ANF Formation Protéger, valoriser
et faire vivre un logiciel libre
1-3 Juillet 2024, Frejus

Smilei)

Community scaling and challenges

Francesco Massimo (z10f )
CNRS-LPGP-Université Paris Saclay
francesco.massimo@universite-paris-saclay.fr

Development team

Charles Prouveur* *
Mathieu Lobet*
Julien Derouillat
Haïthem Kallala
Juan Jose Silva Cuevas

Olga Abramkina*

Arnaud Beck*
Imen Zemzemi
Guillaume Bouchard (now at CEA-DAM)

Francesco Massimo*

Mickael Grech*
Frédéric Pérez*
Tommaso Vinci*
Anna Grassi*
Marco Chiaramello

*Code architect *Permanent staff

Acknowledgements

Contributing labs, institutions & funding agencies




Particle in Cell (PIC) code concept

Sample Plasma with Macro-Particles
(1 Macroparticle = position, momentum, charge, …)

+

Discretize space with computational grid

Define E, B, ρ, J on the grid cells

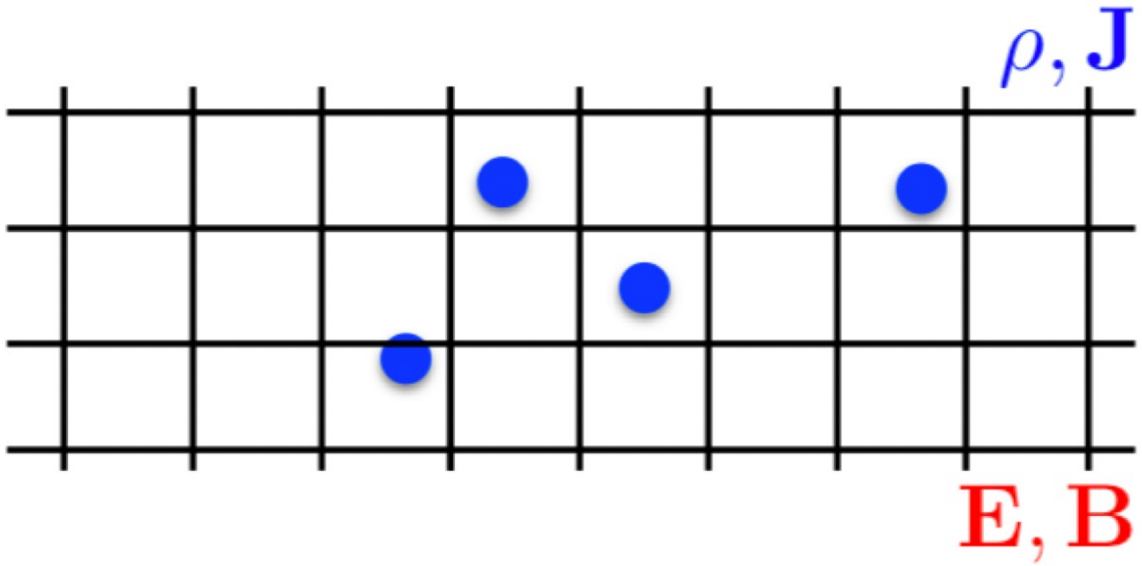

PIC codes self-consistently model the plasma behaviour

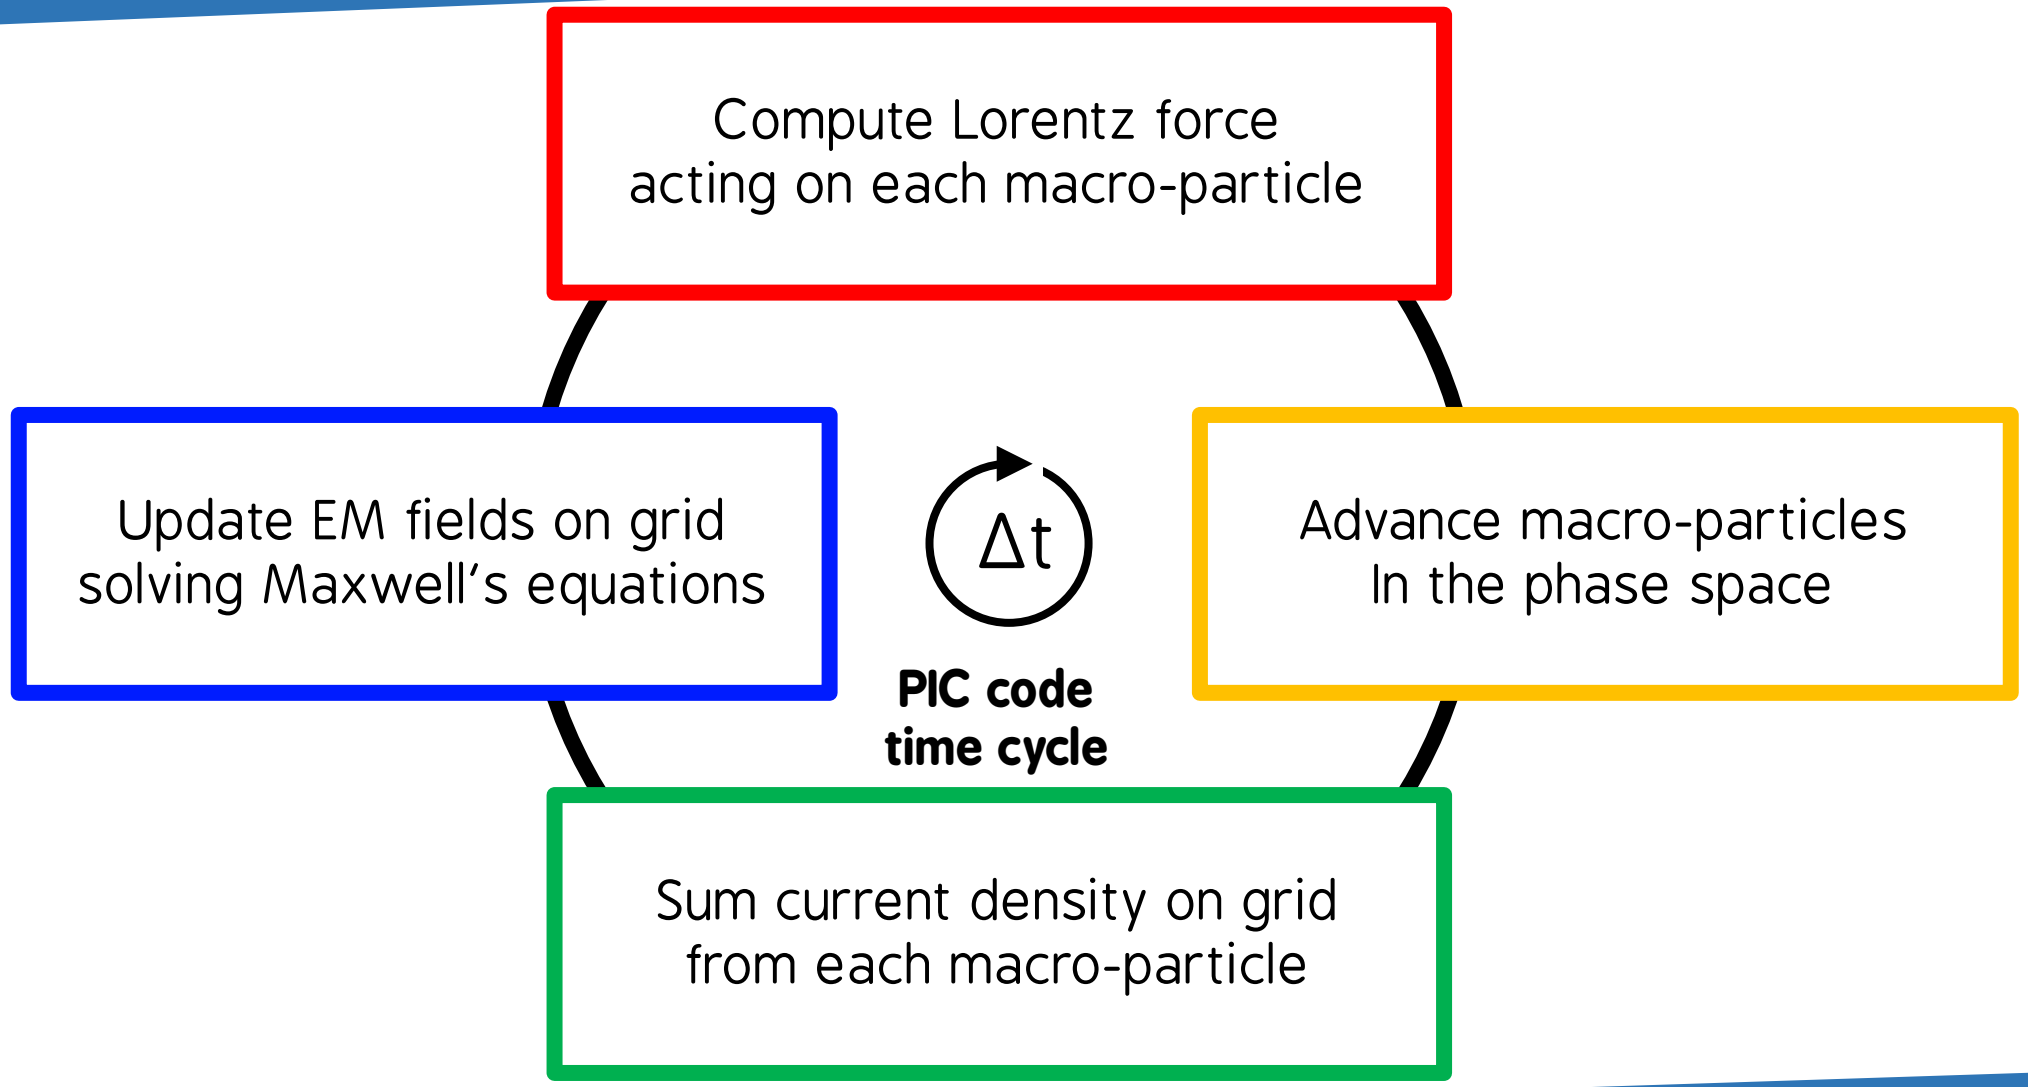

Examples of applications from the last Smilei workshop participants

High-Harmonic Generation

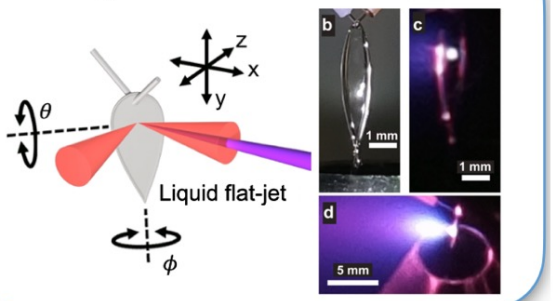

Radiation sources from DLT

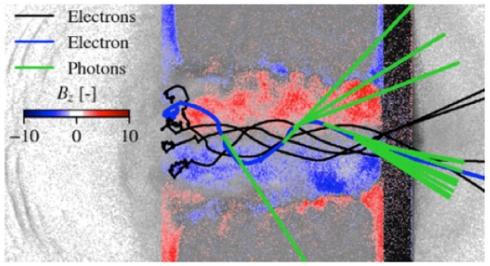

Solid to plasma transition

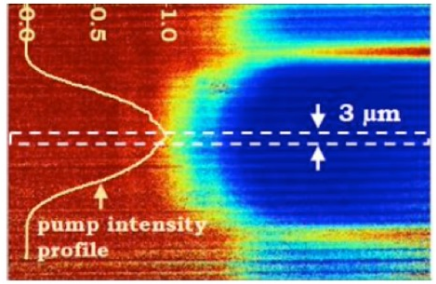

Solar radio-burst

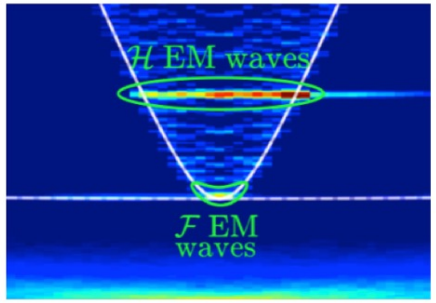

Comet boundaries

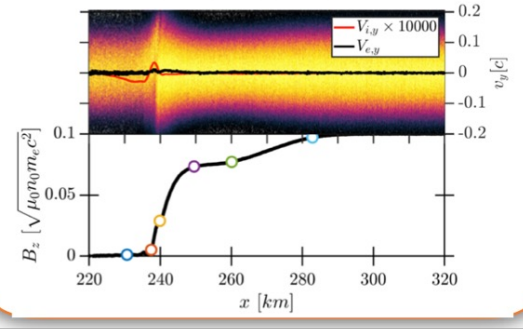

Collisionless shocks & Dark Matter

PIC method advantages:

- Wide range of physics applications
- Conceptually simple
- Efficiently implemented on small or massively parallel supercomputers

History of the project

2013
Start of the project*

2014
GitLab release to co-dev

