI • OpenMP • Taskification • GPU support • Vectorization • ARM support

Towards fast and reliable simulations

1) Numerical methods and High Performance Computing are key.
2) Combining all features is the challenge => **Software engineering**

Smilei) HPC Developments

2019 - 2022	New architectures	• GPU : OpenACC + Thrust • ARM
2019 - 2022	Asynchronism	• Asynchronism, overlapping, heterogeneous computing • OpenMP task, Eventify (JSC)
2017 - 2018	Vectorization	• Data structure adaptation • OpenMP SIMD pragmas • Better cache efficiency
2015 - 2017	Load balancing	• Hilbert curve patch exchange
2015 - 2017	Multithreading	

Balance the workload between processors

Domains automatically adapt
to the simulation evolution

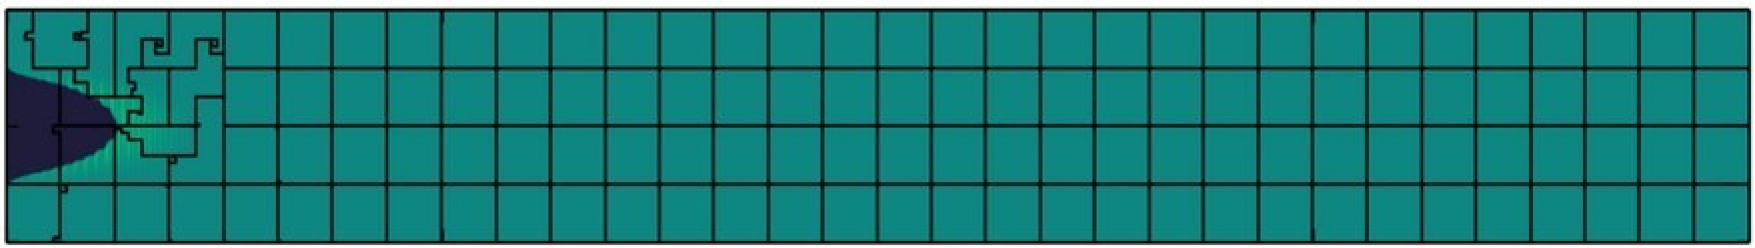

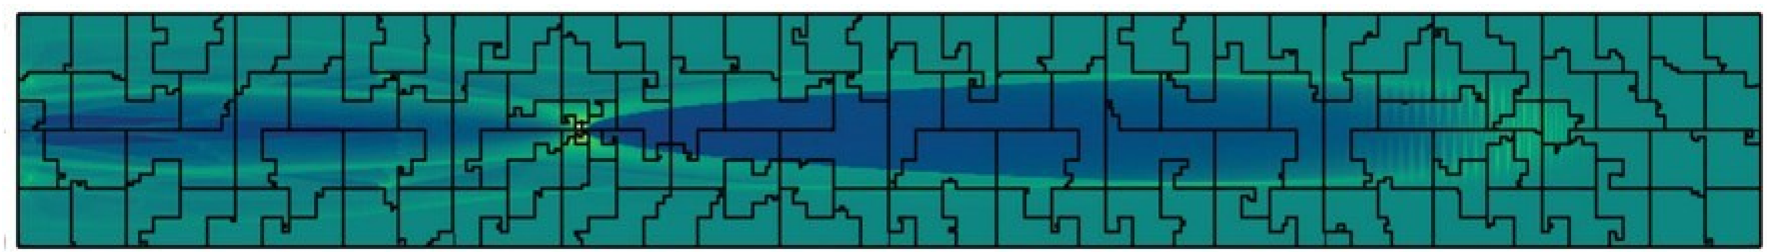

Laser wakefield simulation ~ 2x faster

GPU

Graphical Processing Unit: compute intensive accelerator. A probable path to exascale.

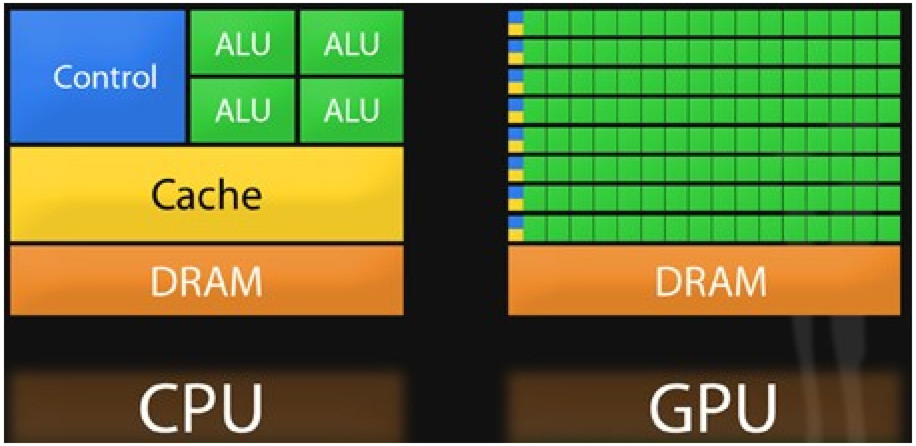

+ Compute intensive, power efficient

 - Bad for logic and data movement

 - Needs to be addressed specifically: SMILEI uses openACC + Thrust

ARM based architecture

- ARM was initially designed for embebded and mobile applications. Europe and Japan are now pushing the development of ARM super-computers.
- On-going work to make SMILEI efficient on ARM-based architecture (FUGAKU and soon in European exascale).
- Optimization and vectorization efficiency enabled and ensured with many compilers (GNU, LLVM, ARM compiler, Fujitsu compiler).
- Public release early 2022

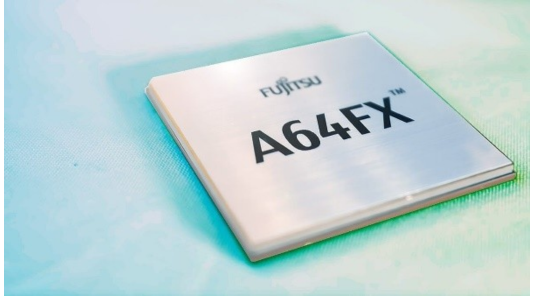


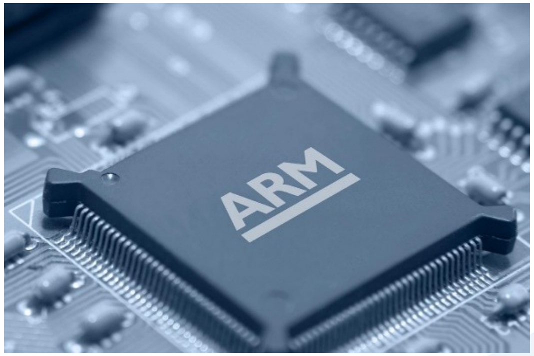

Task Based Asynchronism

Task = Instructions + Data scope + Dependencies (if any)

Task dependency graph:

