EMUZAG

Time-domain electromagnetics for wakefields and impedance

Baseline EM/FDTD + TETM stationary and moving schemes

Igor Zagorodnov

Demonstration and overview · 11 August 2026

Development runtime: EMQ42 0.1.1

We want the wake left behind a relativistic bunch

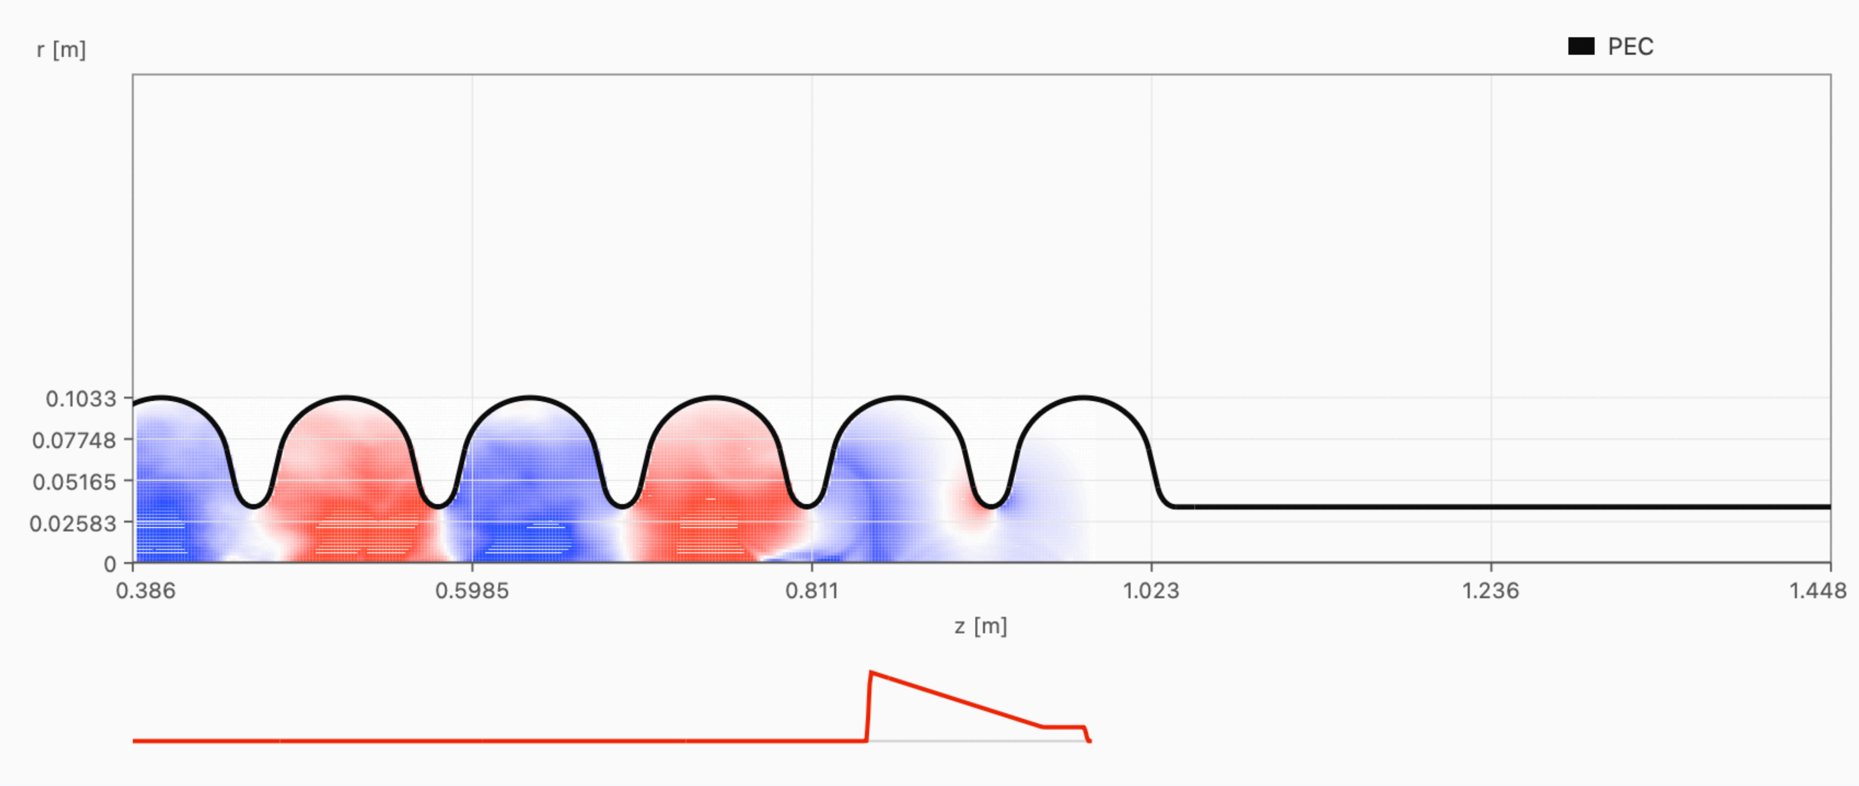

Time-domain fields → wake potential W(s) → coupling impedance Z(ω)

PHYSICS QUESTION

The physical wake length selects the computational frame

SHORT-RANGE WAKE

Long structure, short region near the bunch

Memory follows the bunch

LONG-RANGE WAKE

Compact structure, wake persists far behind

The whole structure remains

Two solver families cover different numerical needs

EM

- Stationary Yee-type FDTD scheme
- Reference path comparable to CST/GdfidL
- Robust baseline for general 3D cases
- Conventional numerical dispersion

TETM

- Stationary and moving formulations
- TETM low-dispersion recurrence
- No beam-direction numerical dispersion
- No extra time-step reduction to suppress it

One product, two numerical tools - selected by physics, not by benchmark identity.

Longitudinal phase error accumulates into a wake error

A tiny phase-velocity error per cell becomes visible after thousands of cells.

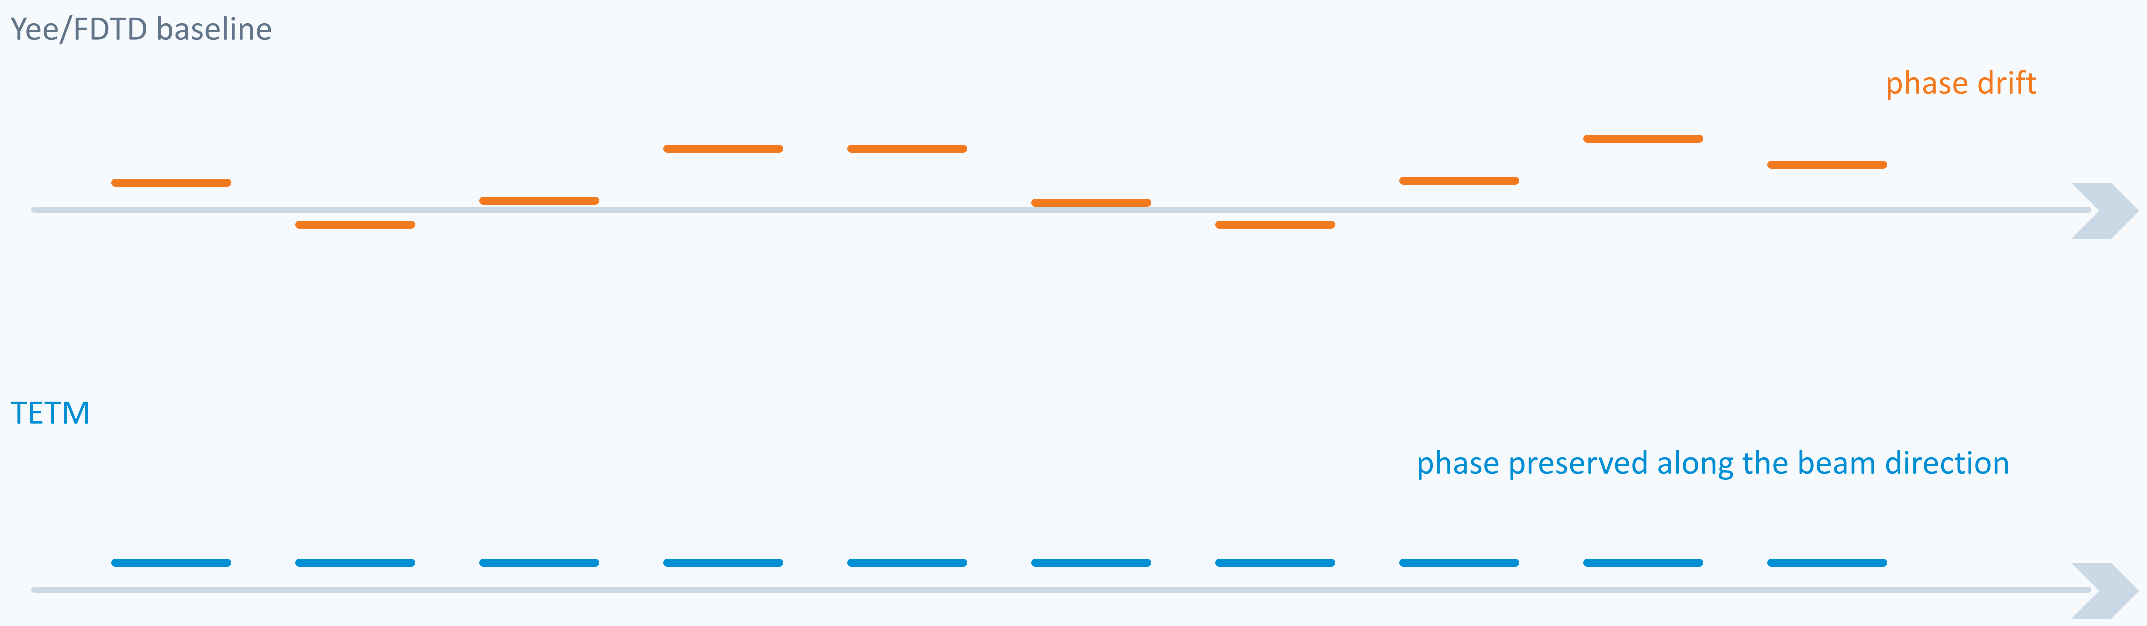

TETM removes the long-structure mesh penalty

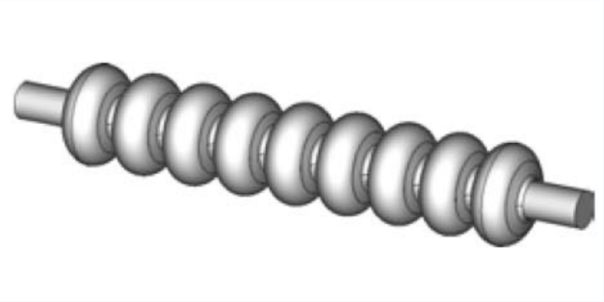

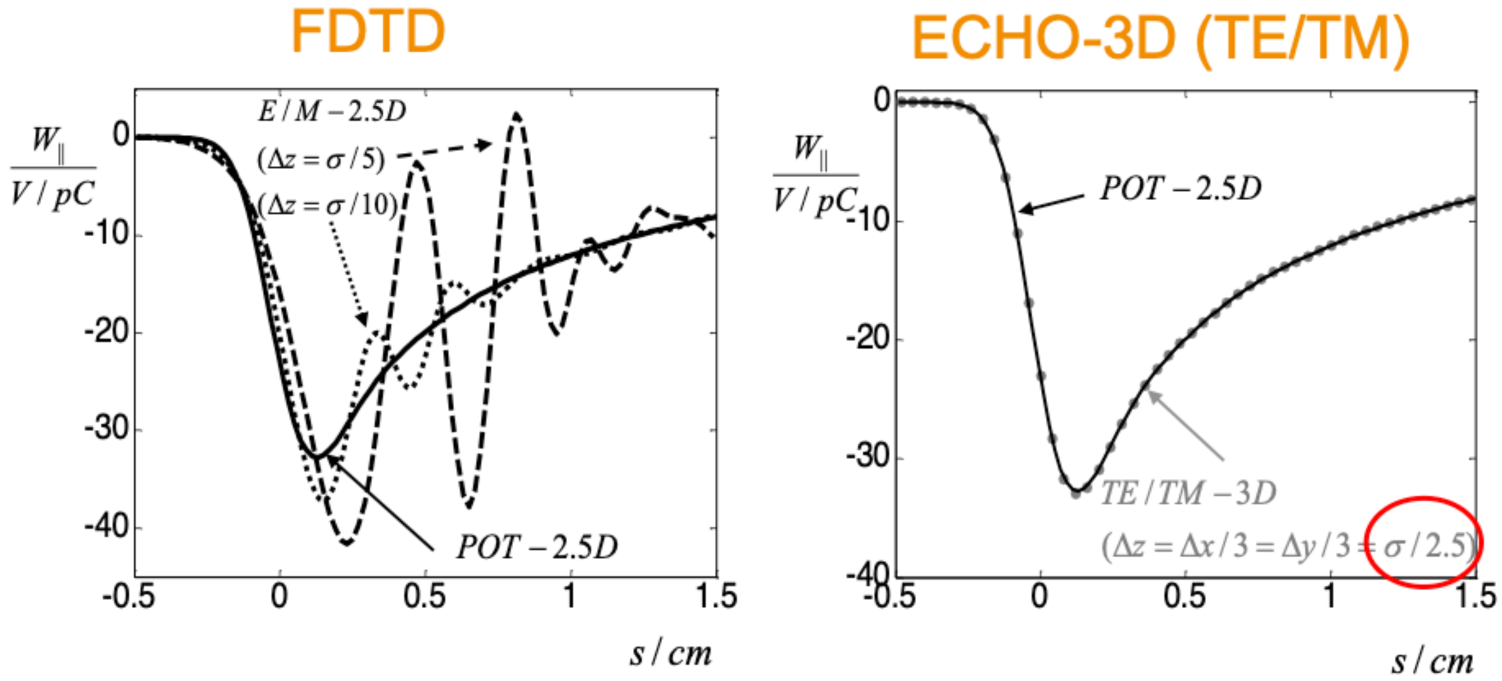

Comparison of the wake potentials for structure consisting of 20 TESLA cells excited by Gaussian bunch with RMS length $\sigma = 1$ mm

E/M splitting $\Delta z \sim \sqrt{\dfrac{\sigma^3}{L}}$

TE/TM splitting $\Delta z \sim \sigma$

Coarse-grid accuracy changes 3D economics

Refining all dimensions by 2 is not a 2× decision.

8×

more 3D cells

≈16×

more explicit update work

EMUZAG target

Accuracy on a comfortable grid

- Conformal boundaries
- Low longitudinal dispersion
- Full time step

The method should buy accuracy before hardware buys cells.

