HiPACE++
Examples and features

Severin Diederichs

HiPACE++ workshop, 11.07.2023

Plasma-based positron acceleration is notoriously challenging

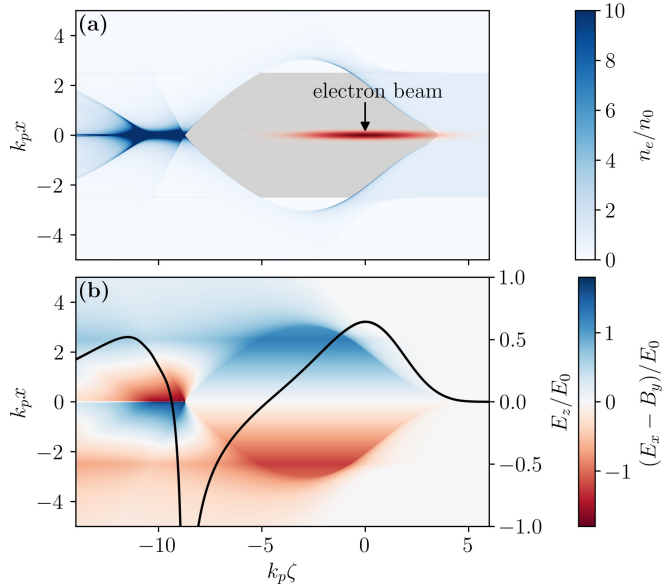

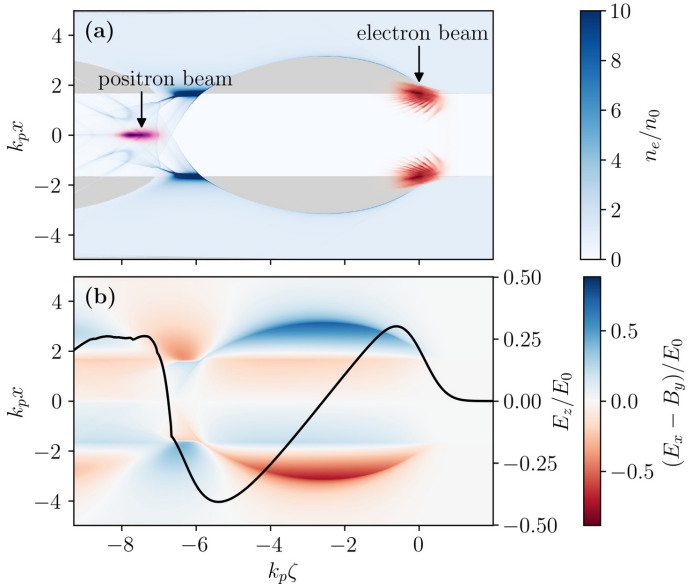

- Often rely on narrow, sharp electron filaments that must be resolved → **high transverse resolution** needed
- Only relevant for collider-like parameters (small beams)
- Require a temperature for convergence → **many plasma particles per cells** are needed

Plasma-based positron acceleration is notoriously challenging

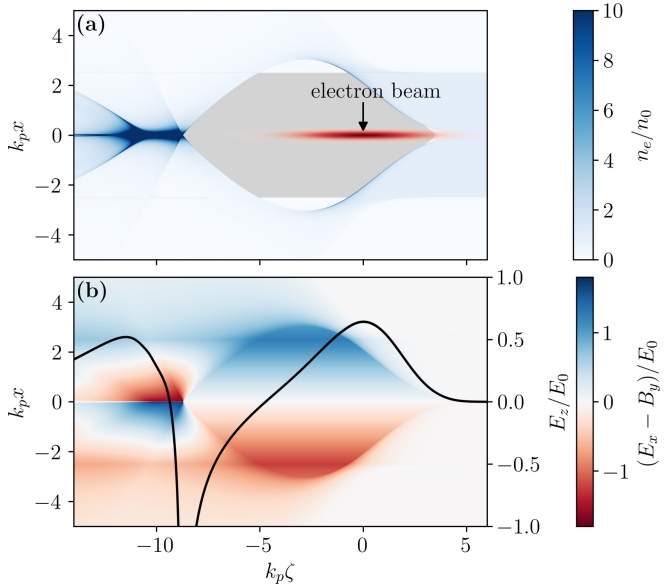

Back in 2019:
Using the predecessor HiPACE
(CPU only)

```
grid points = 2048 2048 10000
time steps = 400
beam particles = 100e6
plasma particles per cell = 25
```

