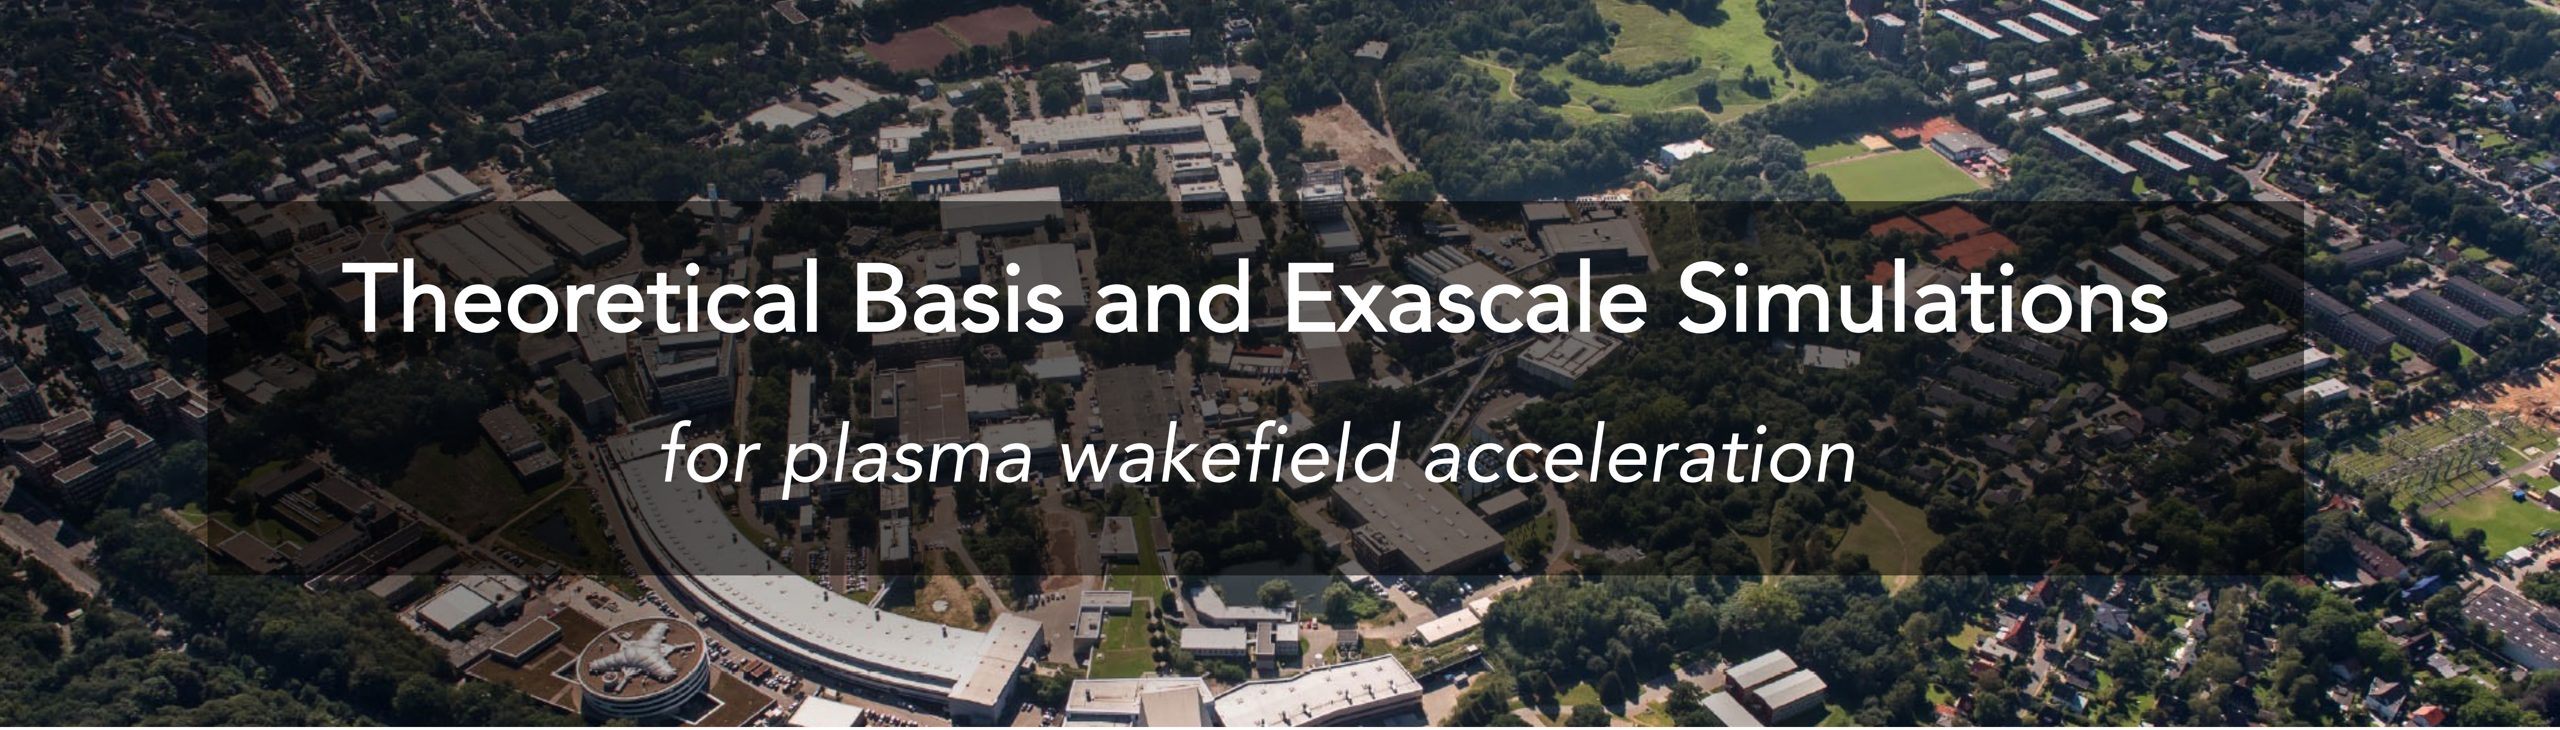

Theoretical Basis and Exascale Simulations

for plasma wakefield acceleration

Maxence Thévenet – DESY

MPA1 team: theory and simulations for plasma acceleration

I. Physics of plasma acceleration

II. The particle-in-cell (PIC) method

III. High-performance computing

IV. Recent activities at DESY

HELMHOLTZ RESEARCH FOR GRAND CHALLENGES

DESY.

I. Physics of plasma acceleration

II. The particle-in-cell (PIC) method

III. High-performance computing

IV. Recent activities at DESY

HELMHOLTZ RESEARCH FOR GRAND CHALLENGES

Particle accelerators are large devices

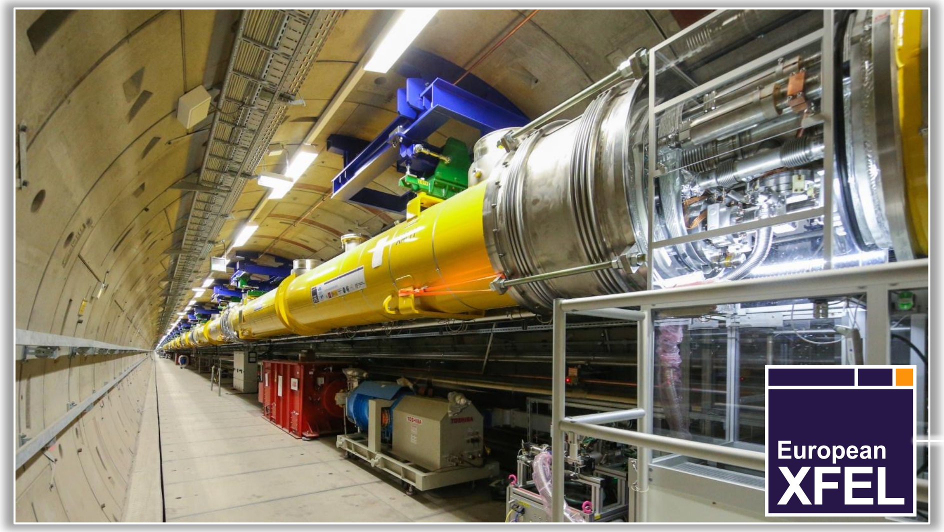

- Kinetic energy: $\mathcal{E}_k \propto L \times E_z$
- Conventional accelerators limited to E_z < 100 MV/m
- Applications from MeV to TeV energy ranges
- **Plasma acceleration** E_z > 10 GV/m → 100x more compact

Plasma acceleration: an alternative to conventional technologies

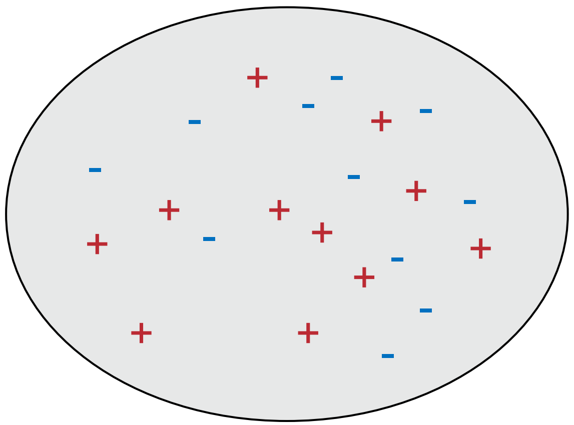

Plasma

- electrons and ions (q/m >1000x larger)
- Sensitive to electromagnetic fields → collective effects
- Electron plasma waves: $\omega_p = \sqrt{\frac{n_e e^2}{m_e \varepsilon_0}}$, $k_p = \frac{\omega_p}{c}$, $\lambda_p = \frac{2\pi}{k_p}$

$$\frac{d\boldsymbol{p}}{dt} = -e(\boldsymbol{E} + \boldsymbol{v}\times\boldsymbol{B})$$