EVIDENCE

Published data confirms the ECHO convergence advantage

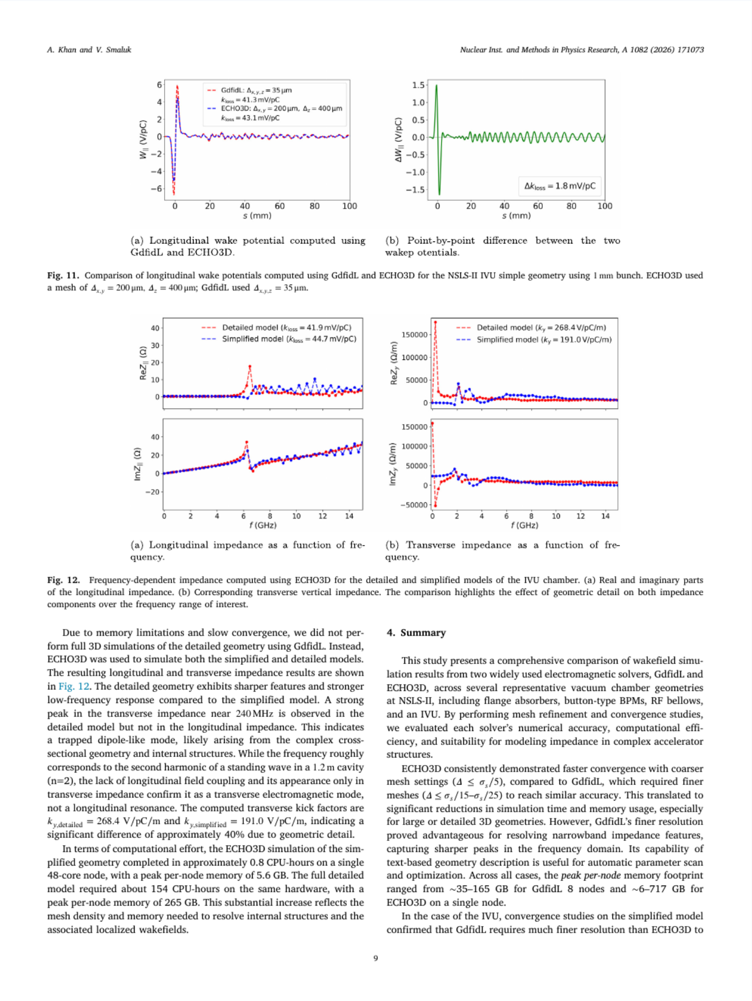

(a) Longitudinal wake potential computed using GdfidL and ECHO3D. (b) Point-by-point difference between the two wakep otentials.

Fig. 11. Comparison of longitudinal wake potentials computed using GdfidL and ECHO3D for the NSLS-II IVU simple geometry using 1 mm bunch. ECHO3D used a mesh of $\Delta_{x,y} = 200\,\mu m$, $\Delta_z = 400\,\mu m$; GdfidL used $\Delta_{x,y,z} = 35\,\mu m$.

(a) Longitudinal impedance as a function of frequency. (b) Transverse impedance as a function of frequency.

Fig. 12. Frequency-dependent impedance computed using ECHO3D for the detailed and simplified models of the IVU chamber. (a) Real and imaginary parts of the longitudinal impedance. (b) Corresponding transverse vertical impedance. The comparison highlights the effect of geometric detail on both impedance components over the frequency range of interest.

Due to memory limitations and slow convergence, we did not perform full 3D simulations of the detailed geometry using GdfidL. Instead, ECHO3D was used to simulate both the simplified and detailed models. The resulting longitudinal and transverse impedance results are shown in Fig. 12. The detailed geometry exhibits sharper features and stronger low-frequency response compared to the simplified model. A strong peak in the transverse impedance near 240 MHz is observed in the detailed model but not in the longitudinal impedance. This indicates a trapped dipole-like mode, likely arising from the complex cross-sectional geometry and internal structures. While the frequency roughly corresponds to the second harmonic of a standing wave in a 1.2 m cavity (n=2), the lack of longitudinal field coupling and its appearance only in transverse impedance confirm it as a transverse electromagnetic mode, not a longitudinal resonance. The computed transverse kick factors are $k_{y,detailed} = 268.4$ V/pC/m and $k_{y,simplified} = 191.0$ V/pC/m, indicating a significant difference of approximately 40% due to geometric detail.

In terms of computational effort, the ECHO3D simulation of the simplified geometry completed in approximately 0.8 CPU-hours on a single 48-core node, with a peak per-node memory of 5.6 GB. The full detailed model required about 154 CPU-hours on the same hardware, with a peak per-node memory of 265 GB. This substantial increase reflects the mesh density and memory needed to resolve internal structures and the associated localized wakefields.

4. Summary

This study presents a comprehensive comparison of wakefield simulation results from two widely used electromagnetic solvers, GdfidL and ECHO3D, across several representative vacuum chamber geometries at NSLS-II, including flange absorbers, button-type BPMs, RF bellows, and an IVU. By performing mesh refinement and convergence studies, we evaluated each solver's numerical accuracy, computational efficiency, and suitability for modeling impedance in complex accelerator structures.

ECHO3D consistently demonstrated faster convergence with coarser mesh settings ($\Delta \le \sigma_s/5$), compared to GdfidL, which required finer meshes ($\Delta \le \sigma_s/15$–$\sigma_s/25$) to reach similar accuracy. This translated to significant reductions in simulation time and memory usage, especially for large or detailed 3D geometries. However, GdfidL's finer resolution proved advantageous for resolving narrowband impedance features, capturing sharper peaks in the frequency domain. Its capability of text-based geometry description is useful for automatic parameter scan and optimization. Across all cases, the *peak per-node* memory footprint ranged from ~35–165 GB for GdfidL 8 nodes and ~6–717 GB for ECHO3D on a single node.

In the case of the IVU, convergence studies on the simplified model confirmed that GdfidL requires much finer resolution than ECHO3D to

Across NSLS-II components

$\Delta \le \sigma/5$

ECHO3D convergence guideline

$\Delta \le \sigma/15$–$\sigma/25$

GdfidL resolution needed for similar accuracy

The paper compares both codes as useful tools - but finds ECHO3D reaches stable wakes on much coarser meshes.

A. Khan & V. Smaluk, “Convergence study of wakefield simulations with GdfidL and ECHO3D,” NIM A 1082 (2026) 171073.

Flange absorber: stable wake on coarser grids

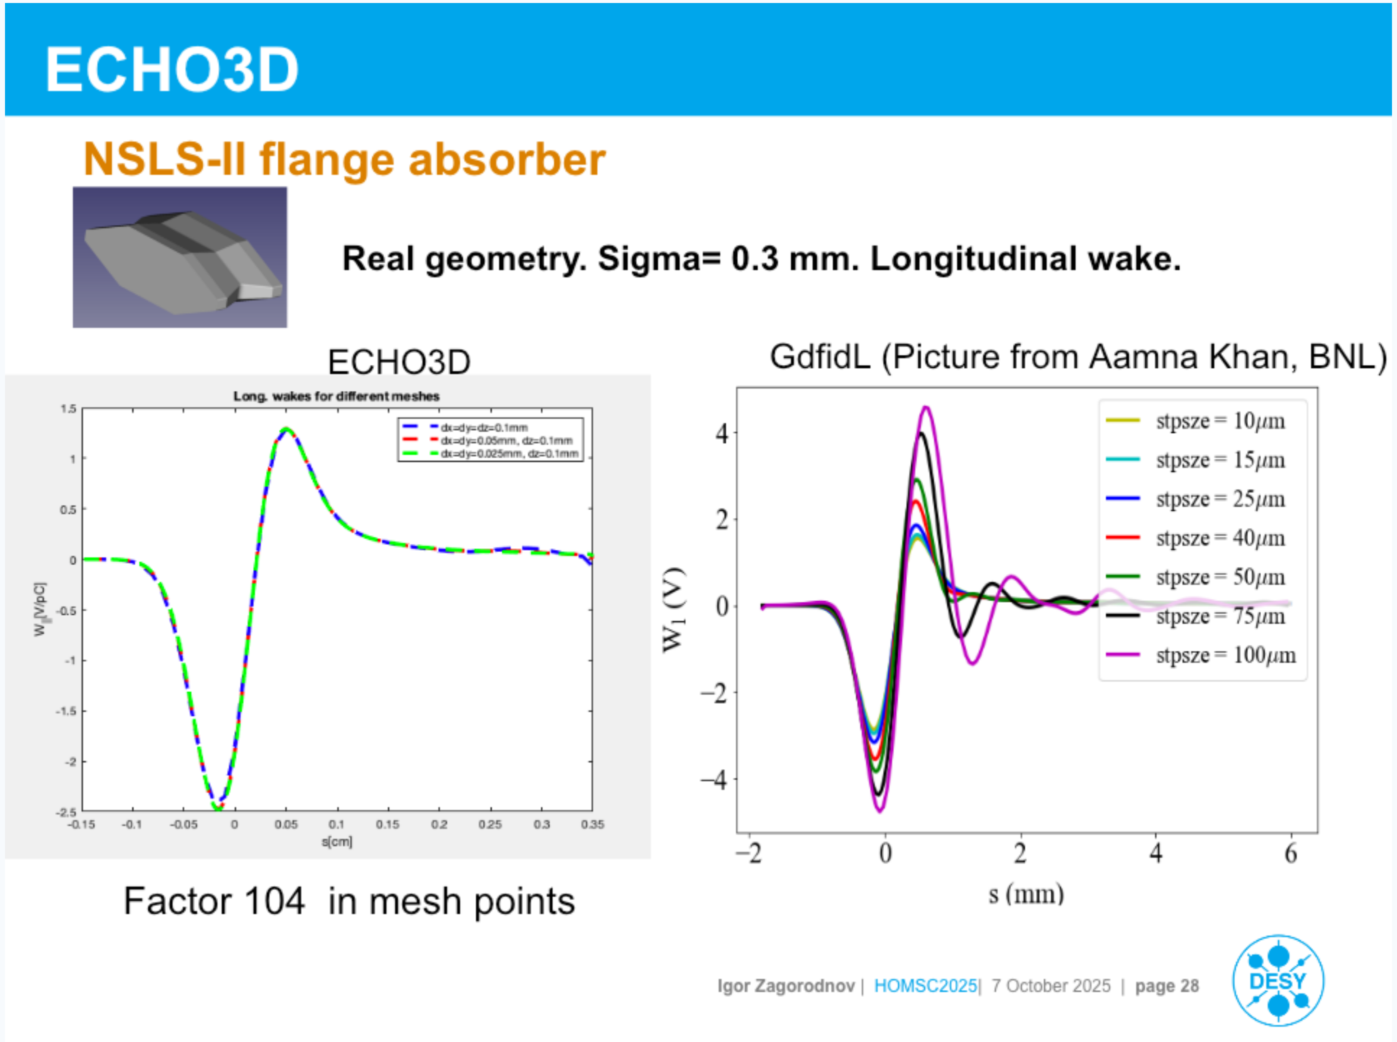

NSLS-II

σ = 0.3 mm

ECHO3D wake curves overlap on coarse anisotropic grids, while GdfidL remains sensitive to uniform refinement.

Published meshes:
ECHO 30×60×60 µm
GdfidL 15×15×15 µm

KEK collimator: 4× larger step, matching dipole wake

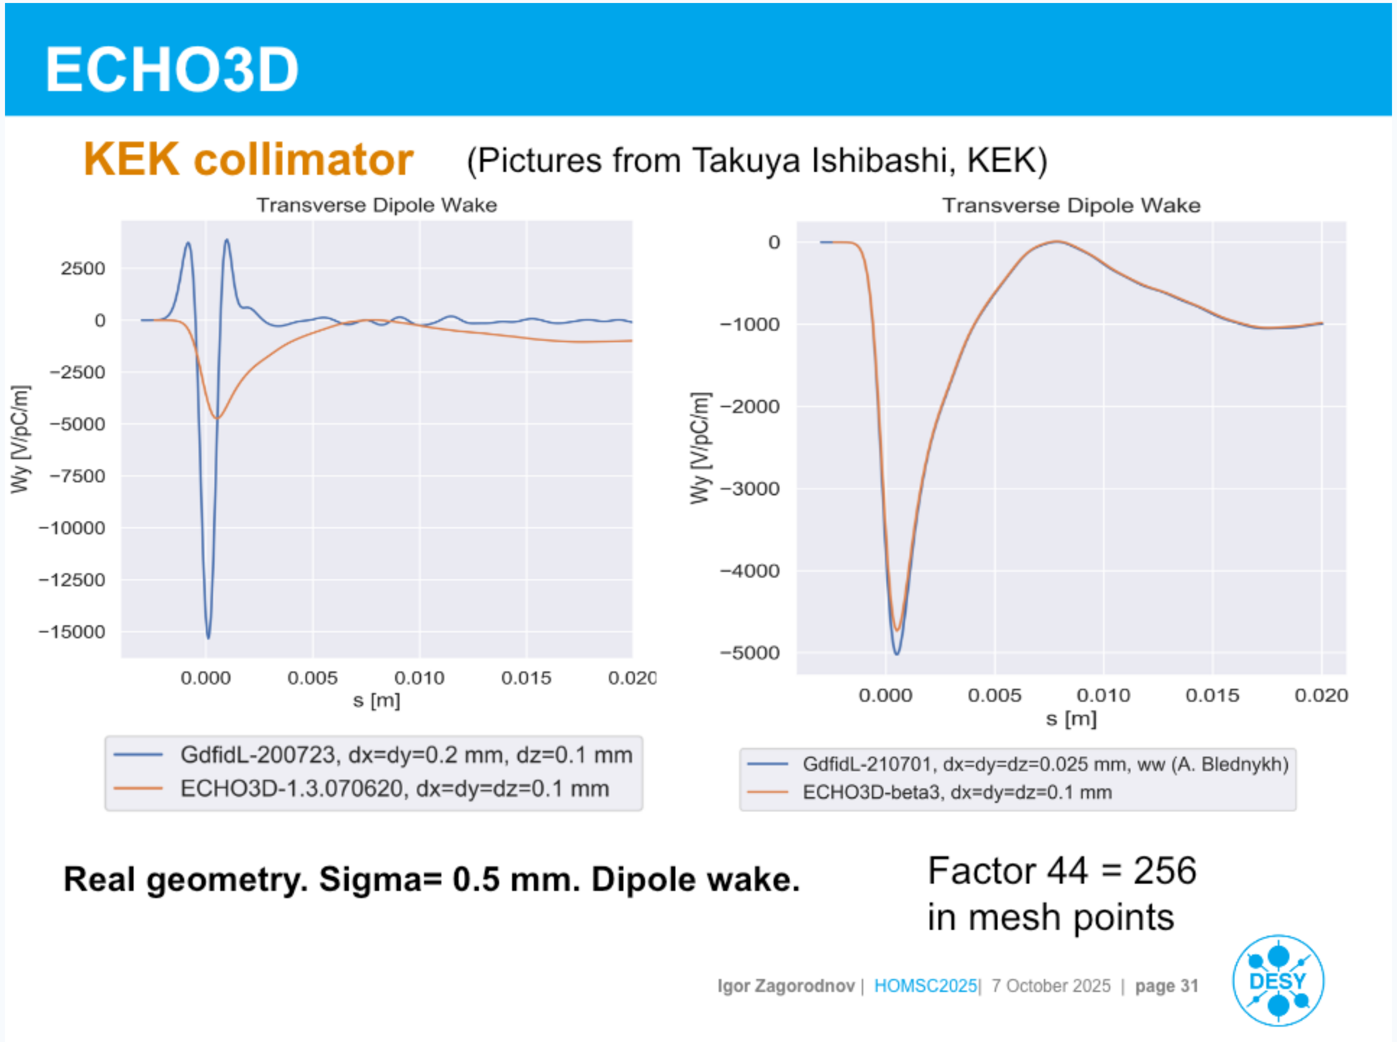

Converged comparison

4×

larger step in each direction

64×

fewer 3D cells at equal volume

With a 4× step-count effect, the update-work ratio reaches ≈256×.