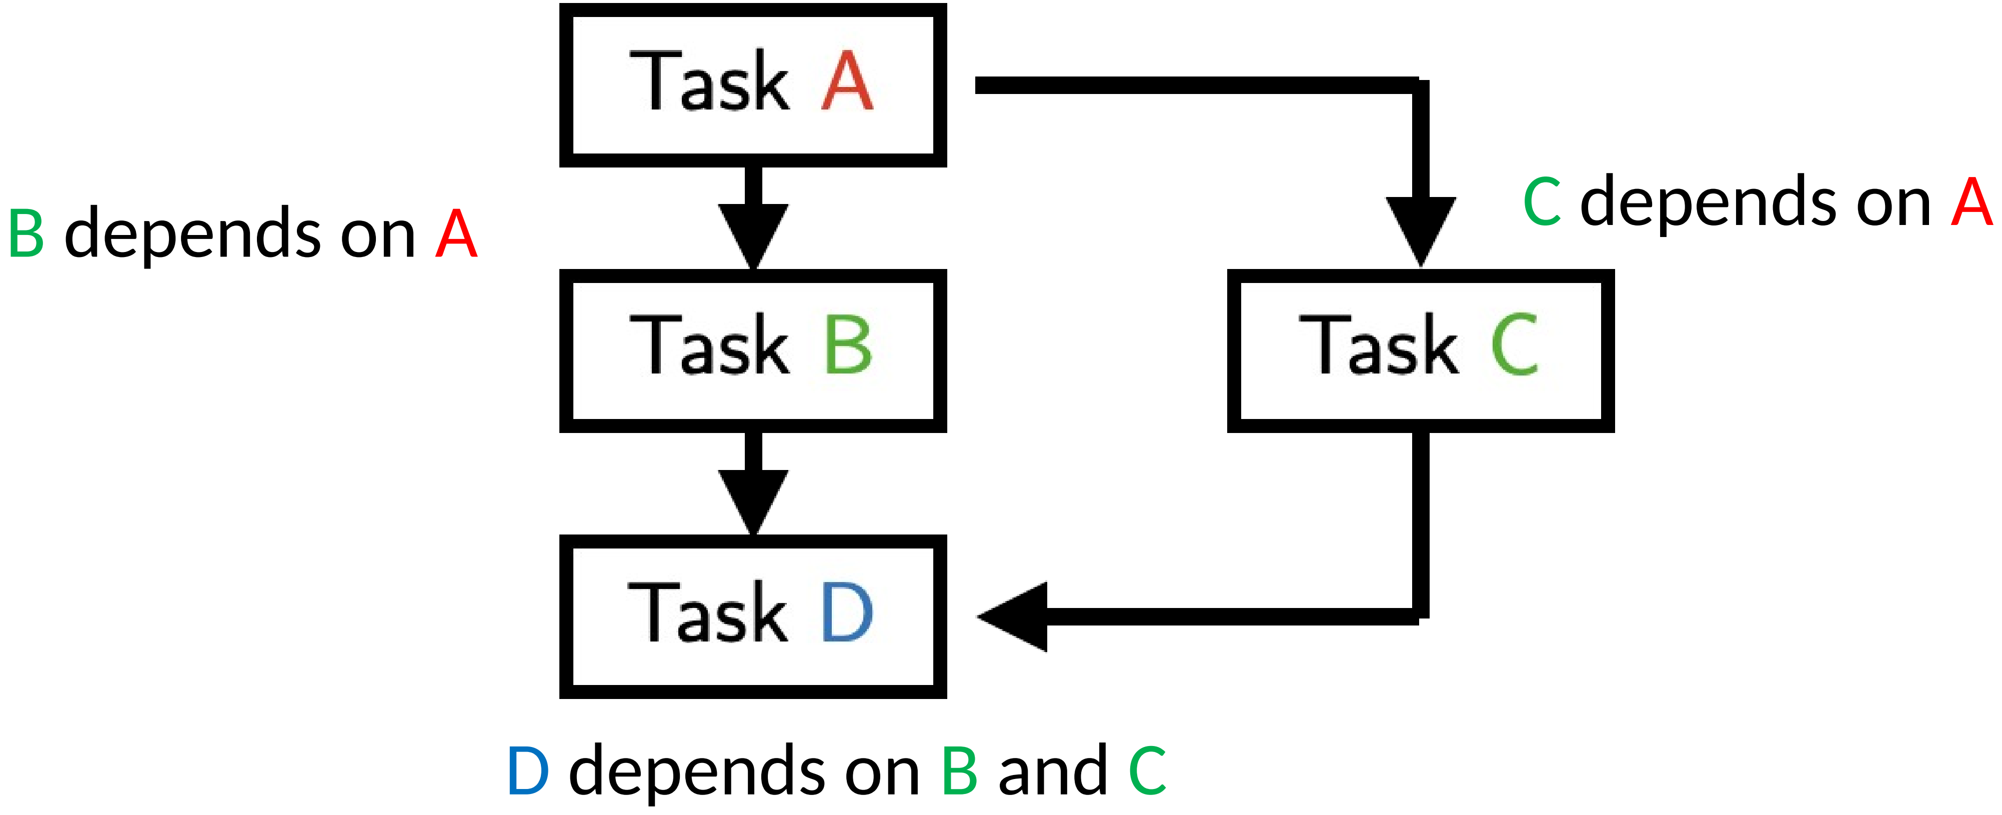

Scheduling Tasks

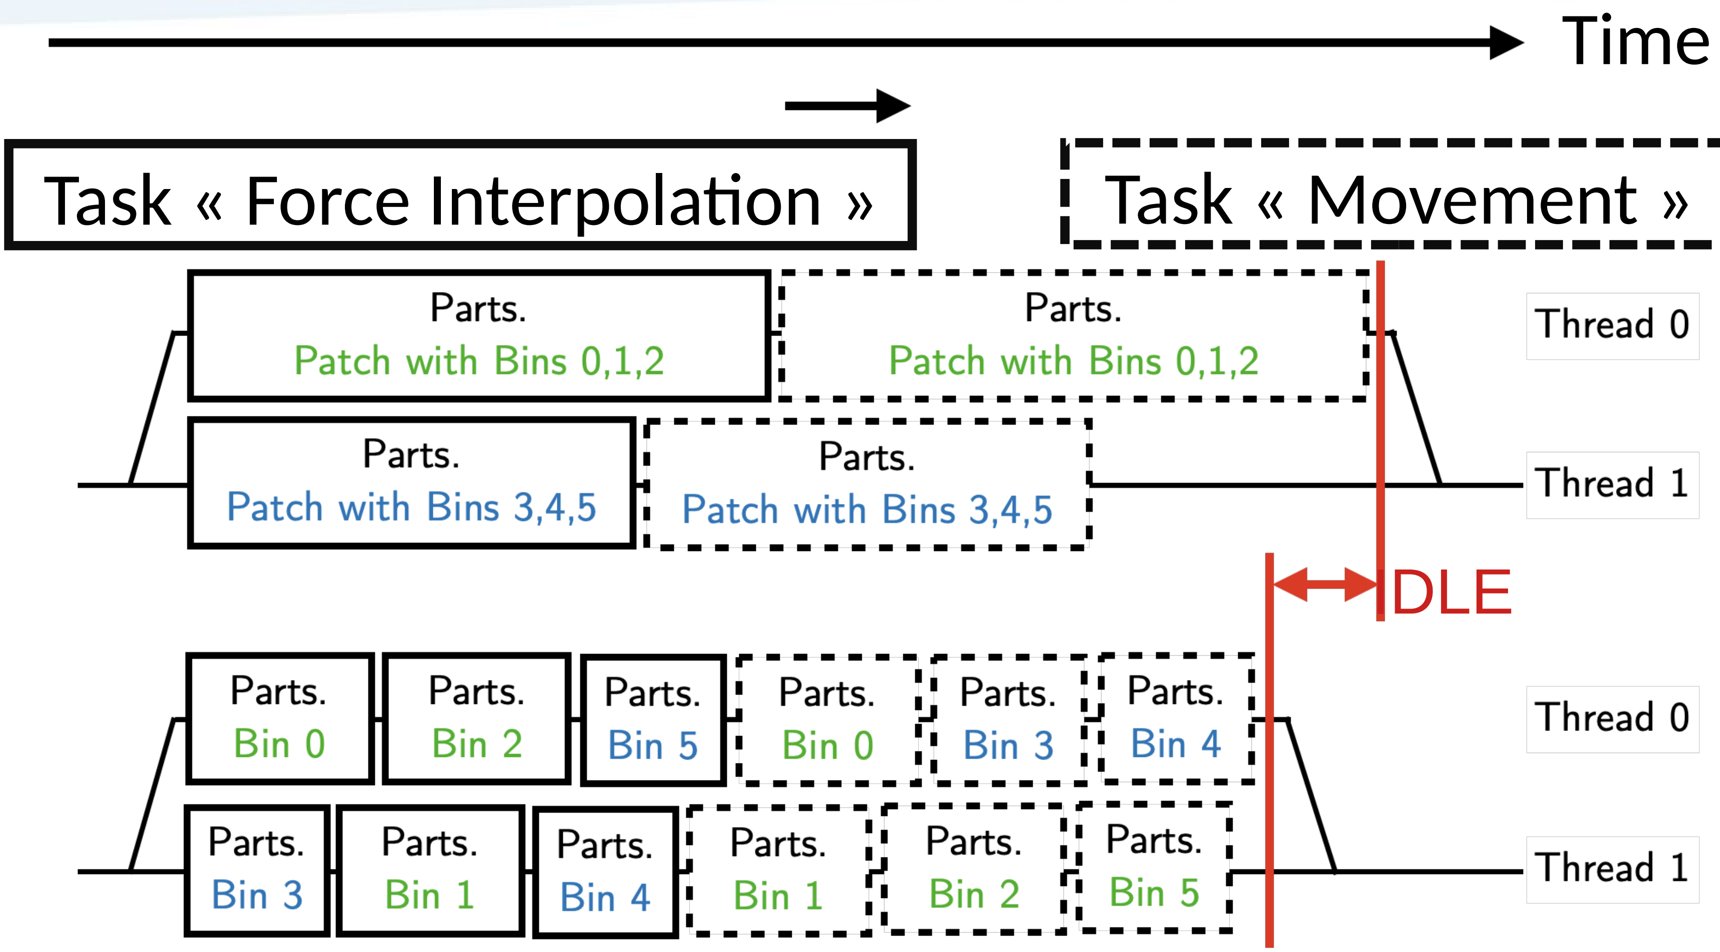

- Improves Load Balance
- Covers latencies
- Helps “Strong scaling” (massive concurrency)

Exascale task scheduler

At the exascale, too many threads and tasks for any developer to explicitly build an optimal graph. The computer must do it at runtime.

- Implemented in SMILEI with openMP tasks.

- It is planned to use the task library "**Eventify**" developed at Jülich Forschungzentrum in the frame of a European cooperation

AI Data Analytics and Scalable Simulations

JÜLICH Forschungszentrum

DE LA RECHERCHE À L'INDUSTRIE cea

- Explore the possibility of tasks on GPU

Other efforts toward exascale

OSIRIS, developed in Portugal and in the U.S. (IST and UCLA)

WarpX, one of 21 apps of the U.S. DOE Exascale Computing Project (ECP)
Development coordinated by LBNL

Similar trends, different technical approaches:

- Task based approach in OSIRIS based on ompSs-2
- Support for ARM architecture in OSIRIS
- Support for GPU in OSIRIS (Cuda) and WarpX (Cuda/HIP/DPC++)
- Mesh Refinement proposed by AMRex implemented in WarpX helps avoiding the memory wall when increasing the resolution

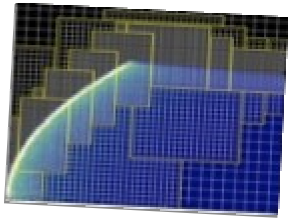

Real Time Simulations

Demand for simulations running at a timescale on par with experiments

- Experimental design
