



EMWSD / Wakis

Electromagnetic and Wake Solver
Development

Meeting #20

Elena de la Fuente, Carlo Zannini, Lorenzo Giacomel, Giovanni Iadarola

Wake solver milestones:



meetings [#17](#), [#18](#), [#19](#), [CEI](#)

0. Wake potential and impedance

1. FIT Maxwell Equations in 3D

2. PEC, Periodic, PMC boundaries

3. Embedded Boundaries: geometry import from **.stl** files

4. Beam injection J_z

5. Materials: ϵ, μ, σ

In progress

x. Documentation

6. on GPU & memory opt.

7. Low beta & SC

8. Open boundaries & PML

9. Numerical tests vs analytic

To Do

10. Leontovich condition for good conductors >100 S/m

11. Staircased geometry: PBA/grid refinement

12. Moving window solver (?)

13. Frequency-dependent materials

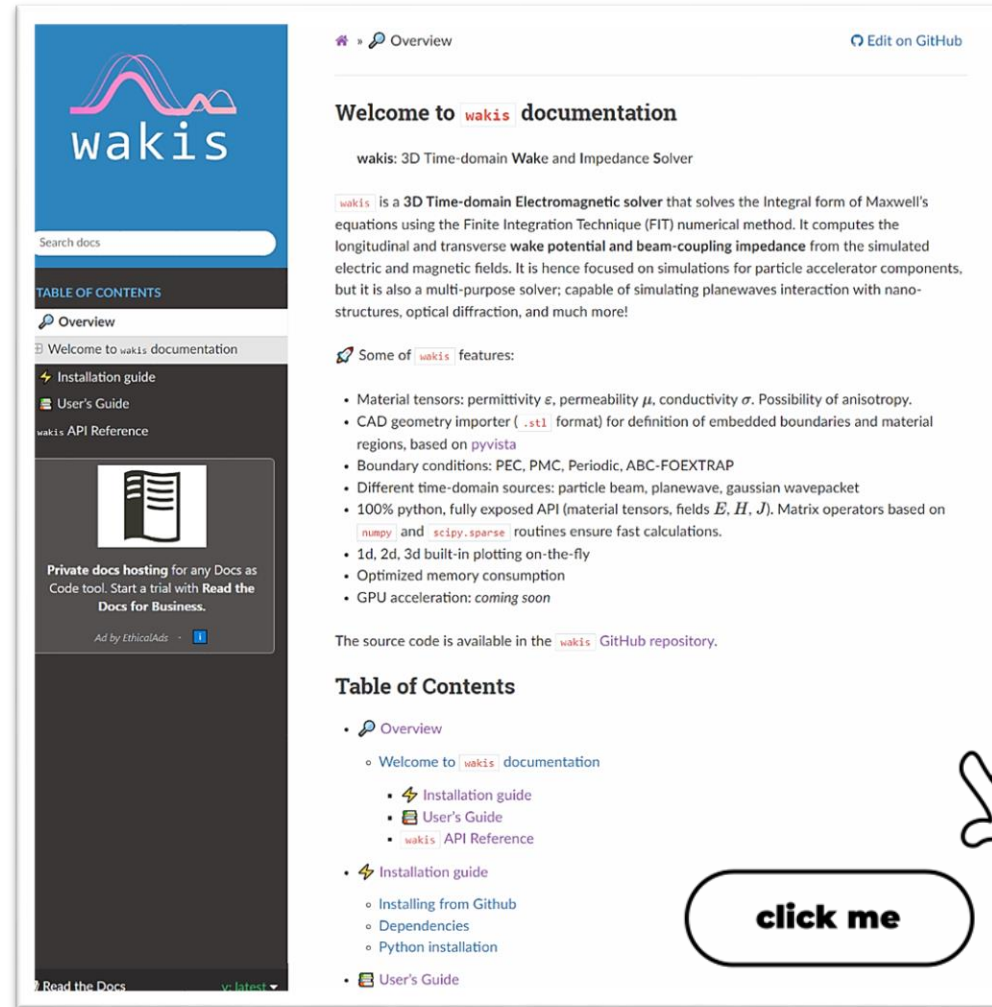
14. ...

Documentation

So far, positive feedback from first users:

- Elena M.
- Josephine P.
- Sebastien J. (BESY)
- PyVista people via Slack 😊

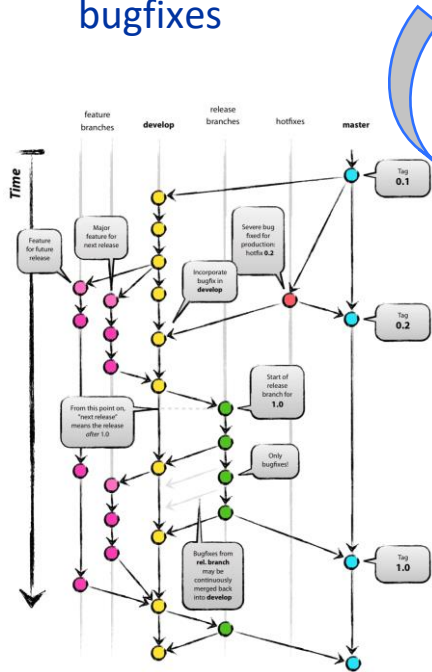
Feedback helped solve installation problems for interactive plots in remote machines. Installation stable for all type of machines in `requirements.txt`



The screenshot displays the 'wakis' documentation website. The left sidebar contains a search bar and a 'TABLE OF CONTENTS' menu with items: Overview, Welcome to wakis documentation, Installation guide, User's Guide, and wakis API Reference. The main content area is titled 'Welcome to wakis documentation' and describes 'wakis: 3D Time-domain Wake and Impedance Solver'. It includes a list of features such as material tensors, CAD geometry importer, boundary conditions, and time-domain sources. A 'Table of Contents' section is also visible, listing 'Overview', 'Welcome to wakis documentation', 'Installation guide', and 'User's Guide'. A hand icon points to a 'click me' button in the bottom right corner of the screenshot.

GitHub strategy

- Stable branch `'main'` for user's to fork
- `'develop'` branch with the latest features
- Individual branches for features
- **cherry-pick** to main for bugfixes

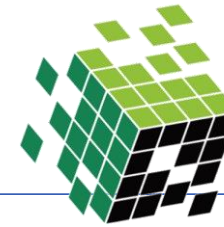


The screenshot shows the GitHub repository page for 'wakis' (Public). The repository has 6 branches and 0 tags. The 'develop' branch is selected. A 'Switch branches/tags' dialog is open, showing a list of branches: main (default), badapple, develop (checked), gpu, noSpeedUp, and speedUp. The commit history is visible, showing a list of commits with their titles and dates. The repository description is: "3D electromagnetic time-domain solver, specialized in wake potential and beam-coupling impedance computation for particle accelerators". The repository has 1 star, 1 fork, and 0 watchers. The repository is updated on Feb 9, 2023.

Old `'wakis'` postprocessor was renamed `'-legacy-wakis'`

The screenshot shows the GitHub repository page for `'-legacy-wakis'` (Public). The repository description is: "Python tool for Wake potential and Impedance calculations". The repository has 1 star, 1 fork, and 0 watchers. The repository is updated on Feb 9, 2023.

GPU implementation



CuPy



NVIDIA

CUDA

User side: Enable `use_gpu=True` in Solver instantiation

Code side:

wakis / field.py

elenafuengar bugfix in inspect3d

Code Blame 430 lines (348 loc) · 14.9 KB

```
1 import numpy as xp
2
3 try:
4     import cupy as xp_gpu
5     imported_cupy = True
6 except ImportError:
7     imported_cupy = False
8
9 class Field:
10     '''
11     Class to switch from 3D to collapsed notation by
12     defining the __getitem__ magic method
13
14     linear numbering:
15     n = 1 + (i-1) + (j-1)*Nx + (k-1)*Nx*Ny
16     len(n) = Nx*Ny*Nz
17     '''
18     def __init__(self, Nx, Ny, Nz, dtype=float,
19                 use_ones=False, use_gpu=False):
```

```
if self.on_gpu:
    return self.array[key].get()
else:
    return self.array[key]
```

wakis / solverFIT3D.py

elenafuengar small bugfixes in emsolve (close h5) and wakesolve (add_space=0)

Code Blame 1434 lines (1191 loc) · 56.1 KB

```
1 from tqdm import tqdm
2
3 import numpy as np
4 import time
5
6 from scipy.constants import c as c_light, epsilon_0 as eps_0, mu_0 as mu_0
7 from scipy.sparse import csc_matrix as sparse_mat
8 from scipy.sparse import diags, hstack, vstack
9
10 from field import Field
11 from materials import material_lib
12
13 try:
14     from cupyx.scipy.sparse import csc_matrix as gpu_sparse_mat
15     imported_cupyx = True
16 except ImportError:
17     imported_cupyx = False
18
19 class SolverFIT3D:
20
21     def __init__(self, grid, wake=None, cfln=0.5, dt=None,
22                 bc_low=['Periodic', 'Periodic', 'Periodic'],
23                 bc_high=['Periodic', 'Periodic', 'Periodic'],
24                 use_conductors=False, use_stl=False, use_gpu=False,
25                 h5=[1.0, 1.0], verbose=1):
```

Check
Here



```
# Move to GPU
if use_gpu:
    if verbose: print('Moving to GPU...')
    if imported_cupyx:
        self.tDsiDmuiDaC = gpu_sparse_mat(self.tDsiDmuiDaC)
        self.itDaiDepsDstC = gpu_sparse_mat(self.itDaiDepsDstC)
        self.iDeps = gpu_sparse_mat(self.iDeps)
        self.Dsigma = gpu_sparse_mat(self.Dsigma)
    else:
        print('*** cupyx could not be imported, please check CUDA installation')

if verbose: print(f'Total initialization time: {time.time() - t0} s')
```