EVIDENCE

PETRA IV: TETM convergence remains stable on coarse longitudinal grids

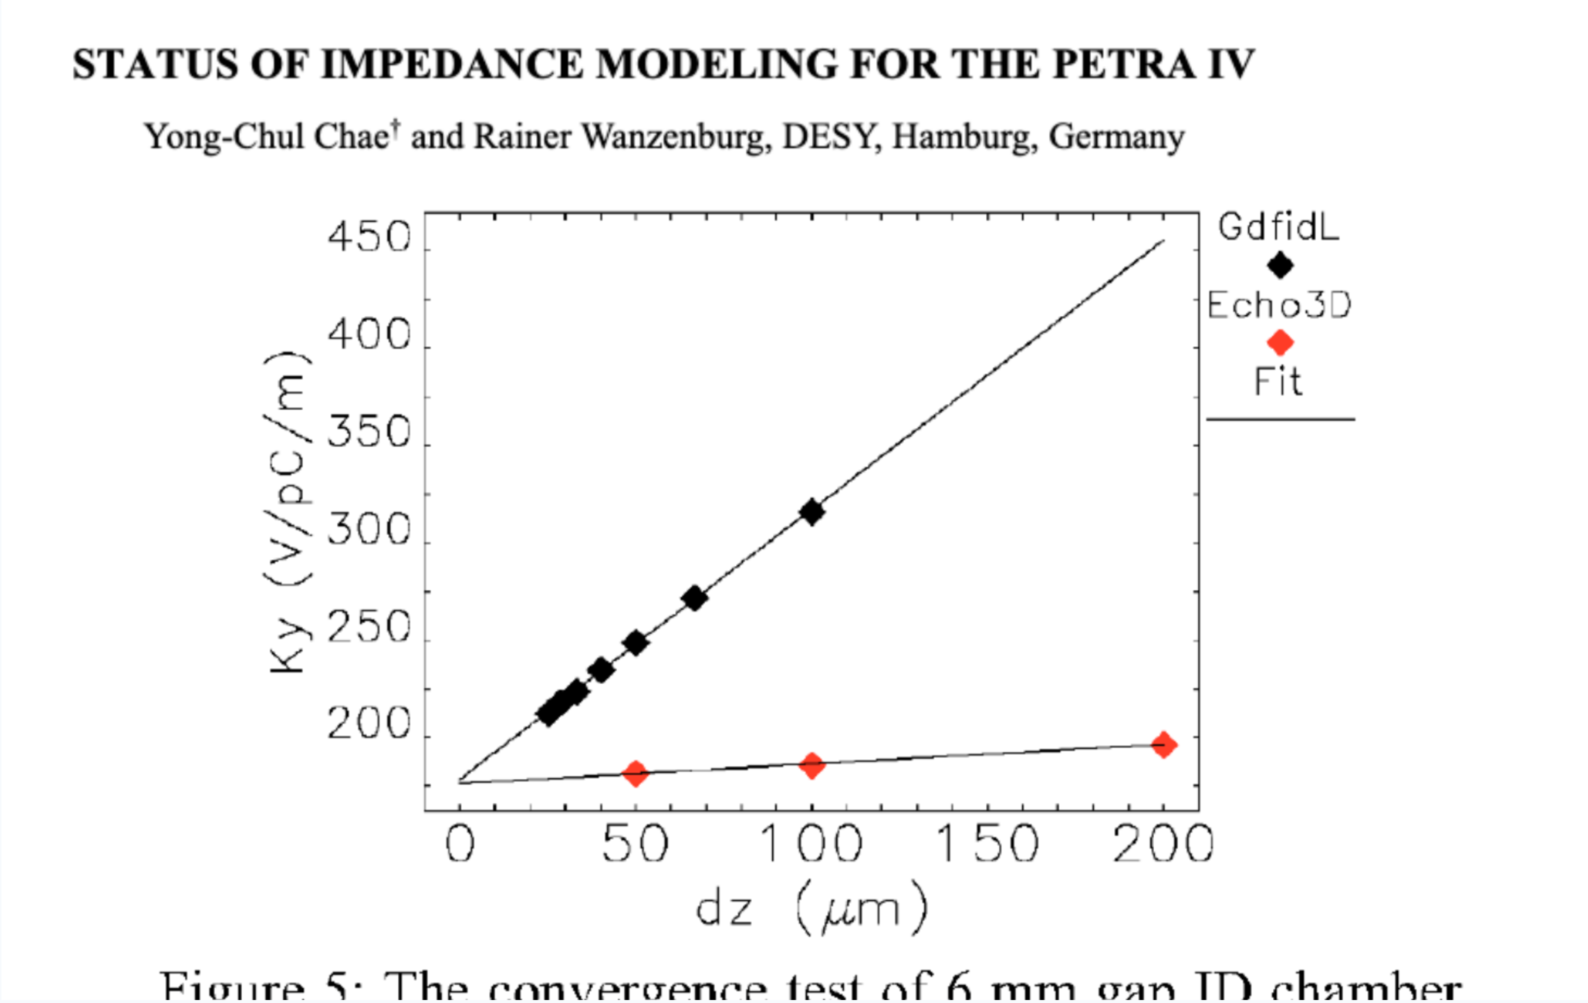

6 mm gap ID chamber

weak

TETM mesh dependence

strong

GdfidL mesh dependence

Independent PETRA IV study

EMUZAG is an independent implementation

Scientific continuity

- Low-dispersion TETM scheme
- Conformal geometry treatment
- Moving-window wake calculation

Independent software implementation

- Written from scratch
- No ECHO3D source code is used
- Developed by Igor Zagorodnov
- Independent engineering practice, registered in 2025

The TETM ideas continue; the EMUZAG code base is entirely new.

EMUZAG turns the old methods into a production code

The numerical core

- Low-dispersion TETM scheme
- Conformal geometry treatment
- Moving-window wake calculation

EMUZAG adds the missing product envelope

- EM/Yee plus stationary and moving TETM
- PML, indirect integration and material losses
- MPI + threads on Cartesian bricks
- CAD, assembly, mesh analysis and postprocessor

Goal: TETM accuracy with GdfidL-class scalability and a modern workflow.

Open boundaries are part of the physical model

Outgoing radiation must leave the domain without returning to the wake monitor.

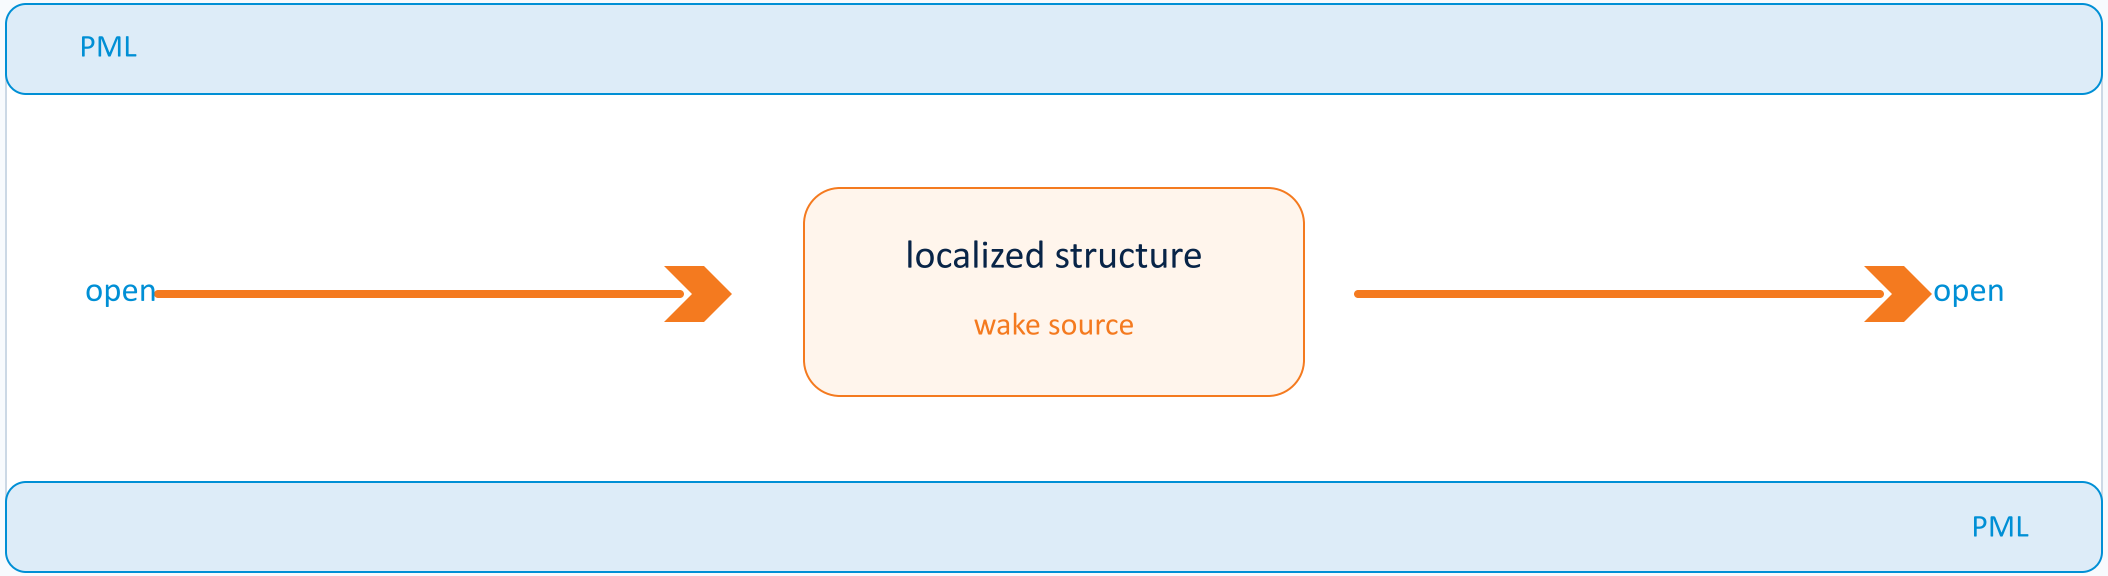

Stationary: longitudinal and transverse PML where required

Moving: transverse PML + causal open head/tail stream

Indirect integration works in both computational frames

The wake can be accumulated directly — and completed from Huygens history.

STATIONARY MESH

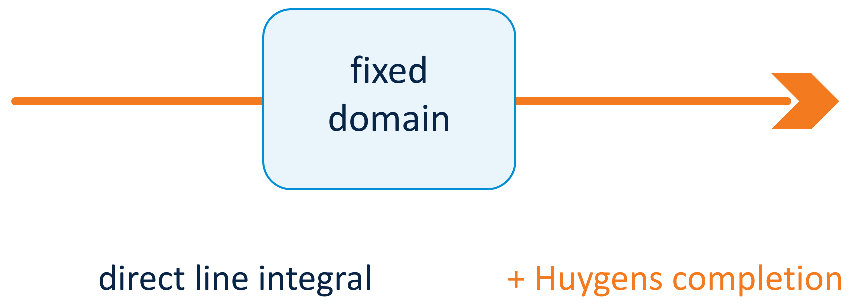

MOVING MESH

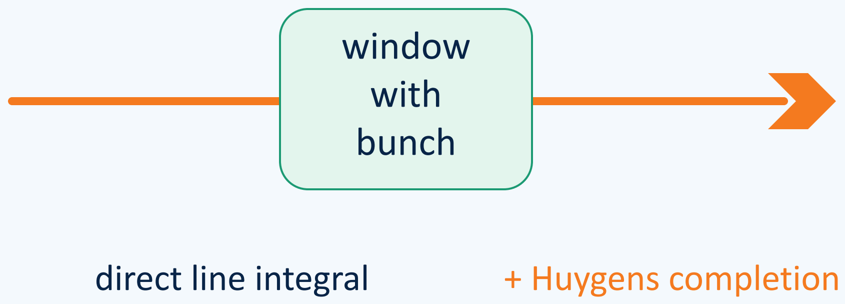

Correct open-space initial field → scattered-field history → indirect wake

A localized Iris can be simulated in open space

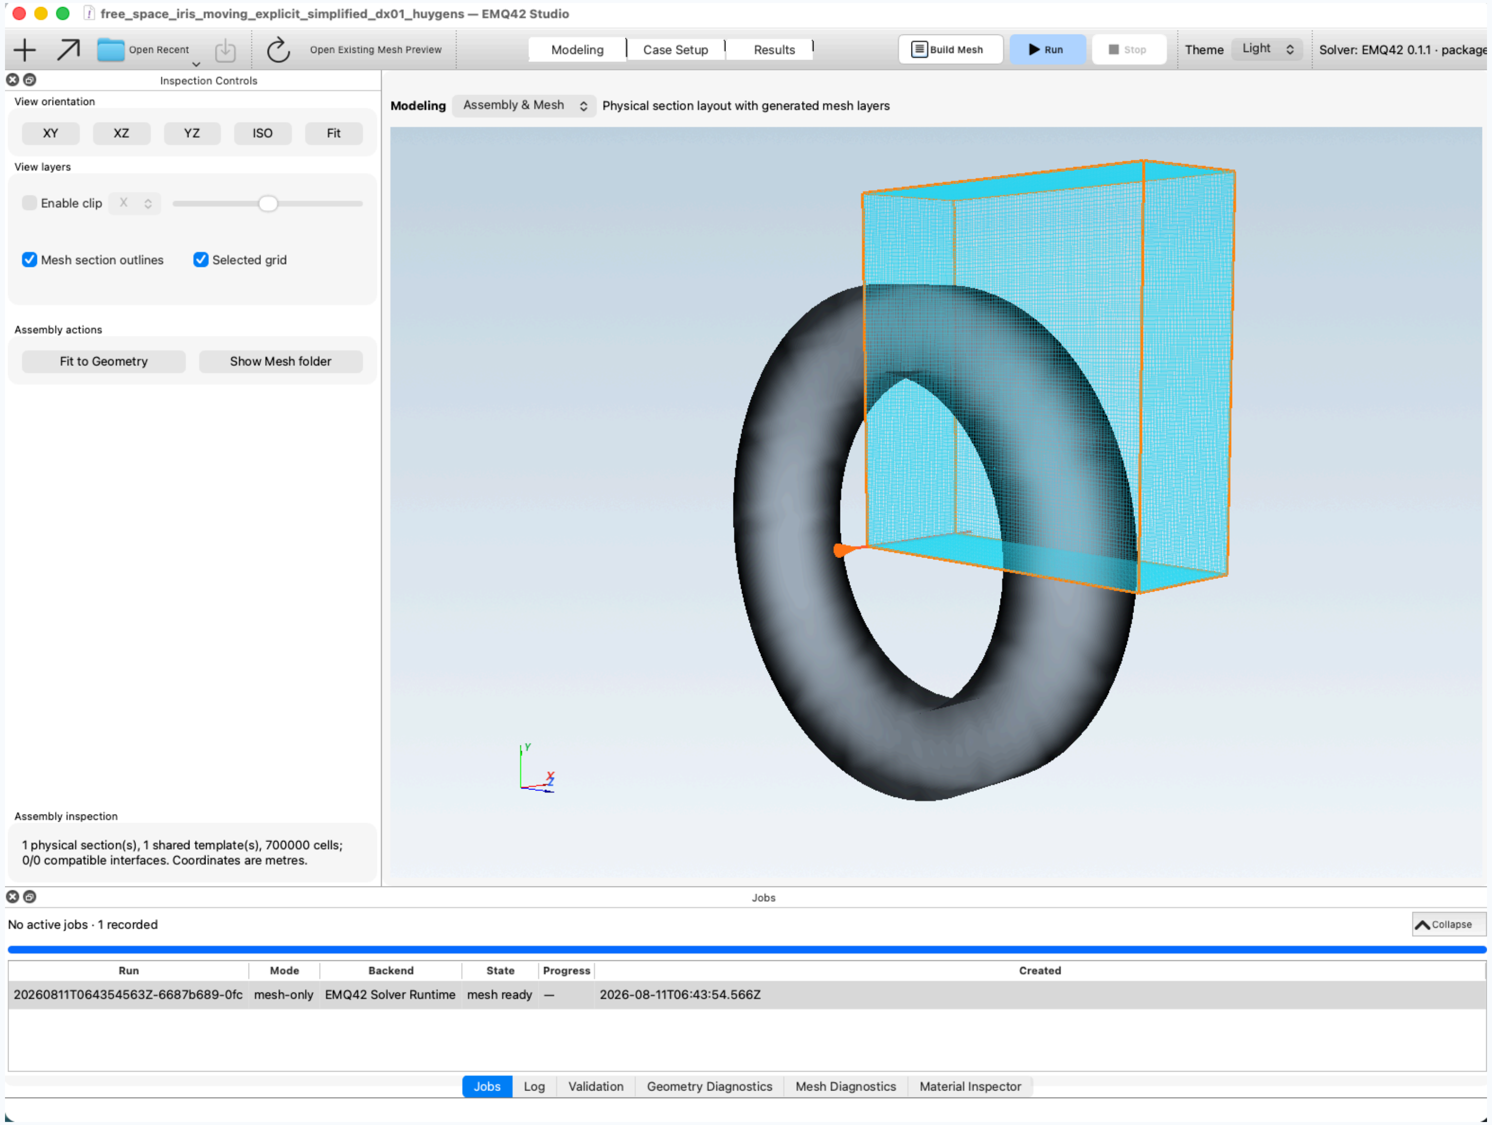

OPEN-SPACE IRIS

80 mm

Huygens-corrected wake

- open/open in x
- transverse PML
- open-space initial field
- stationary EM + TETM
- moving TETM

All three corrected reference outputs match.