For the beam: $\boldsymbol{v_z} \sim c$

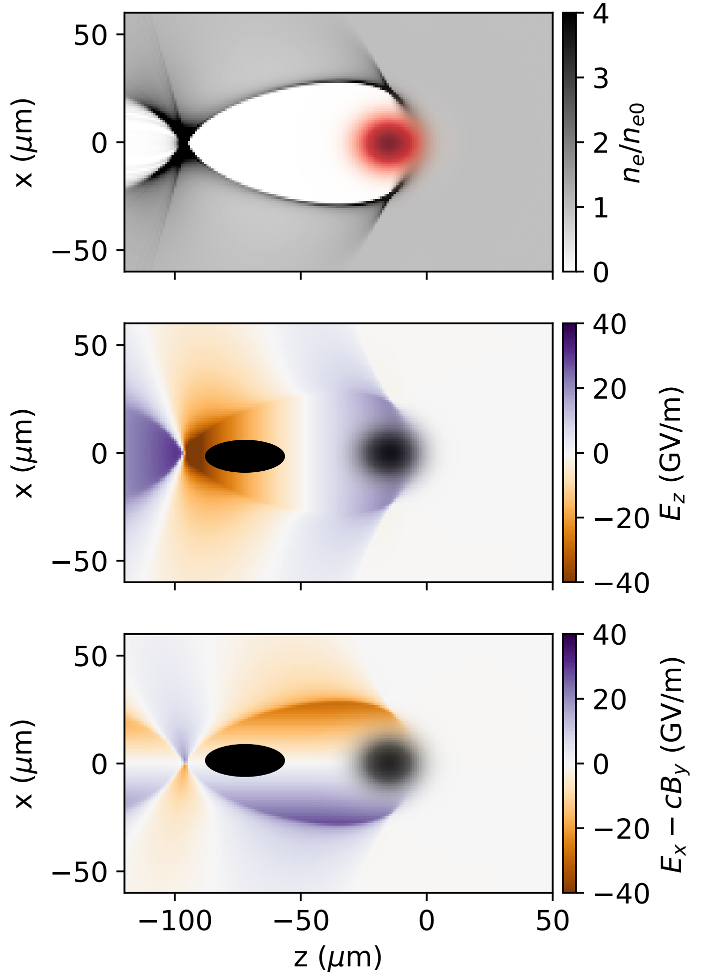

Wake ~ λ_p = 10s – 100s μm
Propagation ~ mm-m distance

T. Tajima, J. M. Dawson. *PRL* 43.4 (1979)

The wake can be driven by a particle beam or a laser pulse

Electron beam

- $L < \lambda_p$
- $\boldsymbol{v} \sim c\boldsymbol{e_z}$
- $\rho, \boldsymbol{J} \rightarrow \boldsymbol{E}, \boldsymbol{B}$
- Fields depend on the beam current

Laser pulse

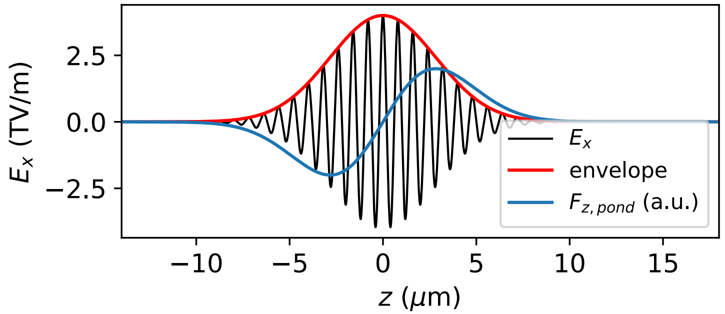

- Length $L < \lambda_p$, width $w_0 < \lambda_p$, wavelength $\lambda = 0.8$ μm
- $a_0 = \frac{eE_0}{m_e \omega c}$; non-relativistic $a_0 \ll 1$
- Ponderomotive force: $\boldsymbol{F_p} = -\frac{1}{2m_e\bar{\gamma}} \boldsymbol{\nabla} \overline{|qA_x^2|}$

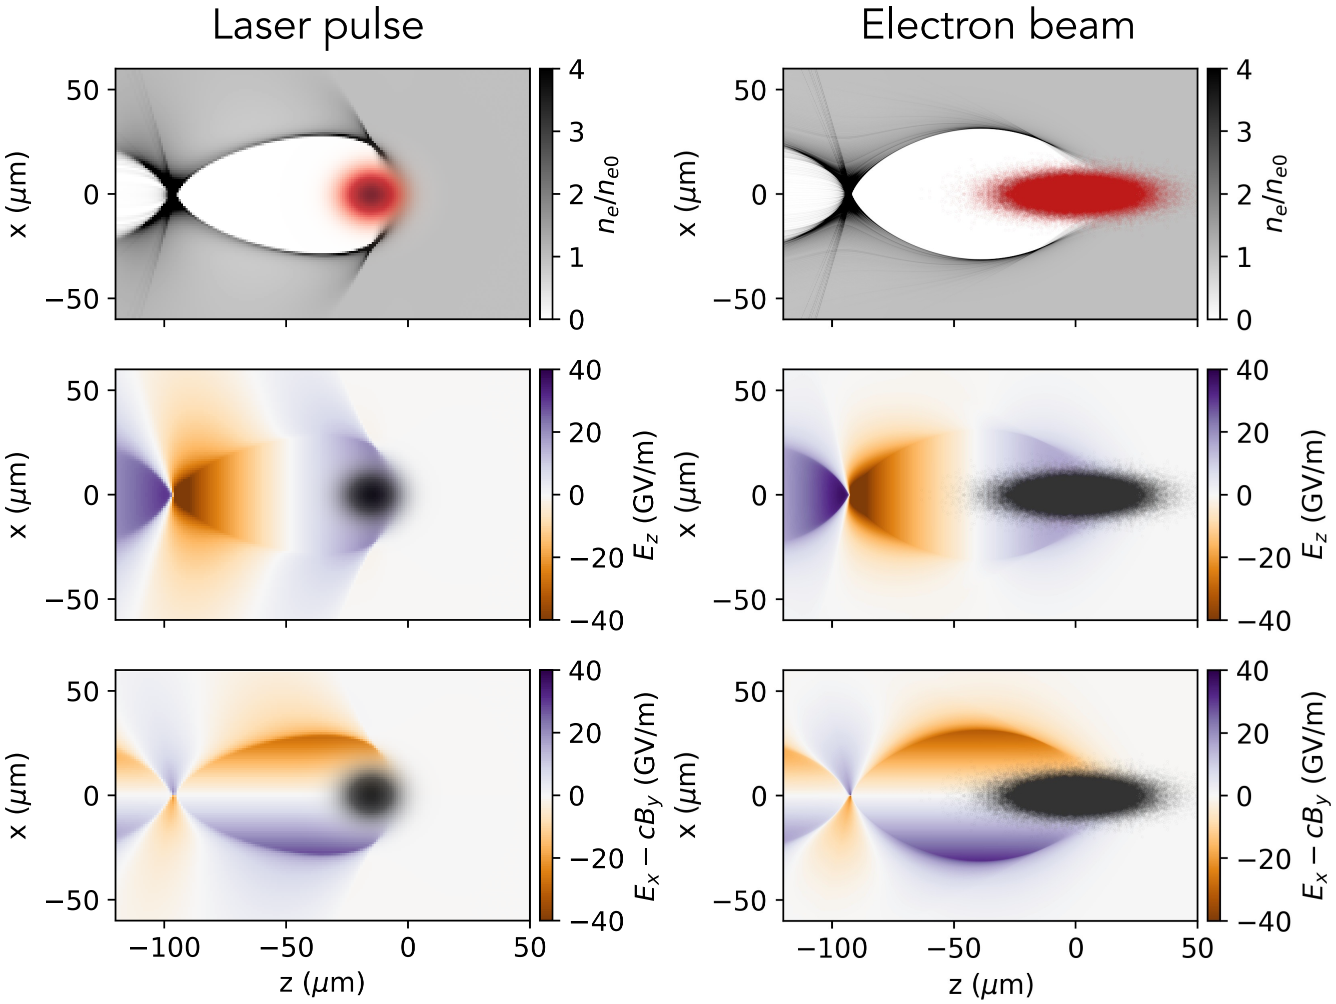

Wakefield acceleration is a complex dynamic process

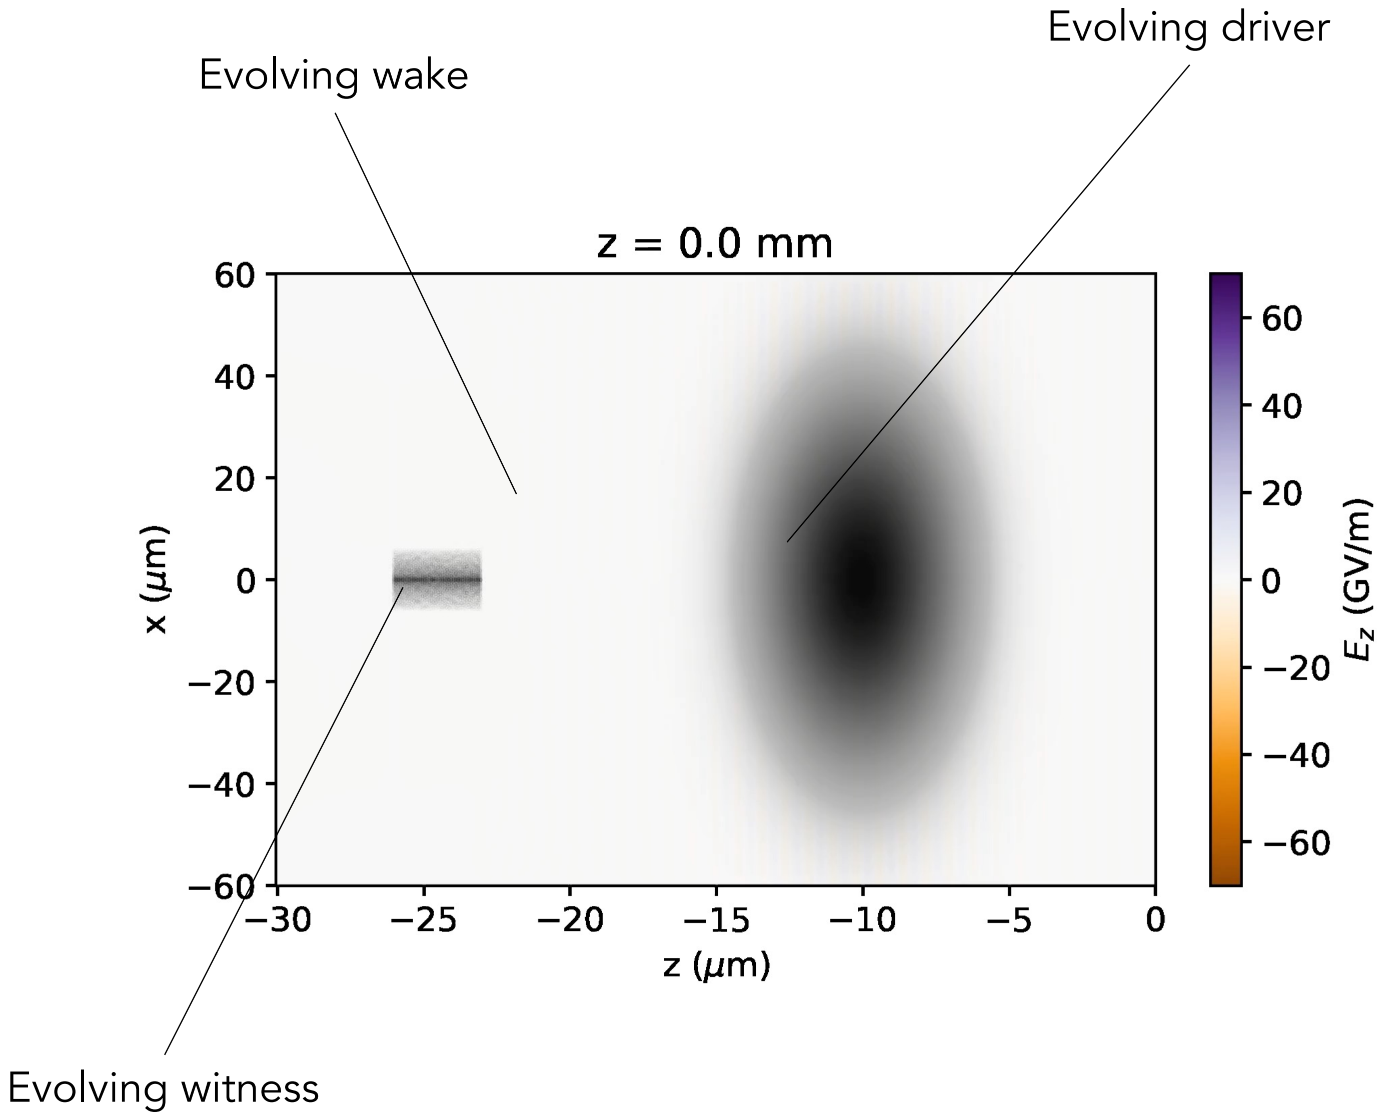

Wake excitation 1/2: linear regime

- Small perturbation of electron density
- Co-moving coordinate $f(z,t)$ → $f(\zeta = z - ct, t)$
- Quasi-static approximation: Neglect some time derivatives
- Electrons behaves as a laminar fluid → cold fluid theory