- Experiment monitoring

Today this is approached with reduced models.

Single stage simulation roughly on the order of 10 min.

Access to exascale systems can increase both speed and accuracy of these simulations provided strong scaling is achieved.

High Fidelity Simulations

- Ultra High resolution necessary for accurate simulations.
- Detailed features of the laser and target.
- Advanced physics (betatron radiation, QED).
- Danger of hitting the memory wall (refinement might be required).

Many stages

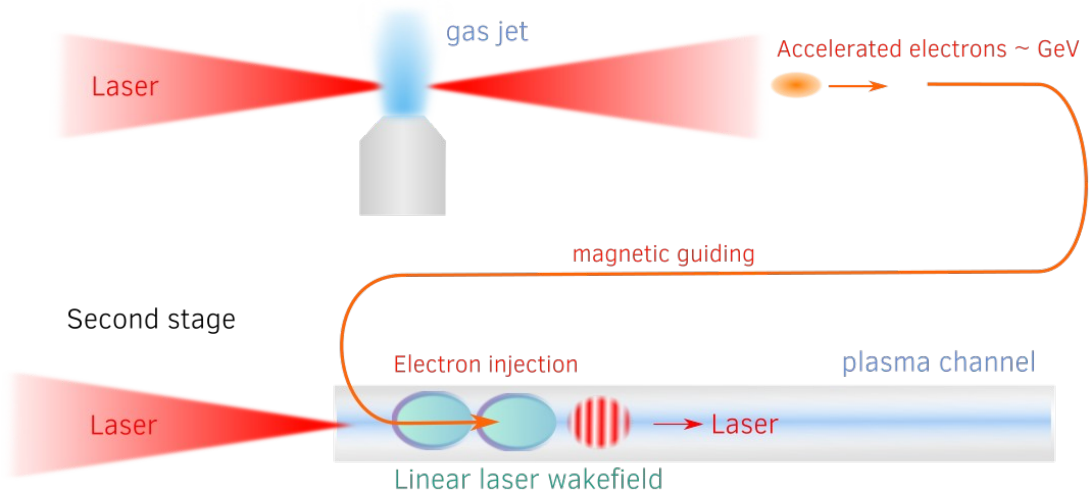

Stages are sequential so the problem size is constant

=> the problem qualifies for exascale provided strong scaling again.

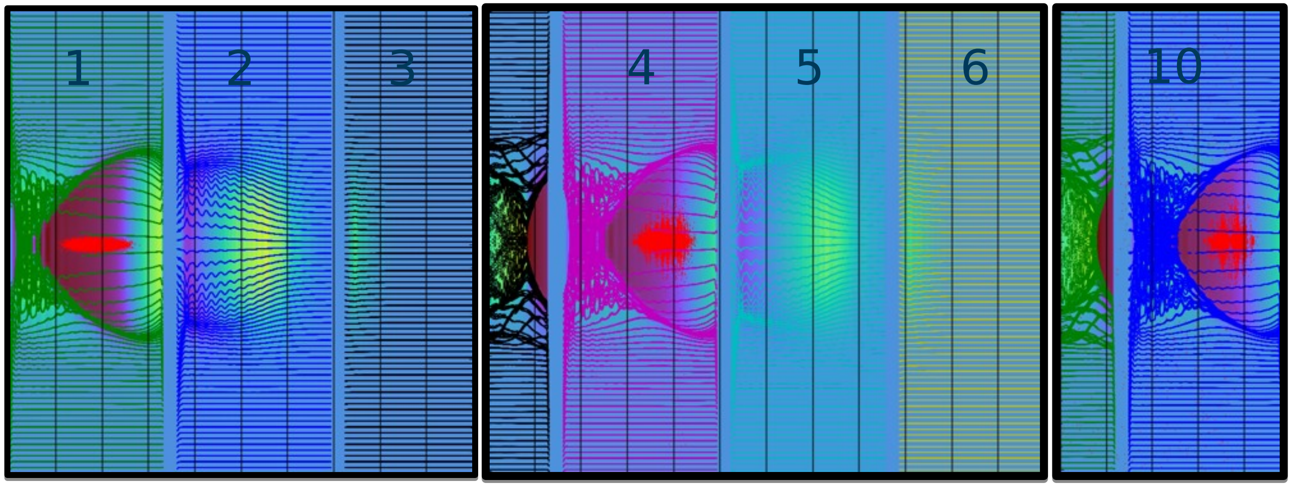

10 stages simulation by WarpX

D. Amorin, M. Thevenet, A. Huebl

Conclusion

- Exascale is coming.
- Codes and algorithms adaptations are mandatory.
- High gradient accelerator simulations are good enough candidates for it.
- New developments are underway to support this new incoming technology.
- Software engineering is the key challenge here in order to benefit from it.
- The U.S are a step ahead of Europe thanks to the ECP project so let’s not miss the euroHPC opportunity.
- We see SMILEI as a contribution to a European joint effort. We hope to see you at our 3rd User & Training Workshop in March 2022 !