GPU implementation (II)



Tested it in **pcbe-abp-gpu001** (thanks Gianni 😊):

- Pushing the limits, we can go up to 20,000,000 cells, takes 20 minutes
- New bottleneck: **memory!**
- Scales linearly with mesh (e.g., 5M cells: 3,5 kMB → 20M cells: 11.5k MB)

```
In [1]: run condcubcavitymm.py
Starting simulation with N_cells = 20000000
Generating grid...
Importing stl solids...
Assembling operator matrices...
Applying boundary conditions...
Adding material tensors...
Pre-computing ...
Moving to GPU...
Total initialization time: 73.22341299057007 s
Running electromagnetic time-domain simulation...
2%|█
```

| 412/24470 [00:19<18:30, 21.66it/s]

```
Every 2.0s: ps -axo cpuid,pcpu,pmem,user,pid,tim... pcbe-abp-gpu001: Fri Aug 9 13:32:43 2024
CPUID %CPU %MEM USER PID TIME COMMAND
38 1469 0.5 aforlara 877968 07:18:32 python _new_wake_script.py --yaml_path /afs/cern.ch/
16 1350 0.5 aforlara 877965 06:43:08 python _new_wake_script.py --yaml_path /afs/cern.ch/
46 1339 0.5 aforlara 877962 06:39:43 python _new_wake_script.py --yaml_path /afs/cern.ch/
8 101 0.7 aforlara 867375 04:18:37 /afs/cern.ch/work/a/aforlara/public/SPS_studies/from
8 86.7 1.0 anglard 659638 6-16:21:09 python main_fft_study_hparamsCV.py
32 83.0 9.1 edelafue 876875 00:52:23 /home/edelafue/miniconda3/bin/python /home/edelafue/
34 58.3 2.2 elamb 878933 00:02:13 /home/elamb/.vscode-server/cli/servers/Stable-blc0a1
34 47.8 1.2 elamb 878987 00:01:05 /home/elamb/.vscode-server/cli/servers/Stable-blc0a1
32 13.6 4.0 elamb 871864 00:25:34 /home/elamb/clean/executable/2023_apr/miniconda/bin/
38 9.6 1.8 aforlara 866735 00:27:21 /home/aforlara/.vscode-server/cli/servers/Stable-file
11 5.0 0.0 root 3907 1-06:18:42 [nv_queue]
11 3.6 0.0 root 3856 21:39:16 [nv_queue]
```

```
(base) edelafue@pcbe-abp-gpu001:~$ gpustat --show-user
pcbe-abp-gpu001 Fri Aug 9 12:24:29 2024 470.256.02
[0] NVIDIA TITAN V | 32'C, 0% | 9 / 12066 MB | gdm(4M)
[1] NVIDIA TITAN V | 36'C, 99% | 11472 / 12066 MB | edelafue(11455M) gdm(4M)
[2] NVIDIA TITAN V | 32'C, 0% | 622 / 12066 MB | aforlara(613M) gdm(4M)
[3] NVIDIA TITAN V | 32'C, 0% | 77 / 12065 MB | gdm(56M) gdm(15M)
(base) edelafue@pcbe-abp-gpu001:~$
```

Comparison with CPU speed in limit case 20M:

```
In [1]: run condcubcavitymm.py
Starting simulation with N_cells = 20000000
Generating grid...
Importing stl solids...
Assembling operator matrices...
Applying boundary conditions...
Adding material tensors...
Pre-computing ...
Total initialization time: 38.916051149368286 s
Running electromagnetic time-domain simulation...
6%|█
```

| 1357/24470 [57:11<39:08:15, 6.10s/it]

Only 9% of RAM used, but takes 40 h!!!! To complete



Memory opt. and speedup

Run through the **memory profiler** for a test simulation of 1M cells (~700 MB).

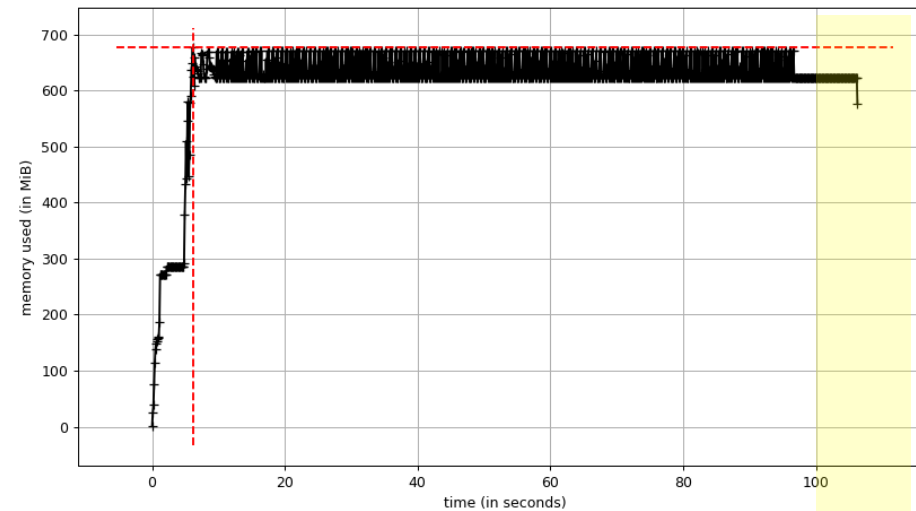
Storing fields E, H, J and tensors σ, ϵ, μ in **numpy (x,y,z)x3d matrices** inside **Field class** caused spikes in memory every timestep...

GPU slowed down or crashing.

Complete refactor of the **Field class** to store fields and tensors in **1x1d array** of length **$3 \cdot N_x \cdot N_y \cdot N_z$** .

Access to data and matrix form through magic methods (setter, getter, `__getitem__`, `__setitem__`, `__add__`, `__mul__`, `__div__`,...), completely transparent to the solver and the rest of the code 😊.

/home/edelafue/miniconda3/bin/python memprofiler.py

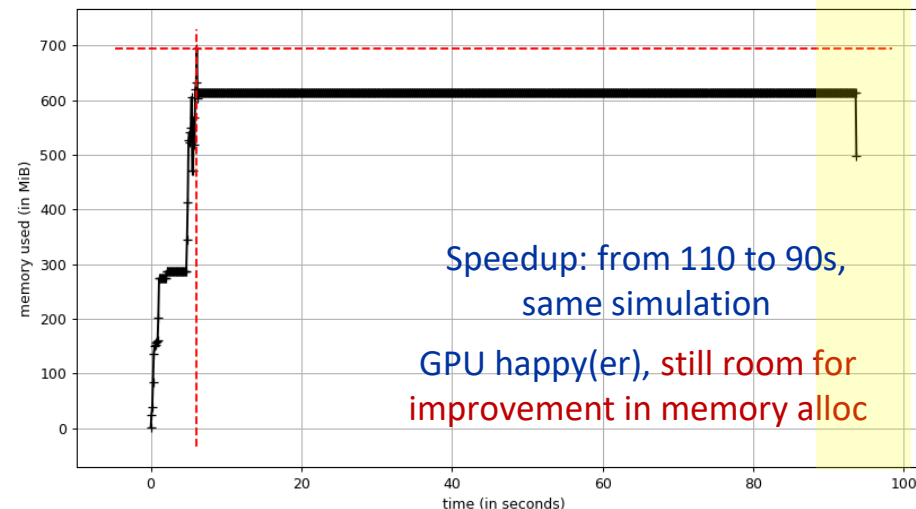


Check
Here



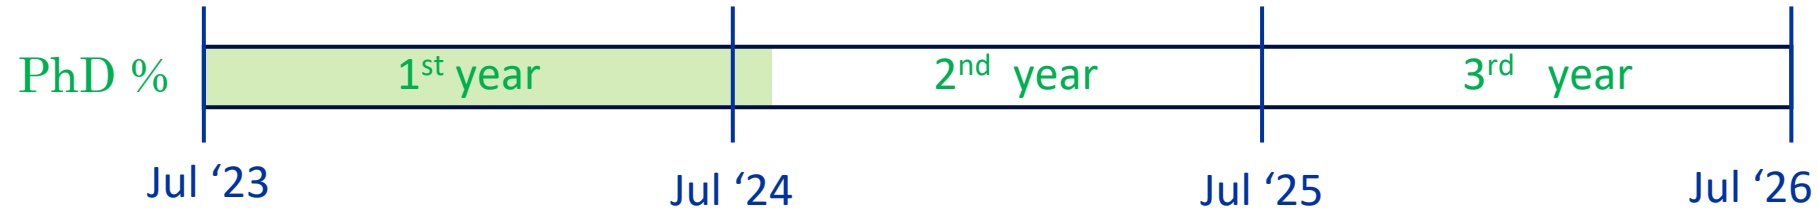
Refactor comparison

/home/edelafue/miniconda3/bin/python memprofiler.py



Speedup: from 110 to 90s,
same simulation
GPU happy(er), still room for
improvement in memory alloc

Wake solver milestones:



meetings [#17](#), [#18](#), [#19](#), [CEI](#)

0. Wake potential and impedance

1. FIT Maxwell Equations in 3D

2. PEC, Periodic, PMC boundaries

3. Embedded Boundaries: geometry import from **.stl** files

4. Beam injection \mathbf{J}_z

5. Materials: ϵ, μ, σ

In progress

x. Documentation

6. on GPU & memory opt.