- Convenient variable: pseudo-potential
  $\psi = \phi - a_z$
  $\boldsymbol{F} = m_e c^2\left(-\partial_\zeta \psi \boldsymbol{e_z} + \boldsymbol{\nabla}_\perp \psi\right)$ for an electron with $\boldsymbol{v} \sim c\boldsymbol{e_z}$

$$\psi = \frac{-k_p}{2}\int_\zeta^{+\infty} \overline{|\boldsymbol{a_l^2}(\boldsymbol{u})|}\,\sin\left(k_p(\zeta - u)\right) du$$

→ Resonant excitation $k_p L \sim 1$
→ Harmonic waves, ¼ plasma wavelength focusing/accelerating

L.M. Gorbunov and V.I. Kirsanov, Sov. Phys. JETP 66 (1987)

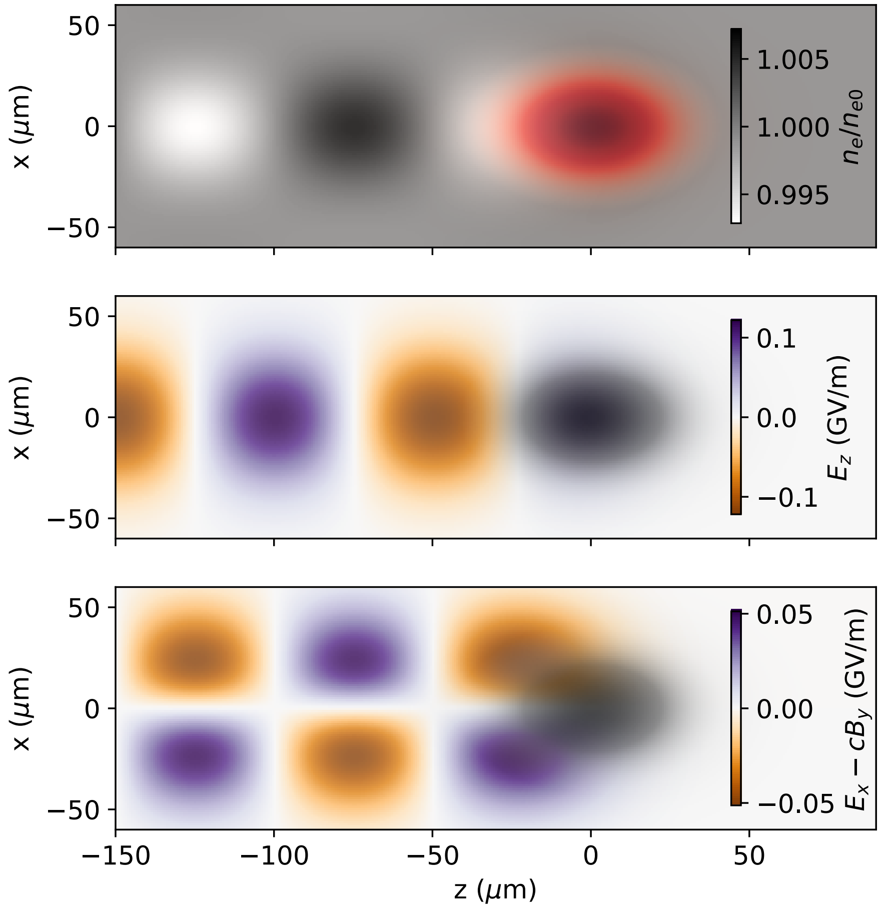

Condition for (quasi-)linear regime
Laser for $k_pL \sim 1$ & $k_p w_0 \leq 1$, $\boldsymbol{a_0} < \boldsymbol{1}$
Beam for $k_pL \sim 1$ & $k_p R \leq 1$, $\frac{\boldsymbol{n_b}}{\boldsymbol{n_e}} < \boldsymbol{1}$

Wake excitation 2/2: blowout regime

- Fluid description is not accurate
- Cold non-relativistic wave-breaking field [1] $E_0 = \frac{m_e \omega_p c}{e}$
- In the ion cavity
  - $E_r - cB_\theta = \frac{E_0 k_p r}{2}$ [2] Linear and independent on z
  - E_z independent on r
- Semi-analytic models [3,4] for E_z

→ Excellent properties for accelerating a beam

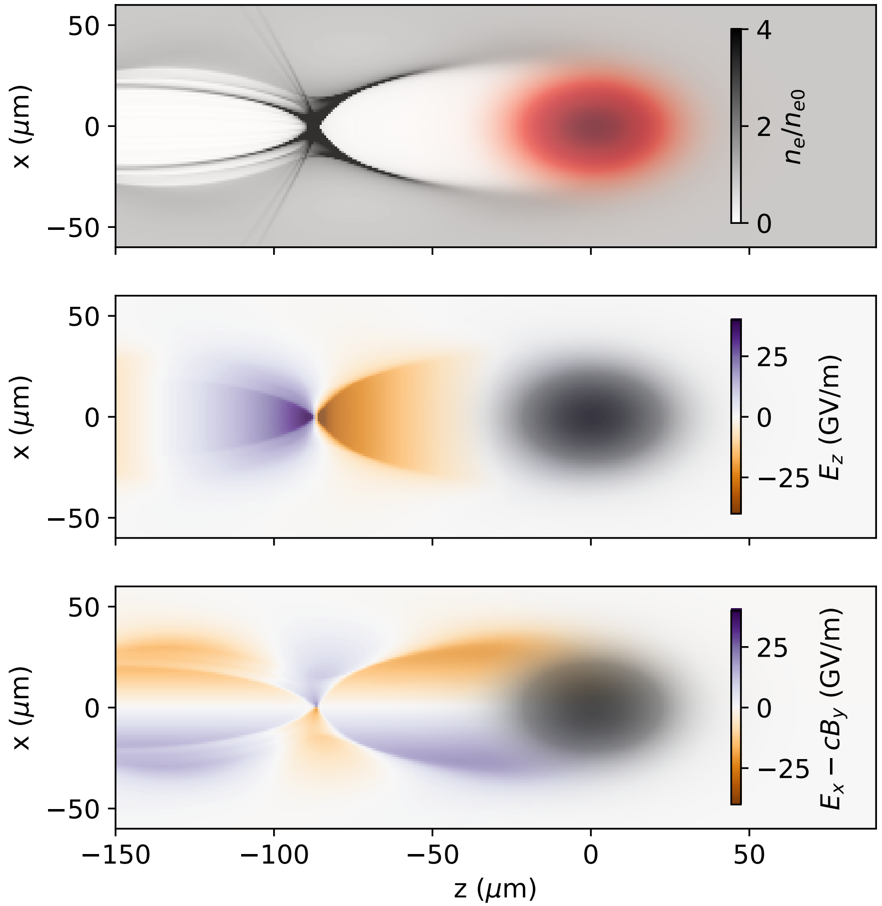

Condition for the blowout regime

Laser for $k_p L \sim 1$ & $k_p w_0 \leq \sqrt{a_0}$, $\boldsymbol{a_0 \gg 1}$

Beam for $k_p L \sim 1$ & $k_p R \leq 1$, $\boldsymbol{\frac{n_b}{n_e} \gg 1}$

[1] J. M. Dawson Phys. Rev. **113** (1959)
[2] J. B. Rosenzweig et al., PRA 44 (1991)
[3] A. Pukhov et al., PPCF 64 (2004)
[4] W. Lu et al., PRL 96 (2006)

Laser pulses and particle beams are affected by the plasma

Laser driver

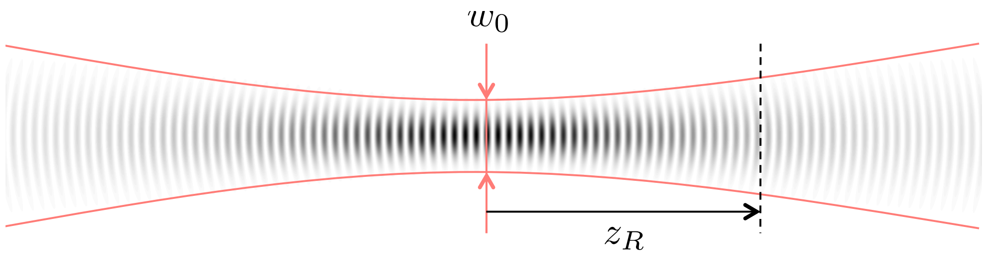

In vacuum

- $v_\varphi \sim c, v_g \sim c$
- Diffraction over the Rayleigh length $z_R = \frac{\pi w_0^2}{\lambda}$

In a plasma: refractive index $\eta_r = 1 - \frac{\omega^2}{2\gamma}$

- $v_{wake} = v_g = \left(1 - \frac{\omega_p^2}{2\omega^2}\right) c < c \rightarrow$
- Relativistic self-focusing $P_c = 2$

Numerous effect are hard to evaluate analytically

→ Quantitative predictions require numerical simulations

Driver and witness particle beams

- Decelerating driver & accelerating witness → Transformer ratio
- Betatron oscillations $\omega_\beta = \frac{\omega_p}{\sqrt{2\gamma}}$
- Trapping of a witness beam

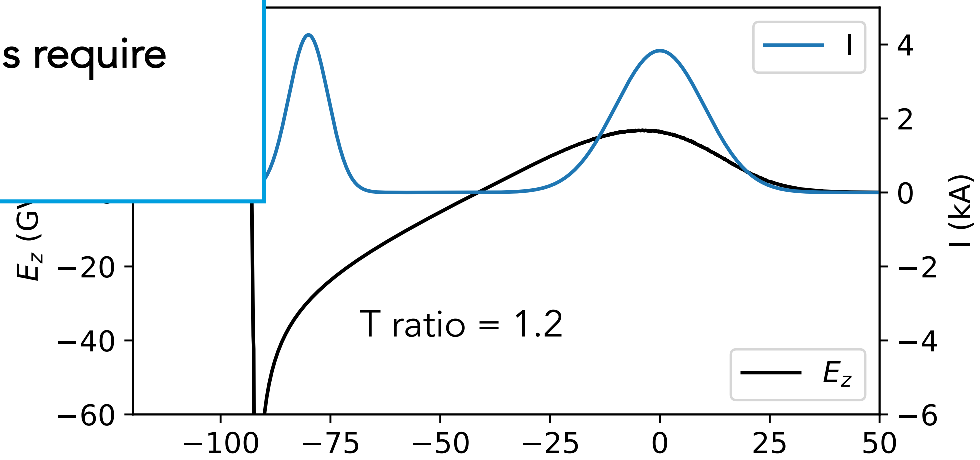

Chen, PRL (1986)