A localized Iris can be simulated in open space

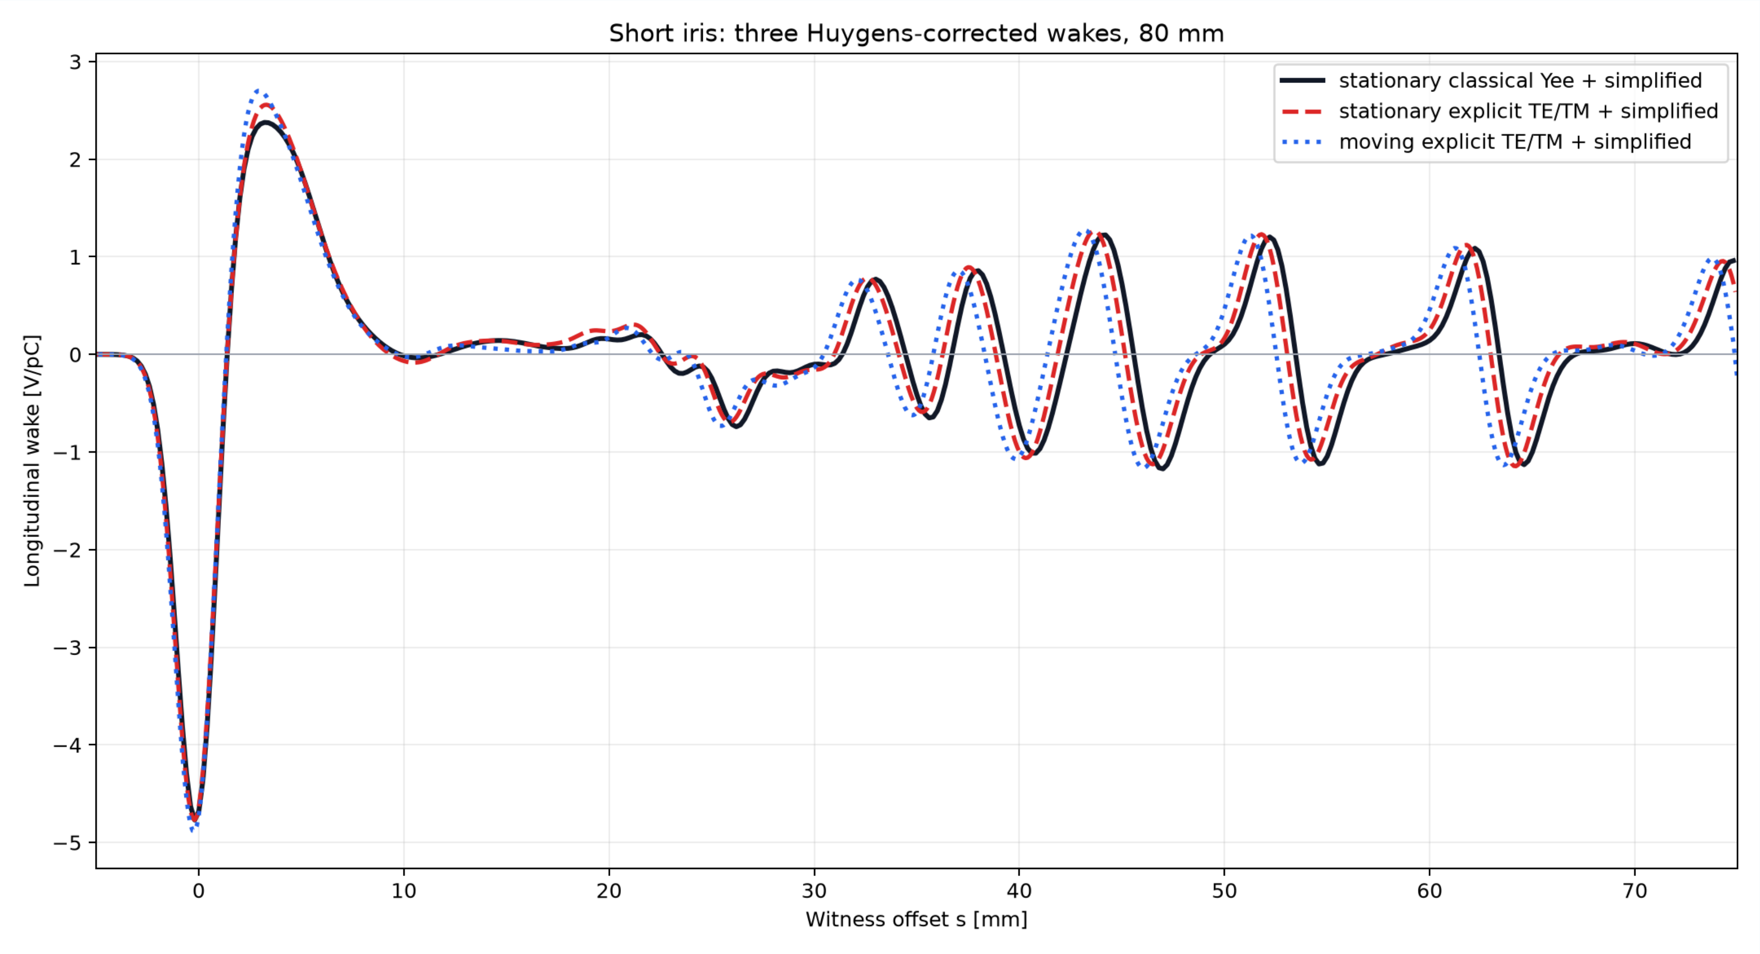

OPEN-SPACE IRIS

80 mm

Huygens-corrected wake

- open/open in x
- transverse PML
- open-space initial field
- stationary EM + TETM
- moving TETM

All three corrected reference outputs match.

Lossy materials belong inside the solve

Wake, impedance and heating must come from the same field history.

LOSSY VOLUME

εr · μr · σbulk

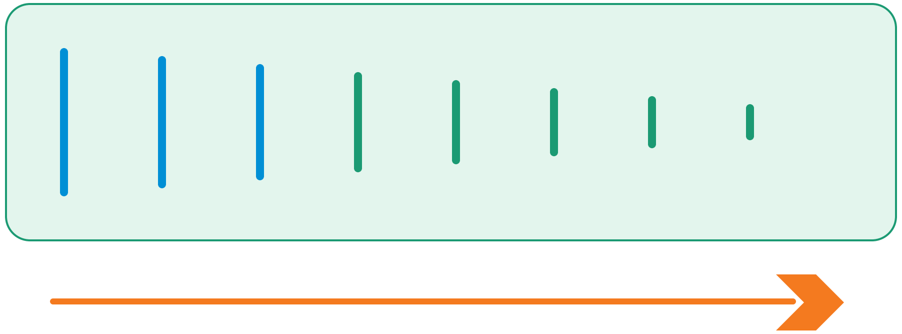

field energy → Joule heat in the volume

HIGH-CONDUCTIVITY WALL

σwall · Zs(ω)

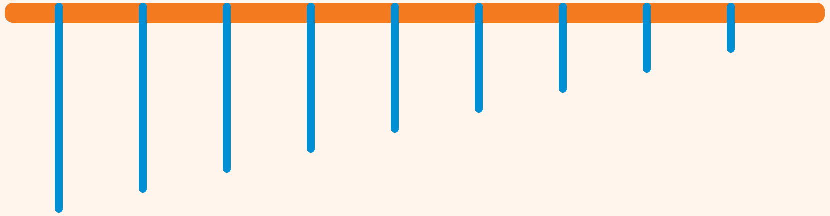

surface response → wall power deposition

Cartesian bricks scale memory and work

Every production run uses the same 3D ownership model.

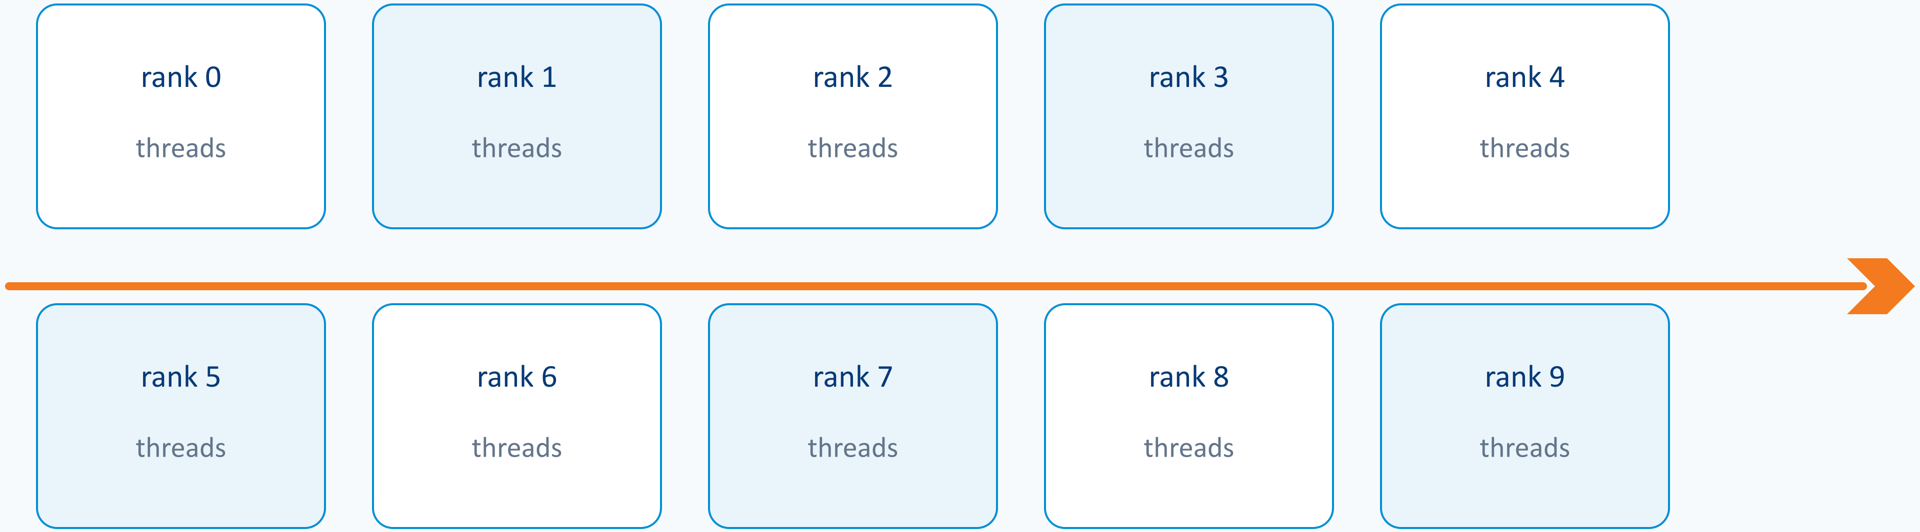

MPI exchanges brick faces · OpenMP works inside each rank · memory follows owned volume

Tapered collimator benchmark: one path across hybrid layouts

Tapered collimator

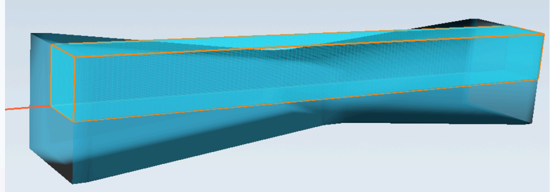

Pure MPI reaches 3.90× at P8

N6 Tapered Collimator · stationary explicit TETM SC · 13.97M field nodes · 2316 steps · b7defce7

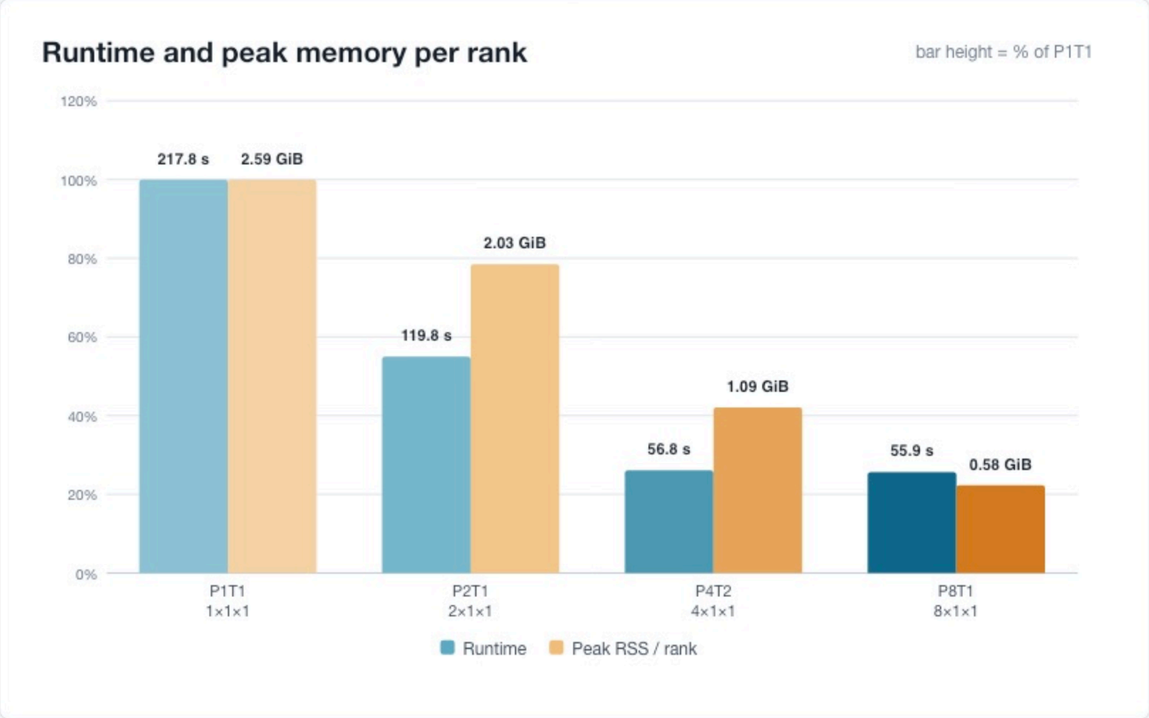

3.90×
MPI speedup
P8T1 8×1×1 vs P1T1

-78%
peak RSS per rank
2.59 → 0.58 GiB

Preparation peak reduced
P8T1 peak falls 0.66 → 0.58 GiB after the fix.
It remains 1.7% faster than P4T2.

