

PSI Center for Scientific Computing, Theory and Data

Performance-Portable Space Charge Methods in OPALX

For Hadron and Lepton Accelerator Modelling

A. Adelmann, R. Ammann, Y. Bi, J. Biddiscombe, P. Calvo, S. Klapproth, A. Liemen, S. Muralikrishnan, C. Rogers, J. Thompson, D. Winklehner, **M. Zampetakis***

Space Charge 2026, 7th ICFA BD Mini-Workshop, Oxford, UK, 2 October 2026

*michail.zampetakis@psi.ch

Overview

OPALX: Motivation and Computational Framework

Space-Charge Methods and Verification

- PIC, P^3M, Landau damping and disorder-induced heating

SwissFEL Application

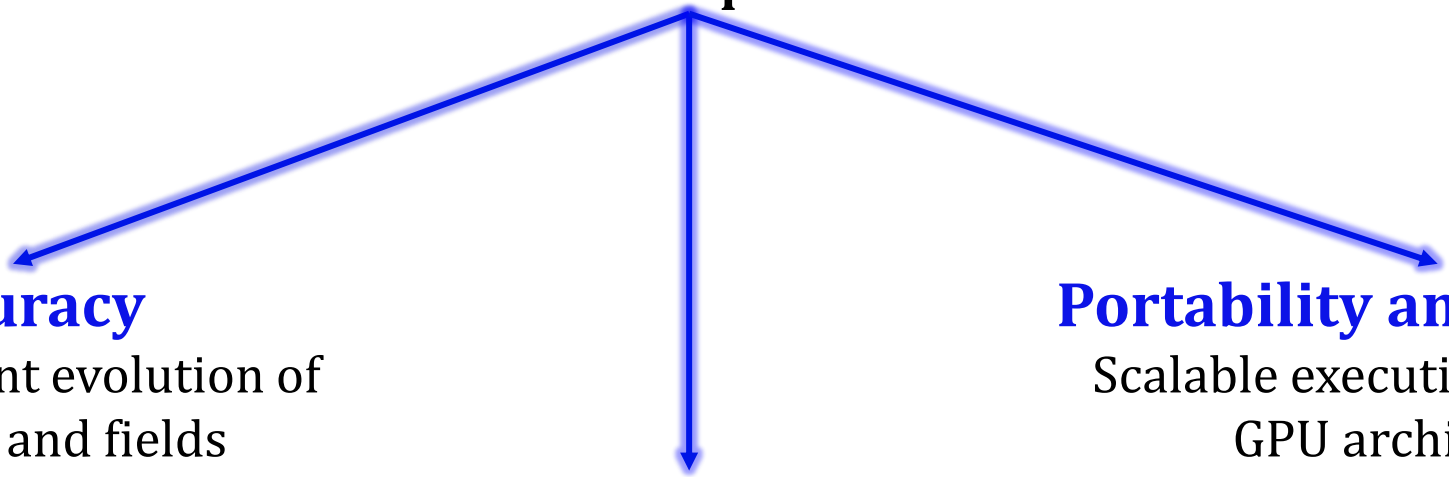
- OPAL comparison and CPU/GPU scaling

Conclusions and Outlook

Introduction

Higher charge densities ⟹ Stronger space-charge effects

Simulation requirements:



Accuracy
Self-consistent evolution of particles and fields

Modelling scope
A common framework for different accelerator types and beam regimes

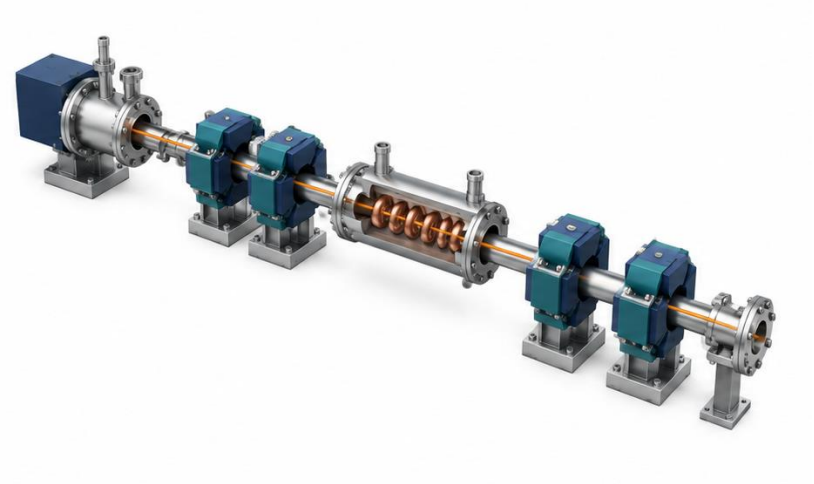
Portability and Scalability
Scalable execution on CPU and GPU architectures

OPAL – Object Oriented Parallel Accelerator Library

An established, open-source framework for parallel beam dynamics simulations, including 3D space charge.[1]