Particle-in-Cell: self-consistent plasma & fields description

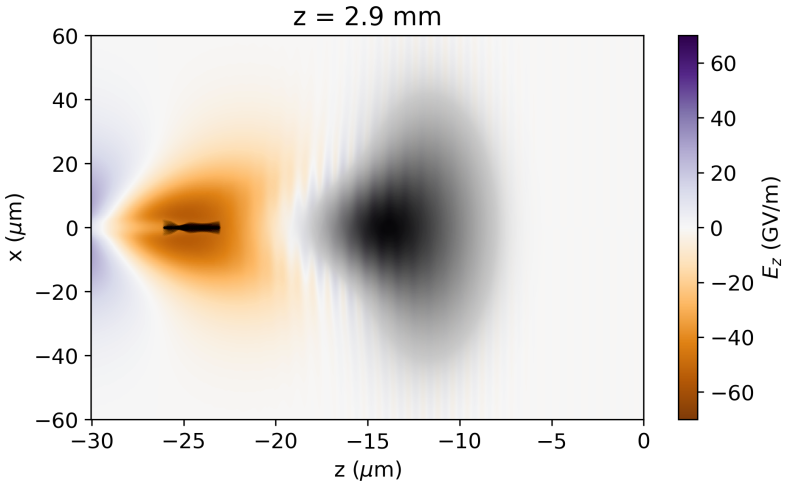

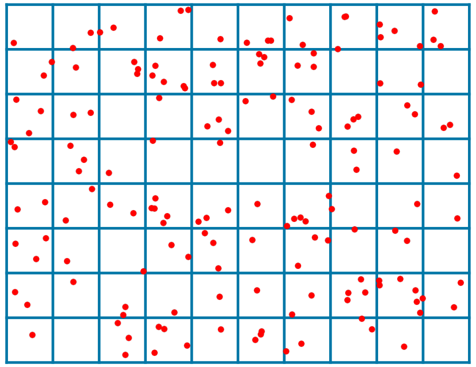

Regular mesh
Macroparticles

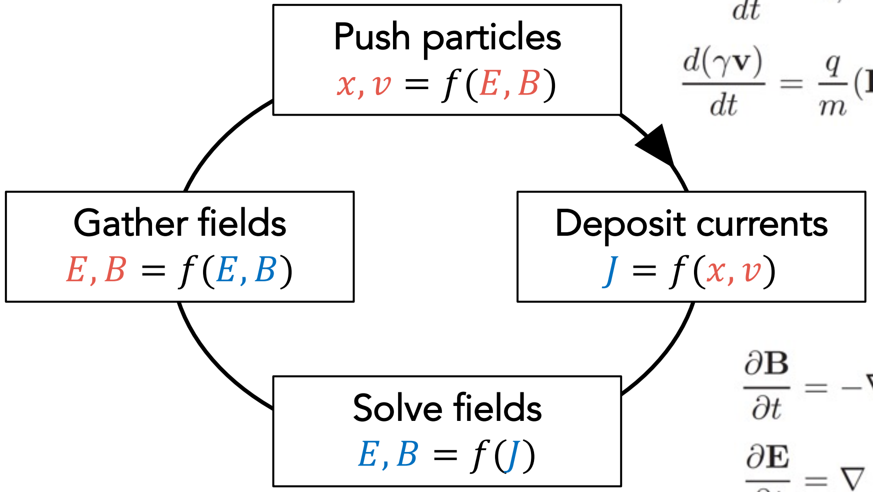

$$\frac{d\mathbf{x}}{dt} = \mathbf{v},$$

$$\frac{d(\gamma\mathbf{v})}{dt} = \frac{q}{m}(\mathbf{E} + \mathbf{v} \times \mathbf{B}),$$

$$\frac{\partial \mathbf{B}}{\partial t} = -\nabla \times \mathbf{E},$$

$$\frac{\partial \mathbf{E}}{\partial t} = \nabla \times \mathbf{B} - \mathbf{J},$$

Lagrangian description of plasma
Eulerian description of fields

→ 3D simulations of plasma acceleration are very expensive

Several methods & approximations are well-suited to PIC

➢ Reduced geometry: Azimuthal decomposition or RZ

- CALDER-circ, OSIRIS, QPAD, FBPIC, WarpX

→ Quasi-cylindrical problems

Lifschitz, A. F., et al. *JCP* 228.5 (2009)

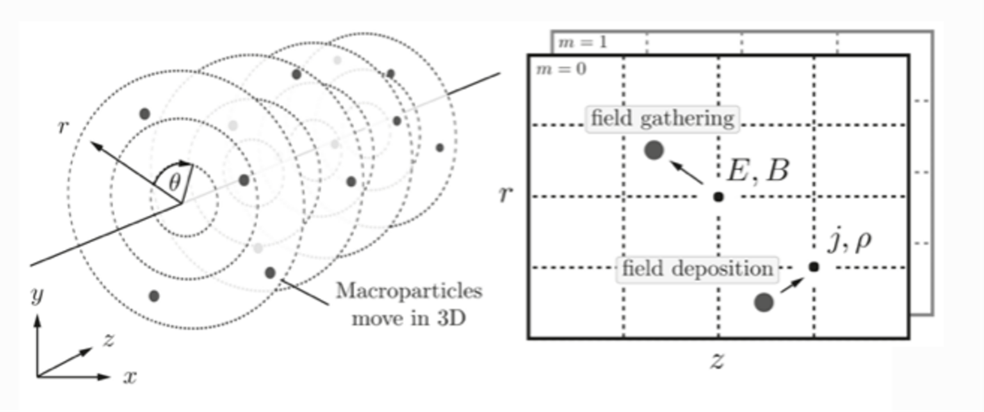

- **The fields are decomposed into azimuthal modes**

$$F(r, z, \theta) = Re\left[\sum_{m=0}^{N_m - 1} \hat{F}_m(r, z) e^{im\theta}\right]$$

m=0: purely cylindrical mode
m=1: dipole mode
m=2: quadrupole mode

- **Each azimuthal mode is represented by a 2D r-z grid**

https://fbpic.github.io/overview/pic_algorithm.html #cylindrical-grid-with-azimuthal-decomposition

Several methods & approximations are well-suited to PIC

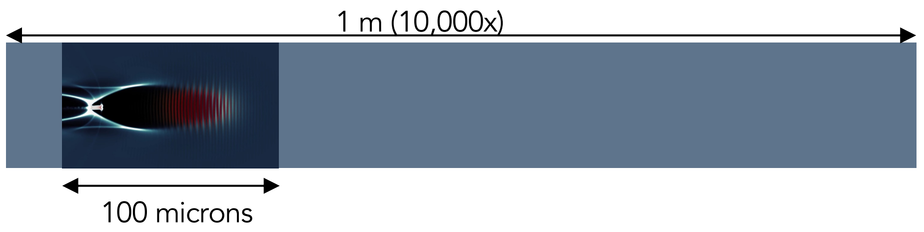

Problem: the CFL condition limits the time step to $\Delta t < c\Delta z$

- **Boosted frame method** J.-L. Vay PRL **98,** 130405 (2007)

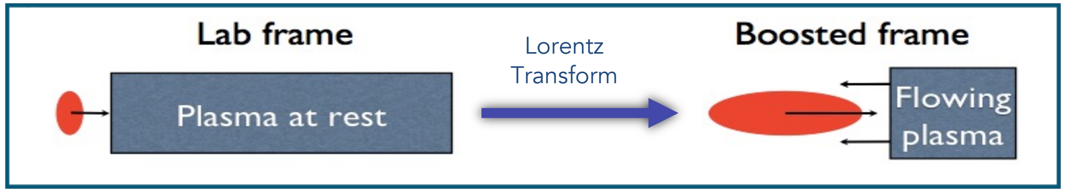

- Reduces number of time step by orders of magnitude
- Prone to Numerical Cherenkov Instability (NCI)
- Methods exist to mitigate NCI (PSATD + Galilean transform [1,2], RIP [3])

[1] R. Lehe et al., PRE 94 (2016)
[2] M. Kirchen et al., Phys. Plasmas 23 (2016)
[3] A. Pukhov *JCP* 418 (2020)

- **Reduced model: quasi-static PIC (QS-PIC)**

- Beam & wake: $\boldsymbol{v} \sim c\boldsymbol{e_z}$
- Quasi-static approximation
- From the driver at a given time, calculate the plasma response → no history