Stationary explicit TETM + stationary

The benchmark checks time and memory over Cartesian rank grids - not a special small-case backend.

Same numerical path from 1 rank to cluster

Tapered collimator benchmark: one path across hybrid layouts

Tapered collimator

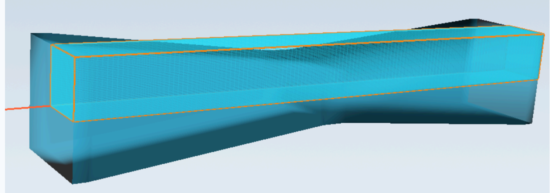

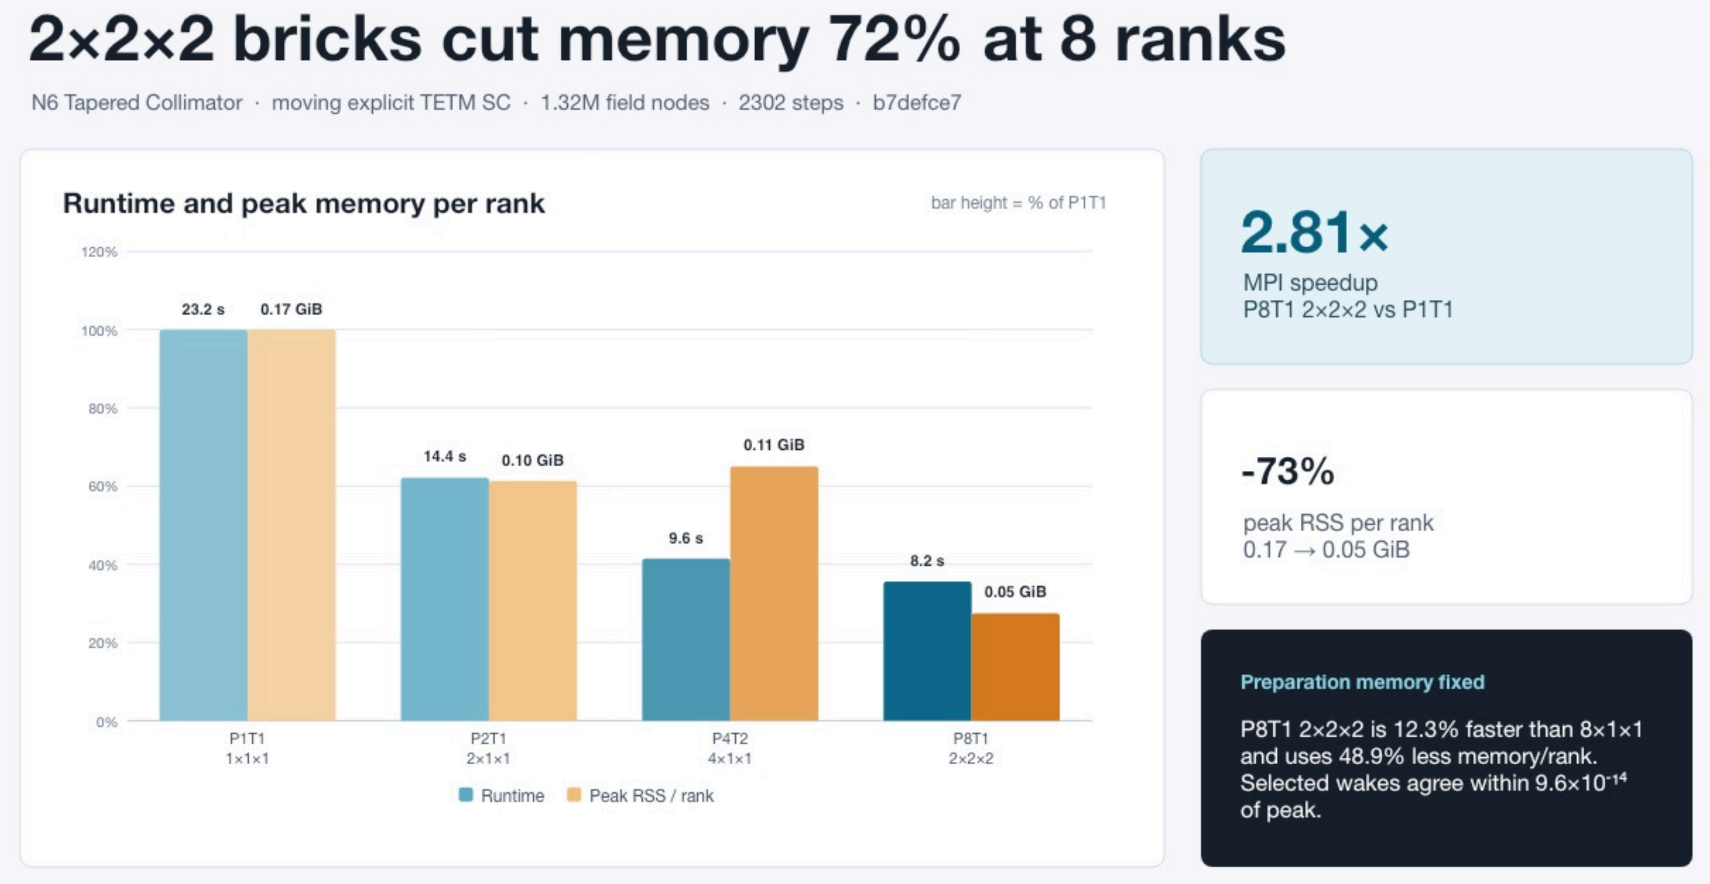

Stationary explicit TETM + moving

The benchmark checks time and memory over Cartesian rank grids - not a special small-case backend.

Same numerical path from 1 rank to cluster

Tapered collimator benchmark: EMUZAG vs ECHO3D

Tapered collimator

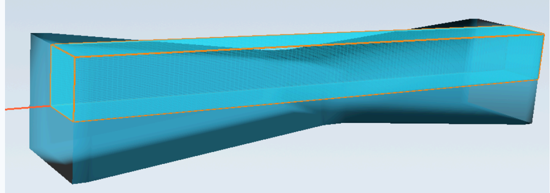

Stationary explicit
TETM + moving

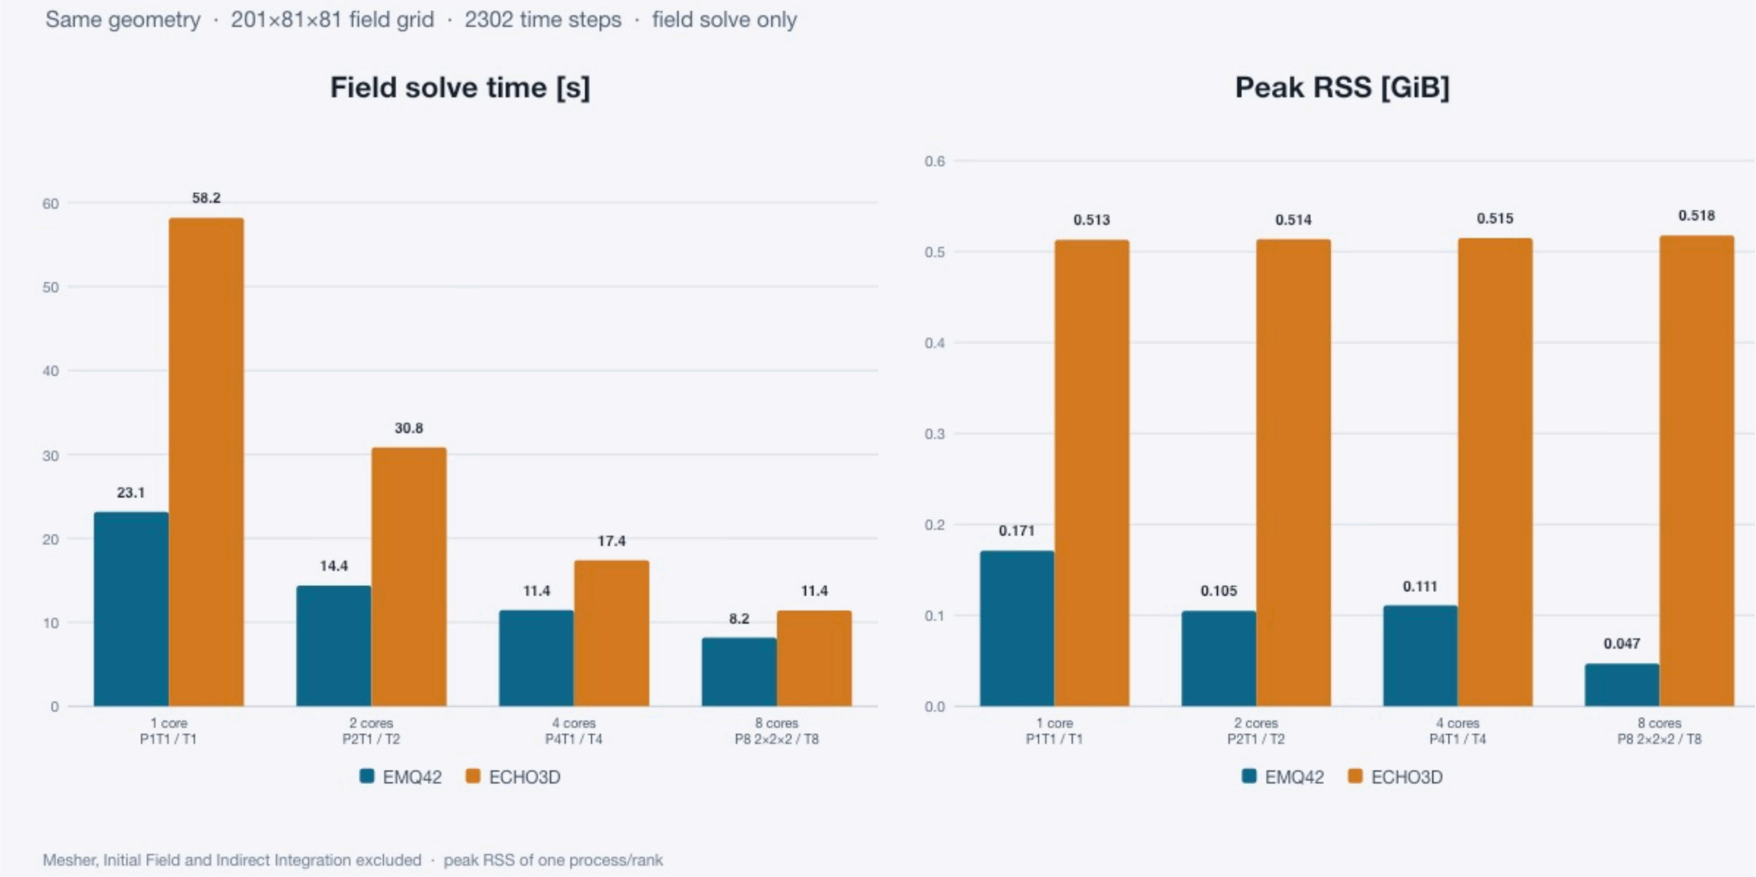

The differentiator is the combination

Capability	CST	GdfidL	ECHO3D	EMUZAG
EM/Yee-type stationary baseline	✓	✓	—	✓
Zero beam-direction dispersion	—	✓	✓	✓
Stationary + moving formulations	stationary only	✓	moving only	✓
Conductivity / lossy materials	✓	✓	—	✓
MPI + threads	✓	✓	threads only	✓
Integrated CAD / setup / results	✓	text/STL	—	✓
Conformal mesh	✓	—	✓	✓

One geometry can contain several materials

VACUUM	DIELECTRIC	LOSSY VOLUME	PEC	CONDUCTIVE WALL
εr = 1, μr = 1	εr, μr	εr, μr, σbulk	ideal conductor	σwall + surface model

Material priority resolves overlaps; the selected coefficients travel with the same Cartesian ownership as the fields.

Ceramic Break is a genuine multi-material 3D case

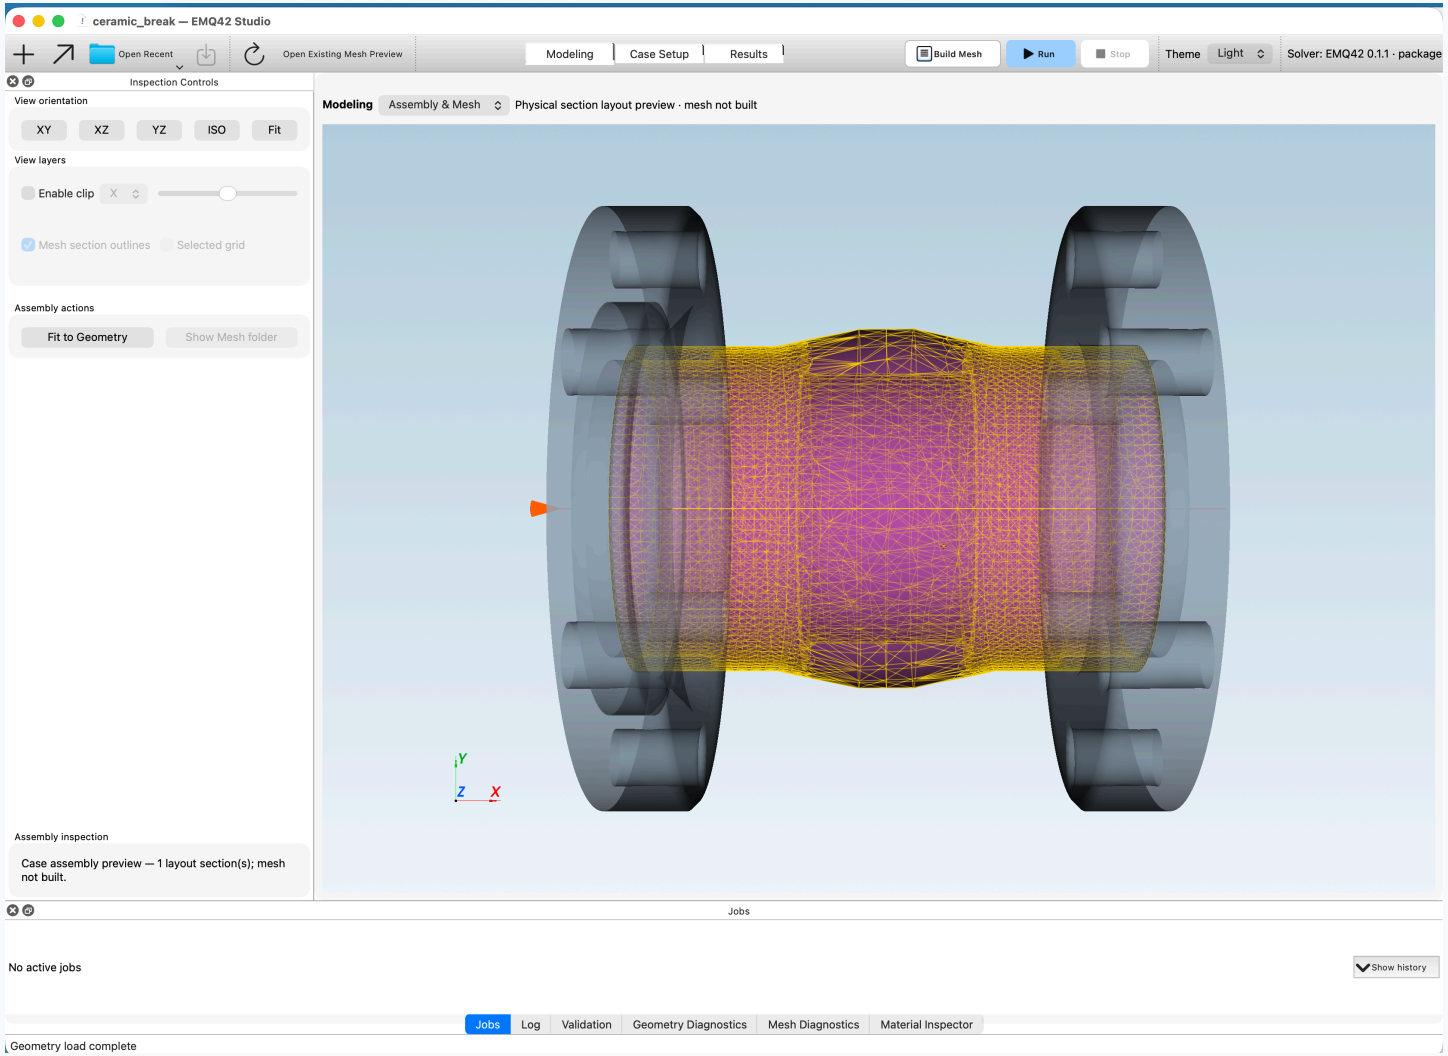

Ceramic body

$\varepsilon r = 11$

Two PEC flanges

moving TETM
150 mm wake

Beam direction →

Ceramic Break separates geometry, material law and solver policy

GEOMETRY

ceramic_body.stl
top_flange.stl
bottom_flange.stl

→

MATERIAL MAP

ceramic: εr = 11
flanges: PEC
background: vacuum

→

SOLVE

moving TETM
implicit + SC
P1 × T4

Add σbulk to the ceramic definition: the passive update damps the wake and accumulates positive Joule energy - without changing CAD or solver policy.