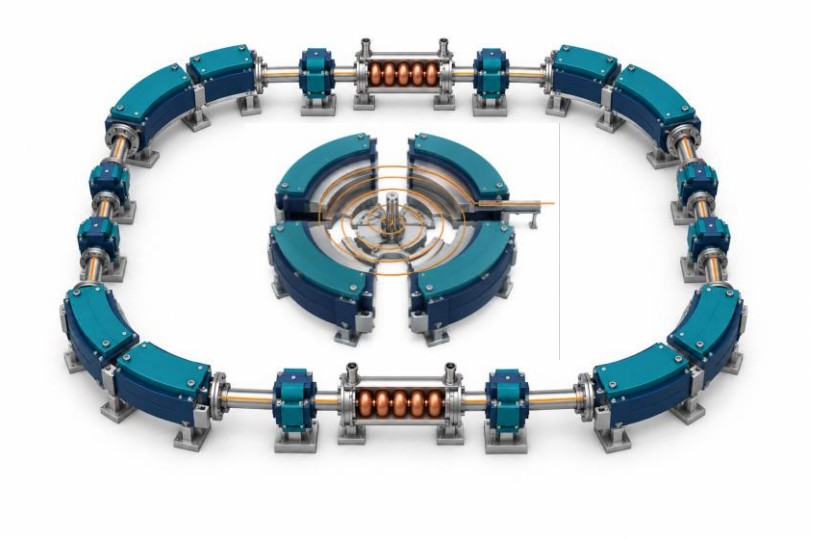
OPAL–T

Beamlines, LINACs, Photoinjectors and XFEL beamlines



OPAL–CYCL

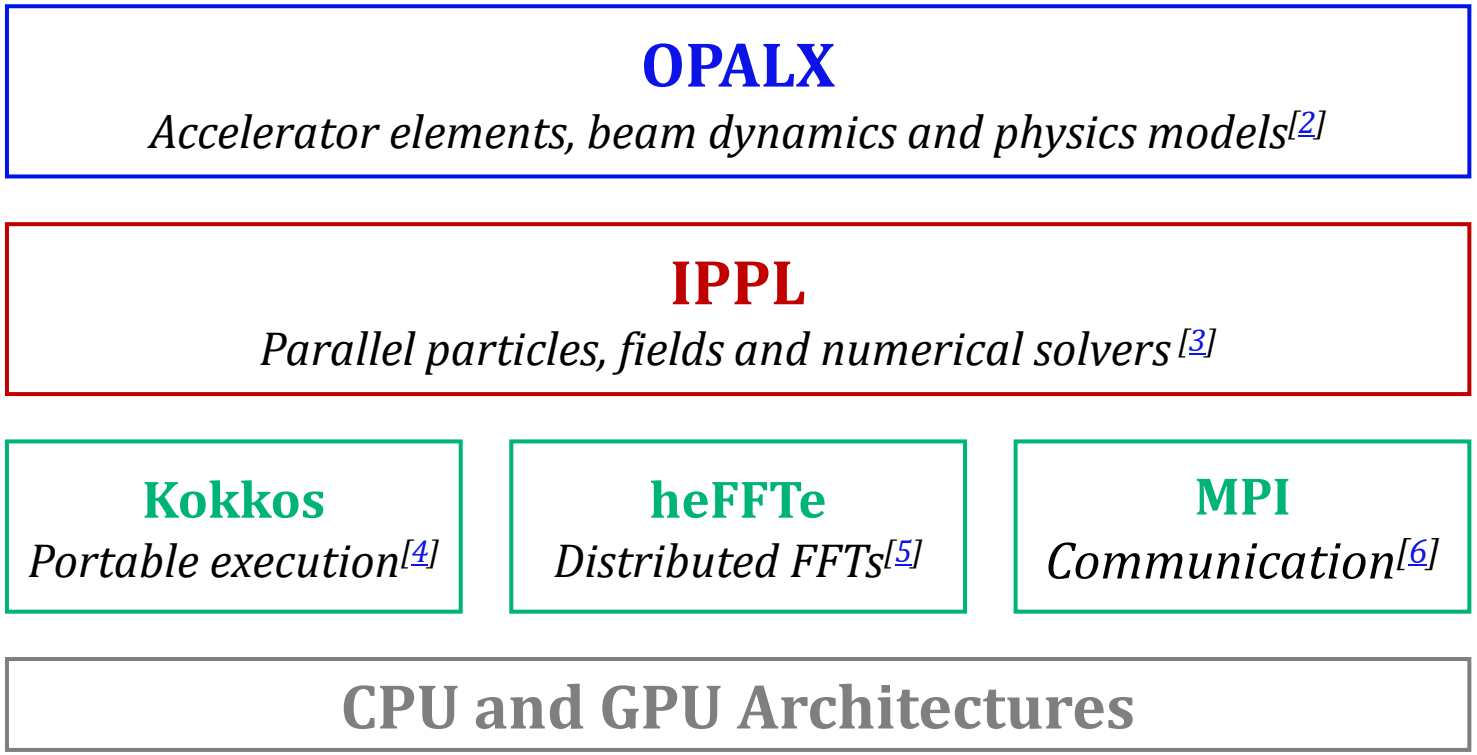
Cyclotrons, FFAs and general ring lattices.



Building on OPAL's modelling experience, OPALX targets portable execution on CPU and GPU architectures.

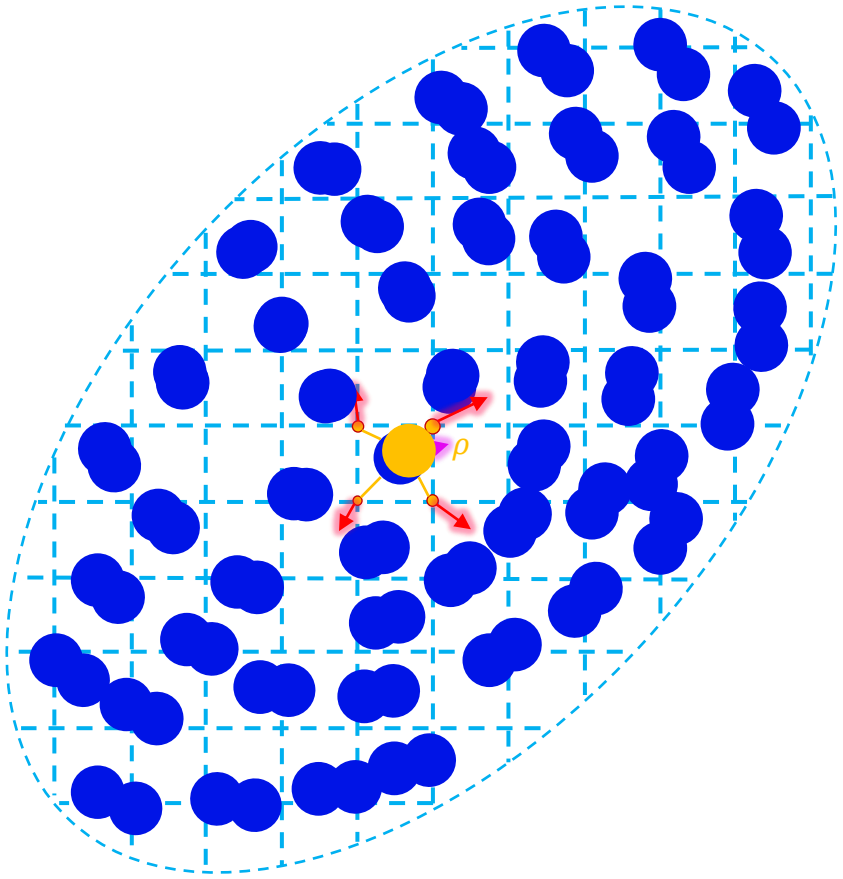
OPALX: Performance-Portable Beam Dynamics