Multiple geometries

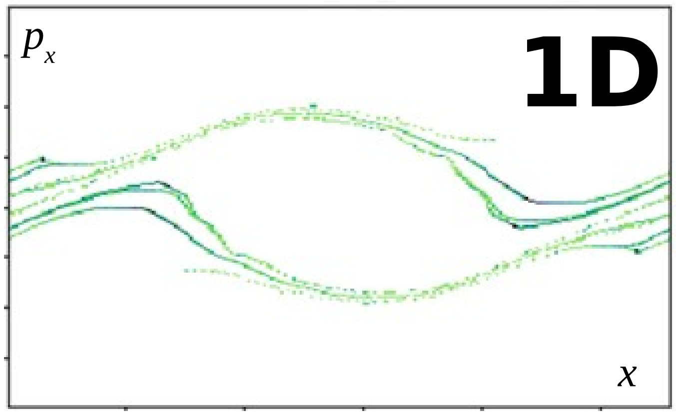

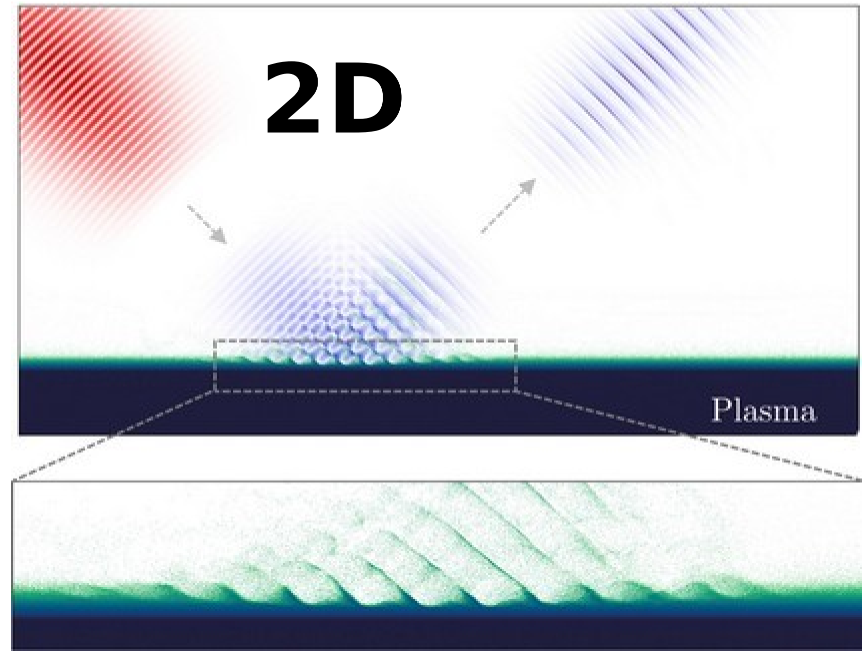

Cylindrical azimuthal modes
(3D particles)

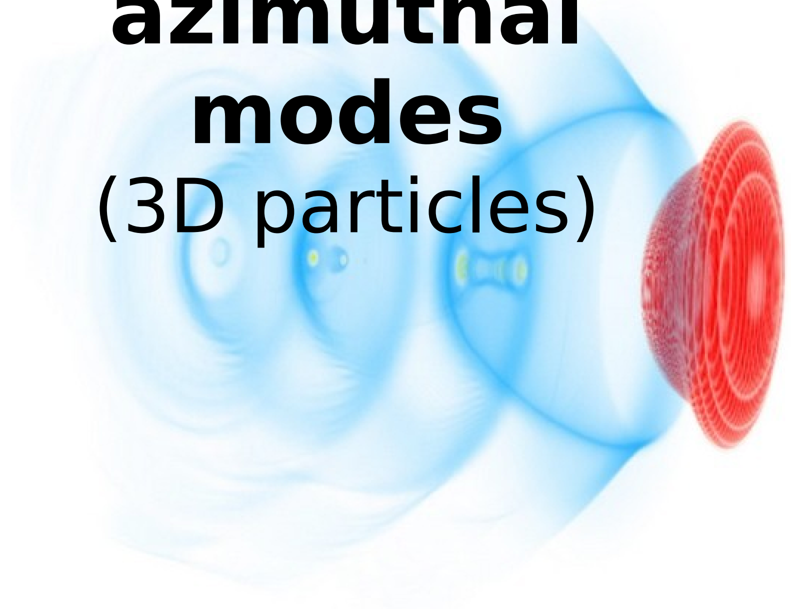

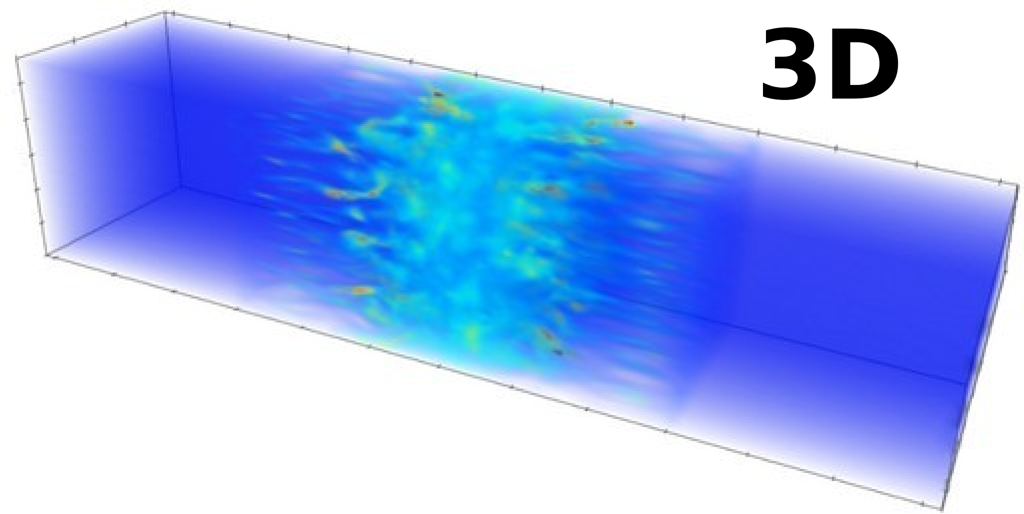

Additional physics

Radiation reaction

+ photon production

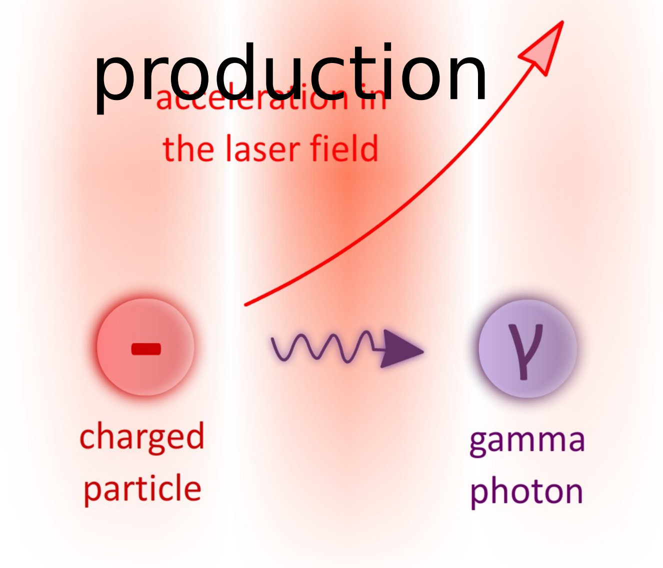

Non-linear Breit-Wheeler

pair cascades

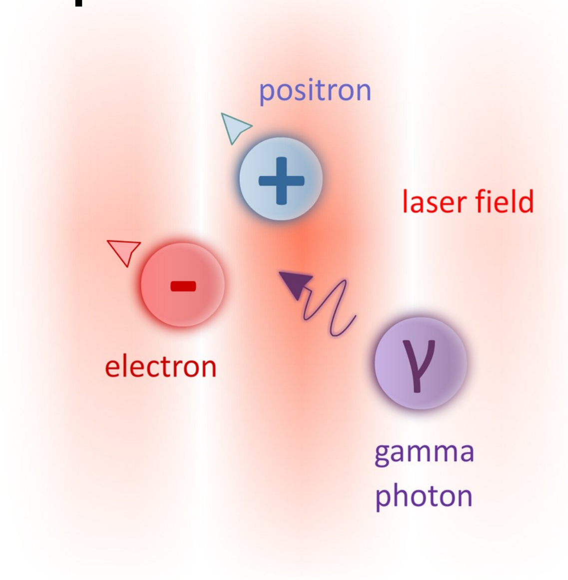

Lobet et al., JPCS (2016)
Niel et al., PRE *97* ***(2018); PPCF*** *60* ***(2018)***