7. Low beta & SC

8. Open boundaries & PML

9. Numerical tests vs analytic

To Do

10. Leontovich condition for good conductors >100 S/m

11. Staircased geometry: PBA/grid refinement

12. Moving window solver (?)

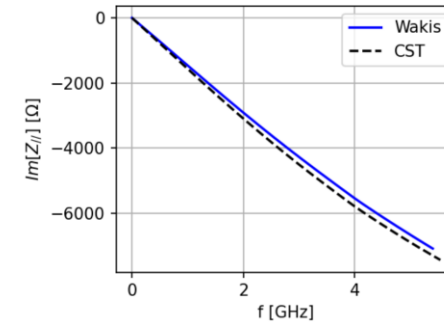
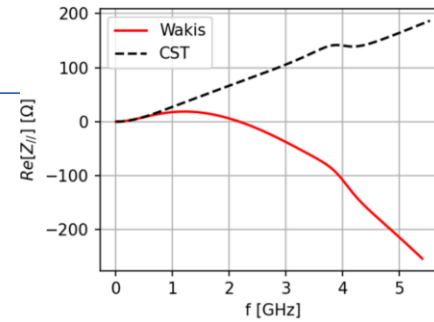
13. Frequency-dependent materials

14. ...

Low-beta simulations

Thanks to [Elena Macchia's work with wakis](#), we noticed some disagreement in the Real part of the impedance that worsen with lower β :

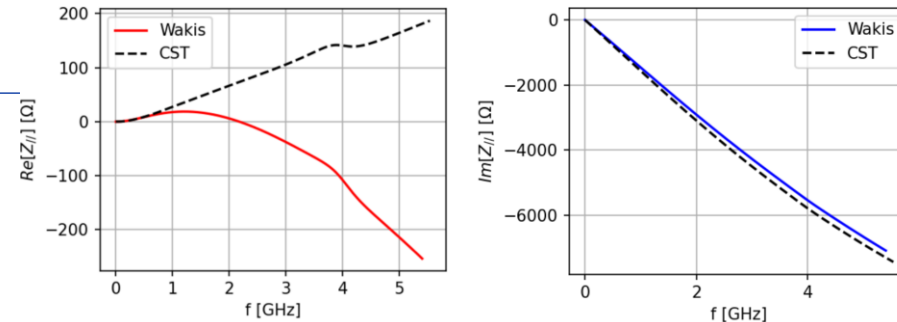
As β decreases, the real parts of Z in the two solvers start to drift apart.



Low-beta simulations

Thanks to [Elena Macchia's work with wakis](#), we noticed some disagreement in the Real part of the impedance that worsen with lower β :

As β decreases, the real parts of Z in the two solvers start to drift apart.

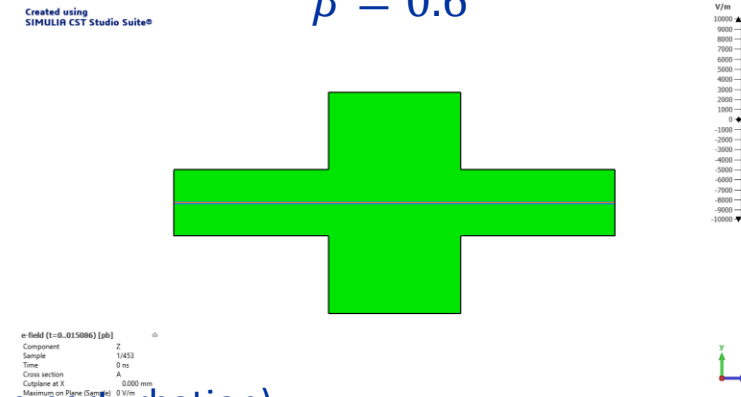
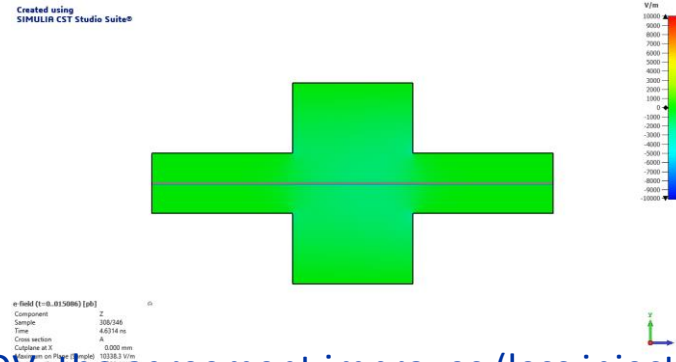
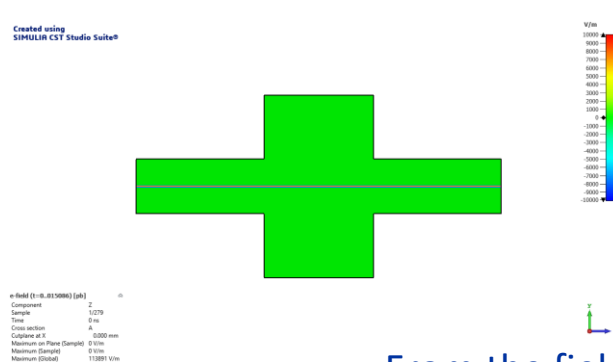


$\beta = 1$

$\beta = 0.8$

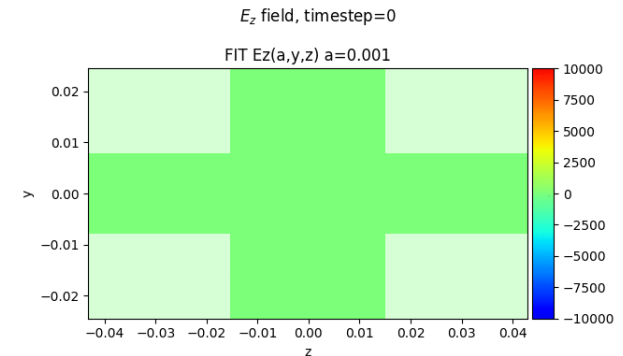
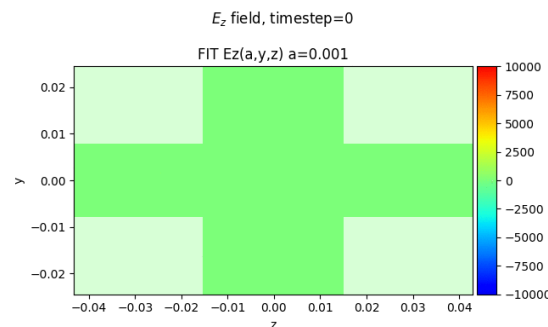
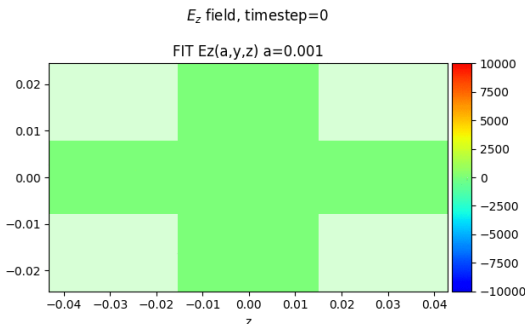
$\beta = 0.6$

CST



From the field POV, the agreement improves (less injection perturbation)

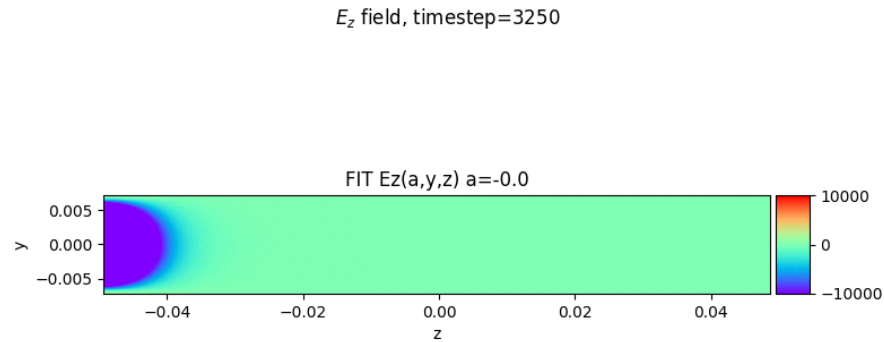
wakis



Low-beta simulations (II)

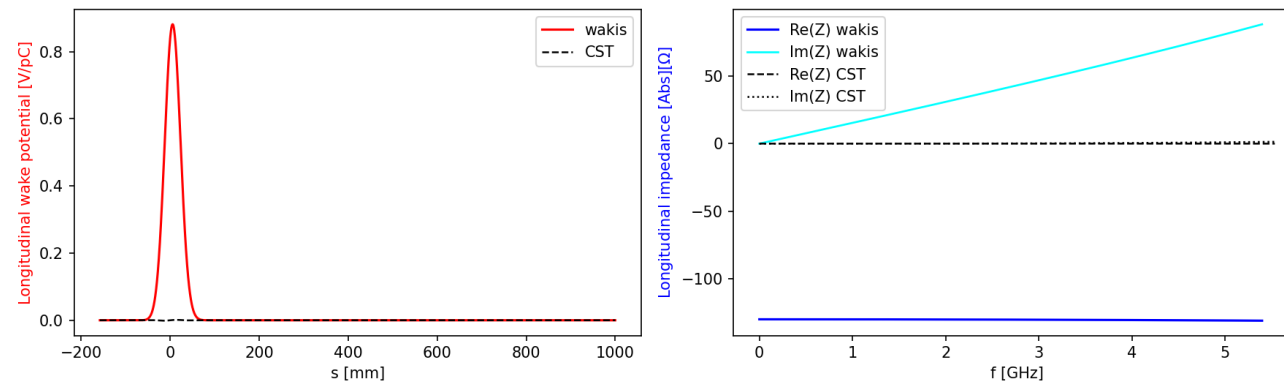
Testing with a simple squared pipe in CST: need $\text{Re} = 0$ and only reactive (Im) impedance, increasing with β . Quite difficult to converge even in CST (need 80 cells per wavelength)