→ No CFL condition, large time step for the beam
→ Cannot capture injection

3D QS-PIC simulations rely on 2D PIC loop in ζ from head to tail

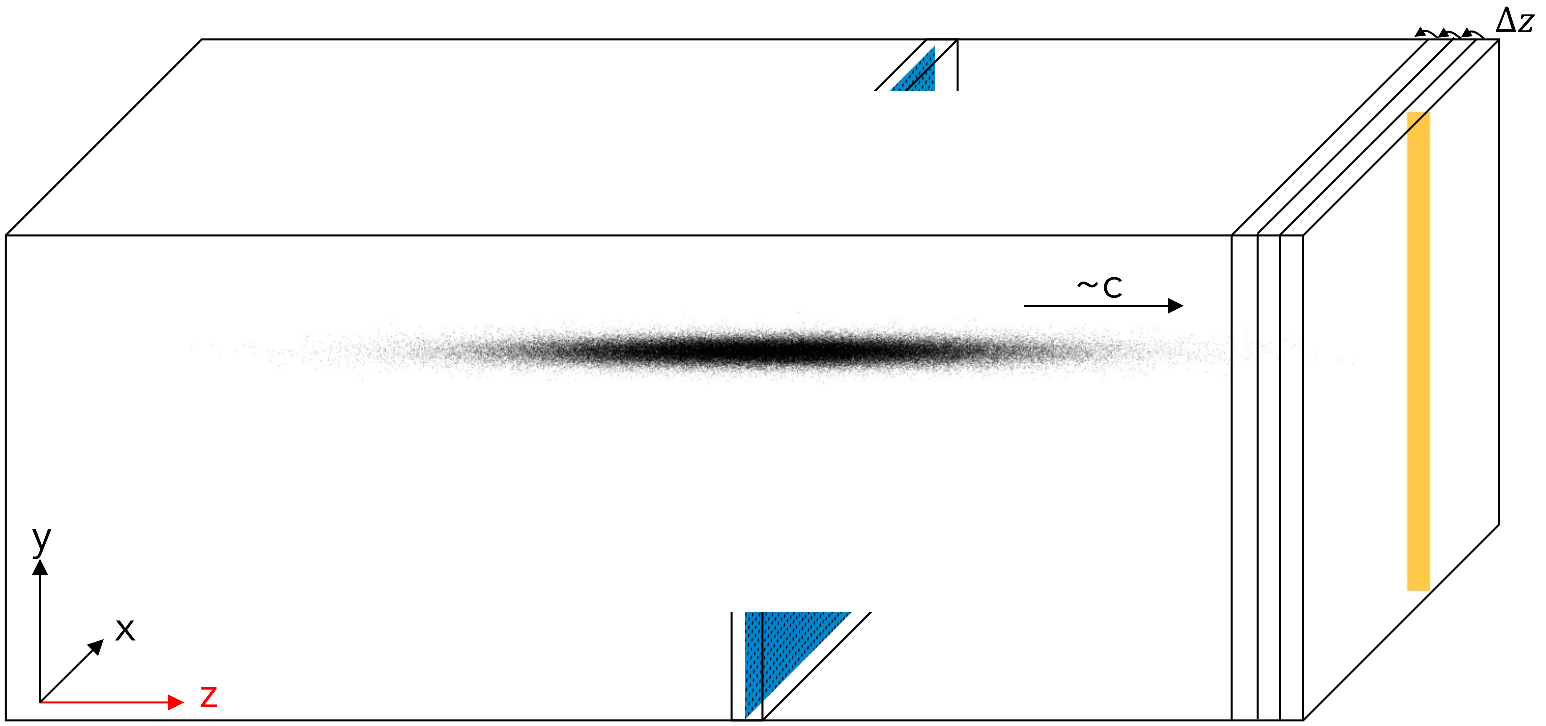

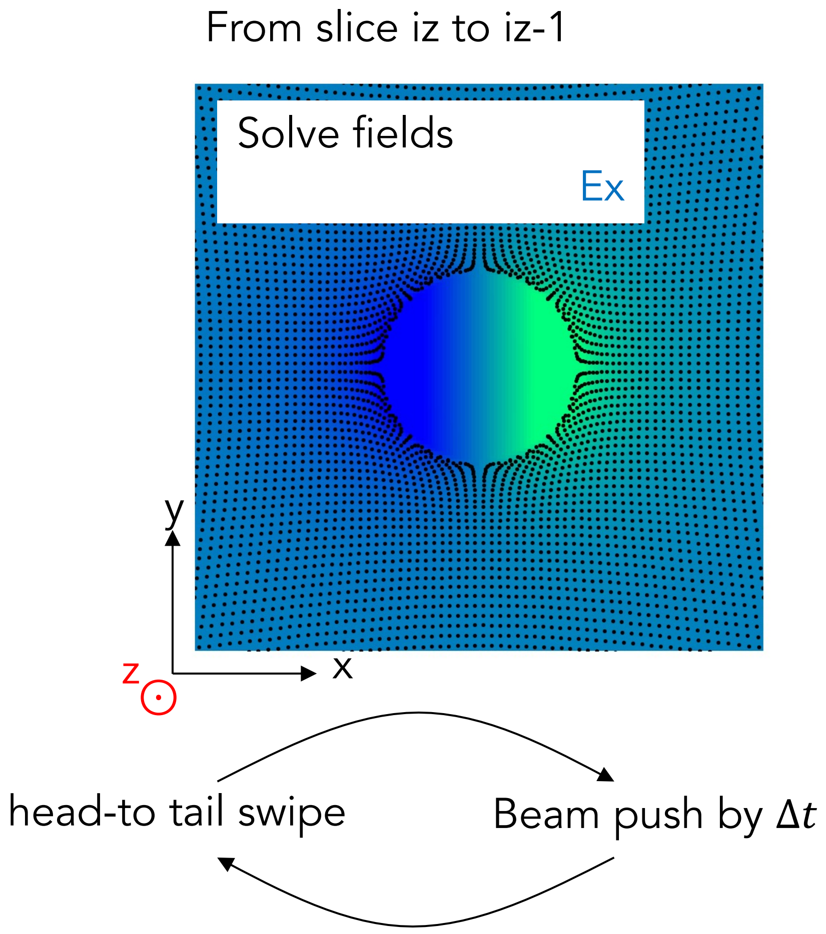

- Only 1 slice of fields and plasma particles
- Head-to-tail swipe to calculate fields on each slice (N_z)
- For each slice, 2D PIC iterations (in z) with Poisson solves
- Can be parallelized transversally and longitudinally

Comparison EM-PIC & QS-PIC

	EM-PIC	QS-PIC
Algorithm	Fields and particles advanced in time $\partial_t f = \cdots$	Beam particles advanced in time Plasma particles advanced in space ζ Fields: 2D Poisson equation $\Delta_\perp f = s$
Data	$n_x n_y n_z (10 + ppc * 10)$	$n_x n_y \ (10 + ppc * 10)$
Operations	$\boldsymbol{n_t}$ **PIC iterations** $\boldsymbol{n_x n_y n_z}$	$\boldsymbol{n_t n_z}$ **PIC iterations** $\boldsymbol{n_x n_y}$
Existing codes	CALDER, EPOCH, FBPIC, OSIRIS, PIConGPU, Smilei, Vsim, WarpX, …	HiPACE++, INF&RNO, LCODE, QuickPIC, QPAD, WAND-PIC, …
Advantages and limitations	General but expensive	Large time step but injection not always captured

- Languages (Python, C++, C, Fortran, etc.)
- Capabilities (boosted frame) & methods
- Physics (collisions, ionization, QED)
- Geometries (1/2/3D, quasi-cylindrical)
- Open-source, supported platforms

I. Physics of plasma acceleration

II. The particle-in-cell (PIC) method

III. High-performance computing (HPC)

IV. Recent activities at DESY

HELMHOLTZ RESEARCH FOR GRAND CHALLENGES

Supercomputers accelerate parallel applications

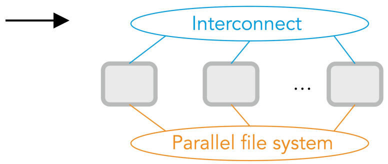

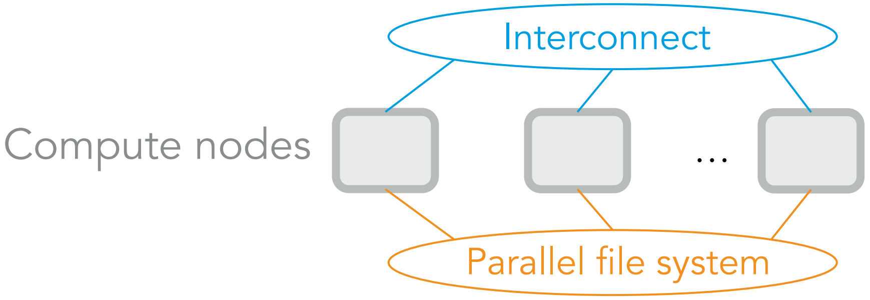

In practice for PIC

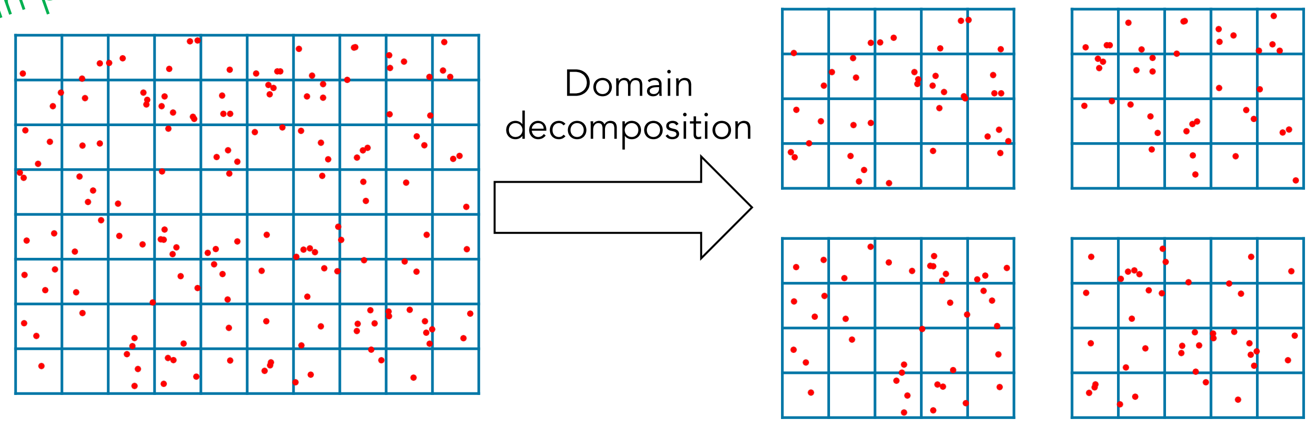

WarpX

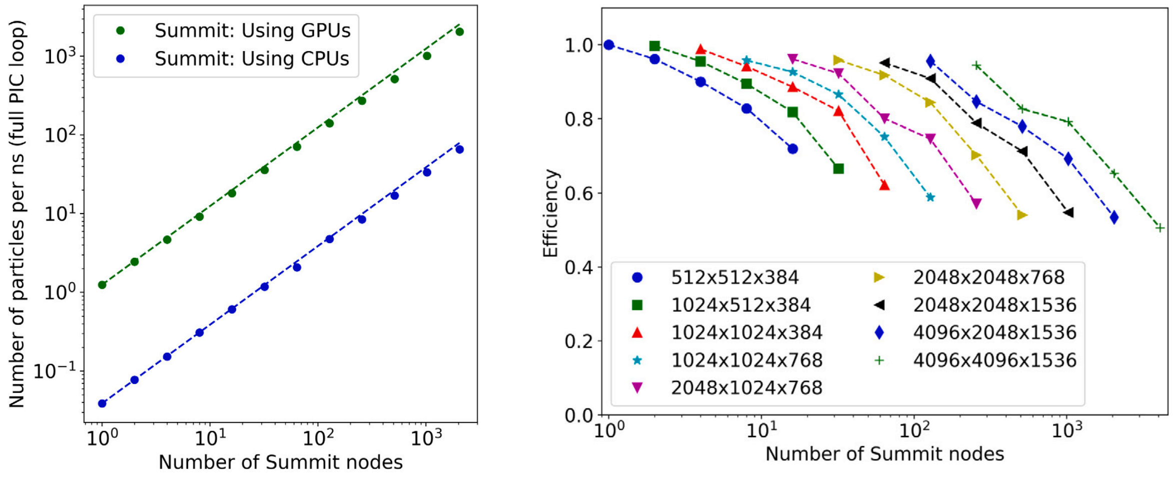

A. T. Myers et al., Parall. Comp. 108 (2021)

Data must be exchanged at each iteration

- Fields: 1 (or more) boundary "*ghost*" cells
- Particles: changing sub-domain

→ **Halo exchange**

Strong scaling (total problem size = cst)
Weak scaling (problem size per node = cst)

DESY.