Advanced numerical methods

- Full-PIC = resolve the laser wavelength
- Approximation : reduced equations on laser envelope

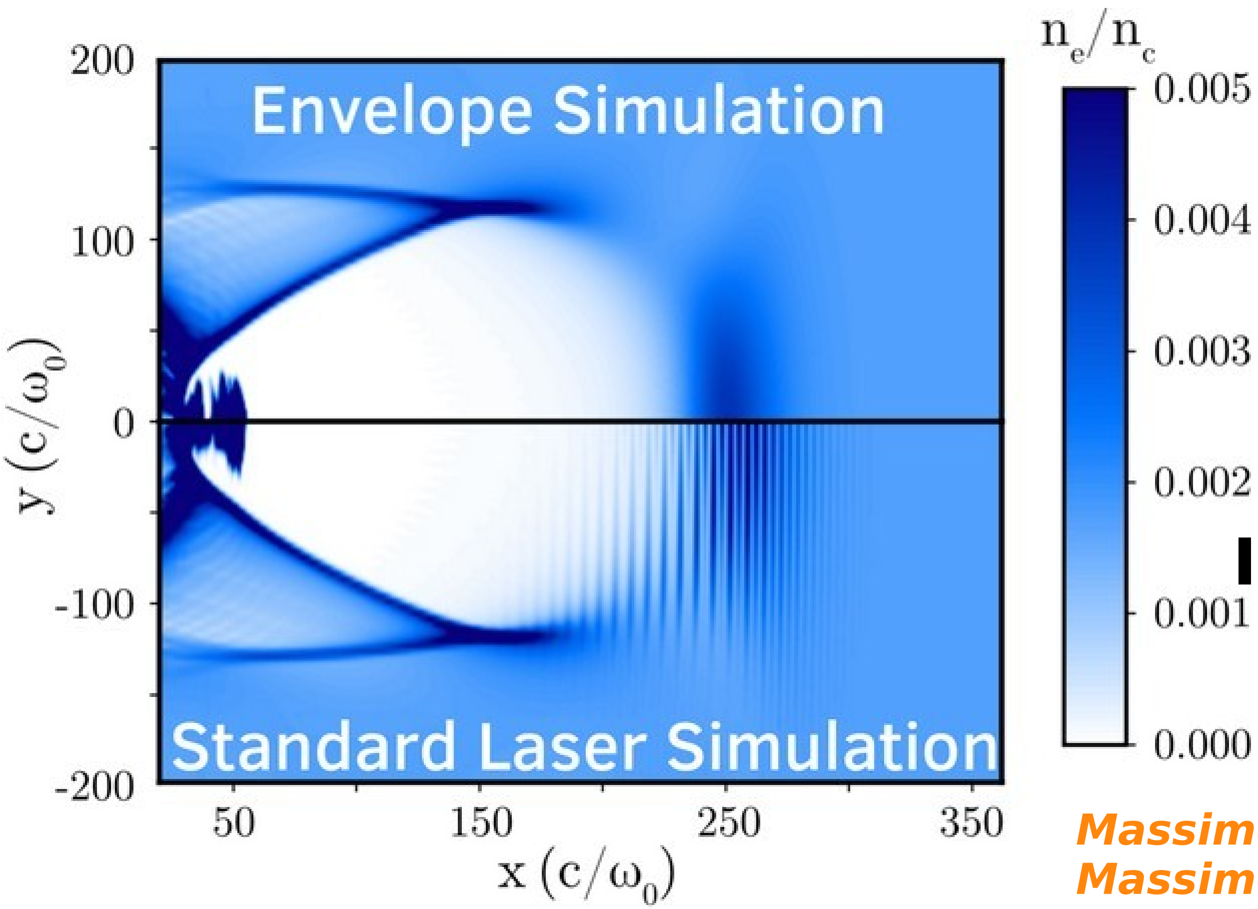

Envelope simulation:

20 x faster in 3D

100 x faster in AM

Ionization compatible

Massimo et al., PPCF (2019)
Massimo et al., IOP Proceedings (2019)
Massimo et al., PRE (2020)

Parallel computing

Many nodes
+ network

MPI protocol

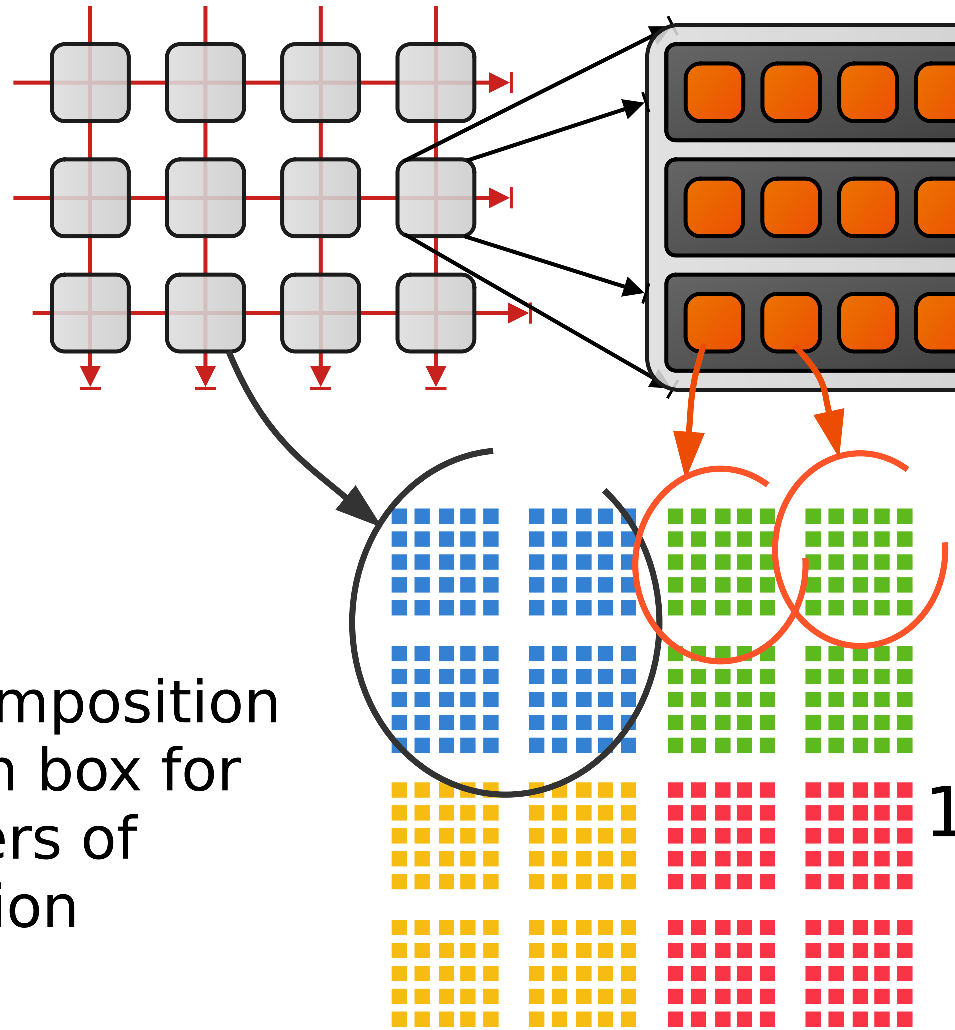

Many cores
per node

**OpenMP
directives**

Hierarchical decomposition
of the simulation box for
multiple layers of
parallelization

Good scaling to
100 000s of cores

Smilei)

Is part of the French national
benchmark for supercomputing

Happi post-process

- The repository includes a *python* module

```
$ ipython
In [1]: import happi; S = happi.Open("simulation_directory")
```

- Plot results

```
In [2]: rho = S.Probe(0,"Rho")

In [3]: rho.plot(timestep=180000, vmin=-0.02)
```