For $\beta = 1$



WP and impedance coming from injection perturbation

Benchmark with CST Wakefield Solver

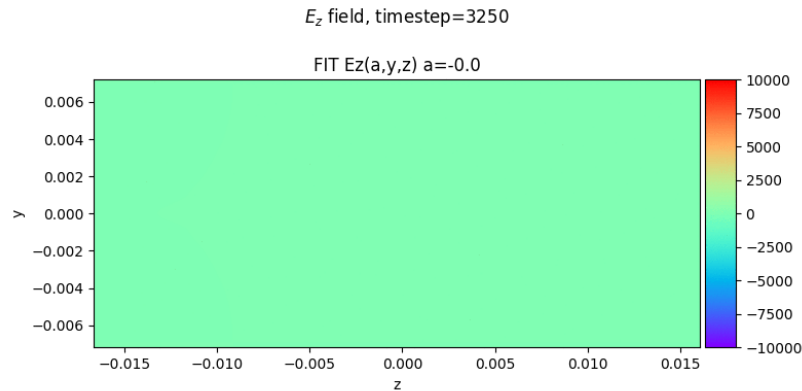


add_space = 0

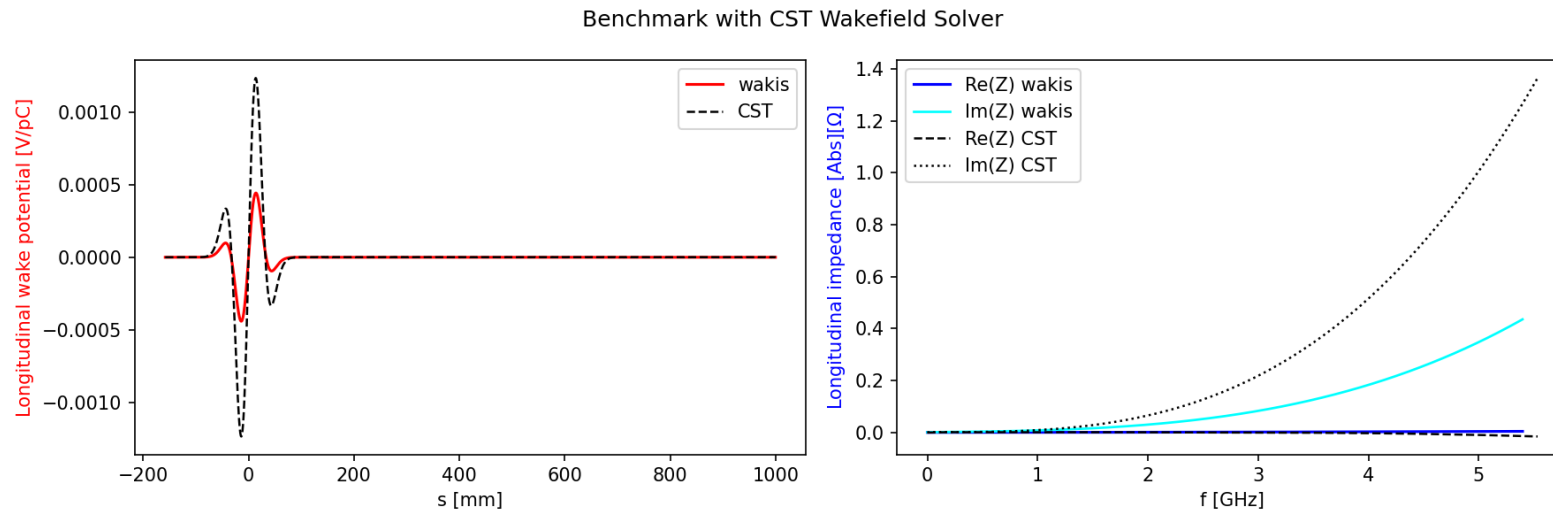
Low-beta simulations (II)

Testing with a simple squared pipe in CST: need $\text{Re} = 0$ and only reactive (Im) impedance, increasing with β . Quite difficult to converge even in CST (need 80 cells per wavelength)

For $\beta = 1$



Removing the first and last 50(!) cells of the domain, the perturbation vanishes, finding expected WP and Z



Is this what CST does behind the scene??

add_space = 50



Low-beta simulations (III)

Testing with a simple squared pipe in CST: need $\text{Re} = 0$ and only reactive (Im) impedance, increasing with β . Quite difficult to converge even in CST (need 80 cells per wavelength)

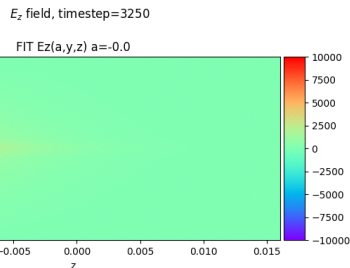
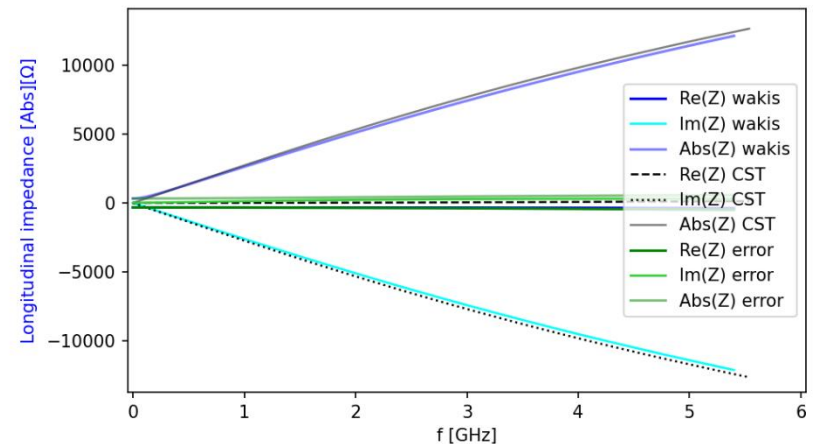
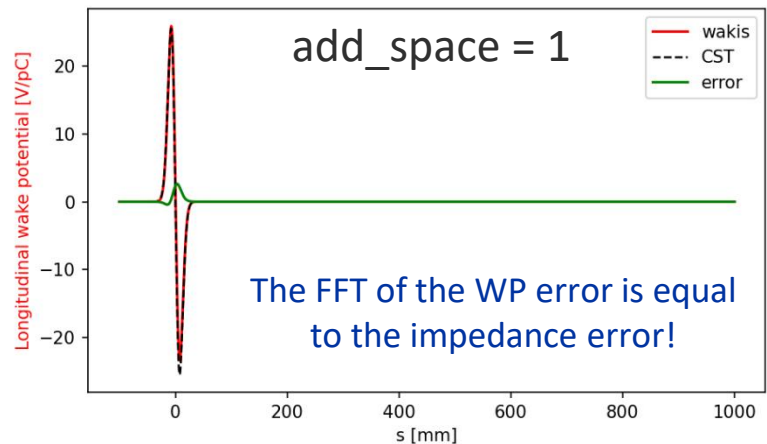
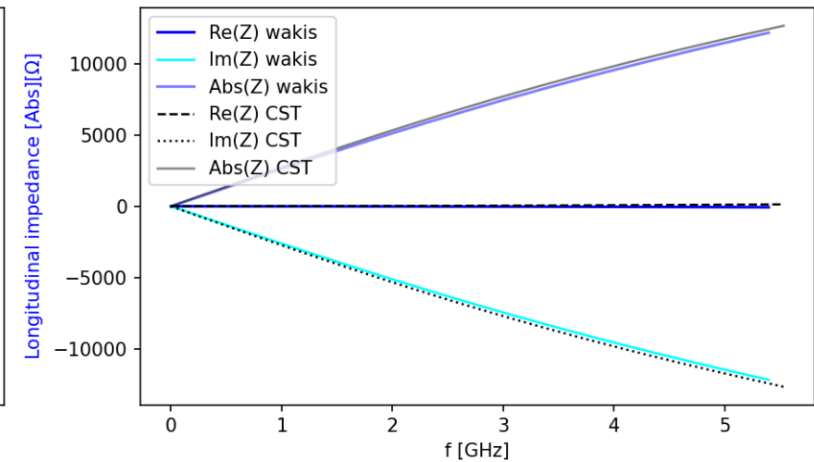
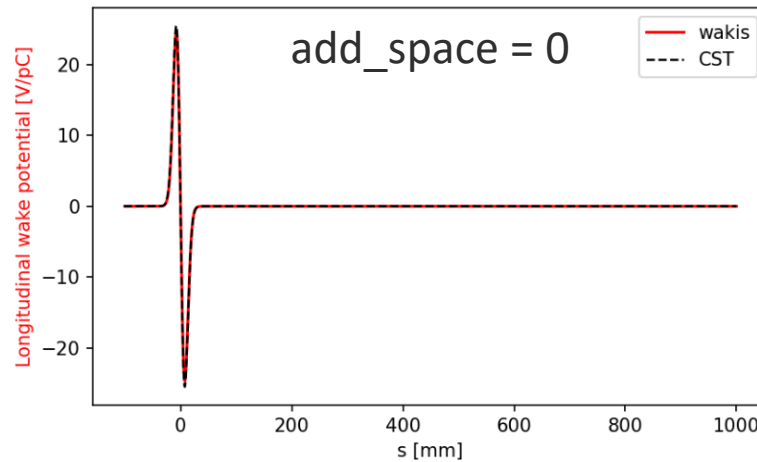


We could make a model to remove this SC from the final WP, it only depends on the mesh