256 nodes, 12000 CPU cores
run time: 1-2 days
costs: 10 000 node hours (CPU)

Simulation costs were prohibitively high

Plasma-based positron acceleration is notoriously challenging

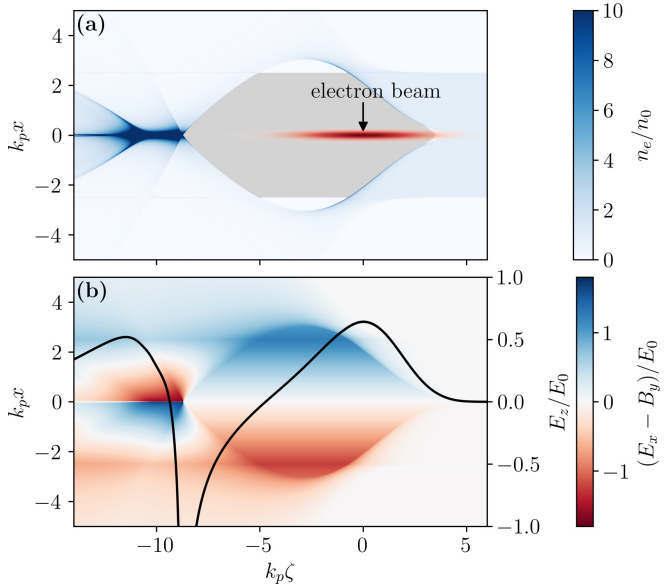

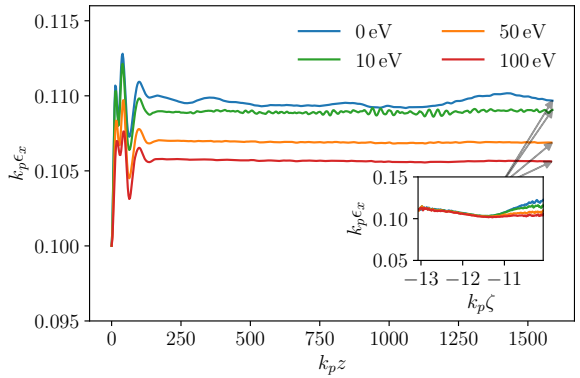

Back in 2019:
Using the predecessor HiPACE
(CPU only)

```
grid points = 2048 2048 10000
time steps = 400
beam particles = 100e6
plasma particles per cell = 25
```

256 nodes, 12000 CPU cores
run time: 1-2 days
costs: 10 000 node hours (CPU)

Now:
Using HiPACE++
(on GPU)

Explicit field solver saves 5x

Simulations are cheap,
can do parameter scans

```
grid points = 4095 4095 1920
time steps = 600
beam particles = 175e6
plasma particles per cell = 49
```

32 nodes, 128x A100 GPUs

run time: 19min
costs: 10 node hours (GPU)*

*Our allocation: 30 000 node hours
"large allocation": 156 000 node hours

Moderate runs even possible on laptops

Example:

e^+ acceleration with a hollow-core driver Jain et al., PRL 2015

2015: 3D simulations too expensive

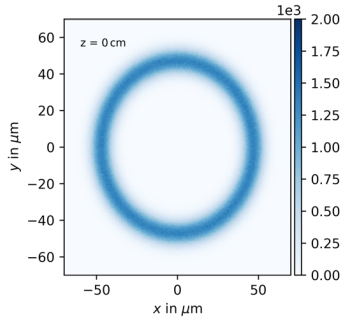

Donut-shaped beam pushes plasma into center

→ allows for positron acceleration

Moderate runs even possible on laptops

Example:
e^+ acceleration with a hollow-core driver Jain et al., PRL 2015

2015: 3D simulations too expensive

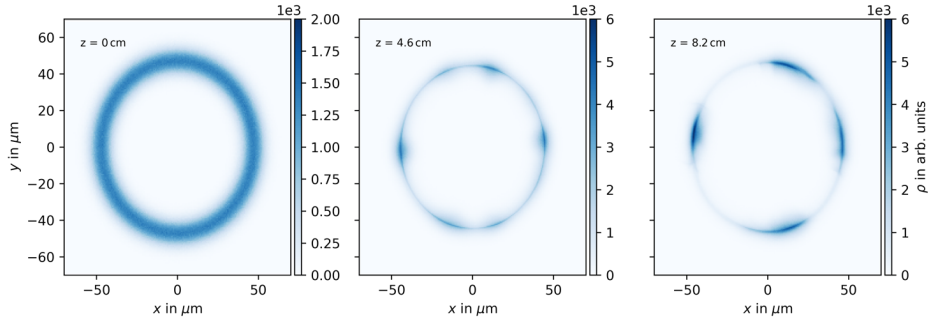

Now:
1 hour on a laptop equipped with an NVIDIA RTX2070
in *single precision*

```
grid points = 1023 1023 1024
time steps = 300
beam particles = 10e6
plasma particles per cell = 1
```

Donut-shaped beam pushes plasma into center
→ allows for positron acceleration

Automated beam loading algorithm: SALAME

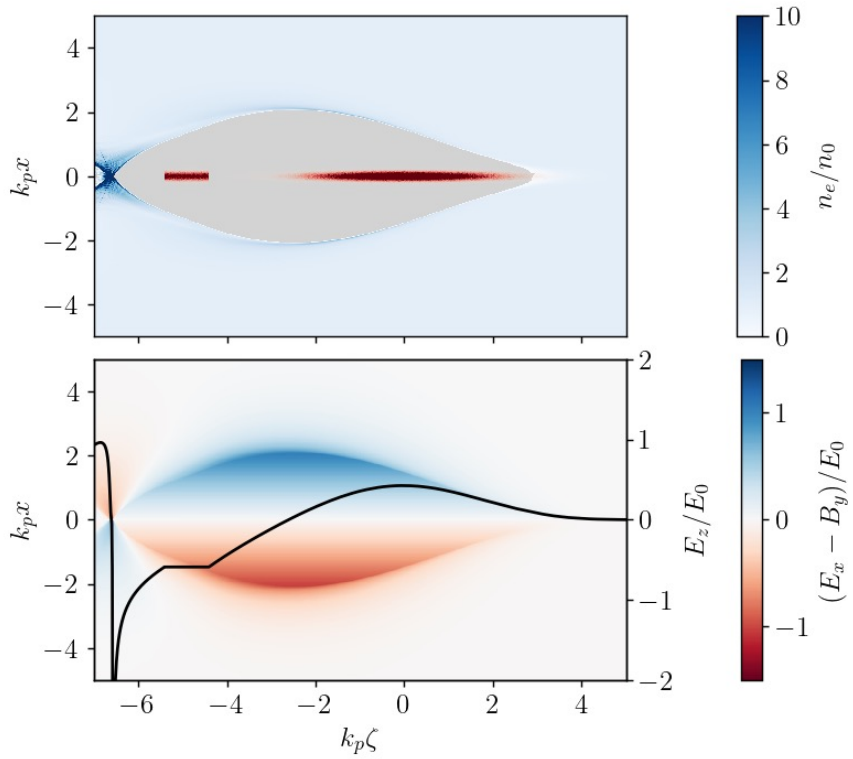

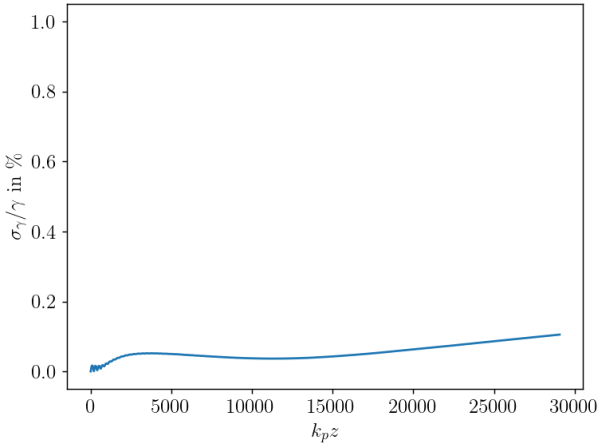

Automated beam generation to flatten an arbitrary wakefield*

```
beam.do_salame = 1
```

```
grid points = 1023 1023 1024
time steps = 1000
beam particles = 20e6
plasma particles per cell = 1
```