2016
First physics studies
and large-scale simulations

2018
Reference paper

2023
Prix Science Ouverte
du logiciel libre de la recherche, catégorie scientifique et technique

*Objective: develop the first open-source PIC code harnessing new paradigms of high-performance computing

Open-source and community oriented
documentation • chat • online tutorials • post processing & visualization training workshops • summer school & master trainings • issue reporting

Multi-Physics & Multi-Purpose
advanced physics modules: geometries, collisions, ionization, QED; broad range of applications: from laser-plasma interaction to astrophysics

High-performance
C++/Python • MPI/OpenMP/OpenACC/CUDA/HIP • SIMD • HDF5 designed for the latest architectures

A project anchored to French and European HPC landscape…

Integration in the French & European HPC landscapes

- Running on all super-computers in France and many in Europe
- 10s millions computing hours every year via GENCI & PRACE / EuroHPC
- GENCI technological survey
- French Project NumPEX, Exascale project

GENCI
Le calcul intensif au service de la connaissance

Special / early access to various machines

2015.2015 IDRIS / Turing BlueGene-Q
2016.2016 CINES / Occigen
2018.2018 TGCC / Irene-Joliot-Curie
2019.2019 IDRIS / Jean Zay
2021.2021 RIKEN / Fugaku
2022.2022 CINES / Adastra (GPU)

… and a teaching platform

Python interface and "user-friendliness"
helped in using Smilei to teach plasma physics

- at the Master/doctoral levels in Europe in various winter/summer schools
- in user & training workshops
- via online tutorials
- EUR Plasma - Ecole polytechnique – Université Paris Saclay