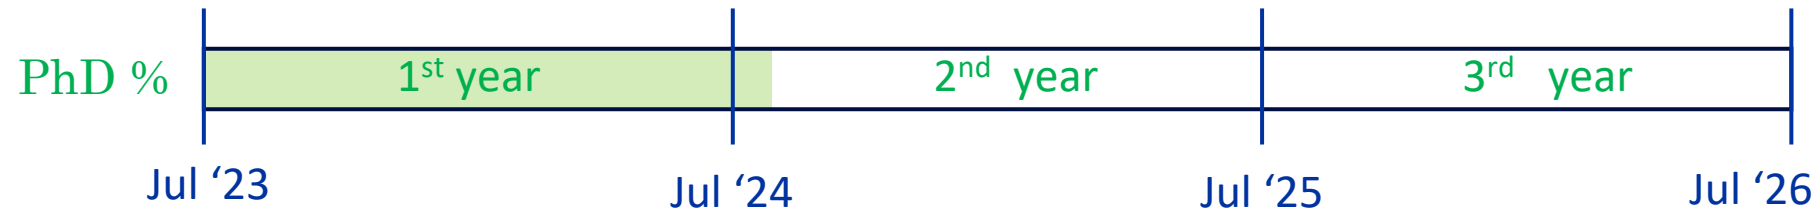
Benchmark with CST Wakefield Solver

For $\beta < 1$

We find the opposite behavior... the more cells we skip, the worse agreement



Wake solver milestones:



meetings [#17](#), [#18](#), [#19](#), [CEI](#)

In progress

To Do

0. Wake potential and impedance

1. FIT Maxwell Equations in 3D

2. PEC, Periodic, PMC boundaries

3. Embedded Boundaries: geometry import from **.stl** files

4. Beam injection \mathbf{J}_z

5. Materials: ϵ, μ, σ

x. Documentation

6. on GPU & memory opt.

7. Low beta & SC

8. Open boundaries & PML

9. Numerical tests vs analytic

10. Leontovich condition for good conductors >100 S/m

11. Staircased geometry: PBA/grid refinement

12. Moving window solver (?)

13. Frequency-dependent materials

14. ...

PML implementation

- Derivation from Snell and Fresnel laws: [Codimd](#) (needs a dedicated presentation...)
- Summary:
 - Many different types, some affect the curl operator (convolutional) some have multiple poles, etc...
 - For now, the simplest PML consist of adding an anisotropic conductivity that grows exponentially inside the PML

$$\sigma_x = \frac{\epsilon_0}{2\Delta t} \left(\frac{x_{N_{pml}} - x_{0_{pml}}}{N_{pml}\Delta x} \right)^n \text{ for E}$$
$$\sigma_x^* = \mu_0 \sigma_x \text{ for H (optional)?}$$

We add it to the conductivity tensor in the PML region before the time-stepping (all pre-computed)

```
def fill_pml_sigmas(self):
    """
    Routine to calculate pml sigmas and apply them
    to the conductivity tensor sigma
    """

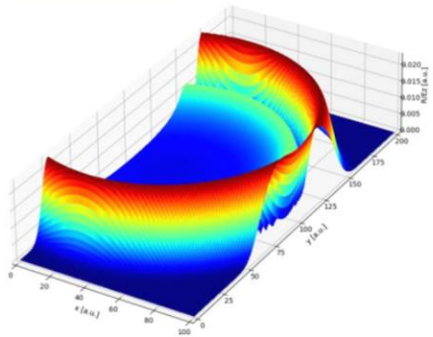
    # Initialize
    sx, sy, sz = np.zeros(self.Nx), np.zeros(self.Ny), np.zeros(self.Nz)
    pml_exp = 2

    # Fill
    if self.bc_low[0].lower() == 'pml':
        sx[0:self.npml] = eps_0/(2*self.dt)*((self.x[self.npml] - self.x[:self.npml])/(self.npml*self.dx))**pml_exp
        for d in ['x', 'y', 'z']:
            for i in range(self.npml):
                self.sigma[i, :, :, d] = sx[i]
```

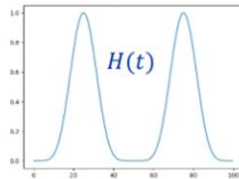
Test PML reflection, sounds familiar...

PML study in Warp

Test domain $t = 200\delta t$

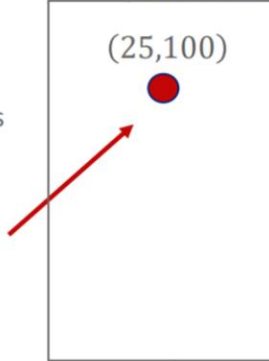


$100\delta x \times 200\delta y$
Simulate 200 t steps
 $E_z(25,100) = H(t)$



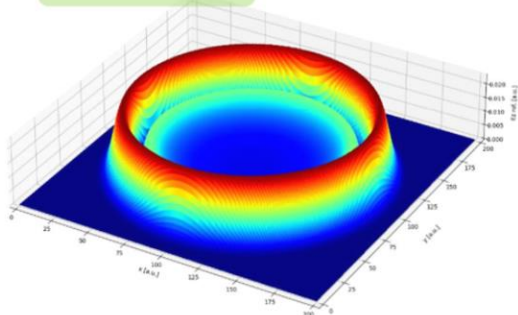
100 cells

(25,100)



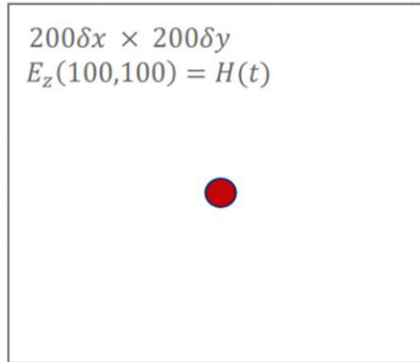
200 cells

Ref domain $t = 200\delta t$



$200\delta x \times 200\delta y$
 $E_z(100,100) = H(t)$

200 cells



200 cells

[EMWSD #15](#)

2D simulation in Warp
(never replicated in WarpX)

$$\mathbf{R}(0:75, 0:200) = \mathbf{E}_{z, test}(0:75, 0:200) - \mathbf{E}_{z, ref}(100:175, 0:200)$$

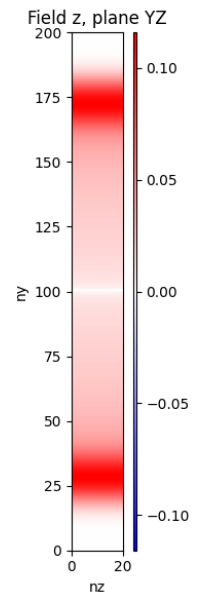
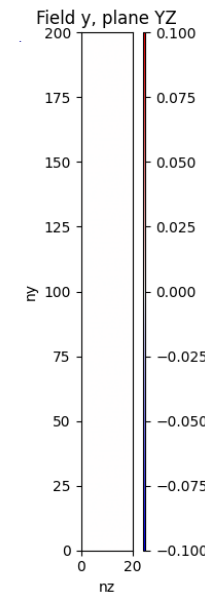
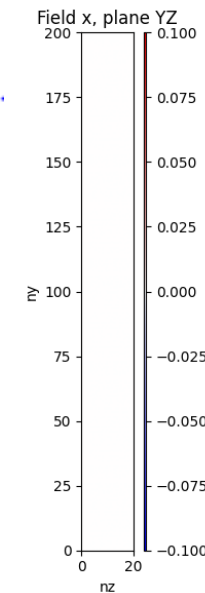
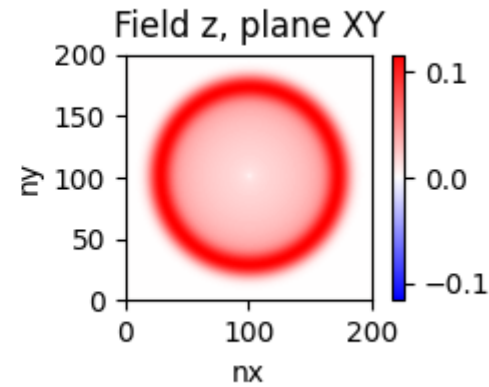
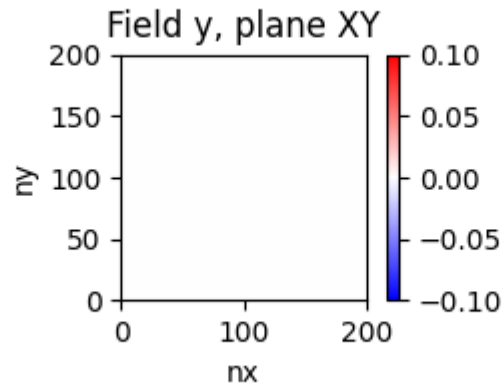
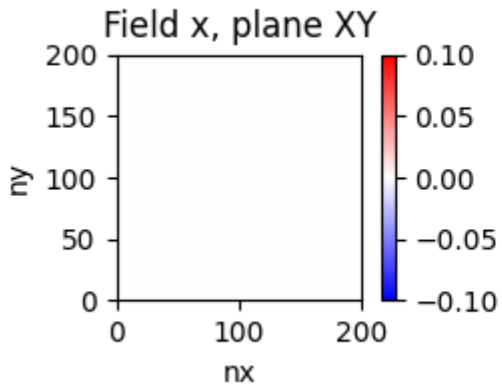
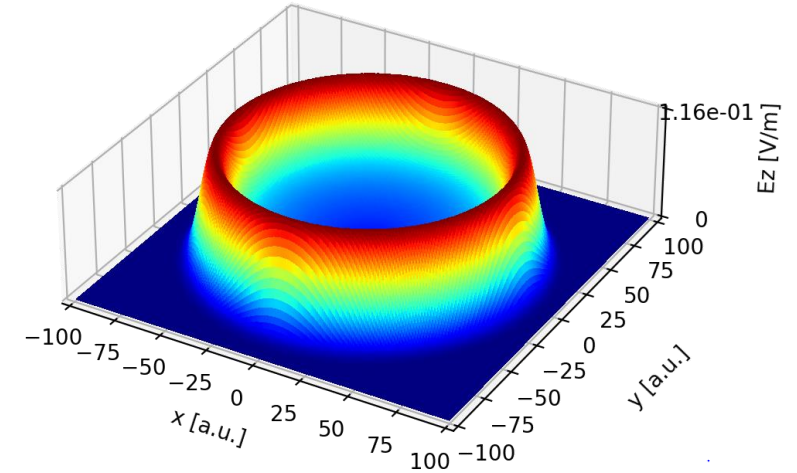
*Original test from [Berenger paper](#) presenting the PML formulation