Run time: 97 s
Costs: 0.4 node hours using on 16 nodes on Perlmutter (64 A100 GPUs)

*more information on the algorithm
Diederichs et al., PRAB 2020

Laser envelope solver

Implemented the INF&RNO[1] laser envelope model[2]
also implemented in Wake-T

Benchmark vs Wake-T

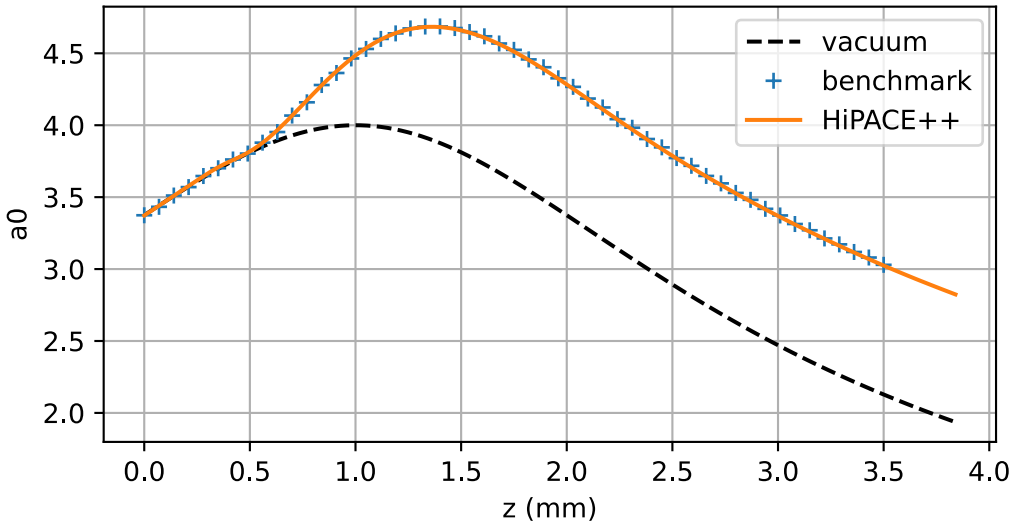

[1]Benedetti et al. AIP Conf. Proc. 2010
[2]Benedetti et al. PPCF 2018