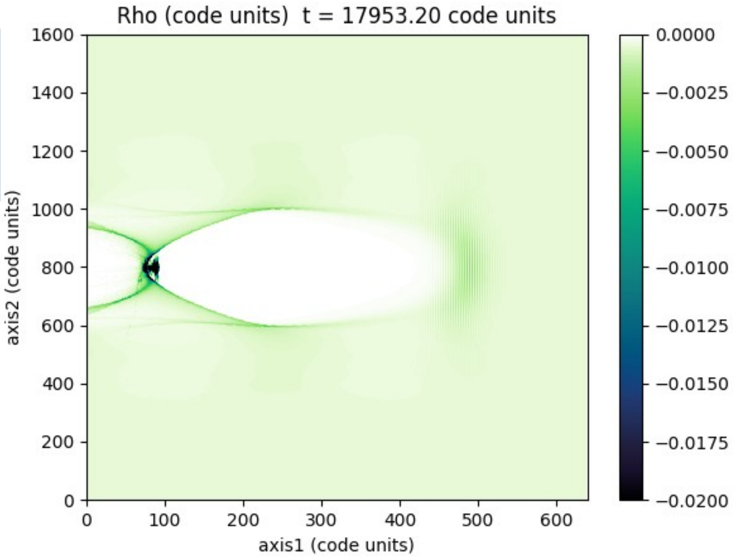

- Data manipulation

```
In [4]: data_array = rho.getData()
In [5]: rho.toVTK()
```

Smilei)

Open-source
PIC code

M. Grech
F. Perez
T. Vinci

M. Lobet
F. Massimo
J. Silvacuevas

MAISON DE LA SIMULATION

A. Beck
G. Bouchard

… and many more

maisondelasimulation.fr/smilei

github.com/SmileiPIC/Smilei

app.element.io/#/room/#Smilei-users:matrix.org

Derouillat* et al., *CPC* 222 *(2018)

High Performance

Vectorization activated only where there are many particles per cell

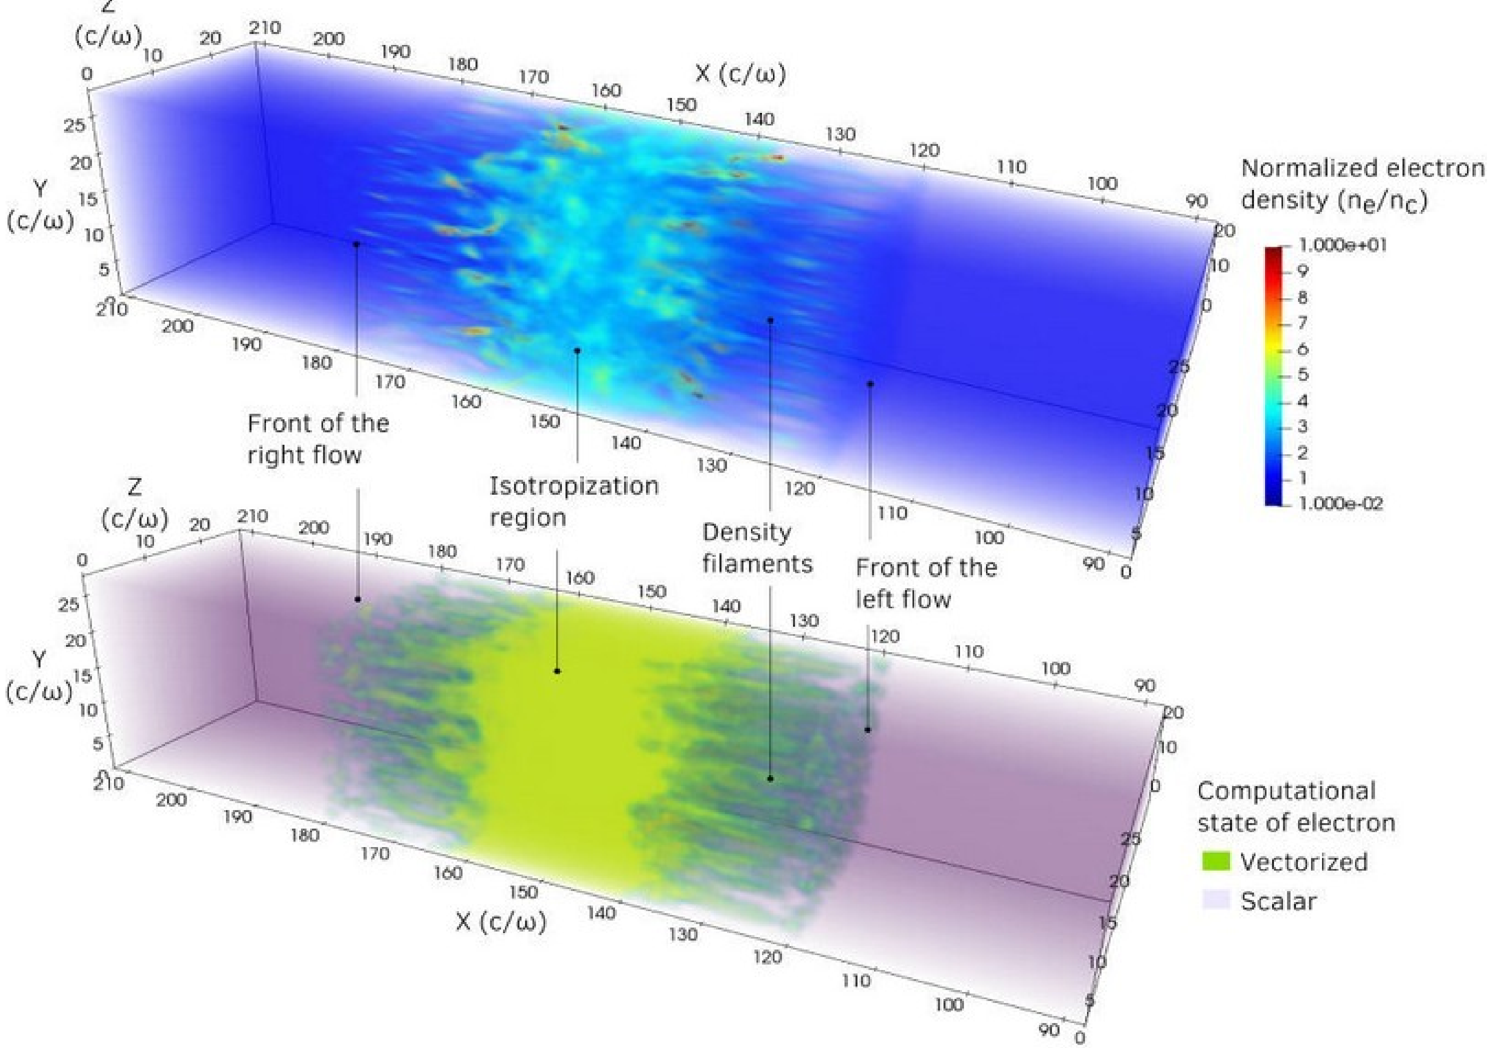

Beck et al., CPC 244 *(2019)*

Many options for solvers

- Charge-conserving current deposition

  Esirkepov, CPC** 135 **(2001)

- Orders of interpolation:

  2 or 4

  (3 or 5 points)

- Several FDTD schemes:

  “Yee”, “Cowan”, “Lehe”

  Nuter et al., EPJD** 68 **(2017)

- Spectral solver available *via* ***PICSAR*** *(beta)*

  picsar.net

Ionization by fields

- Monte-Carlo
- Multiple events in 1 timestep
- May define a custom ionization rate