Conductive walls have two complementary models

Both avoid meshing the microscopic skin depth.

ADE SURFACE IMPEDANCE

Auxiliary Differential Equation

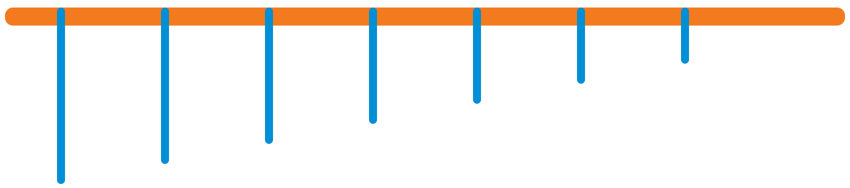

passive fit of Zs(ω) from σ and the bunch spectrum

1D CONDUCTOR LINES

Resolved only through wall depth

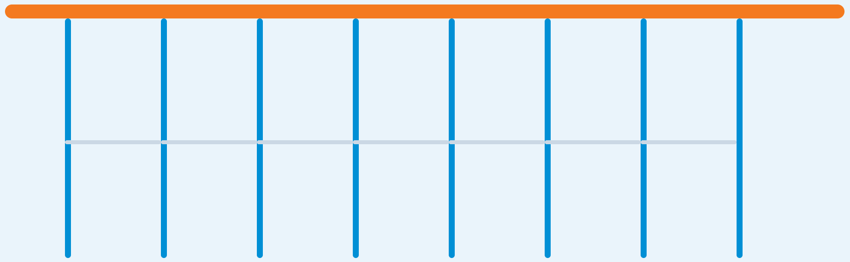

homogeneous or layered wall; conductivity and thickness remain explicit

+ roughness
+ /AC conductivity
+ / Anomalous skin effect

and more

Many layers and transparent metallic walls.

Copper bellows: geometry and material map are inspectable

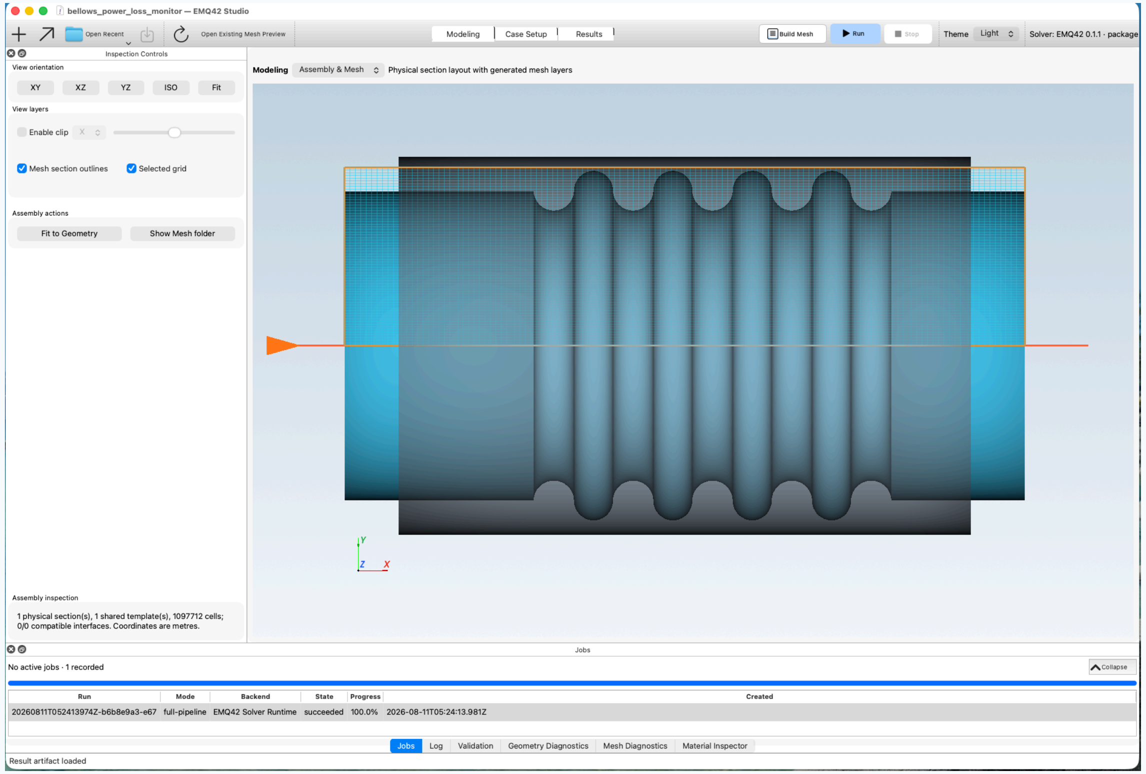

five-valley 3D bellows

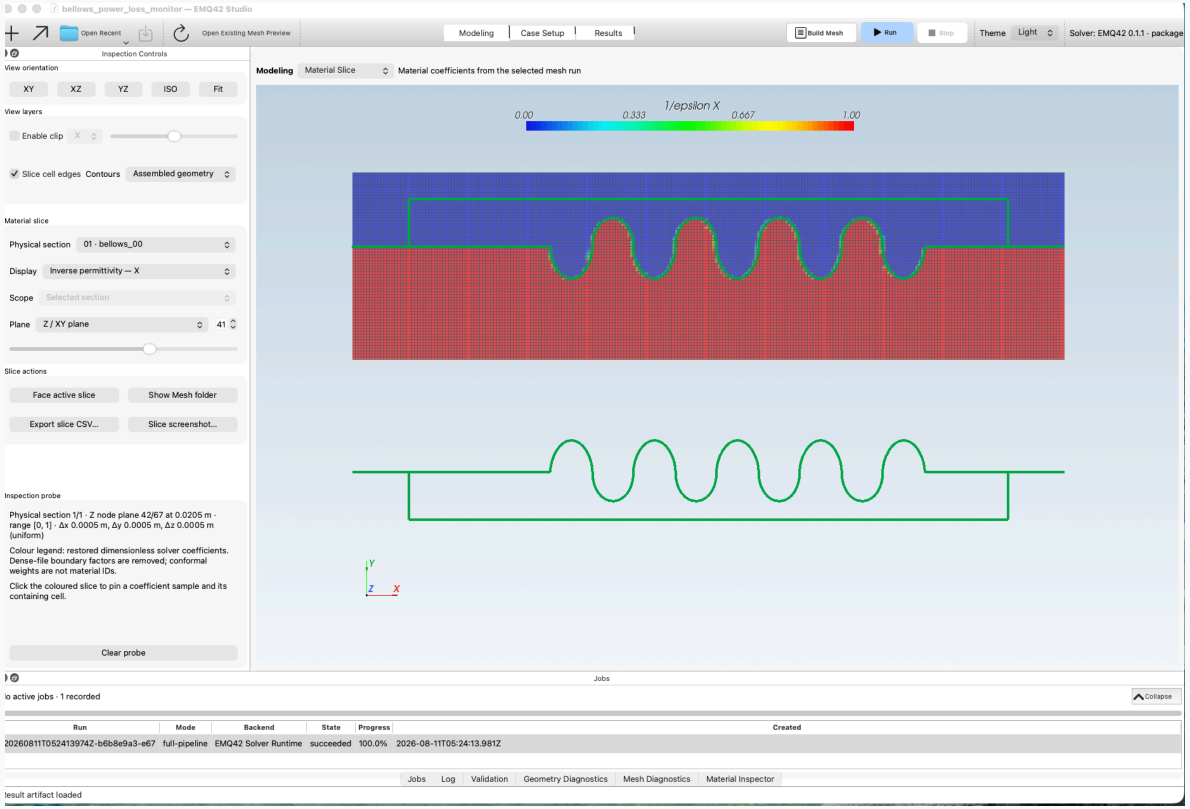

vacuum / conductive-wall interface on the Cartesian mesh

Moving implicit TETM + SC · copper $\sigma = 5.8 \times 10^{7}$ S/m · P1 × T4

Loss Monitor turns wall heating into a traceable result

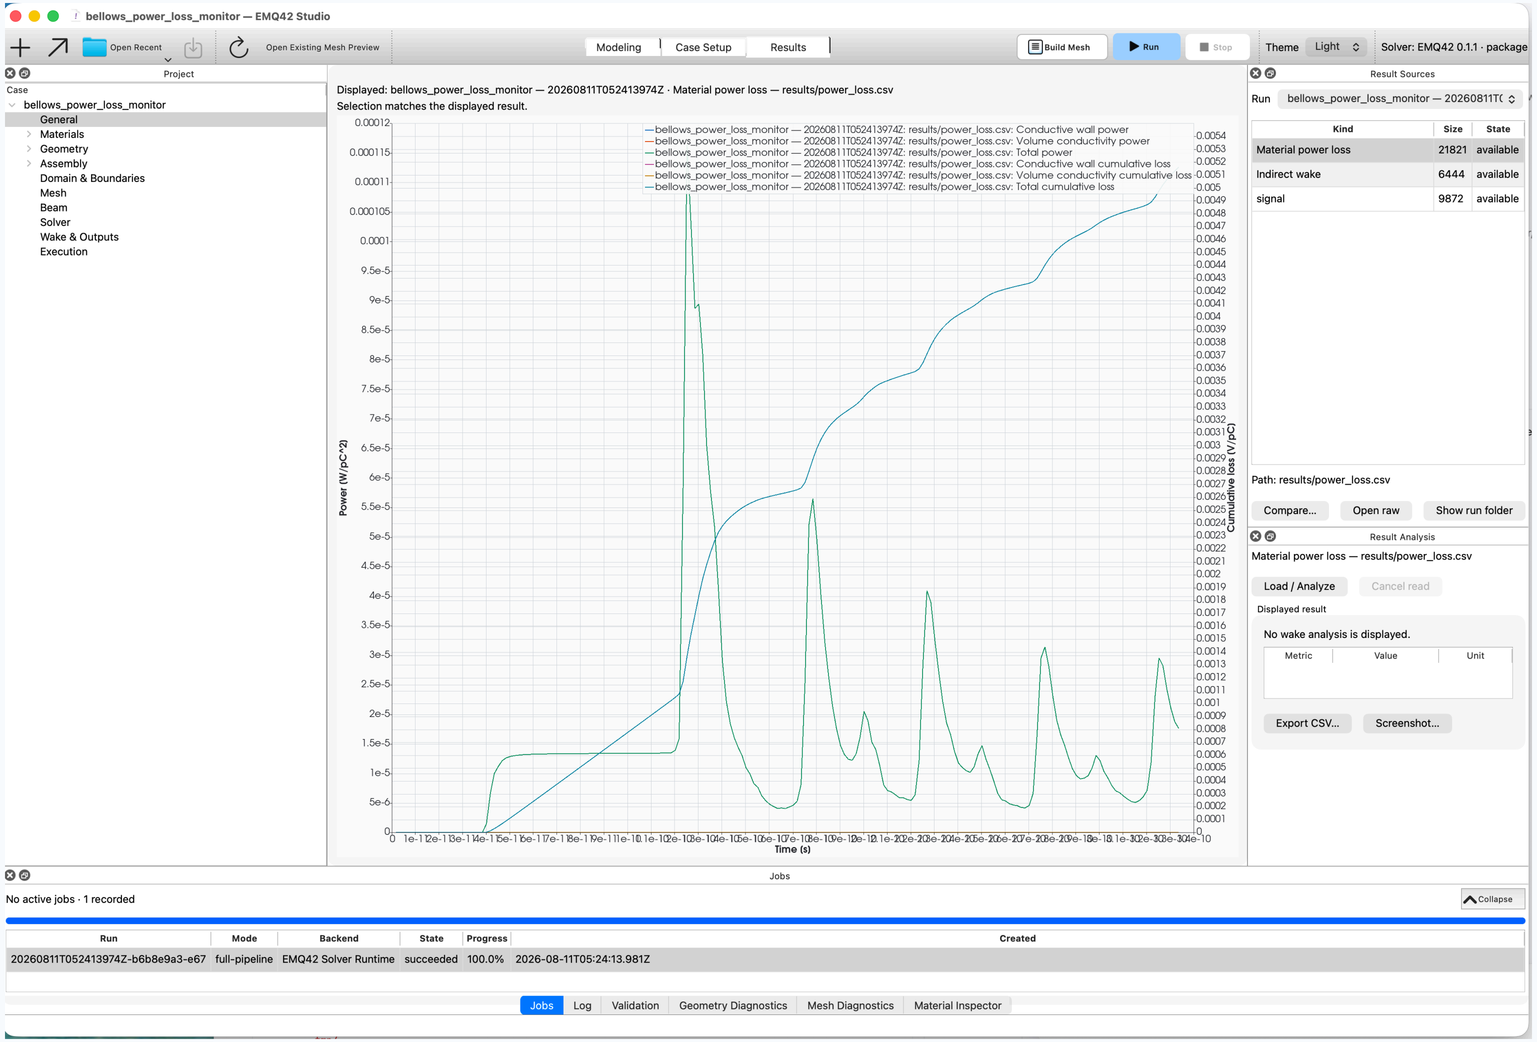

Measured in the run

- instantaneous wall power
- cumulative wall loss
- volume-loss channel
- total power and loss

200 samples
symmetry factor 4

Conformal accuracy is a controlled hierarchy

SC — local diagonal coefficients

iUSC — adds longitudinal weights

full USC — adds transverse weights

→ more nonlocal geometric coupling

Conformal geometry without reducing the time step

New Conformal Scheme* (2002)