Test PML in wakis

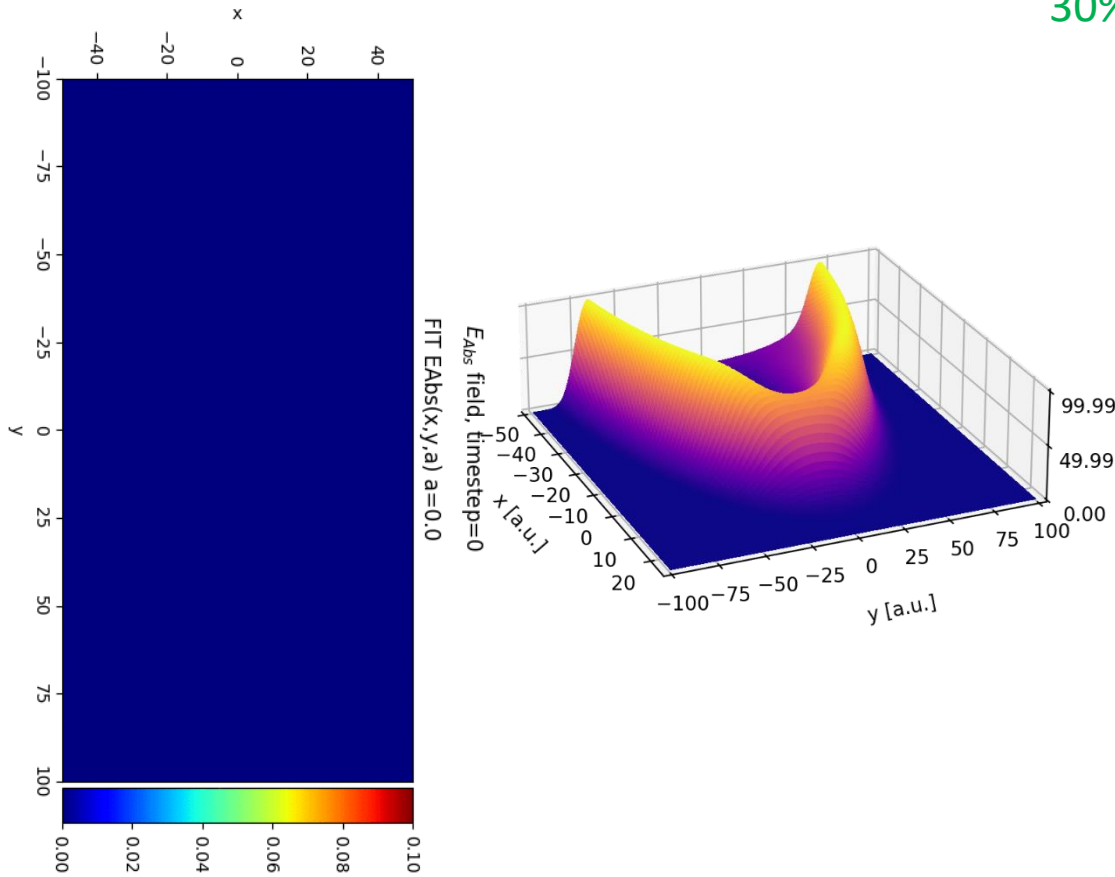
Replicating test and reference 2D simulations in wakis: **successful**

- Created new source class `Pulse` that injects at a given x_s, y_s, z_s
- To have a 2D pulse, $z_s = \text{slice}(0, N_z)$ with PEC BCs
- This Harris pulse gives a very clean positive field pulse in E_z .
- All the other field components are 0

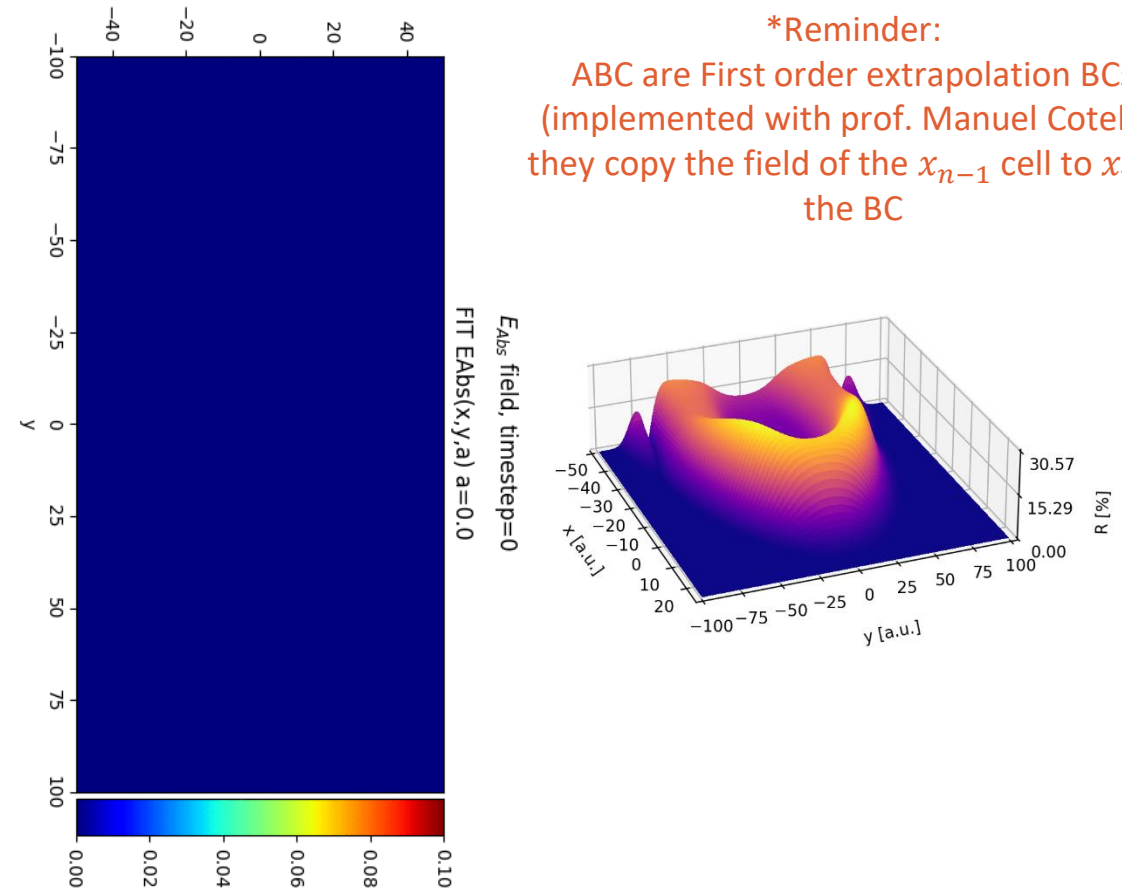


Test PML in wakis (II)

PEC boundaries X, Y
Total reflection



PEC boundaries Y, ABC* boundaries X
Absorption in Xlo, but not in Xhi → bug in ABC
30% reflection xlo, >100% reflection in xhi

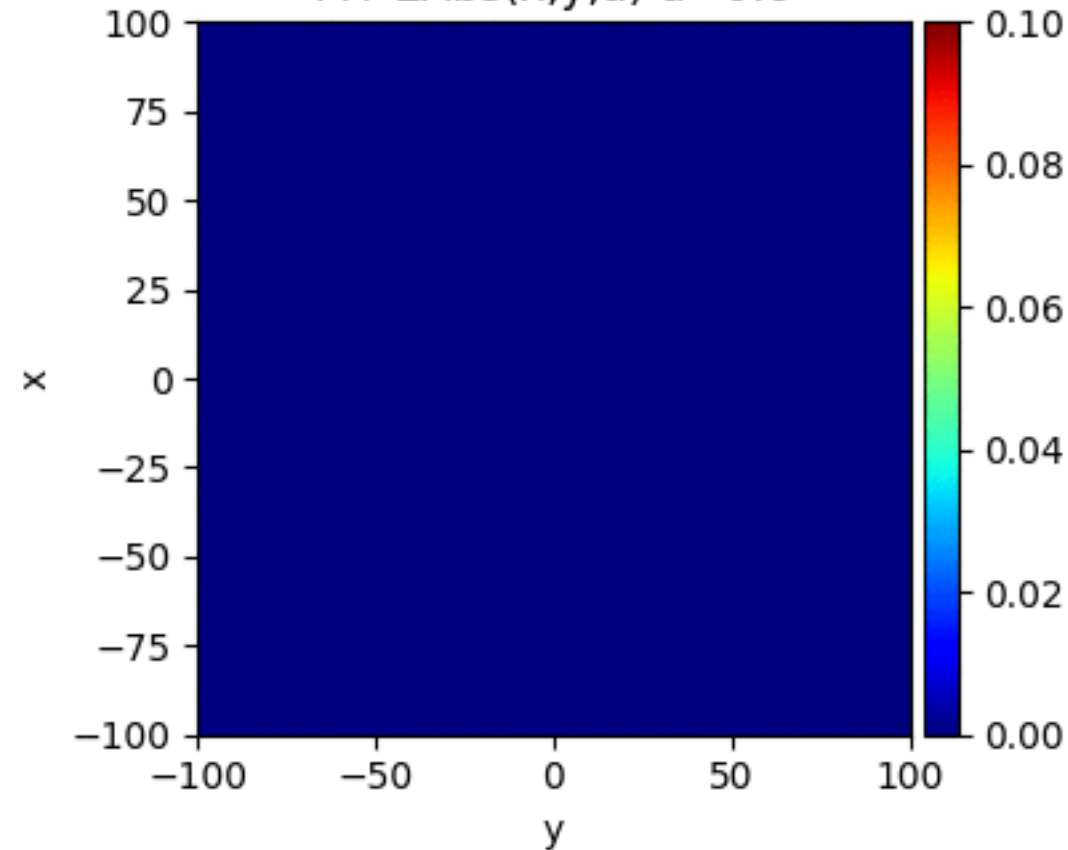


*Reminder:
ABC are First order extrapolation BCs
(implemented with prof. Manuel Cotelo),
they copy the field of the x_{n-1} cell to x_n at
the BC