A modernisation of OPAL for scalable beam dynamics simulations on CPU and GPU architectures.



Focus of this presentation: space-charge methods and their performance in OPALX.

Space-Charge Calculation: The PIC Cycle



(Schematic)

Particles

Macroparticles with positions $\boldsymbol{x}$*, momenta* $\boldsymbol{p}$*, and charges* q

Scatter

Deposit particle charge → mesh

Solve

Calculate the electric field on the mesh

Electrostatic approximation in the beam rest frame:

$$\boldsymbol{\nabla}^2\phi = -\rho/\varepsilon_0\,, \qquad \mathbf{E} = -\boldsymbol{\nabla}\phi$$

Gather

Interpolate the field to the particles

Push

Update particle momenta and positions

Next step

Space-Charge Solvers: IPPL and OPALX

IPPL provides the numerical methods; OPALX integrates selected solvers for accelerator simulations.

Methods available in IPPL	Status in OPALX
FFT-based Poisson solvers	**Implemented**
Particle-Particle Particle-Mesh (P^3M)	**Implemented**
Finite differences with PCG	Integration planned
Finite elements	Integration planned
Particle-in-Fourier	Integration planned

In development within OPALX: 2.5D space-charge solver. *(See the poster by J. Thompson et al.)*

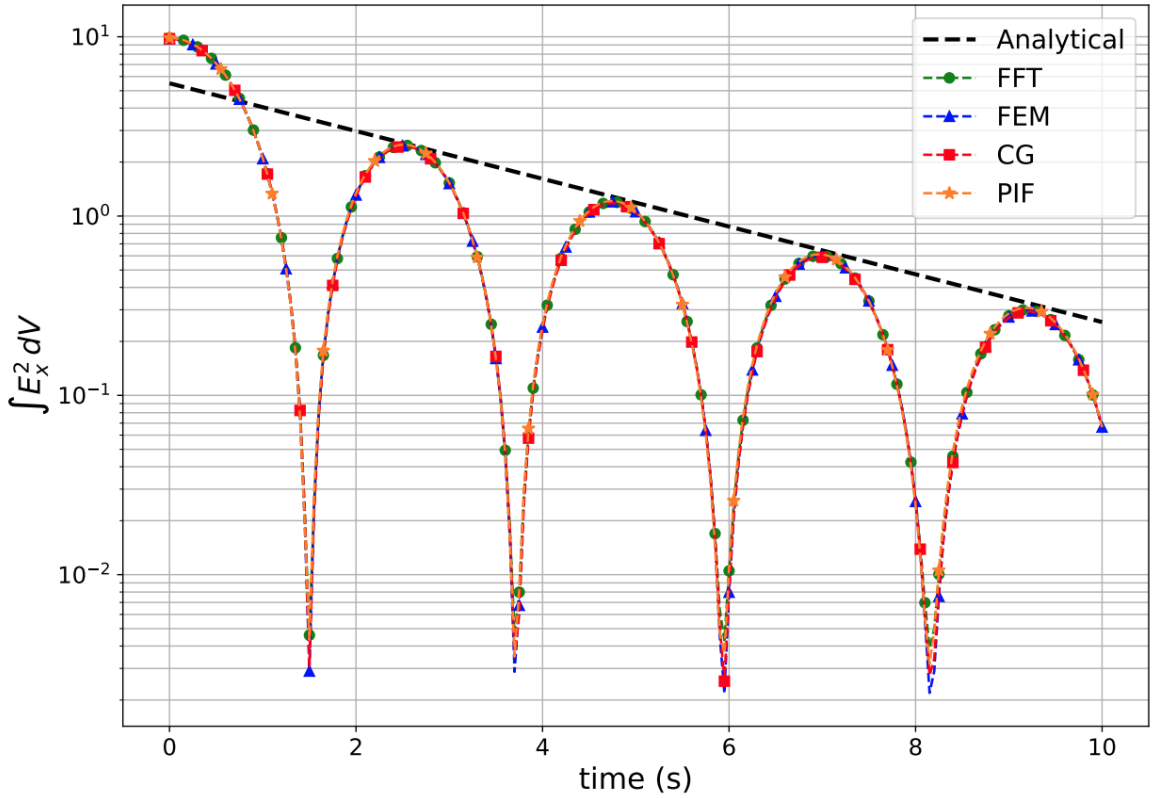
Verification in IPPL: Landau damping

Test: damping of an electrostatic perturbation in a collisionless plasma.

Setup: Initially Maxwellian electrons with a small sinusoidal density perturbation ($a = 0.05$).

Comparison: PIC with FFT, Finite Differences and FEM; Particle-in-Fourier.