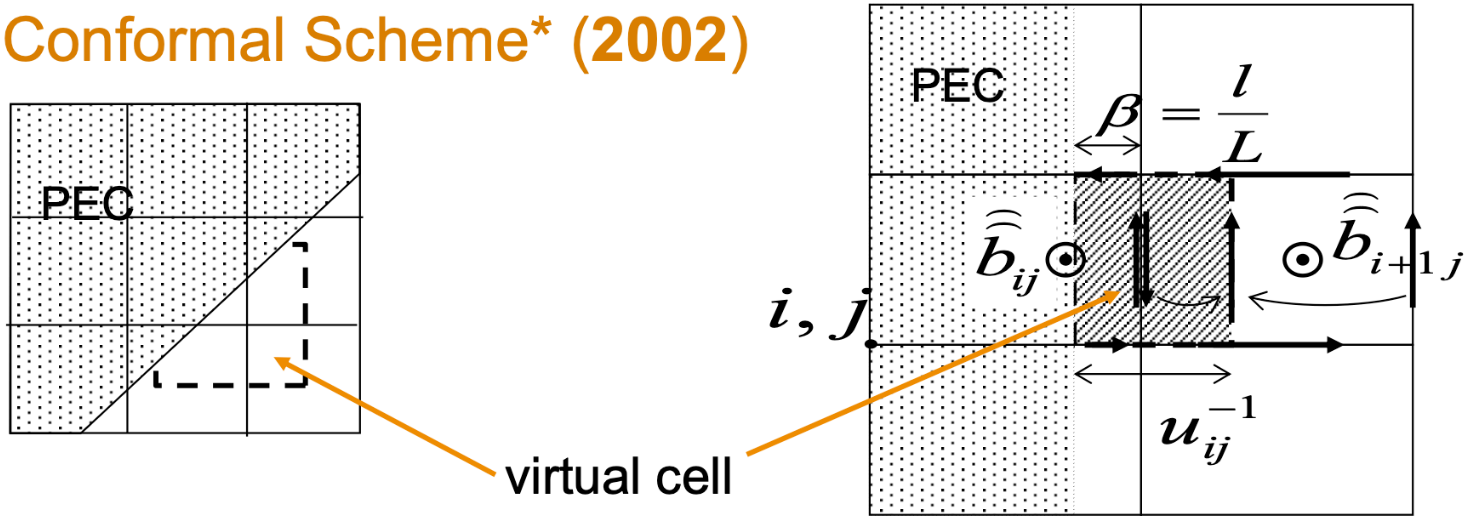

Time step is not reduced!

* Zagorodnov I., Schuhmann R.,Weiland T., *A Uniformly Stable Conformal FDTD-Method on Cartesian Grids*, International Journal on Numerical Modeling, vol. 16, No.2, pp. 127-141 (2003)

Virtual-cell geometry improves curved boundaries while preserving the full stable time step.

Conformal geometry without reducing the time step

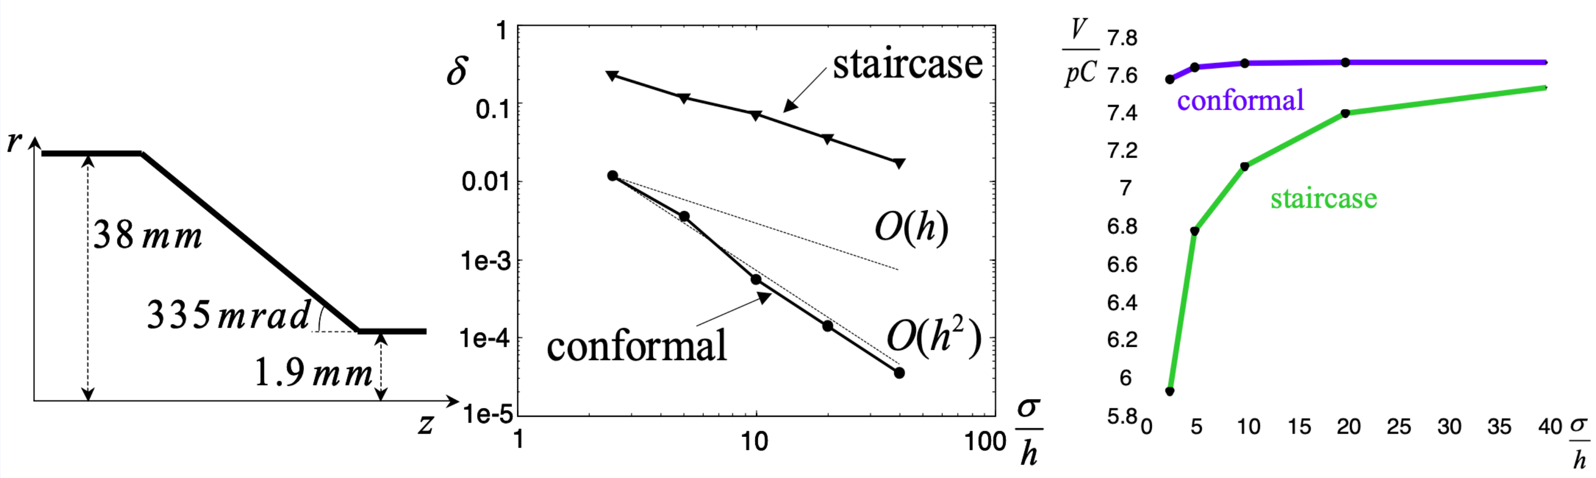

Error in loss factor for a taper $\delta = |L_{calc} - L| L^{-1}$

The error δ relative to the extrapolated loss factor L=-7.63777 Vp/C for bunch with σ=1 mm is shown.

Virtual-cell geometry improves curved boundaries while preserving the full stable time step.

Fields and materials enter only at the moving head

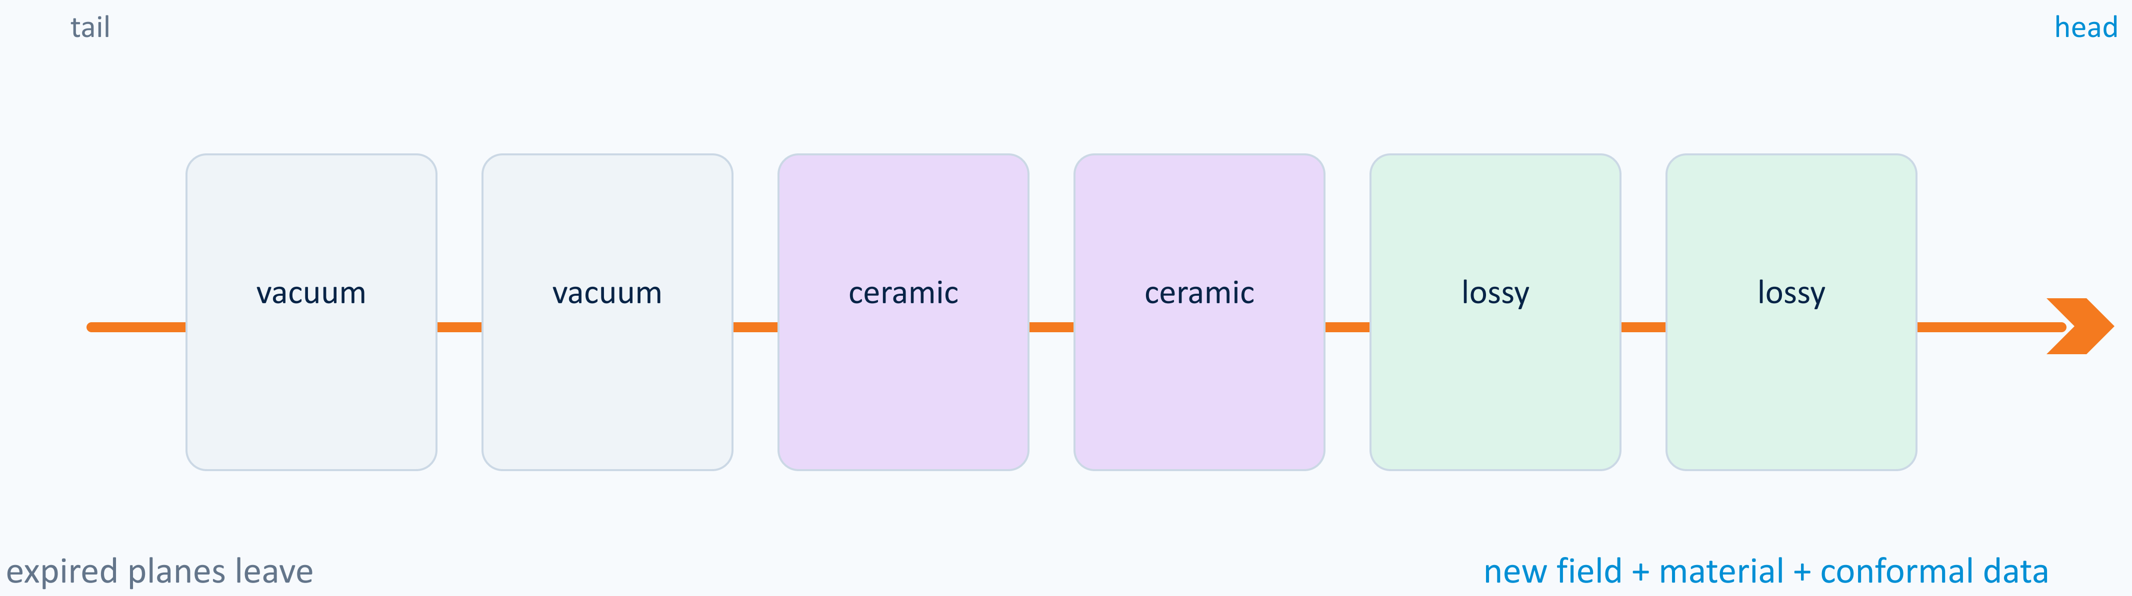

Circular buffering changes storage addresses - never physical causality.

Sections and repeats make long structures practical

A repeated structure should not require repeated authoring.

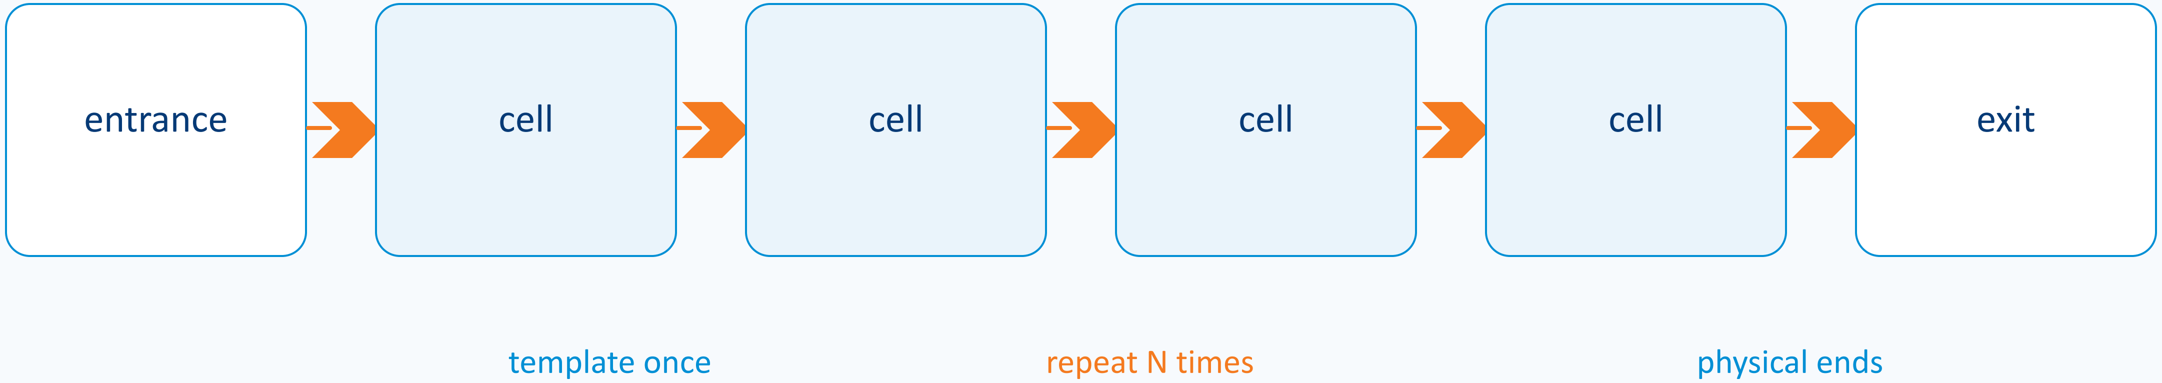

The solver streams or assembles physical sections while preserving material order and interface compatibility.

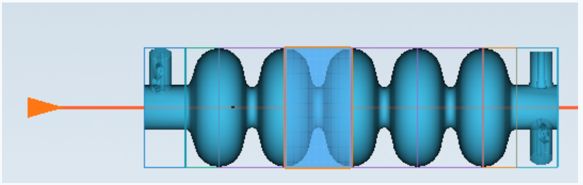

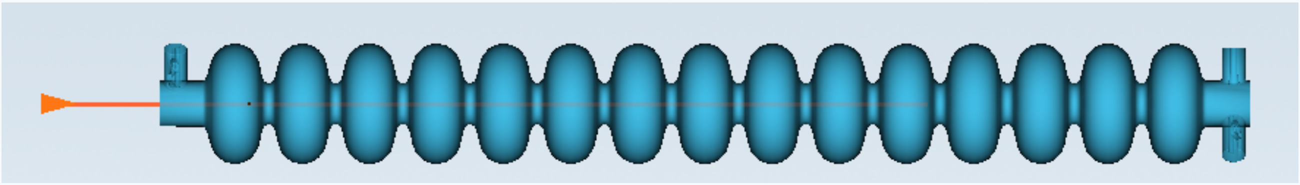

One case targets laptop, workstation and cluster

PERSONAL COMPUTER

macOS · Linux · Windows

MPI ranks + threads

Px × Py × Pz bricks

WORKSTATION

many cores + memory

MPI ranks + threads

Px × Py × Pz bricks

CLUSTER

distributed memory

MPI ranks + threads

Px × Py × Pz bricks

The same parallel model of all

Studio makes the workflow repeatable

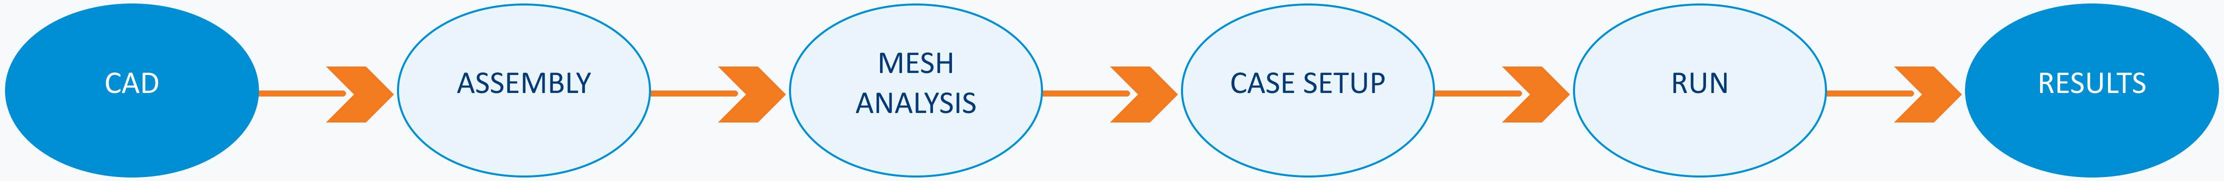

Model once

Inspect geometry and mesh before spending cluster time.