CPU & GPU: different approaches to HPC

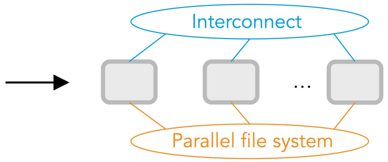

CPU: central processing unit
GPU: graphics processing unit (GPGPU)

Rank	Machine	Node architecture
1 (26)	Fugaku (Japan)	CPU (Arm)
2 (28)	Summit (USA)	CPU + GPU
3 (32)	Sierra (USA)	CPU + GPU
4 (53)	Sunway TaihuLight (China)	CPU + GPU
5 (7)	Perlmutter (USA)	CPU + GPU
6 (16)	Selene (USA)	CPU + GPU
7 (111)	Tianhe-2A (China)	CPU
8 (11)	JUWELS Booster (Germany)	CPU + GPU
9 (25)	HPC5 (Italy)	CPU + GPU
10 (180)	Voyager-EUS2 (USA)	CPU + GPU

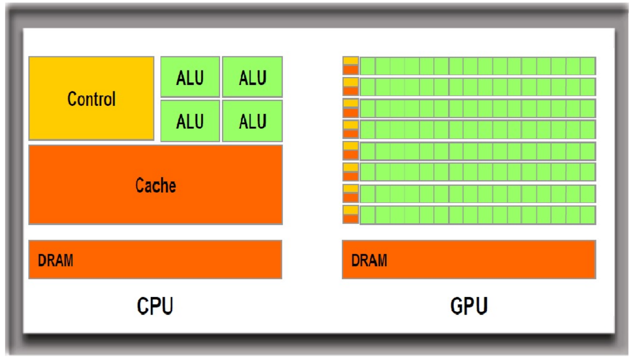

Data movement dominates computation time

Time of data migration = $latency + \frac{message\ size}{bandwidth}$

CPU: low latency → high-speed serial processors

GPU: high bandwidth → high-throughput parallel processors

CPU: 10s of very fast & independent cores
GPU: 1000s of slow cores doing the same operation

https://www.nextplatform.com/2019/07/10/a-decade-of-accelerated-computing-augurs-well-for-gpus/
Kirk, David B., and W. Hwu Wen-Mei. Morgan kaufmann, 2016.
www.top500.org

A portability layer helps support multiple architectures

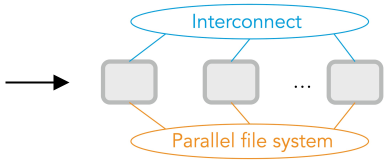

- **Performance-portability**
  - In particular GPU computing
  - Portability layer (Kokkos, Alpaka, RAJA) C++
- **Open Source & Open Repository**
  - Software can be freely used, modified and shared
  - Encourages flexible, modular code
  - Favor good dependency graph rather than duplication

```
for(int i=0; i<N; i++){
    xp[i] += 1.;
}
```

```
CUDA (NVIDIA)
kernel(int* xp) {
    int i = blockIdx.x *
            blockDim.x
            + threadIdx.x;
    if (i<N) xp[i] += 1; }
kernel<<<N, 256>>>(xp);
```

```
Kokkos::ParallelFor( N,
    [=] (int i) {
        xp[i] += 1.;
    }
);
```

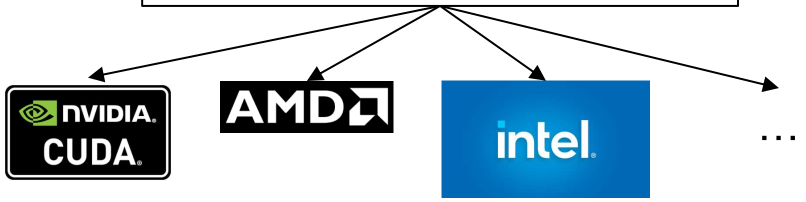

openPMD for I/O in PIC simulations

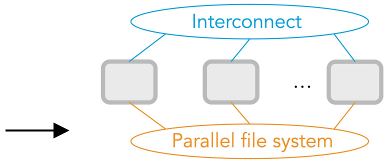

Your favorite code

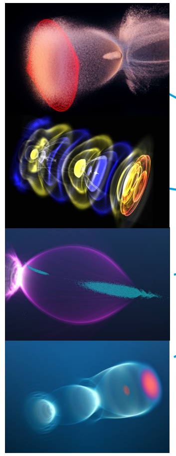

- ✓ Archive (FAIR) & share
- ✓ Analyze & plot (openPMD-viewer, VisualPIC)
- ✓ Interface with other codes
  Beam optics, ICS, FEL, ML → openPMD → …
- ✓ In-situ visualization?

Standard I/O format for particle and mesh data

Pioneered at HZDR, contributors worldwide

standard	https://github.com/openPMD
API	https://github.com/openPMD/openPMD-api
viewer	https://github.com/openPMD/openPMD-viewer

openPMD: high-quality standard for particle & mesh data
- reliable tool for start-to-end simulations
- adopt FAIR principles for longevity
- encourage benchmarks and collaboration in a (reasonably) user-friendly way.
- Good adoption in PIC community
- Wraps around performant file formats (HDF5, ADIOS2)

DESY actively contributes to multiple plasma acceleration codes

And uses more

	Quasistatic	Electromagnetic
Quasi-cylindrical	Wake-T (DESY) → *Conceptual designs (sec-min)* MPA1 Open-source openPMD https://github.com/AngelFP/Wake-T	FBPIC (LBNL + UHH + …) → *LPA RZ with injection* LUX (M. Kirchen, S. Jalas) Open-source GPU openPMD https://github.com/fbpic/fbpic
3D	HiPACE++ (DESY + LBNL) → *3D without injection* MPA1 Open-source GPU openPMD https://github.com/Hi-PACE/hipace	WarpX (LBNL + …) → *3D with injection, QED* Minor contribution Open-source GPU openPMD https://github.com/ECP-WarpX/WarpX or PIConGPU, OSIRIS, …

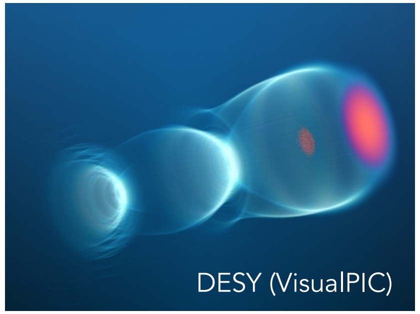

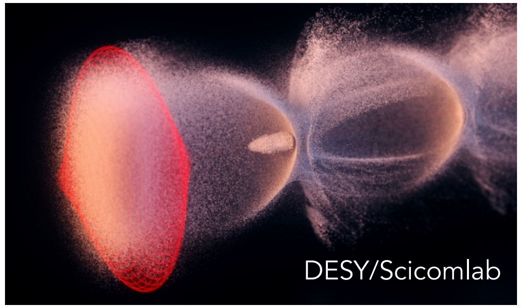

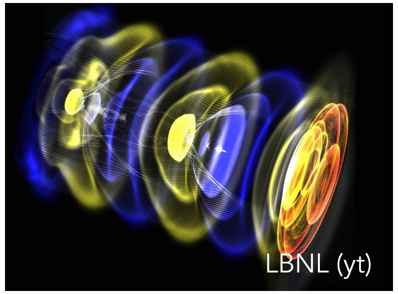

DESY actively contributes to multiple plasma acceleration codes

And uses more

Quasistatic | Electromagnetic

Quasi-cylindrical

Wake-T (DESY)

→ *Conce*

MPA1

Open-sou

https://gith

FBPIC (LBNL + UHH + …)

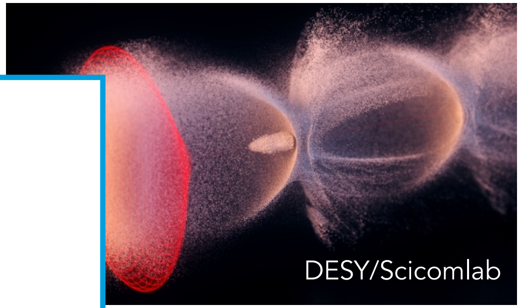

3D

HiPACE+

→ *3D wit*

MPA1

Open-sou

https://github.com/Hi-PACE/hipace

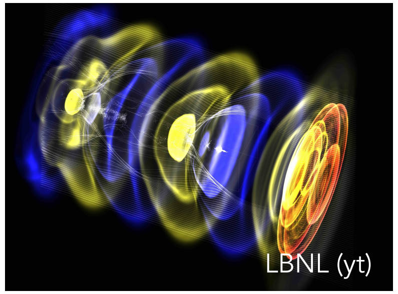

In collaboration with the WarpX team, LBNL

https://github.com/ECP-WarpX/WarpX

or PIConGPU, OSIRIS, …

Beam-driven wakefield acceleration

1 GeV, 5 μm width, 20 μm emittance, 1 nC, 20 μm long, 10^{17} cm^{-3} plasma, $n_b/n_0 = 8$

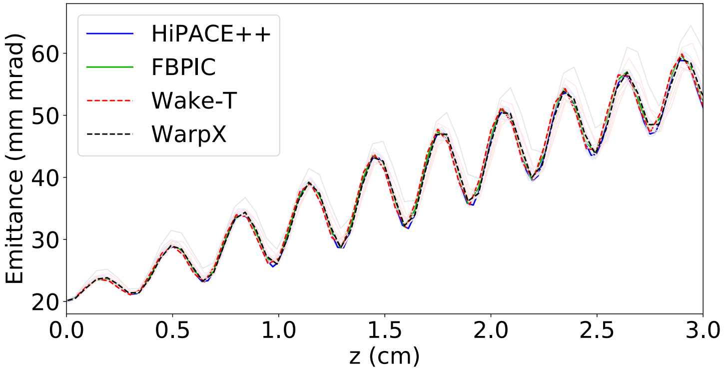

HiPACE++ 1/3 – a quasi-static PIC on GPU (and CPU)

- Collaboration with the ECP WarpX team
- C++, full re-writing of HiPACE (DESY, LBNL)
- GPU porting of the FULL PIC loop for orders-of-magnitude speedup
- Built on top of AMReX & openPMD
  - Data structures and communications
  - Performance-portability (`ParallelFor`)
  - 10000 LOC, 2000 comments
  - Proper HiPACE++ code: 18% compilation time
- HPC programming standards
  - Documented, open-source, open-repository, CMake
  - Continuous Integration
  - Two unit systems (normalized, SI)
  - Single or double precision

→ *Beam-driven used in production, laser-driven work-in-progress*

HiPACE++

3D quasi-static PIC code

Lead dev/PI(s):

DESY + LBNL

https://github.com/Hi-PACE/hipace

Language: C++

Doc: https://hipace.readthedocs.io

2021, just starting

AMD

HiPACE++: a portable, 3D quasi-static Particle-in-Cell code

S. Diederichs,[1,2,3,*] C. Benedetti,[2] A. Huebl,[2] R. Lehe,[2]
A. Myers,[2] A. Sinn,[1] J.-L. Vay,[2] W. Zhang,[2] and M. Thévenet[1]

[1] *Deutsches Elektronen-Synchrotron DESY, Notkestraße 85, 22607 Hamburg, Germany*
[2] *Lawrence Berkeley National Laboratory, 1 Cyclotron Rd, Berkeley, California 94720, USA*
[3] *University of Hamburg, Institute of Experimental Physics, Luruper Chaussee 149, 22607 Hamburg, Germany*
(Dated: September 22, 2021)

arXiv

HiPACE++ 2/3 – the GPU porting strategy exploits fast single-GPU

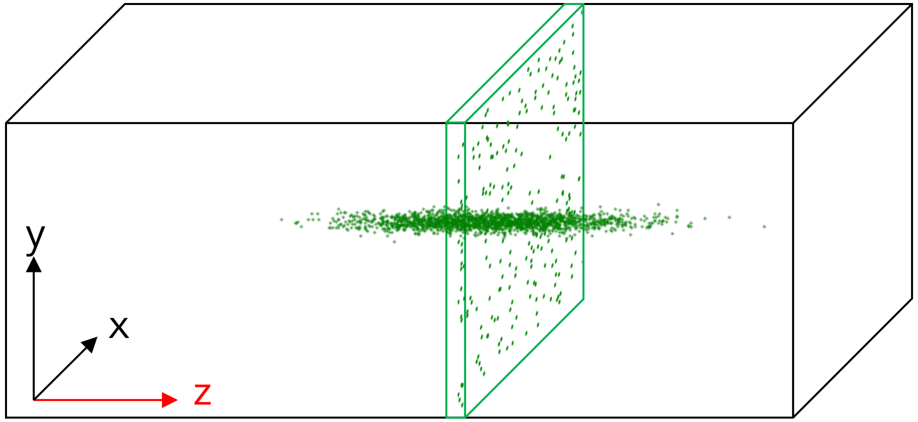

40 CPU cores +
10,000 GPU cores

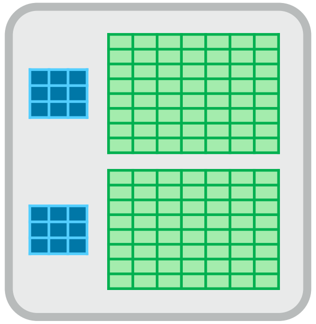

Poster:

- Low memory required for computing (on the device)
  - $N_x \times N_y \times N_{fields}$ grid cells (E, B, J, ρ)
  - $N_x \times N_y \times N_{ppc}$ plasma macro-particles
  - A few millions beam macro-particles

  → All compute data on the device

- Many 2D Poisson solves (FFT) & Helmholtz solves (MG)

  $N_x = N_y = 2048$ is a large problem

  $N_z \times N_{steps} \times 10 = 10^7$ Poisson solves

  → single-GPU cuFFT vs. FFTW

- All other operations (current deposition etc.) work well on GPU

  As demonstrated by WarpX, PIConGPU, etc.