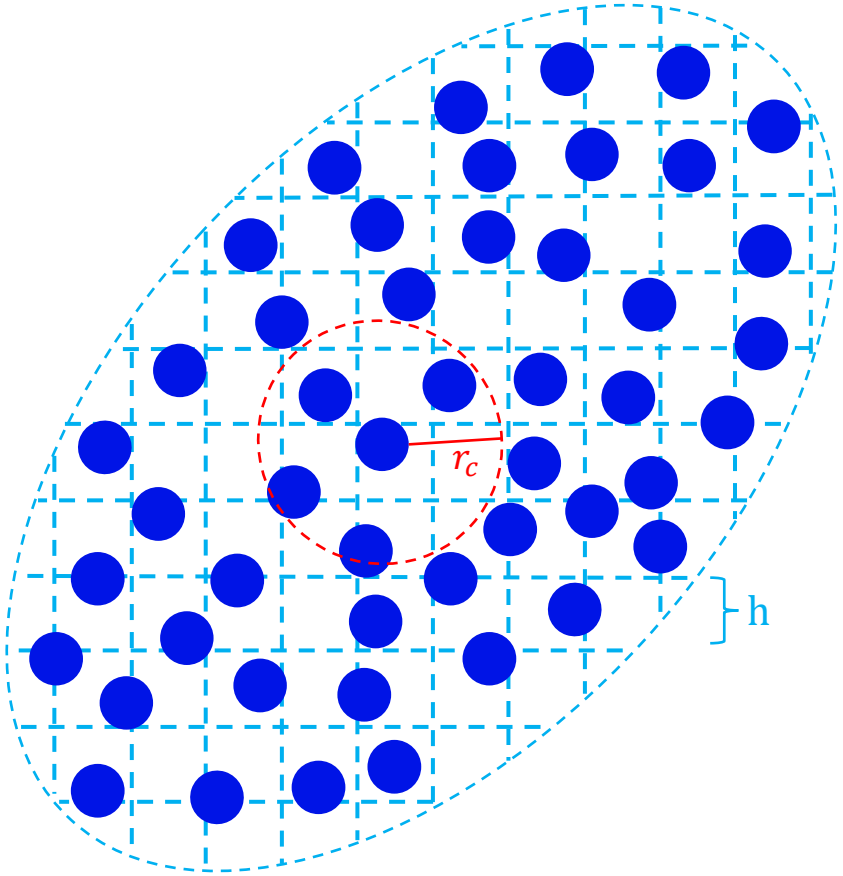
Courtesy of S. Mayani et al., PASC'26.[6]



The methods reproduce the expected damping trend and closely agree over the interval shown.

Particle-Particle Particle-Mesh $\left(\mathbf{P^3M}\right)$

Combine a mesh-based calculation of long-range forces with direct short-range particle interactions.



(Schematic)

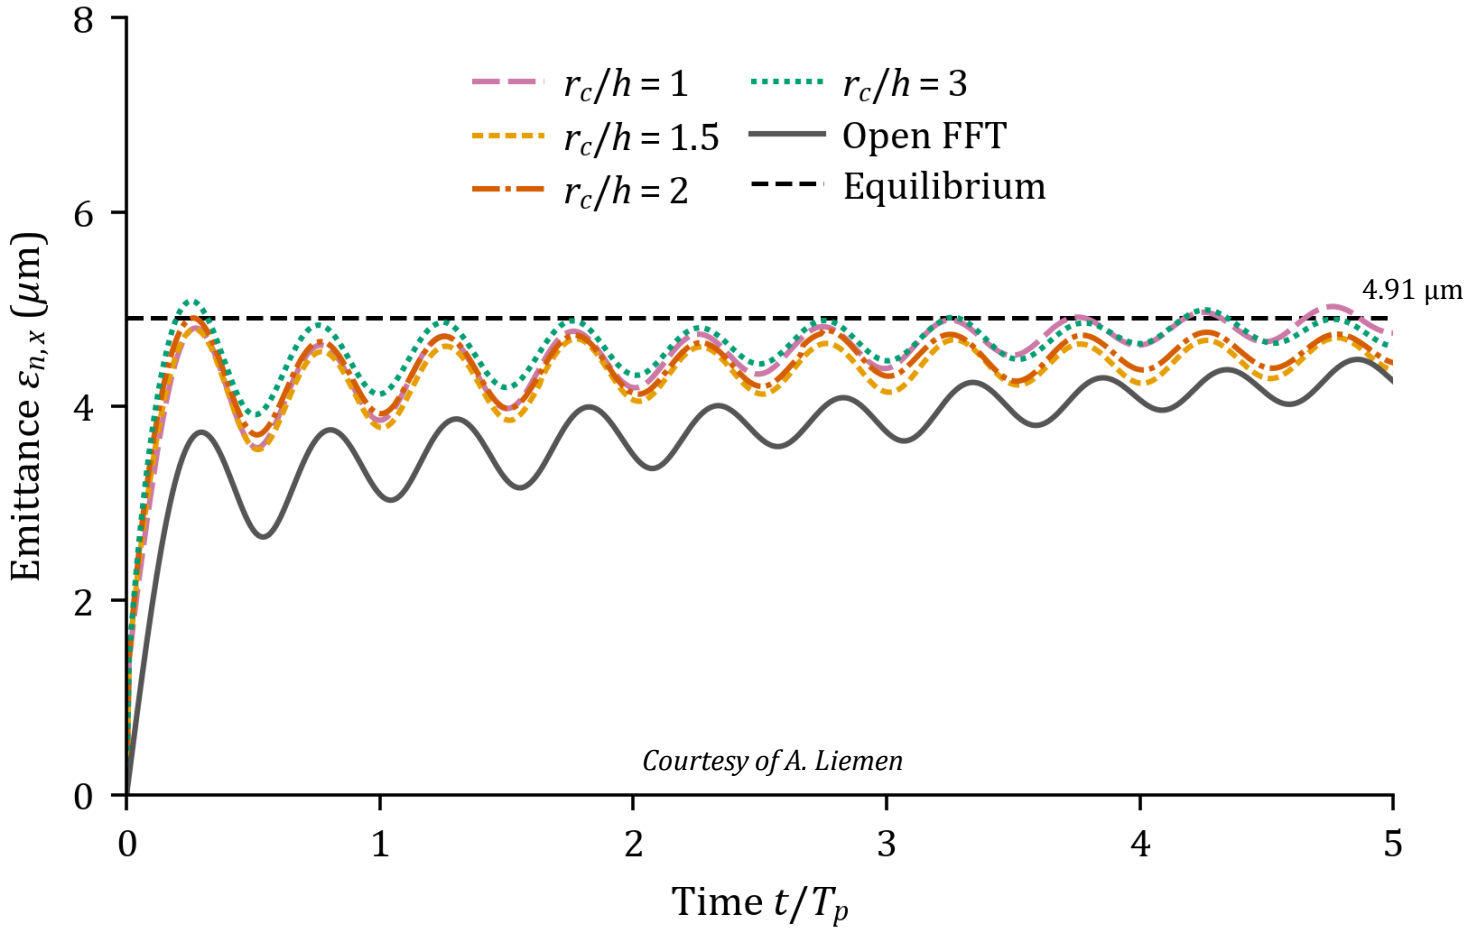
Particle–Mesh: long-range contribution
Calculate the smooth long-range contribution from all particles using an FFT-based mesh solver.