Carbon ionization state *vs* time

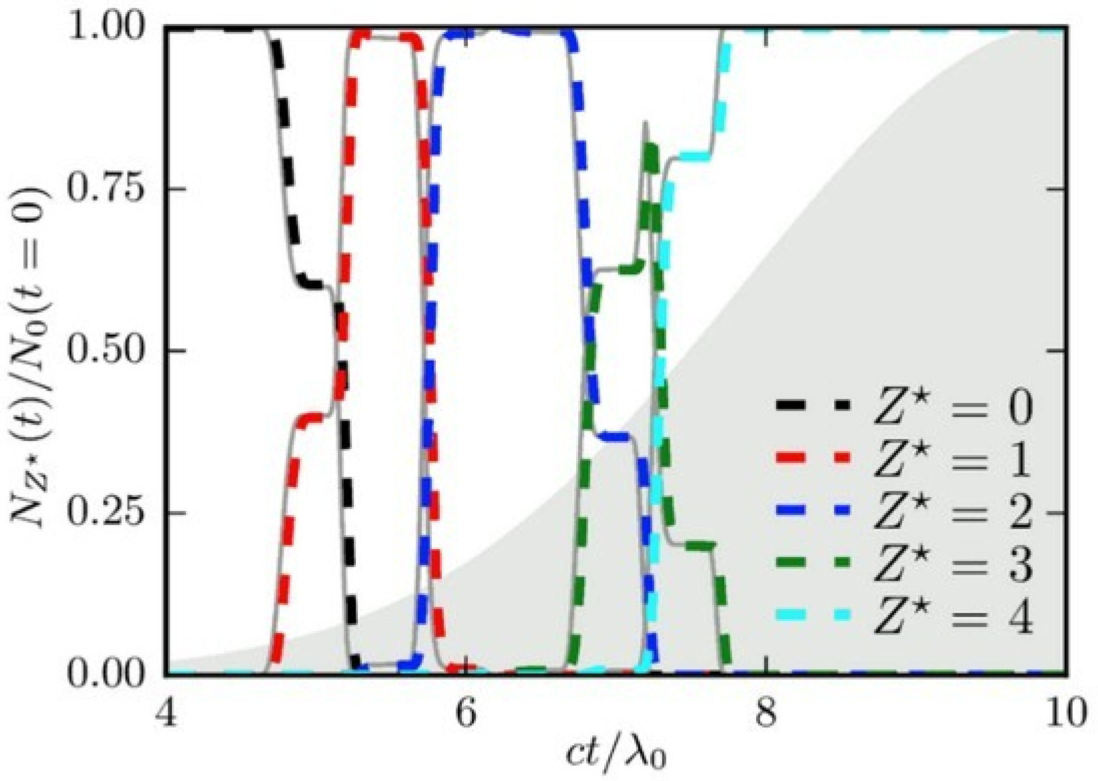

Nuter* et al., *POP 18 (2011)

Processes between pairs of particles

- Collisions
- Collisionnal ionization
- Nuclear reactions (D-D fusion in progress)

Conductivity of Cu
in wide temperature range

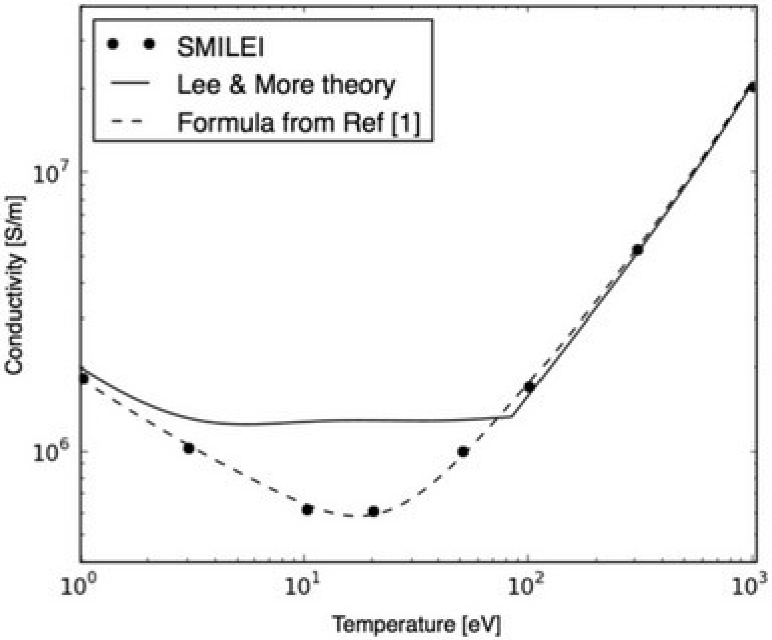

Electron stopping power in Al
elastic + inelastic

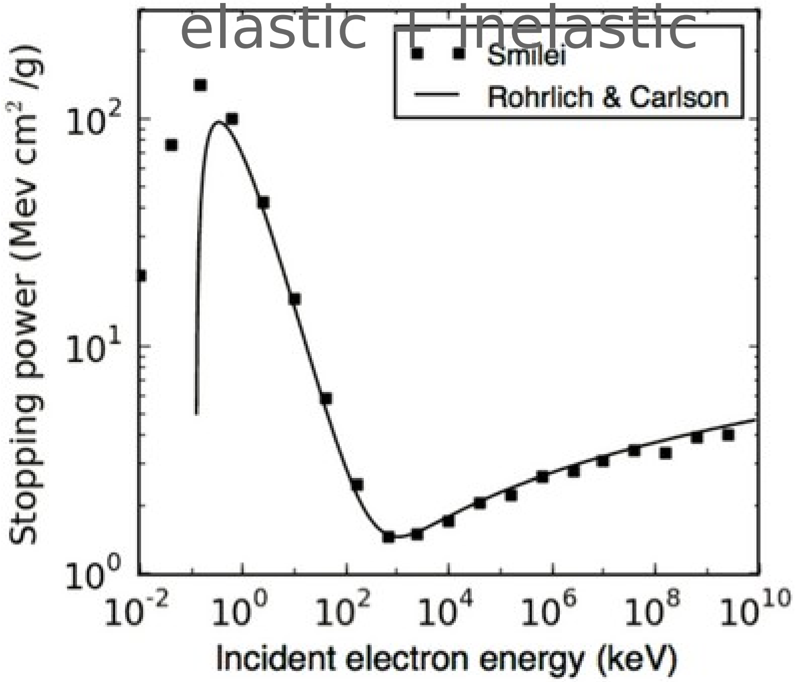

Envelope: wave equation

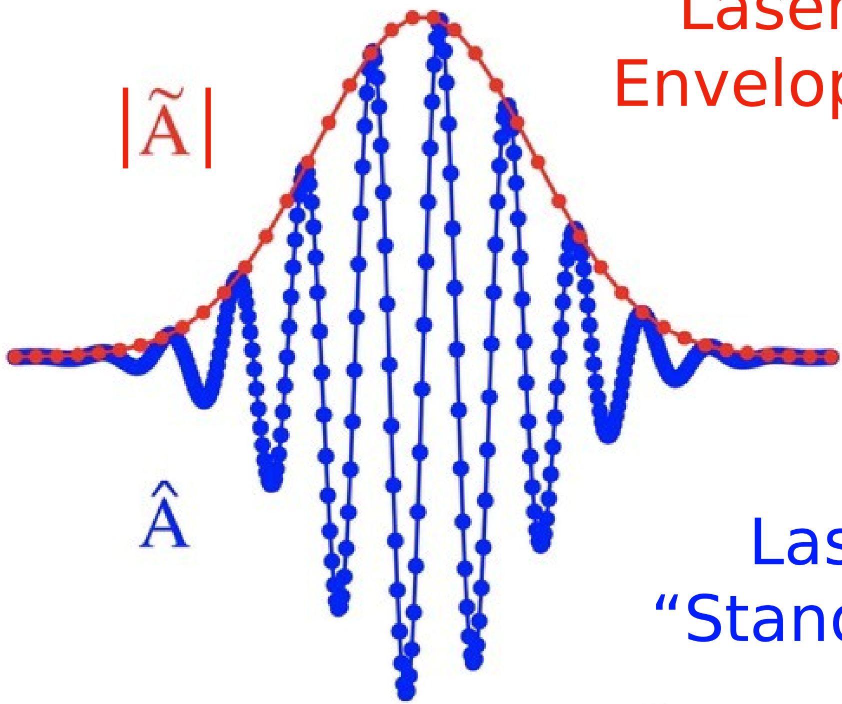

Terzani and Londrillo, CPC (2019)

Laser Complex Envelope

$$\hat{A}(\mathbf{x},t) = Re\left[\tilde{A}(\mathbf{x},t)e^{ik_0(x-ct)}\right]$$

+

D’Alembert Equation:

$$\nabla^2\hat{A} - \partial_t^2\hat{A} = -\hat{J}$$

Envelope Equation:

$$\nabla^2\tilde{A} + 2i\left(\partial_x\tilde{A} + \partial_t\tilde{A}\right) - \partial_t^2\tilde{A} = \chi\tilde{A}$$

Plasma Susceptibility

Envelope: particle motion

Ponderomotive force:
Acts as a radiation pressure on charged particles.
Expels the electrons from high-intensity zones.