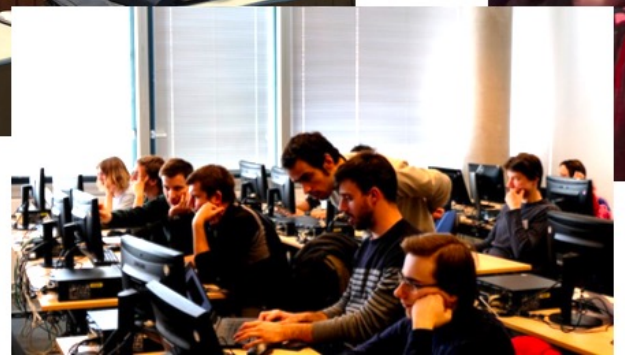

An international, steadily-growing user community

Countries of 1st authors affiliation

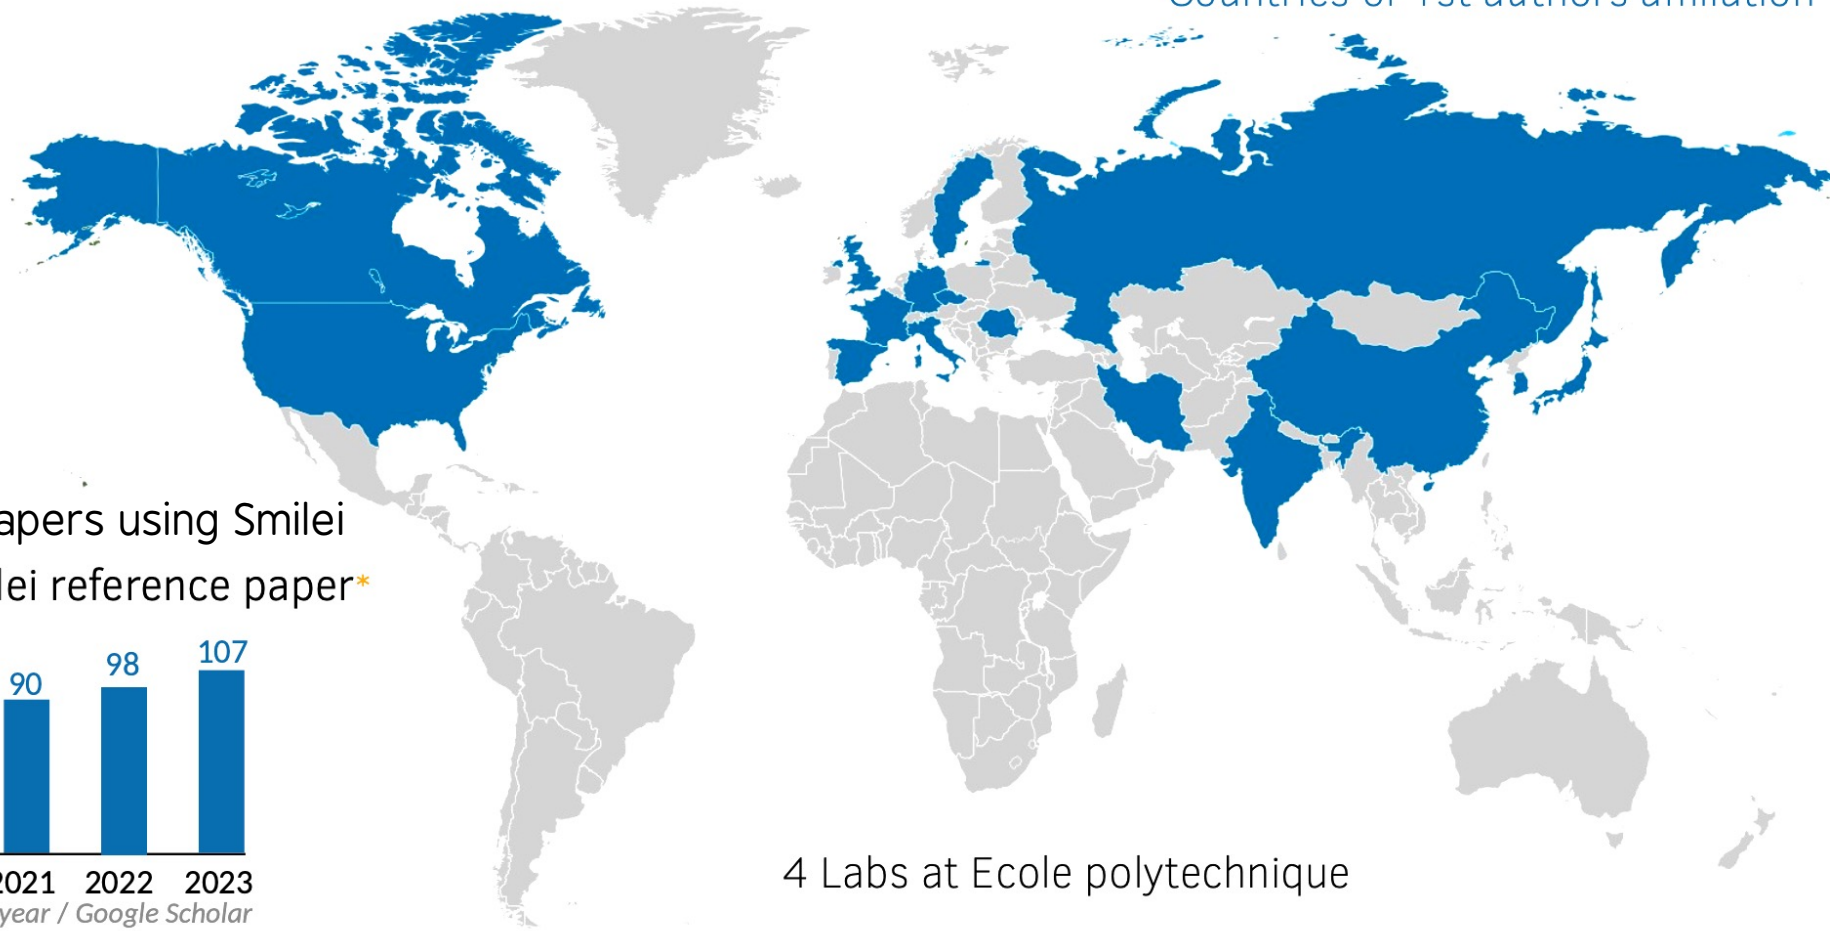

~200 peer-reviewed papers using Smilei

~470+ citations for Smilei reference paper*