Particle–Particle: short-range correction
Evaluate the short-range contribution directly for neighbouring pairs within r_c.

$$\boldsymbol{F}_i = \boldsymbol{F}_i^{\text{long}} + \sum_{\substack{j \neq i \\ r_{ij} < r_c}} \boldsymbol{F}_{ij}^{\text{short}}$$

Disorder-Induced Heating: P^{3}M Verification in OPALX

Short-range particle correlations convert electrostatic potential energy into random motion, increasing beam emittance.



Cold, uniform electron sphere.

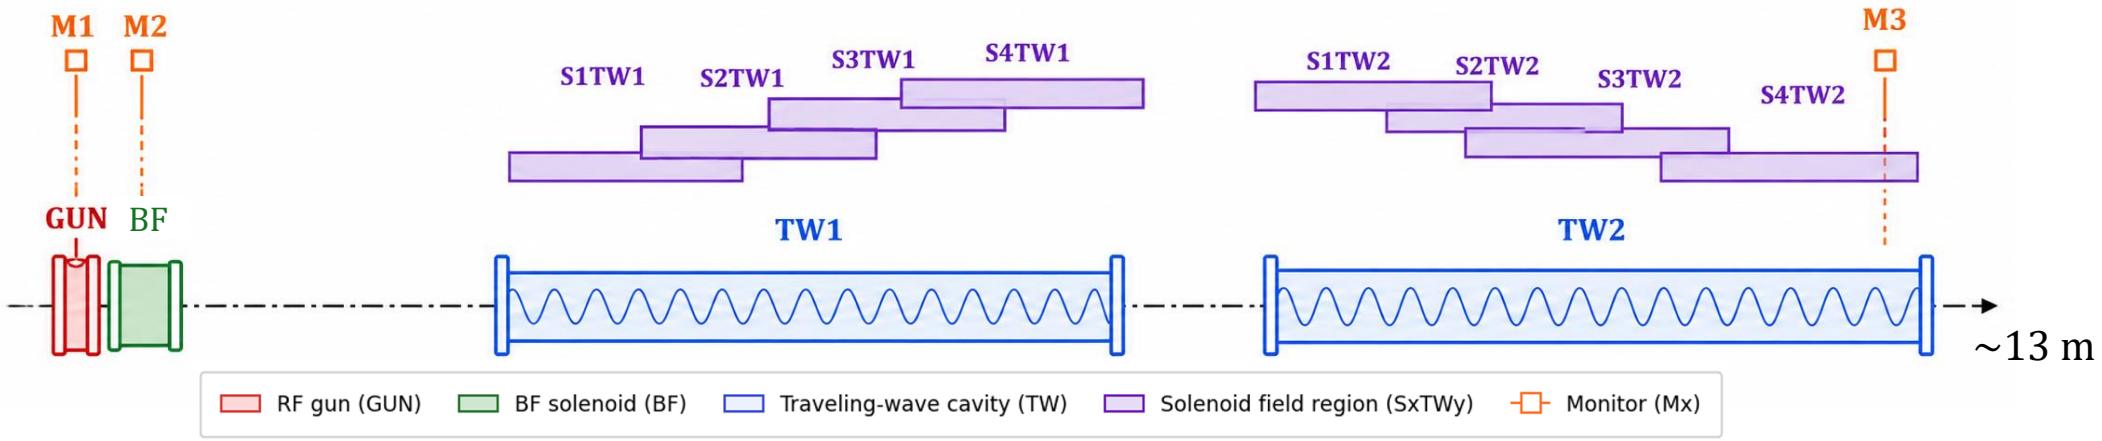
$N_p = 156'055$ macroparticles.

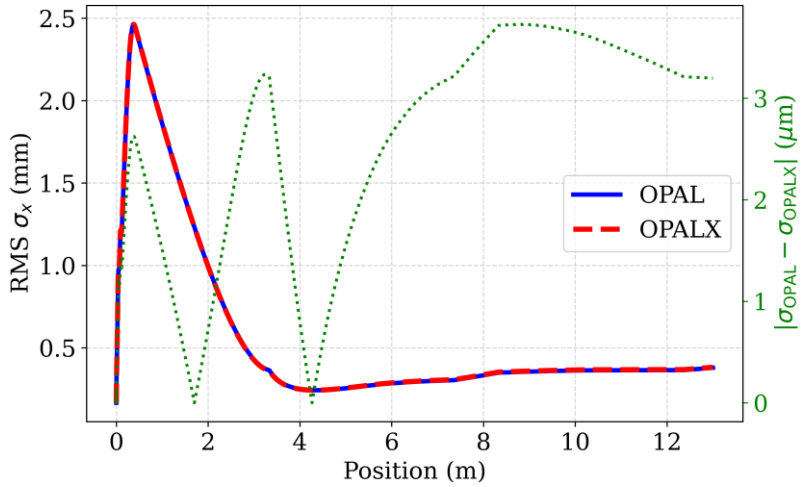
Total charge $Q = -25.0028$ nC.