Test PML in wakis (II)

E_{Abs} field, timestep=0

FIT EAbs(x,y,a) a=0.0



PEC boundaries Y
Periodic X
No reflection at xlo
(but not suitable for impedance)

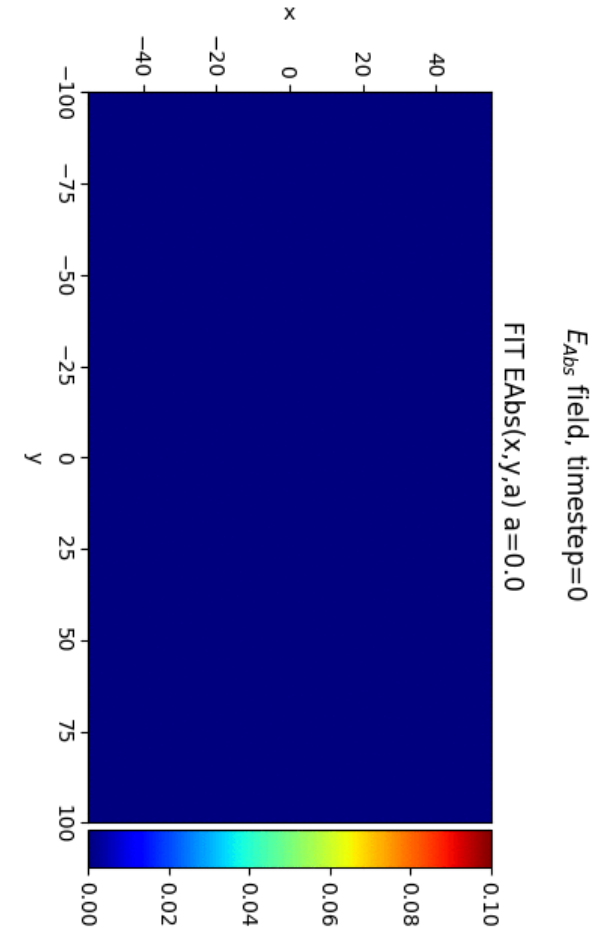
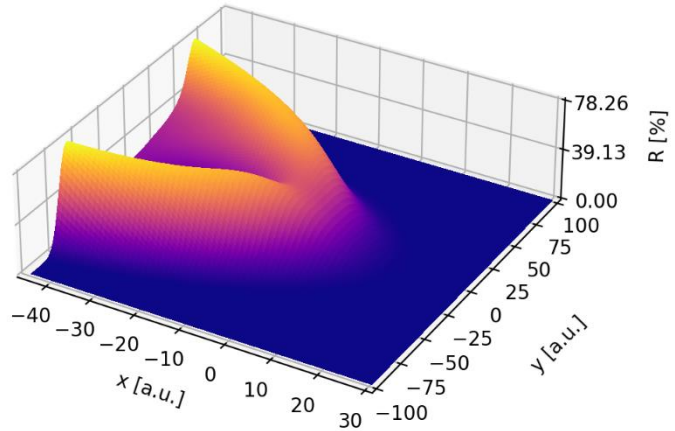
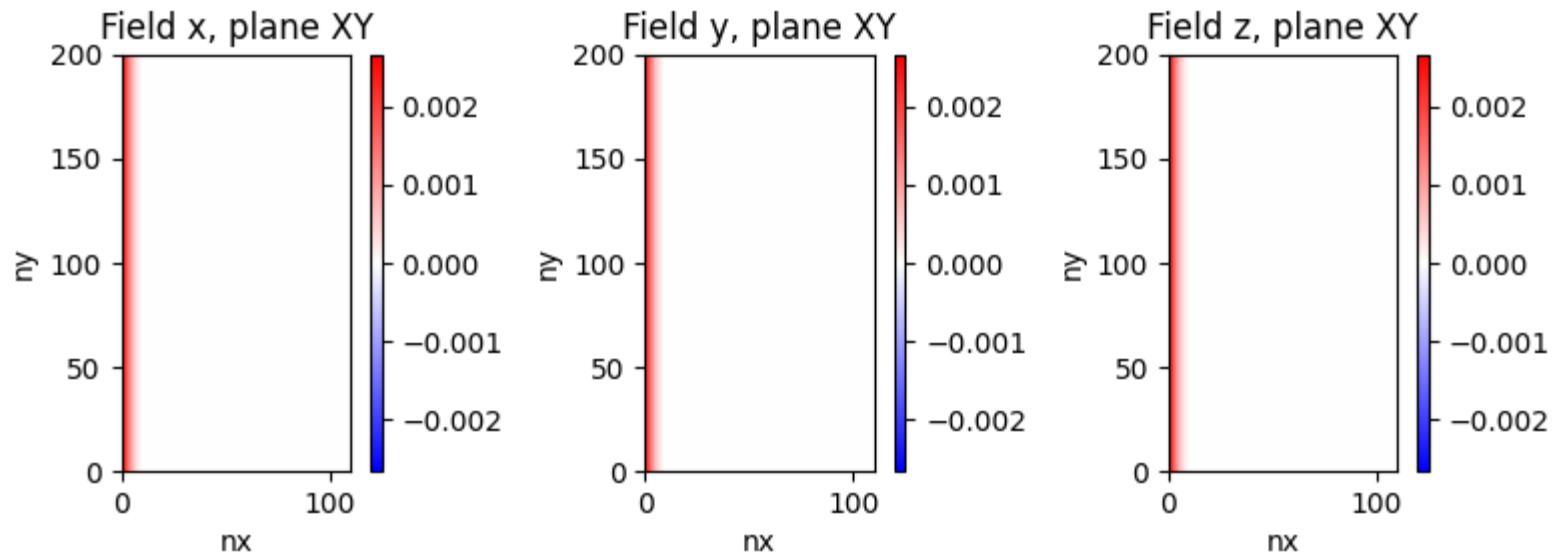
Test PML in wakis (II)

PML boundaries X, Y

70% reflection, depends on N_{pml} layers and max conductivity

$$\sigma_x = \frac{\epsilon_0}{2\Delta t} \left(\frac{x_{N_{pml}} - x_{0_{pml}}}{N_{pml}\Delta x} \right)^2$$

σ conductivity tensor with 10 PML layers in x



Test PML in wakis (II)

PML boundaries X, Y

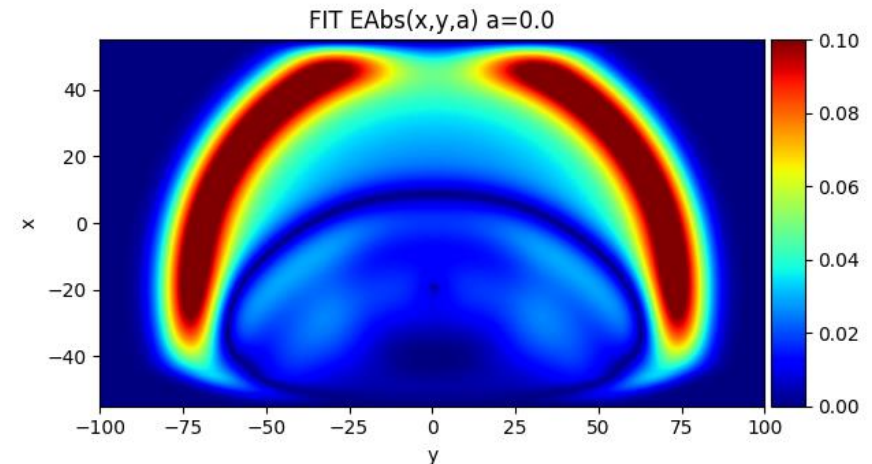
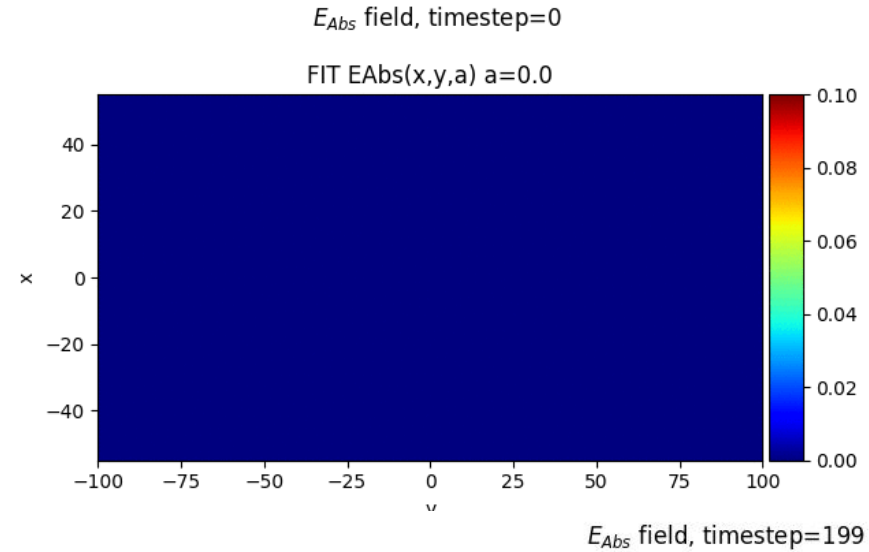
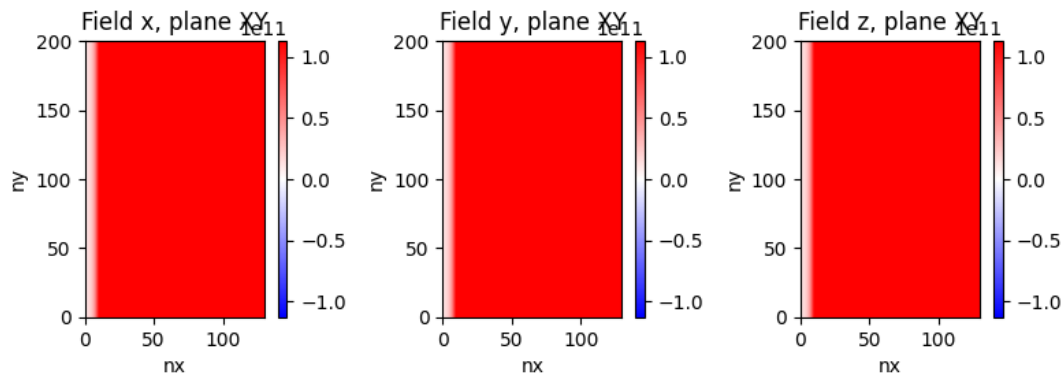
70% reflection, depends on N_{pml} layers and max conductivity