HiPACE++ 3/3 – benchmarks & performance

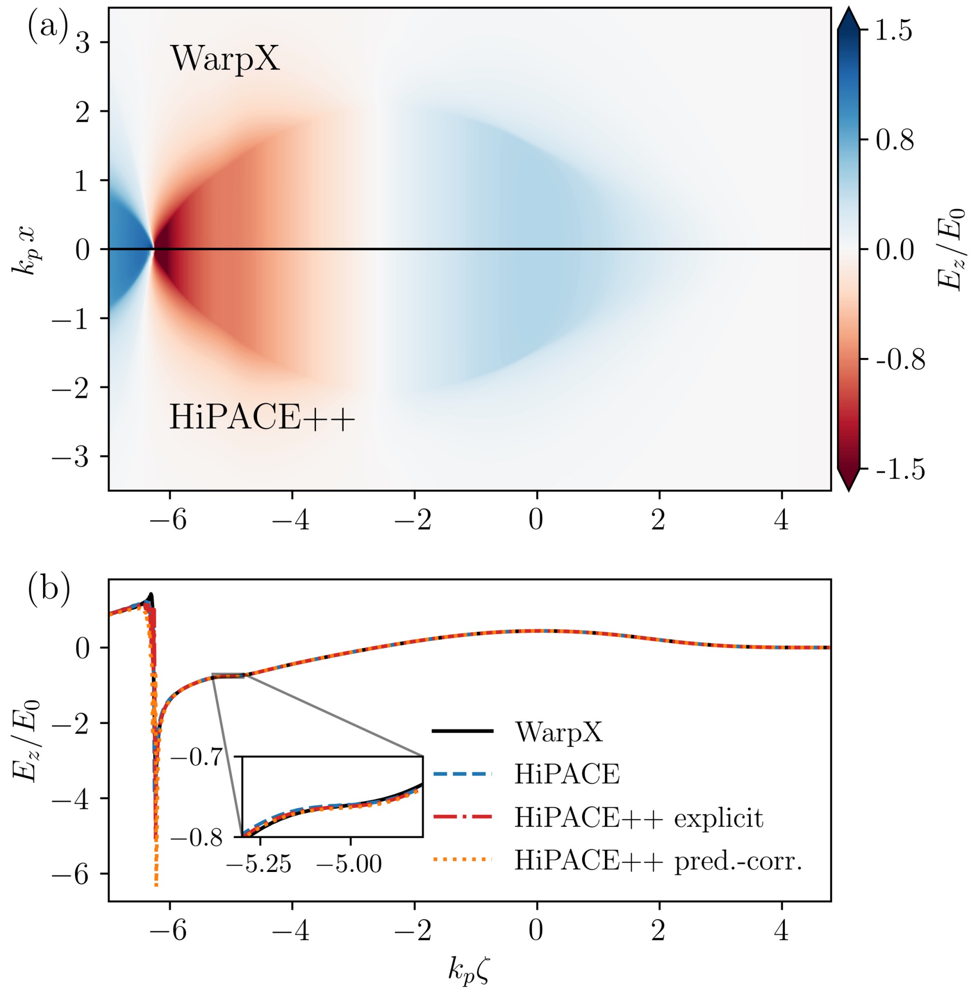

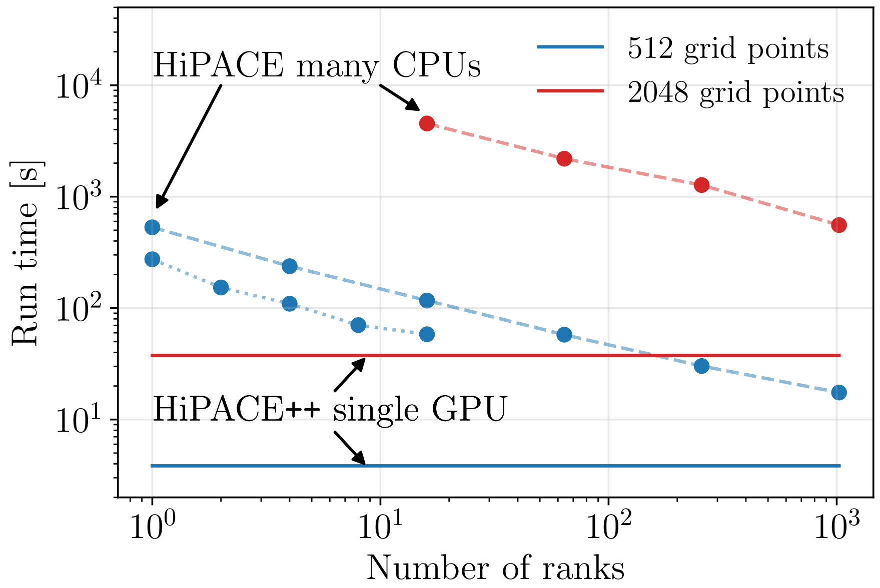

➢ Excellent scaling to hundreds of GPUs

HiPACE++ 3/3 – benchmarks & performance

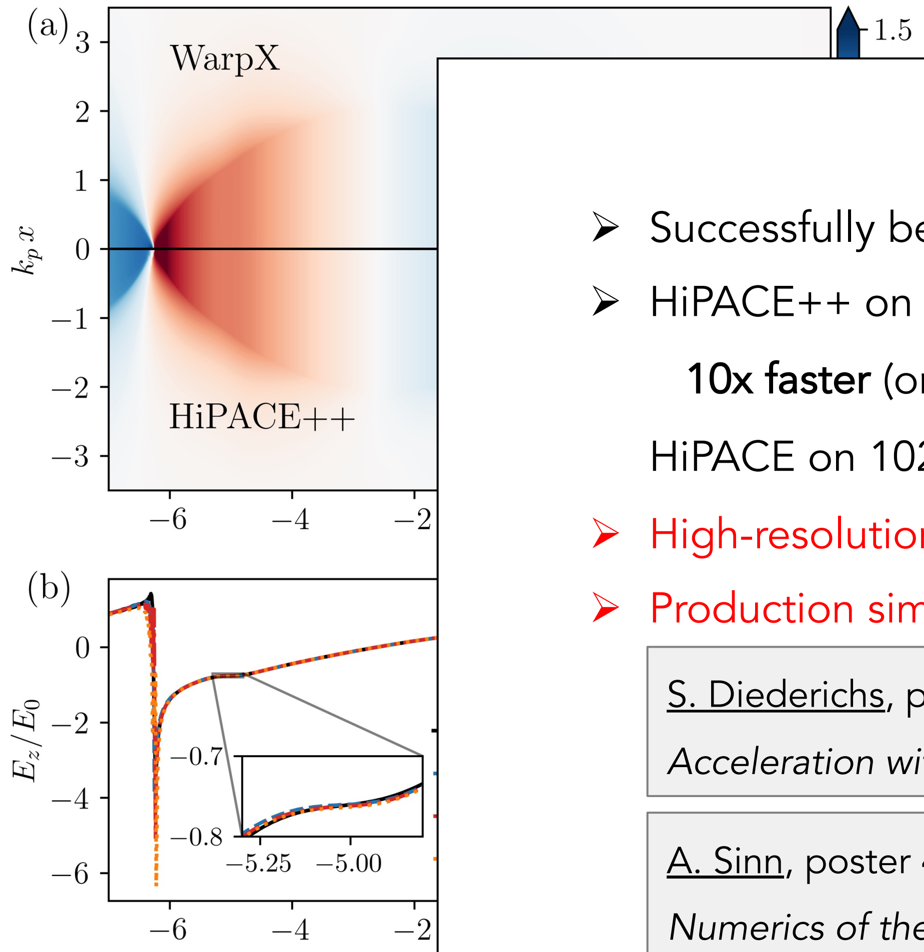

- Successfully benchmarked
- HiPACE++ on 1 GPU (<u>1/4 node</u>)

  10x faster (on the JUWELS Booster) than

  HiPACE on 1024 CPU codes (<u>22 nodes</u>)
- High-resolution simulations within minutes
- Production simulations from laptop `float` to HPC

<u>S. Diederichs</u>, poster 12, *Modelling Positron Acceleration with HiPACE++*

<u>A. Sinn</u>, poster 47, *Improving Performance and Numerics of the Quasi-static PIC Code HiPACE++*

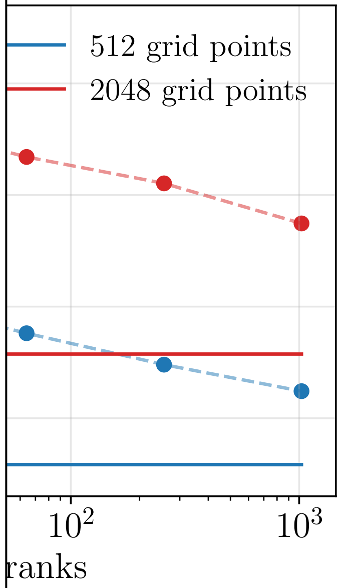

- Excellent scaling to hundreds of GPUs

Conceptual design study: a plasma injector for PETRA IV (PIP4)

The team

I. Agapov, S. Antipov, R. Brinkmann, A. Ferran Pousa, S. Jalas, L. Jeppe, M. Kirchen, W. P. Leemans, A. R. Maier, A. Martinez de la Ossa, J. Osterhoff, M. Thévenet

Petra IV [1] is the upgrade of the Petra III storage ring for synchrotron radiation (2.3 km, 6 GeV), proposing orders-of-magnitude increase in X-ray brightness.