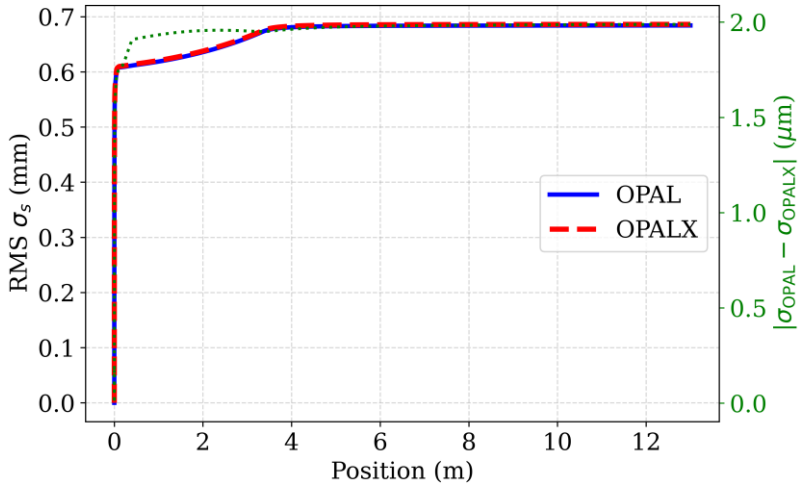
$N = 128^3$ mesh.

Five plasma periods, $T_p = 43.115$ ps.

SwissFEL Booster: OPALX – OPAL Comparison



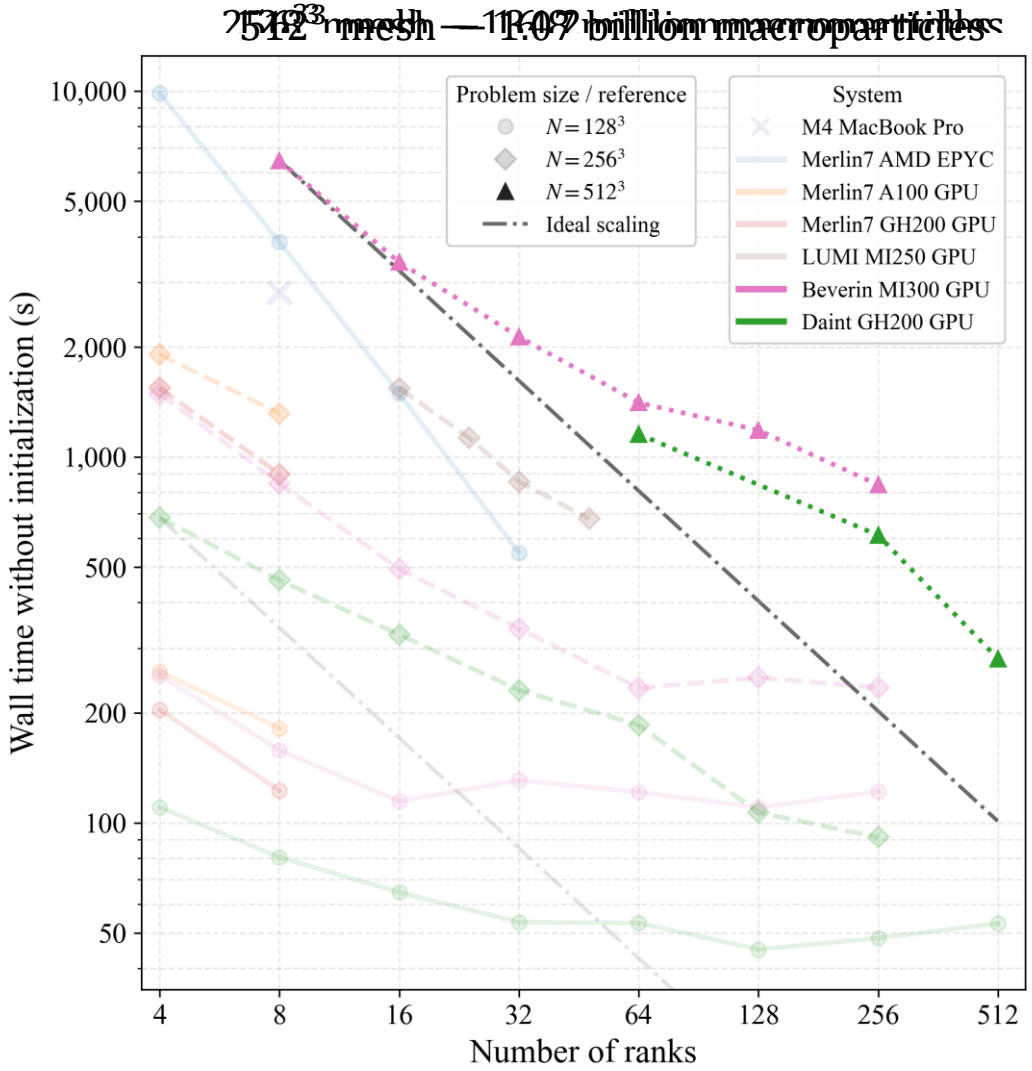




Mesh: 128^3
Macroparticles: 16.8 million

OPALX closely reproduces OPAL's RMS beam-size evolution through the SwissFEL booster **with space charge.**

SwissFEL Booster: CPU and GPU Scaling



- **1,000 timesteps**
- **8 macroparticles per mesh cell**
- **Space-charge solver:** Open FFT
- **M4 MacBook Pro:** CPU, 8 MPI ranks, SERIAL backend
- **CPU:** 4 MPI ranks/node × 32 cores
- **GPU:** 4 GPUs/node, 1 MPI rank/GPU

Conclusions and Outlook

Conclusions:

- **OPALX** combines **accelerator modelling** with **portable execution** on **CPU** and **GPU** architectures.
- **Physics benchmarks test complementary capabilities:** Landau damping in IPPL and disorder-induced heating with P^3M in OPALX.
- **SwissFEL results show close OPAL–OPALX agreement in RMS beam size evolution.** Performance studies demonstrate execution across CPU and GPU systems, including a billion-macroparticle configuration.