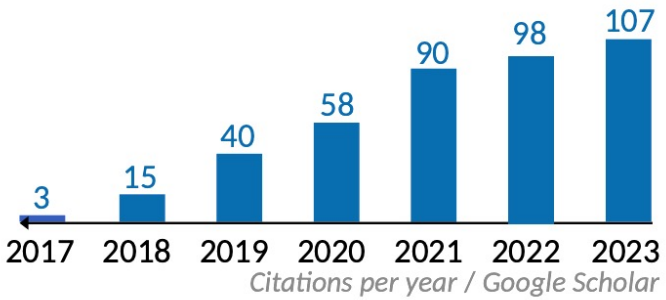

4 Labs at Ecole polytechnique

*Dérouillat et al., Comp. Phys. Comm. 222, 351 (2018)

Resources for the code community: documentation, tutorials and chatroom

An extensive documentation + tutorials

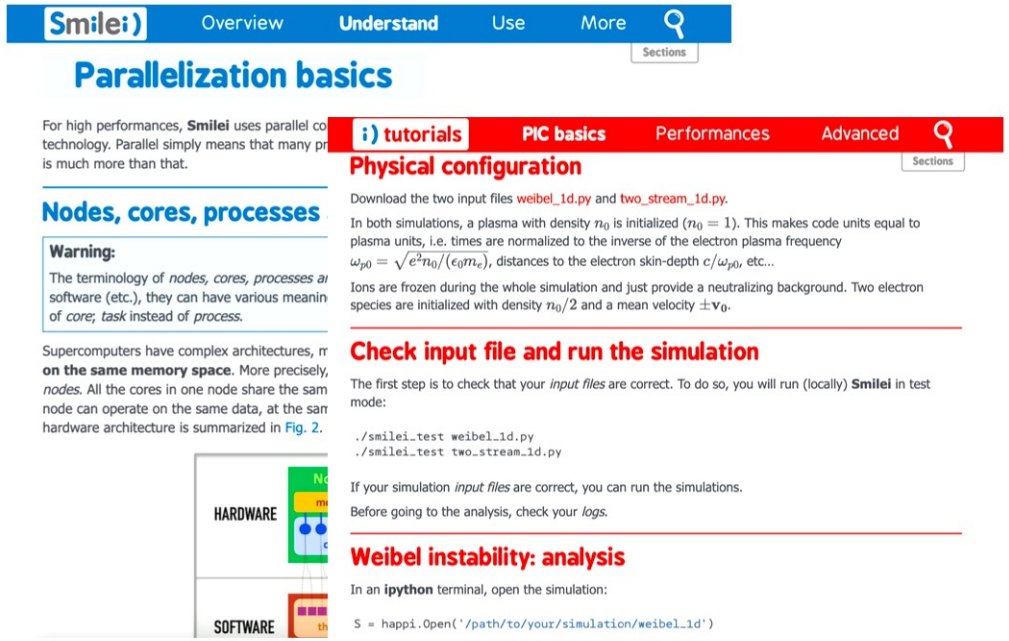

Code chatroom:

contact the developers and other users

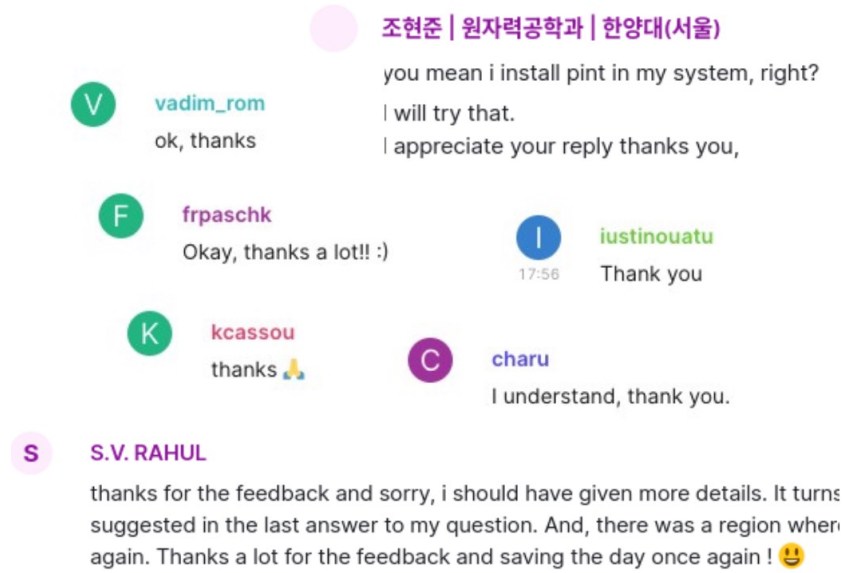

Interaction platforms for a collaborative community

Github
sources, issues, discussions

Element
chat

Smilei User and Training Workshops

- Project Review and Perspectives
- Hands-on tutorials
- Presentations by the users
- Participants from 10 countries in the last edition

Example of Contributions

Dynamic load balancing and SIMD (A. Beck)

A. Beck et. al. Computer Physics Communications 244, (2019)

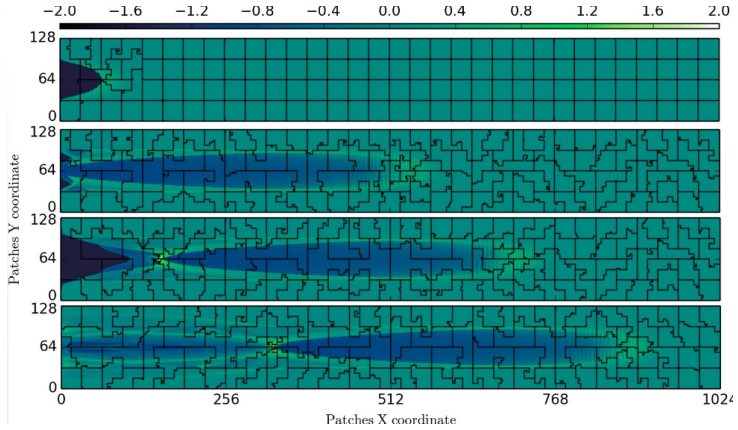

Perfectly Matched Layers (G. Bouchard)