Specs: 6 GeV, > 1 nC/s, 1% energy spread

- LPA based on the LUX design [2]
- 500 MeV prototype [3] & 6 GeV injector
- Novel energy compression concepts required [4]

→ CDR in 2022 (S2E simulations), commissioning in the decade

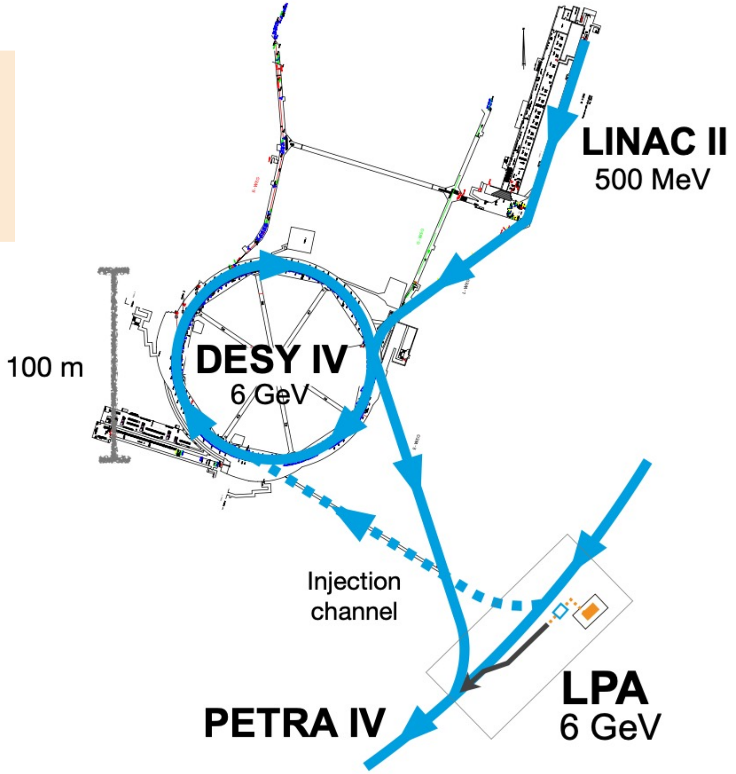

Can the whole injector be replaced by a LPA?

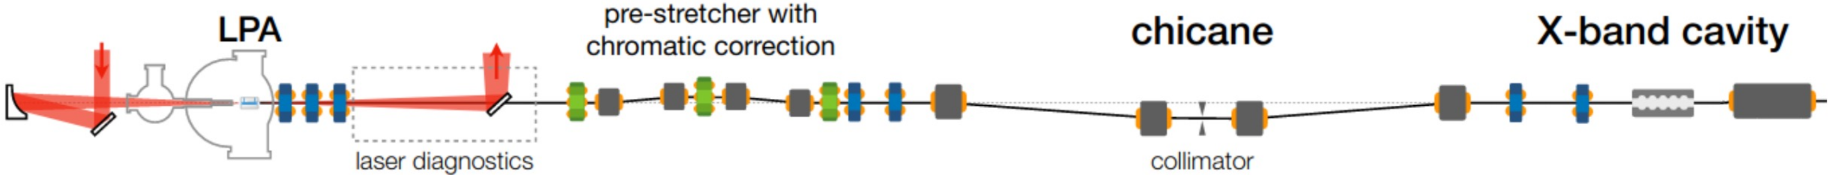

[1] https://www.desy.de/research/facilities__projects/petra_iv
[2] Kirchen, Manuel, *et al. PRL* 126.17 (2021); LUX PI: A. R. Maier
[3] Antipov, S. A., et al. *arXiv preprint arXiv:2106.07367* (2021).
[4]A. Ferran Pousa, *et* al. *arXiv preprint arXiv:2106.04177* (2021).

Other activities in theory and simulations at DESY

Positron acceleration (S. Diederichs)

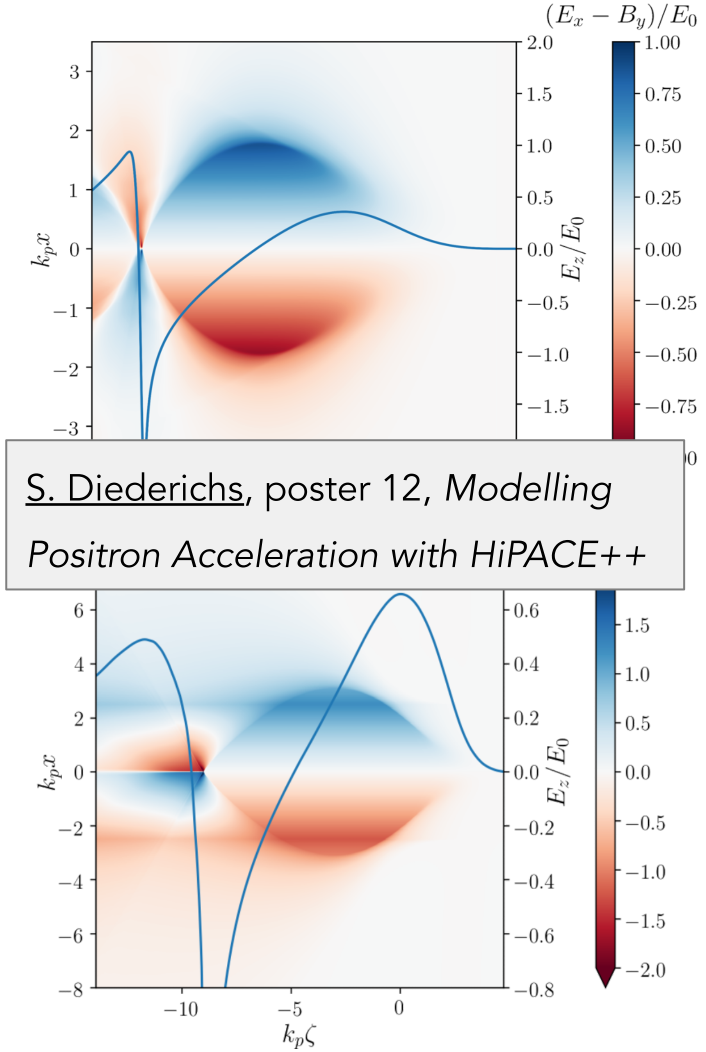

S. Diederichs, poster 12, *Modelling Positron Acceleration with HiPACE++*

Plasma hydrodynamic simulations (M. Mewes)

Capillary discharge

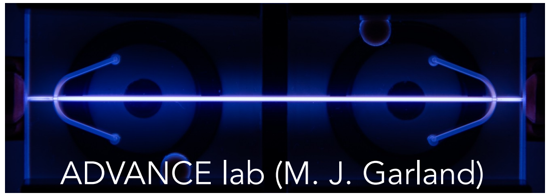
ADVANCE lab (M. J. Garland)

Hydrodynamic Optical Field Ionization (HOFI)

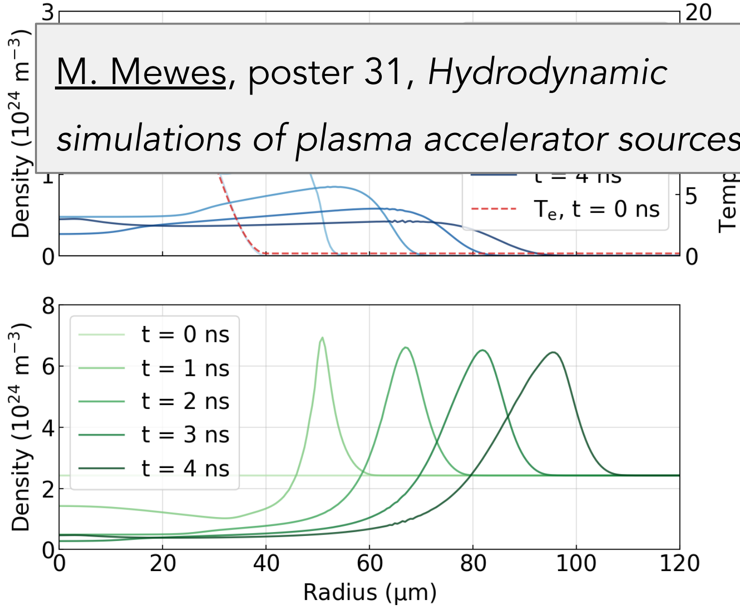

M. Mewes, poster 31, *Hydrodynamic simulations of plasma accelerator sources*

Conclusion

- Basic understanding of plasma acceleration
- Numerical simulations provide invaluable support to the development of plasma acceleration
- PIC provides modelling
- HPC efforts benefit from modern code practices (open-source, dependencies, portability layers)

Our efforts

- Improve simulation capabilities
- Address domain challenges
  - Positron acceleration
  - Plasma Injector for Petra IV (PIP4)
  - Reduce energy spread (dechirper)
  - Investigate plasma sources

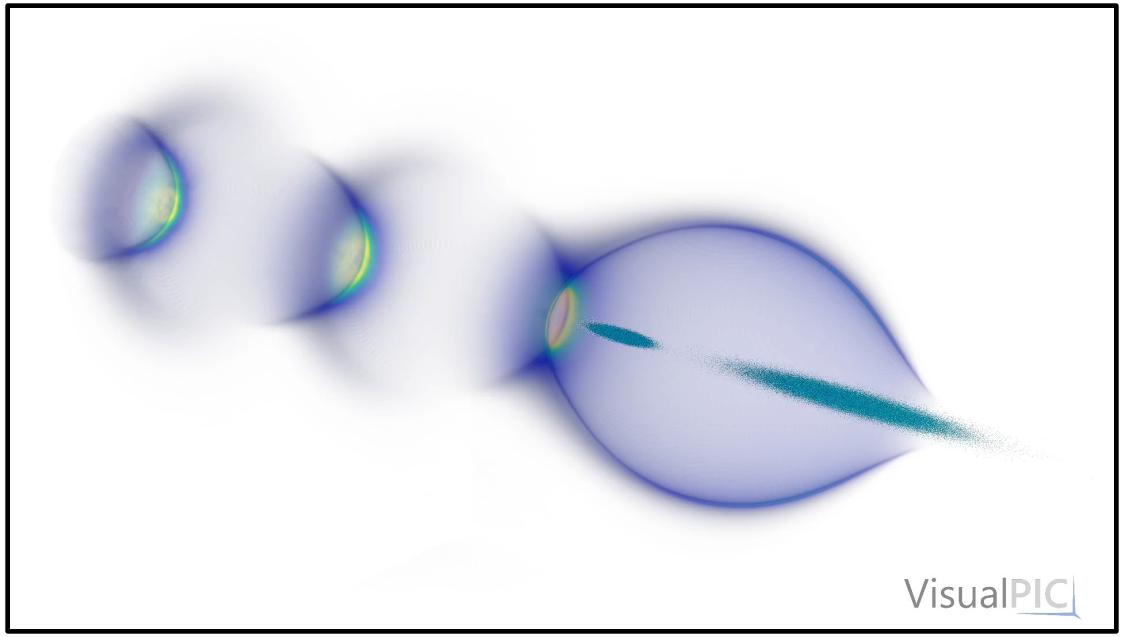

Thank you for your attention

Contributions from:
DESY Severin Dieredichs, Angel Ferran Pousa, Alberto Martinez de la Ossa, Mathis Mewes, Alexander Sinn; LUX team
James Cook University (AU) Gregory Boyle
LBNL The WarpX team (PI: Jean-Luc Vay)
All HiPACE++ contributors

DESY.