Outlook:

- Further performance optimisation on CPUs and GPUs.
- Extend ring-simulation capabilities.
- Continue development and verification of the 2.5D space-charge solver.
- ***OPALX v1.0 release by the end of October — Stay tuned!***

Scan for the OPALX Manual
opalx-project.github.io/opalx-manual/

PSI

Thank you!

OPALX

Scan for the OPALX Manual

opalx-project.github.io/opalx-manual/

References

[1] A. Adelmann *et al.*, “OPAL a versatile tool for charged particle accelerator simulations” (2019), arXiv:1905.06654 [physics.acc-ph].

[2] OPALX Project, OPALX: Object-oriented particle accelerator library for exascale, https://github.com/OPALX-project/OPALX (2026)

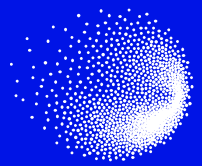
[3] S. Muralikrishnan *et al.*, “Scaling and performance portability of the particle-in-cell scheme for plasma physics applications through mini-apps targeting exascale architectures”, *in Proc.* 2024 SIAM Conference on Parallel Processing for Scientific Computing (PP) (Society for Industrial and Applied Mathematics, 2024) pp. 26–38. Link

[4] C. R. Trott *et al.*, “Kokkos 3: Programming model extensions for the exascale era”, IEEE Transactions on Parallel and Distributed Systems 33, 805 (2022). Link

[5] A. Ayala, S. Tomov, A. Haidar, and J. Dongarra, “heFFTe: Highly Efficient FFT for Exascale”, in *Computational Science – ICCS 2020*, Amsterdam, The Netherlands, 2020, pp. 262–275. doi: 10.1007/978-3-030-50371-0_19.

[6] Message Passing Interface Forum, MPI: A Message-Passing Interface Standard, Version 5.0 (2025). Link

[7] S. Mayani, P. Fischill, S. Muralikrishnan, and A. Adelmann, “A Comparison of Massively Parallel Performance Portable Particle-in-Cell schemes for electrostatic kinetic plasma simulations”. *In Proc.* Platform for Advanced Scientific Computing Conference (PASC '26). Association for Computing Machinery, New York, NY, USA, Article 18, 1–11. https://doi.org/10.1145/3815572.3815754