Laser envelope solver

Implemented the INF&RNO[1] laser envelope model[2]
also implemented in Wake-T

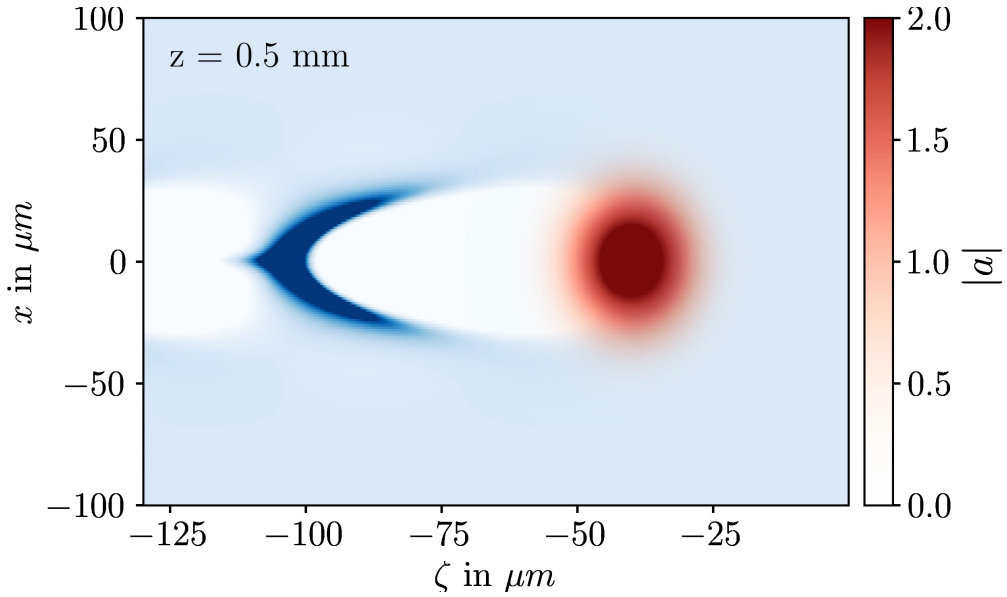

Benchmark vs Wake-T

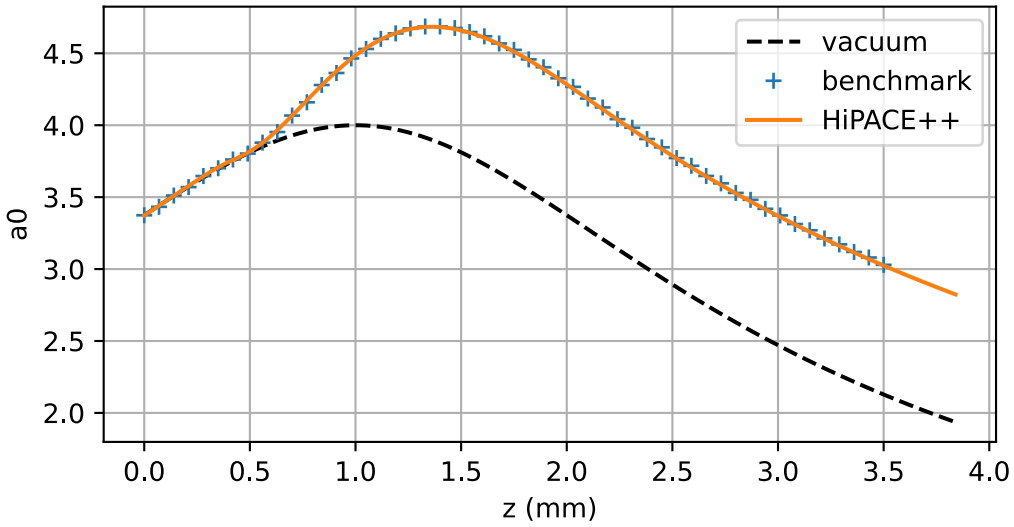

Standard LWFA setting:

```
grid points = 511 511 1000
time steps = 50
plasma particles per cell = 1
```

Run time: 75 s (laser solver 23%)
Costs: 0.02 node hours
(on 1 node on the JUWELS Booster)