Run reproducibly

A case records method, source, boundaries, materials and parallel layout.

Interpret quickly

Wake, impedance, fields and diagnostics stay connected to the model.

Studio or console application

CAD and assembly expose a full cavity with couplers

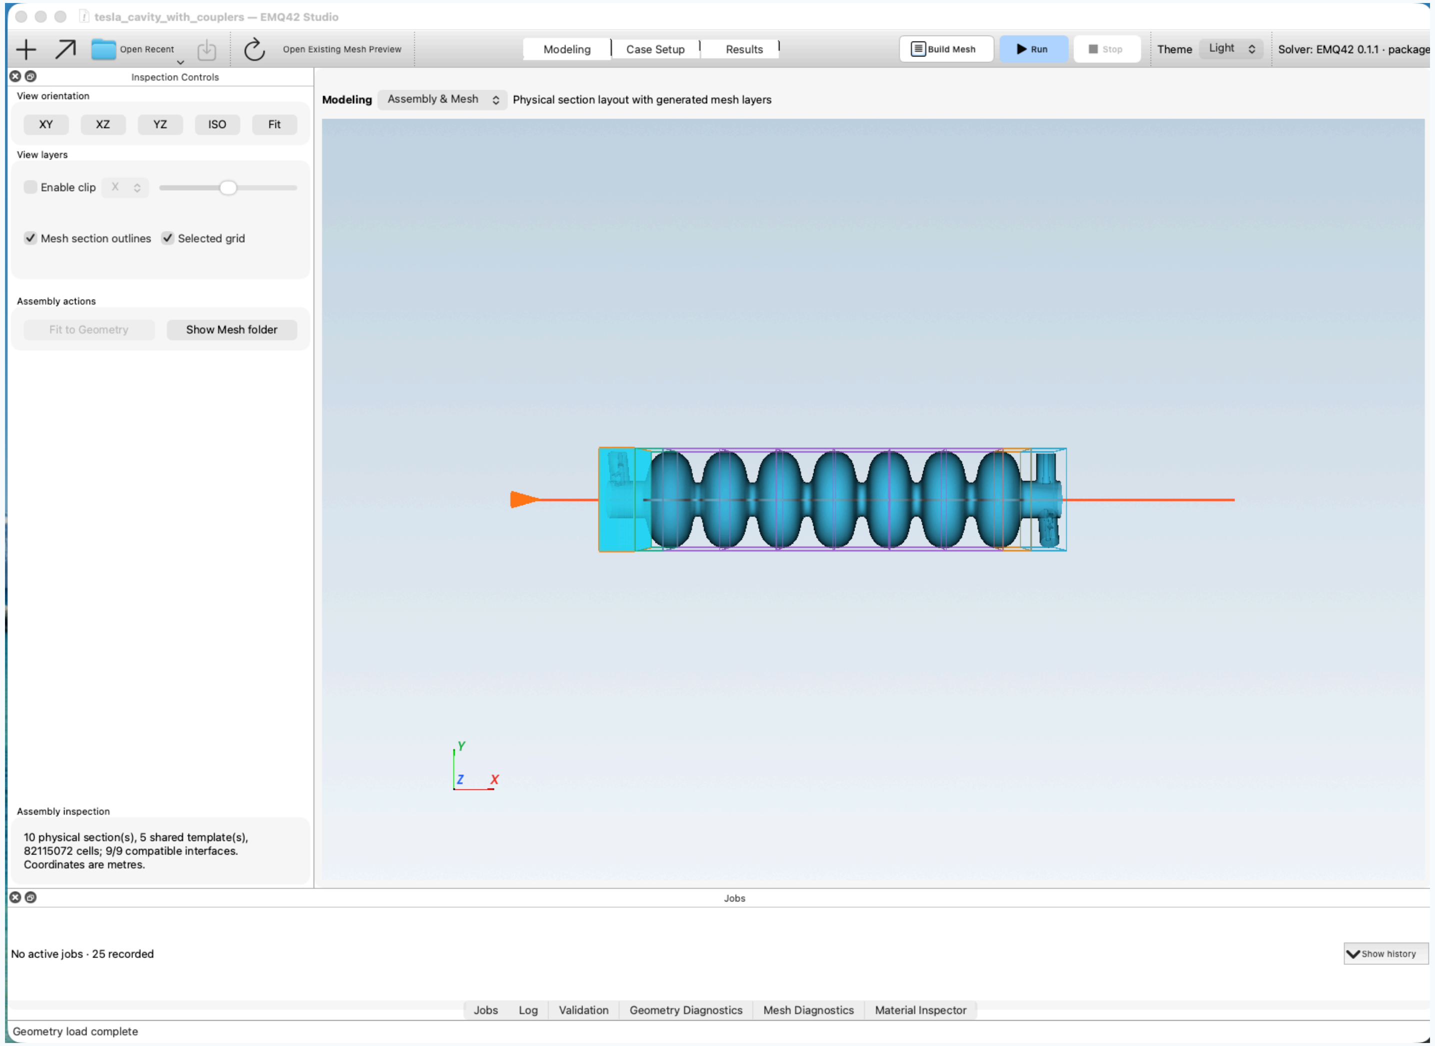

Tesla cavity with couplers

- Imported 3D geometry
- Physical-section assembly
- Mesh-layer compatibility
- Beam axis shown left → right

Case Setup makes the numerical policy explicit

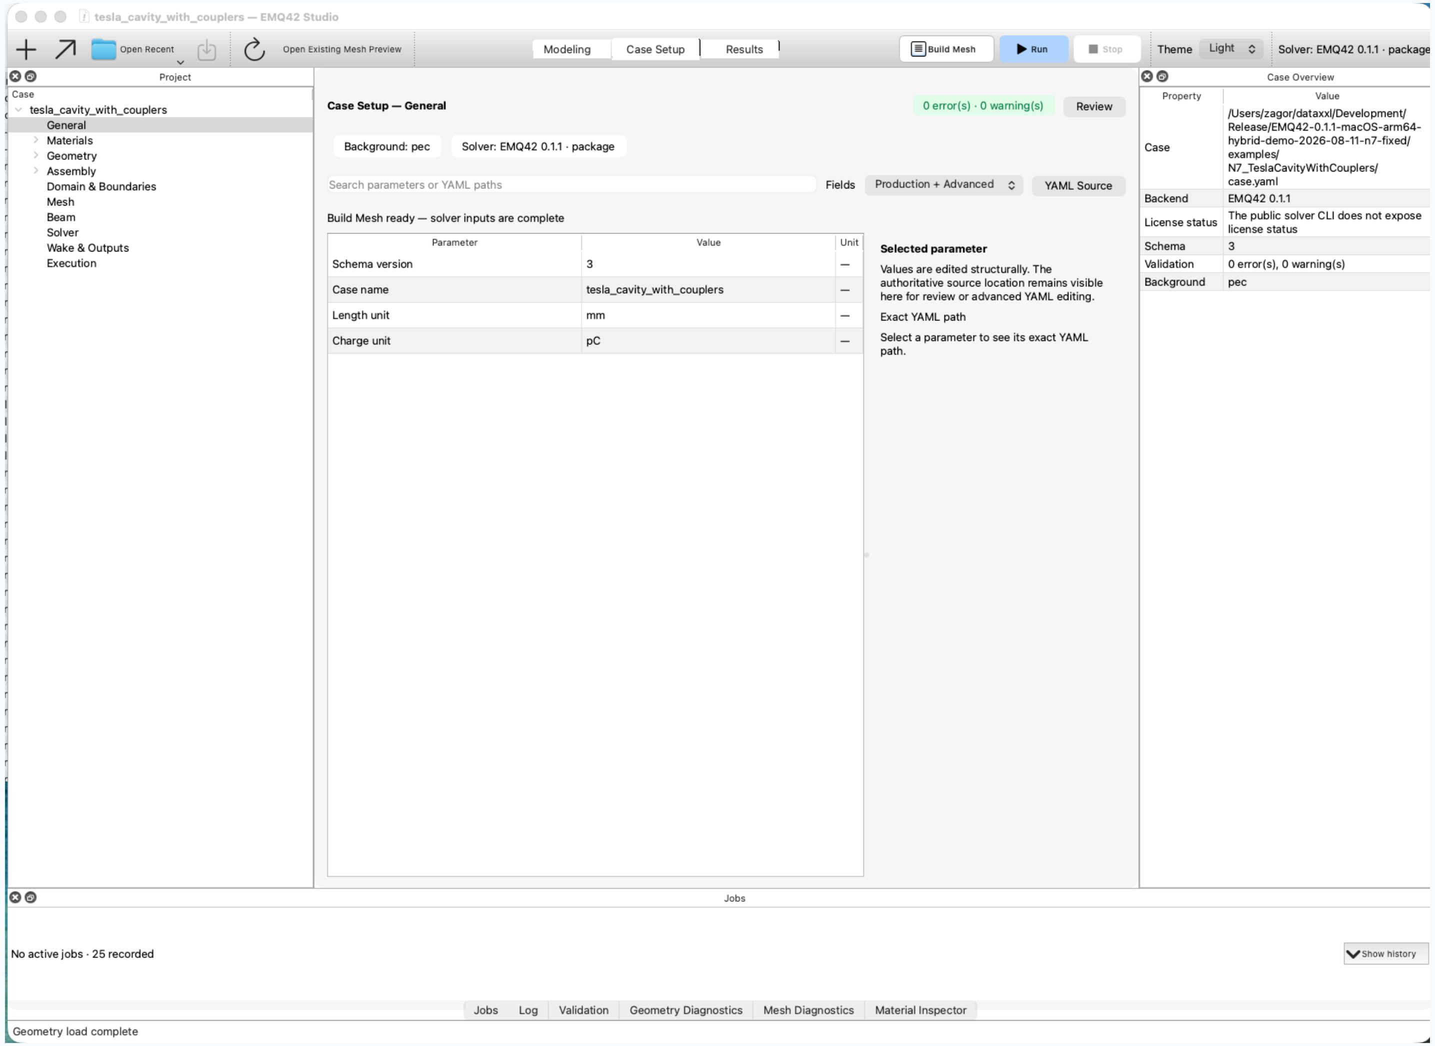

A case records

- EM or TETM method
- stationary or moving frame
- source and wake length
- materials and boundaries
- MPI × thread layout

The postprocessor connects fields to impedance

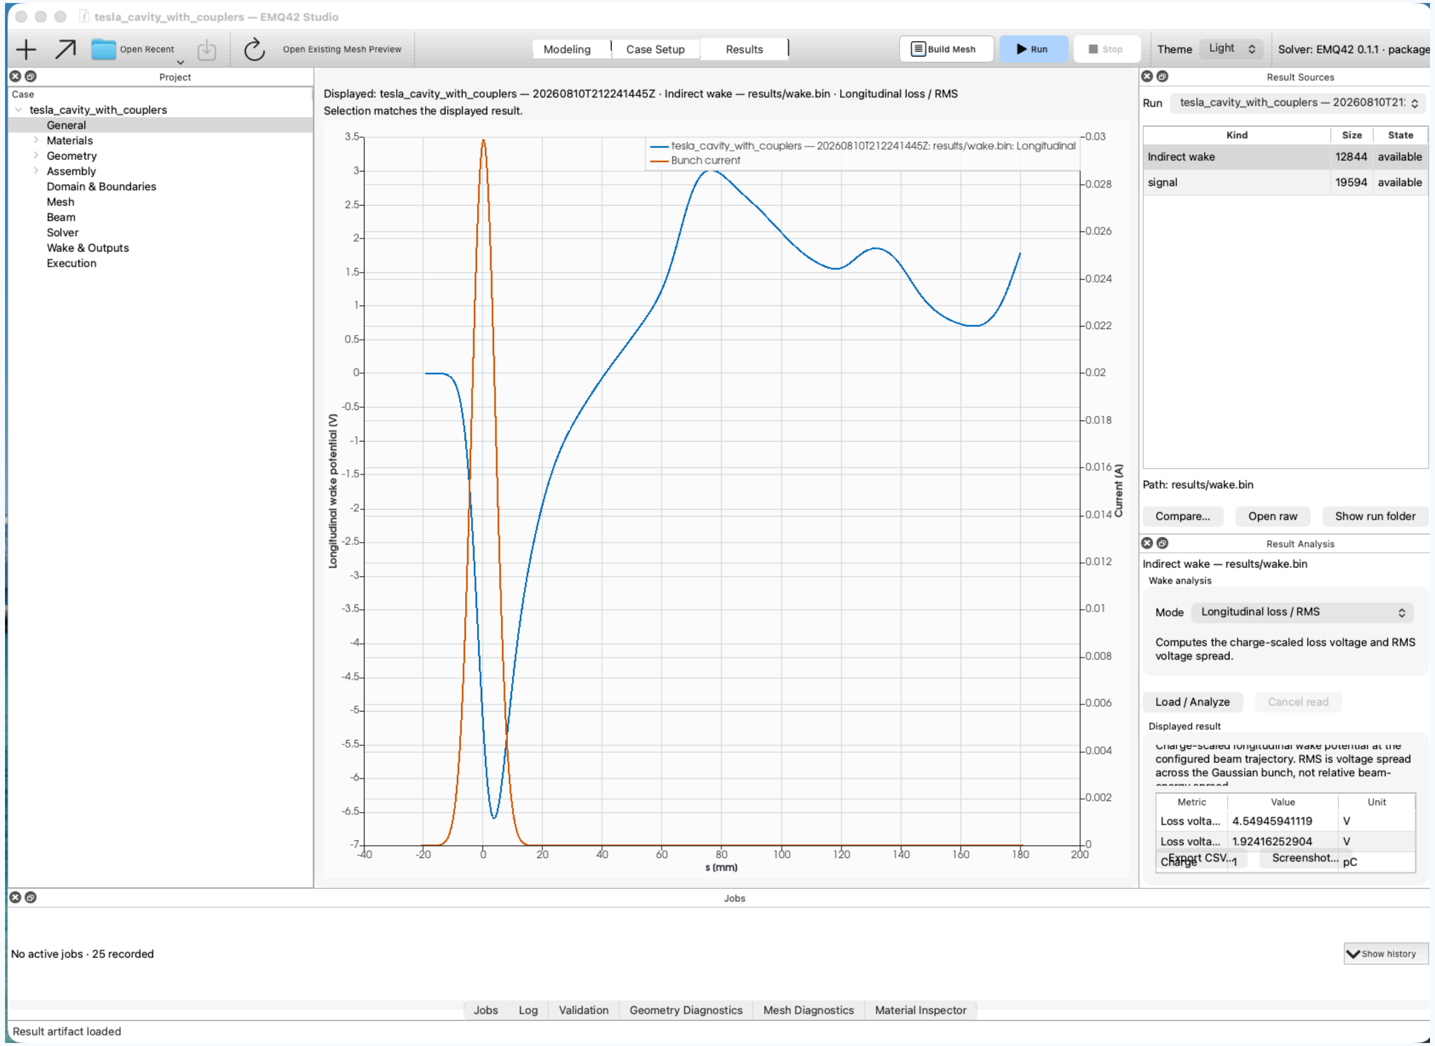

One result context

- Wake potentials
- Longitudinal and transverse impedance
- Field monitors and slices
- Run log and diagnostics

No detached plotting workflow is required for first interpretation.

The evaluation should answer three concrete questions

1 **Accuracy**

Does EMUZAG reproduce trusted wakes on a coarser grid?

2 **Coverage**

Can one workflow cover stationary, moving, materials and full 3D geometry?

3 **Scale**

Does the same case move cleanly from workstation to hybrid cluster?

www.emuzag.de/www.emuzag.com

Request an evaluation license