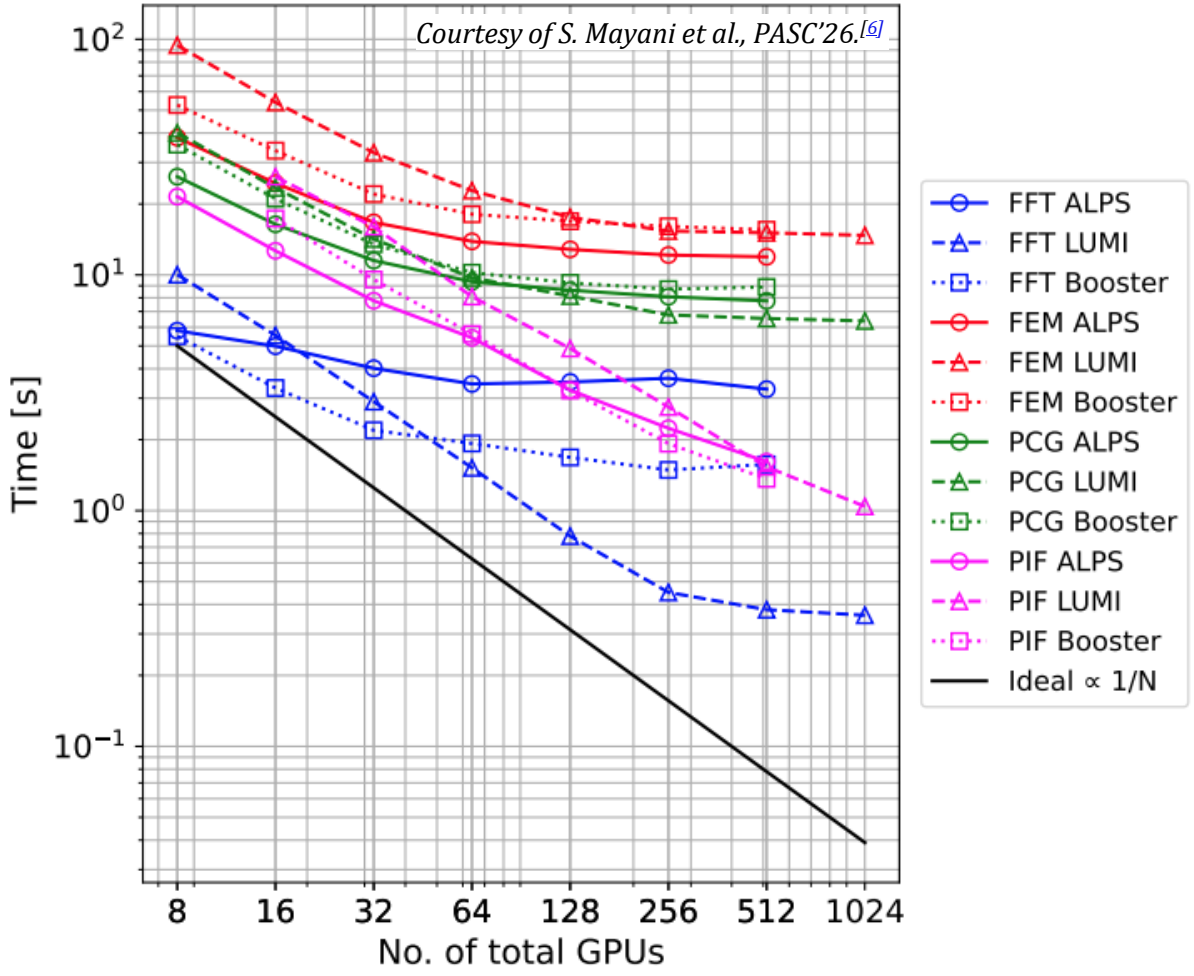


PSI

Backup Slides

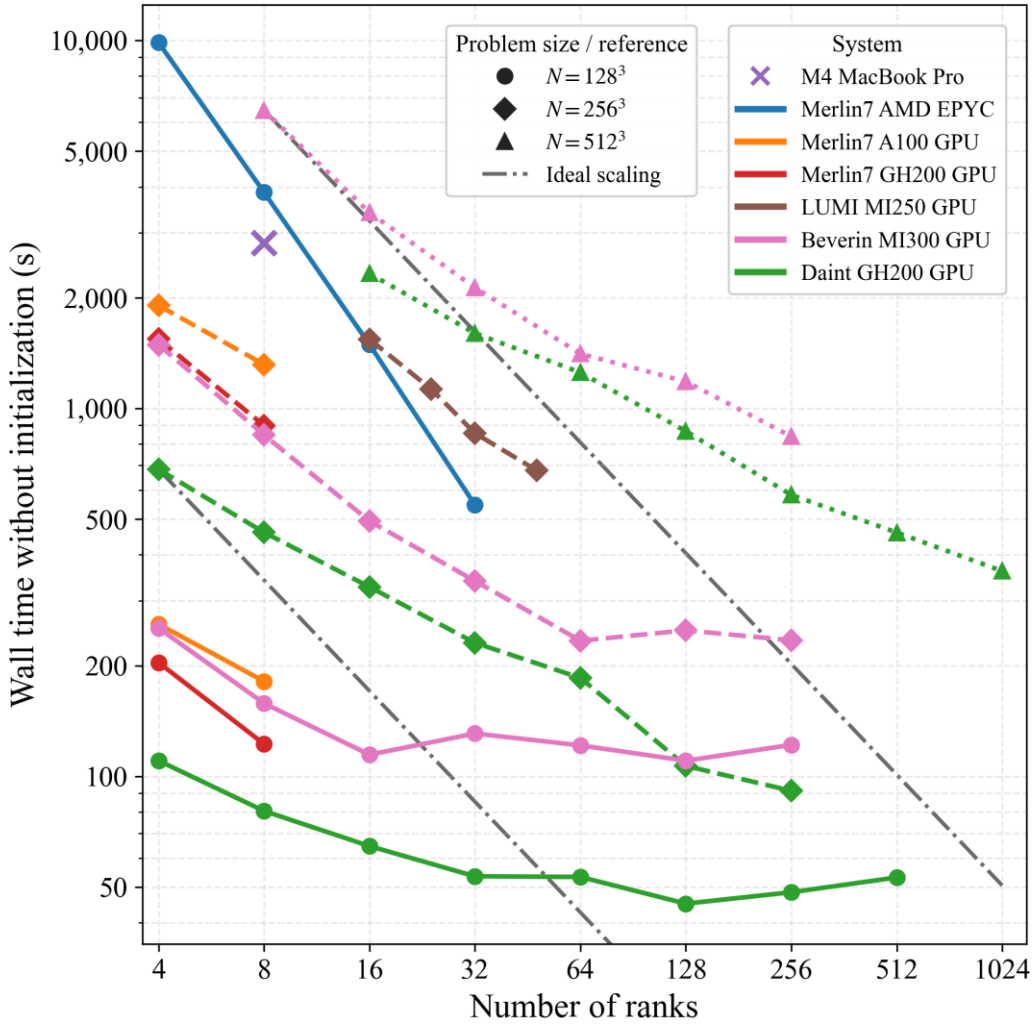
Complete Scaling Plots

IPPL: Landau damping



FFT solver on ALPs had a bug. Updated scaling is much better.

OPALX: SwissFEL booster



DIH Benchmark: Setup and Reference Values

Open-boundary P^3M benchmark, courtesy of A. Liemen

Simulation setup	
Distribution	Uniform electron sphere
Macroparticles	156,055
Sphere radius	1.774 mm
Total charge	−25.0028 nC
Mesh	128^3, initial h = 29.334 µm
Long-range solve	Open Hockney FFT
Cutoff r_c / h	0.5, 1, 1.5, 2, 3
Time step	0.431246 ps
Duration	500 steps ≈ 5 plasma periods
Execution	8 MPI ranks, SERIAL backend

Similarity scaling and reference	
Reference radius	17.74 µm
Length scale S	100
Macro weight	$\approx 10^6$ electrons
Electron density	6.6731×10^{18} m^{-3}
Plasma period	43.115 ps (unchanged)
Reference emittance	0.491 nm
Scaled emittance	$S^2 \times 0.491$ nm = 4.91 µm

Position and velocity scale by S.
Normalized emittance scales by S^2.

Reference: Mitchell and Qiang, IPAC 2015; Bachmann, Chapter 4.

Clusters

Merlin7

- Multi-core CPU node (x86), *2x* AMD EPYC 7742 (x86_64 Rome, 64 Cores, 3.2GHz) with 512GB DRR4 3200Mhz RAM
- Multi-core CPU + GPU node (x86), *2x* AMD EPYC 7713 (x86_64 Milan, 64 Cores, 3.2GHz) with 512GB DRR4 3200Mhz RAM, and *4x* NVidia A100 (Ampere, 80GB)
- Grace Hopper node (arm64), 4x Nvidia Grace CPU (SBSA arm64 , 72 Cores, 3.44GHz) with 128 GB DDR5 6400MHz RAM, and 4x Nvidia H200 (Hopper, 96GB)

LUMI

- 2978 nodes, each node with one 64-core AMD Trento CPU and four AMD MI250X GPUs. A total of 11 912 AMD GPUs.
- The MI250X GPU comes with a total of 128 GB of HBM2e memory, offering over 3.2 TB/s of memory bandwidth.

CSCS Alps

- Daint: 1022 nodes, 4088 combined CPUs and GH200 GPUs
- Beverin: MI300X with 192 GB HBM3 memory and 5.3 TB/s of memory bandwidth