[1]Benedetti et al. AIP Conf. Proc. 2010
[2]Benedetti et al. PPCF 2018

Laser envelope solver

Implemented the INF&RNO laser envelope model

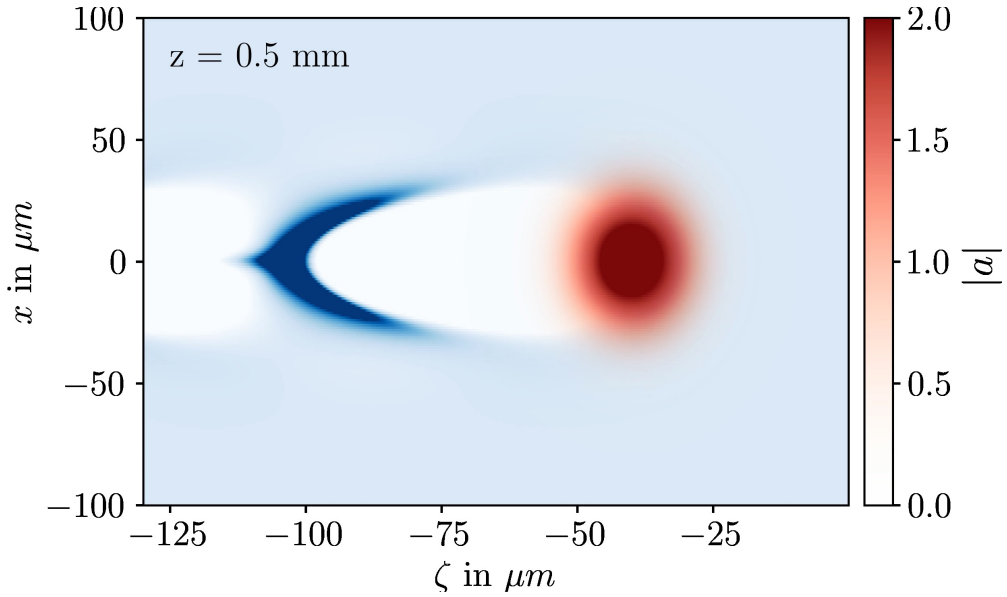

Production run:

```
grid points = 511 511 2048
time steps = 1000
plasma particles per cell = 1
```

Run time: 89 s (laser solver: 14%)
Costs: 0.8 node hours
(on 32 nodes on the JUWELS Booster)

Standard LWFA setting:

```
grid points = 511 511 1000
time steps = 50
plasma particles per cell = 1
```

Run time: 75 s (laser solver 23%)
Costs: 0.02 node hours
(on 1 node on the JUWELS Booster)

Another challenging plasma accelerator: AWAKE

- Very long proton beams (many cm) → huge amount of beam particles required (>3e9)
  beam does not fit on GPU
  very long box required (150000 longitudinal cells)
- long plasma propagation (10m and beyond)
- Still want to resolve small emittance witness beam → high transverse resolution required

Previously:

tilted proton bunch test

```
grid points = 511 511 2048
time steps = 400
beam particles = 500e6
plasma particles per cell = 4
```

Run time: 4 min
Costs: 1 node hour using 16 nodes
(64 A100 GPUs)

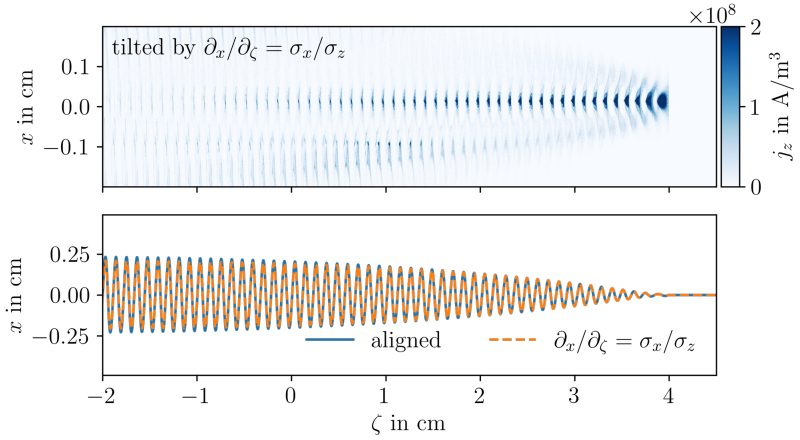

Another challenging plasma accelerator: AWAKE

- Very long proton beams (many cm) → huge amount of beam particles required (>3e9)
  beam does not fit on GPU
  very long box required (150000 longitudinal cells)
- long plasma propagation (10m and beyond)
- Still want to resolve small emittance witness beam → high transverse resolution required

Work in progress:

AWAKE run 1 baseline benchmark

```
grid points = 511 511 150000
time steps = 125
beam particles = 3.5e9
plasma particles per cell = 8
```

Run time: 28 min
Costs: 16 node hours using 32 nodes (128 A100 GPUs)

- Using the latest parallelization (not merged yet)
- 60% of the run time is initialization! This will be fixed soon
- May need to increase resolution
- Mesh refinement will enable high transverse resolution for witness beam

Summary

- HiPACE++ is production ready
- it enables previously unfeasibly simulation settings
- Standard plasma accelerator simulations are very cheap
- We are continually improving the code to cover even the most challenging scenarios via mesh refinement and memory management