G. Bouchard, A. Beck et al., submitted

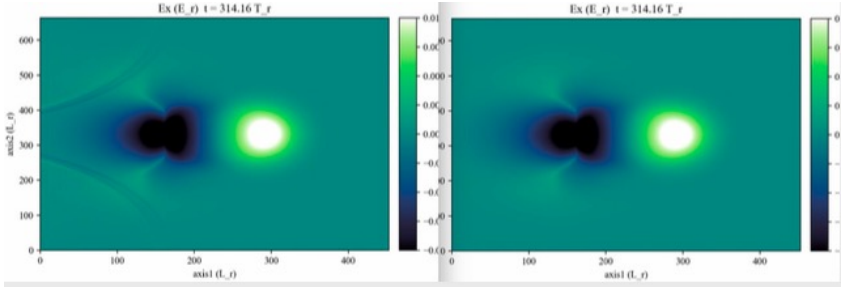

Azimuthal modes geometry (I. Zemzemi)

I. Zemzemi et. al., J. Phys.: Conf. Ser. 1596, 012055 (2020)

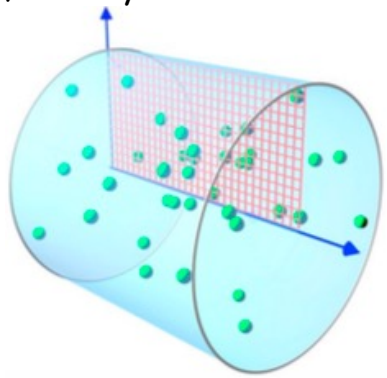

Envelope model (F. Massimo)

Massimo et. al. Plasma Phys. Control. Fusion 61, 124001 (2019)

Massimo et. al. Phys. Rev. E 102, 033204 (2020)

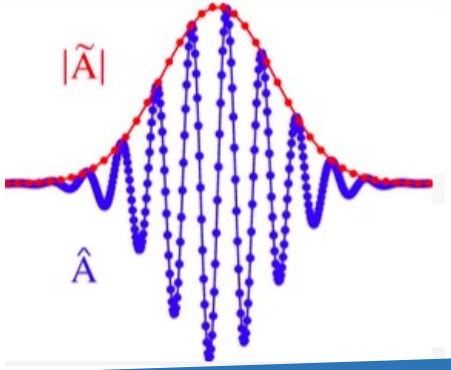

Perspectives

Code & HPC aspects

- GPU porting: more features will be progressively ported and continuous integration in progress
- Advanced IO management (AI approach)

Additional physics modules

- Additional modules for plasma acceleration, atomic processes, more nuclear fusion processes

Keep on building & animating the user community

- Encouraging new developers to join
- Developing an online teaching platform (beyond the tutorial approach)
- Preparing the next Smilei workshop

Governance model

- In Smilei we have a board of trusted maintainers for different sections : GPU, input/output, post-process, vectorization, MPI synchronizations ...

- Ideally at least 2 trusted maintainers should know very well a section of code to crosscheck modifications

- Requires collaborative work at interfaces

Scaling challenges

« As projects scale, work not only increases, but fundamentally changes. »

The Labor of Maintaining and Scaling Free and Open-Source Software Projects, Geiger et. al. Proc. ACM Hum.-Comput. Interact., Vol 5, No. CSCW 1, April 2021

- Support requests skyrocket.
- Maintenance of documentation and collaborative tools
- Maintenance of interfaces with other softwares and/or standards.
- Maintenance of performances in an evolving hardware environment.

Challenges of rapid growth

More maintenance (features, hardware, debugging…)

More communication (conference, workshop, documentation, promotion…)

More support (compilation, use, post-process …)

More integration in house or from contributions (C.I., debugging pull requests)

… leads to overworked maintainers

« Scalar labor » is necessary to maintain a project at its new scale.

Conclusions

- Smilei is an example of 10-year-old open-source project for a multi-purpose physics simulation code
- Some key aspects that helped the project include: development team composed by experts of physics and high performance computing, its scalability from small to large machines
- However the main strength of the code comes from a dynamic user and developer community
- As the size (and community) of an open-source project increases, additional "scalar work" is necessary
- Many resources (documentation, tutorials, …) with different degrees of detail and specificity, and different options to contact the developers (GitHub Issues/Discussions, Element Chat) help both the developers and the users
- The work and resources needed to develop, maintain and address the community's needs will often be underestimated (even by the developers!)