Equations of motion for the macro-particles:

$$\frac{d\bar{\mathbf{x}}_p}{dt} = \frac{\bar{\mathbf{u}}_p}{\bar{\gamma}_p} \qquad r_s = q_s/m_s$$

$$\frac{d\bar{\mathbf{u}}_p}{dt} = r_s\left(\bar{\mathbf{E}}_p + \frac{\bar{\mathbf{u}}_p}{\bar{\gamma}_p} \times \bar{\mathbf{B}}_p\right) - r_s^2\frac{1}{4\bar{\gamma}_p}\nabla\left(|\tilde{A}_p|^2\right)$$

Lorentz Force (plasma fields)

Ponderomotive force (laser envelope)

B. Quesnel and P. Mora, Physics Review E 58, 3719 (1998)

Smilei) Timeline

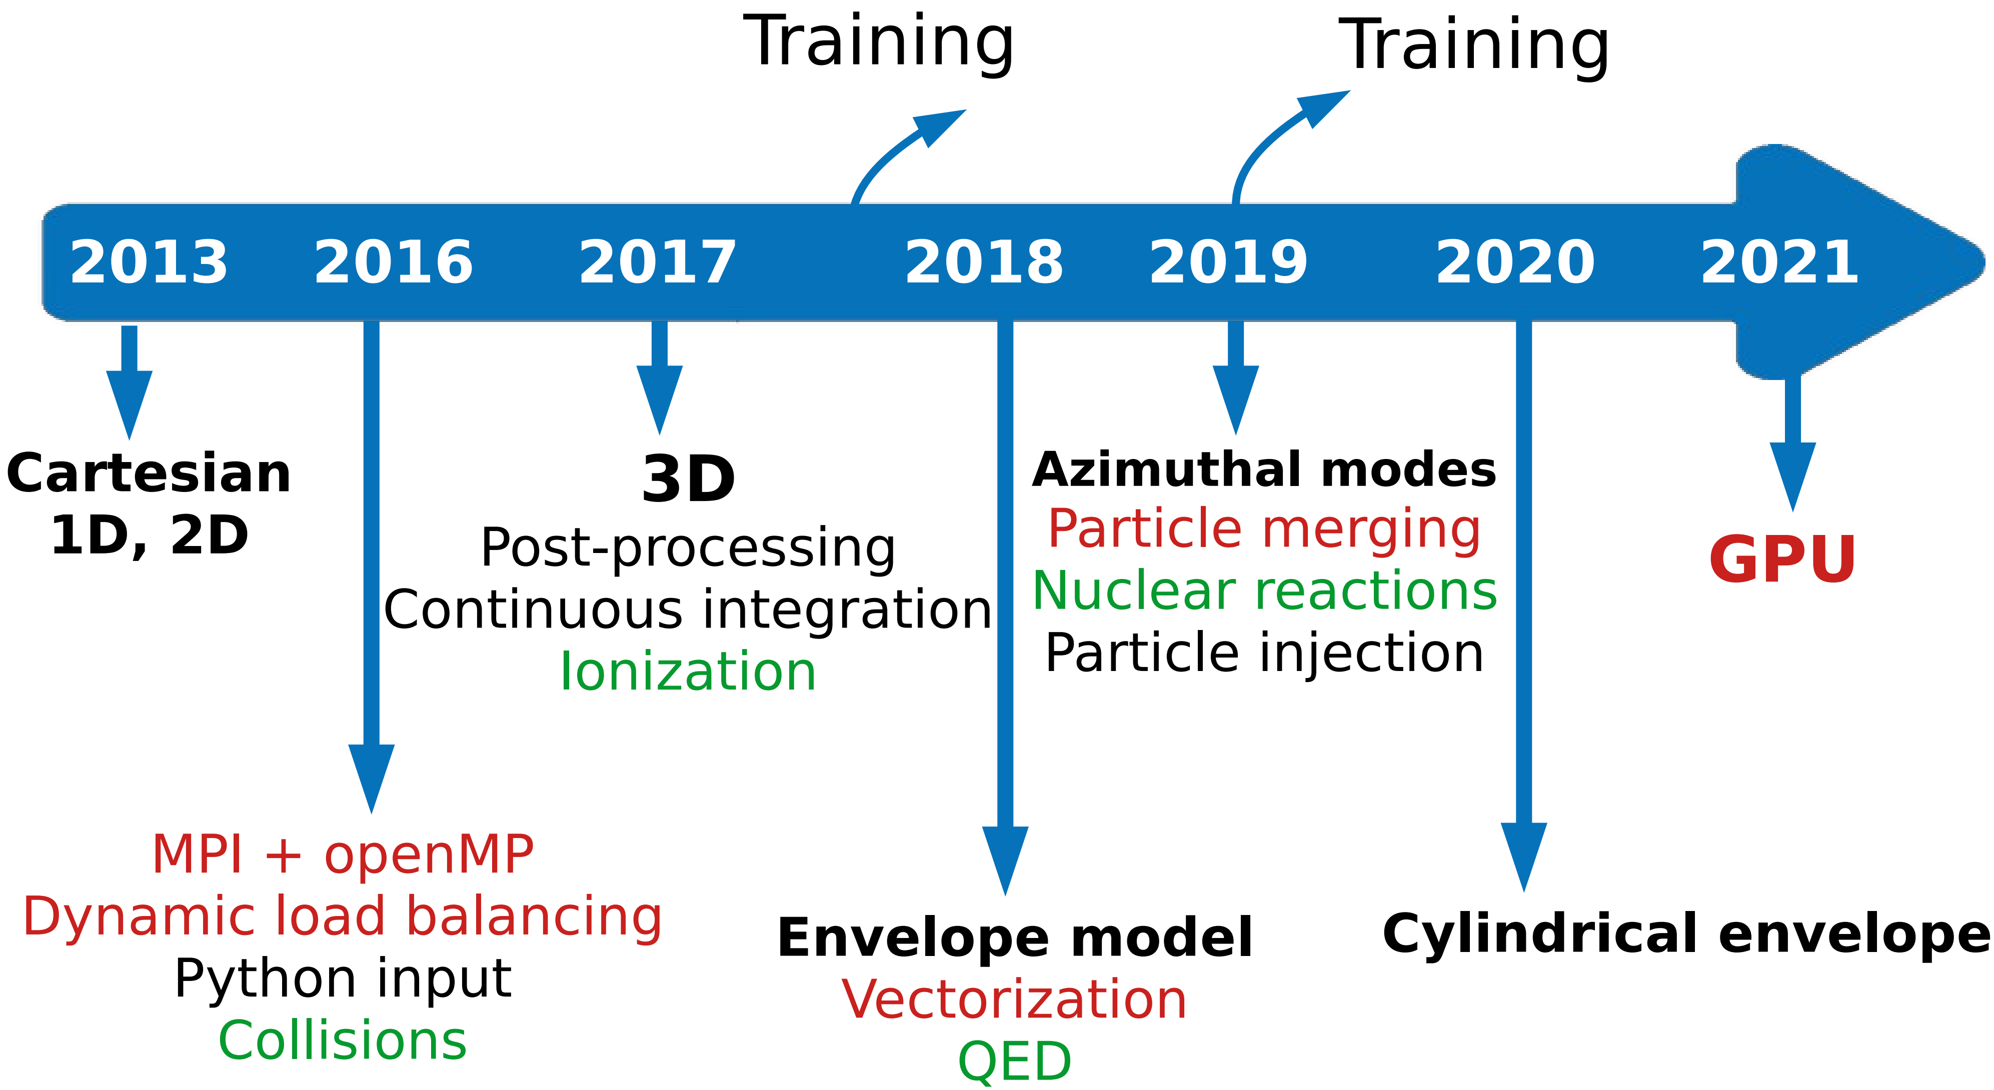

Derouillat et al., CPC 222 (2018)