$$\sigma_x = \sigma_{max} \left(\frac{x_{N_{pml}} - x_{0_{pml}}}{N_{pml} \Delta x} \right)^2$$

If $\sigma_{max} > \frac{\epsilon_0}{2\Delta t}$ simulation unstable

Can be cured modifying $\bar{\epsilon} \rightarrow \epsilon_x = \frac{\sigma_{max}}{2} \epsilon_0$

ϵ^{-1} tensor with 30 PML layers in x

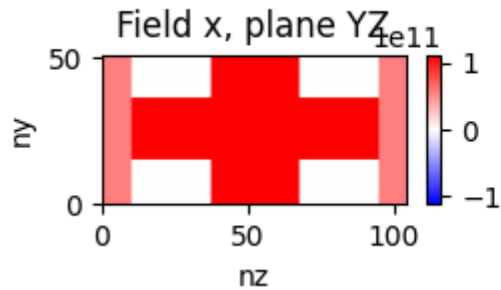


WIP to optimize σ_{max} while keeping simulation stable

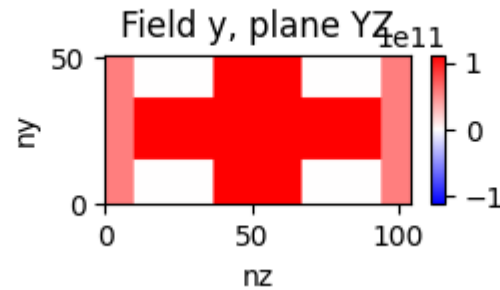
PML on Impedance simulation

Just to try... I tested using this PML in a PEC cubic cavity with modes above cutoff

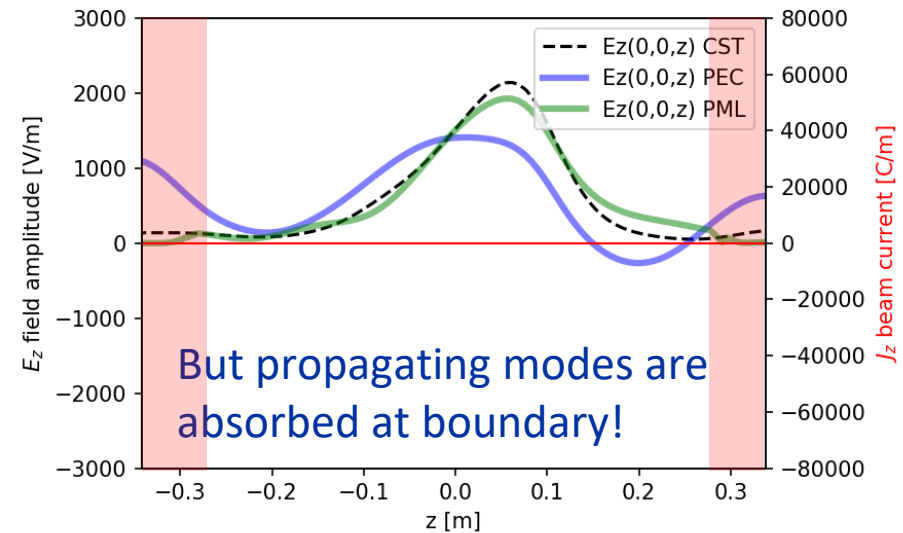
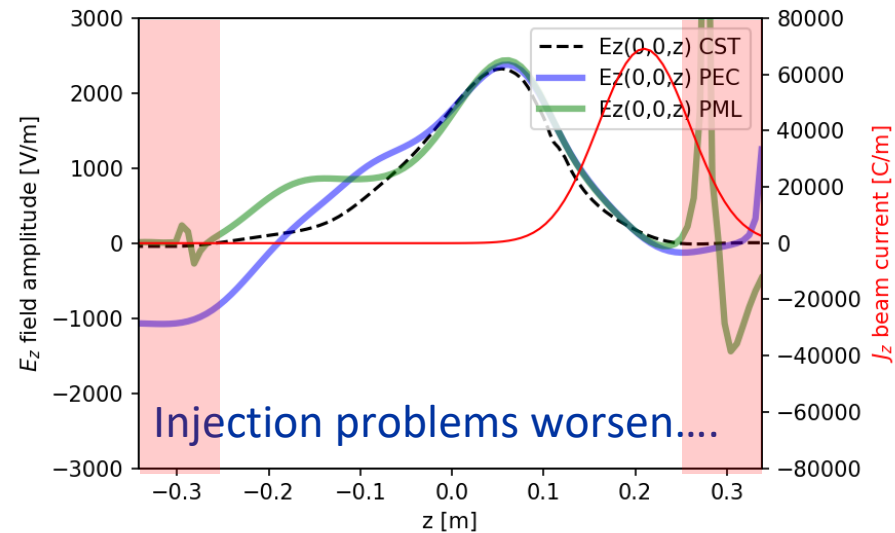
ϵ^{-1} tensor



timestep=930



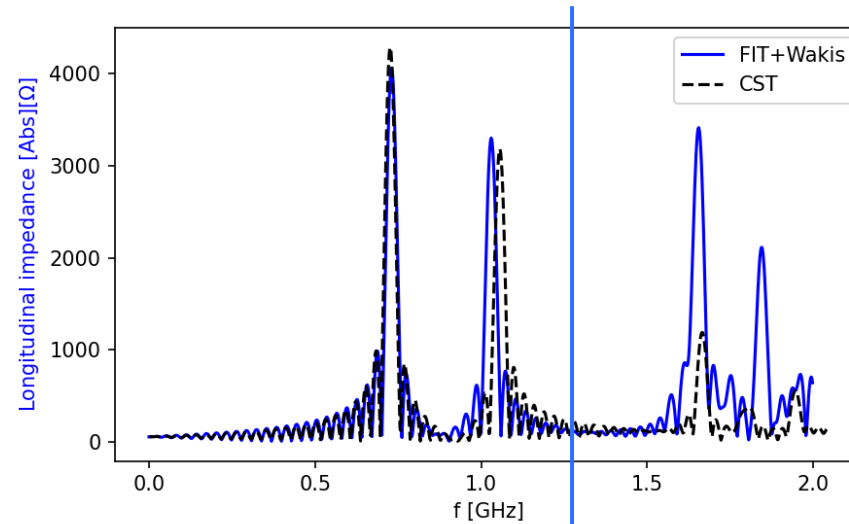
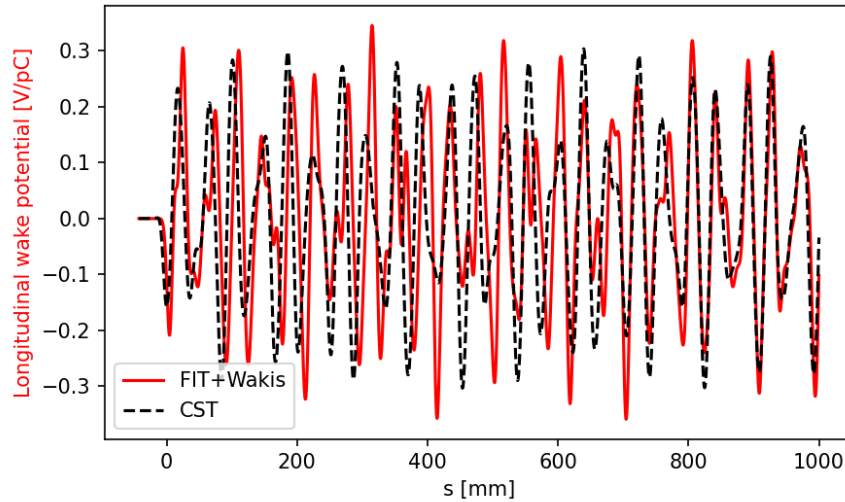
timestep=1740



PML on Impedance simulation (II)

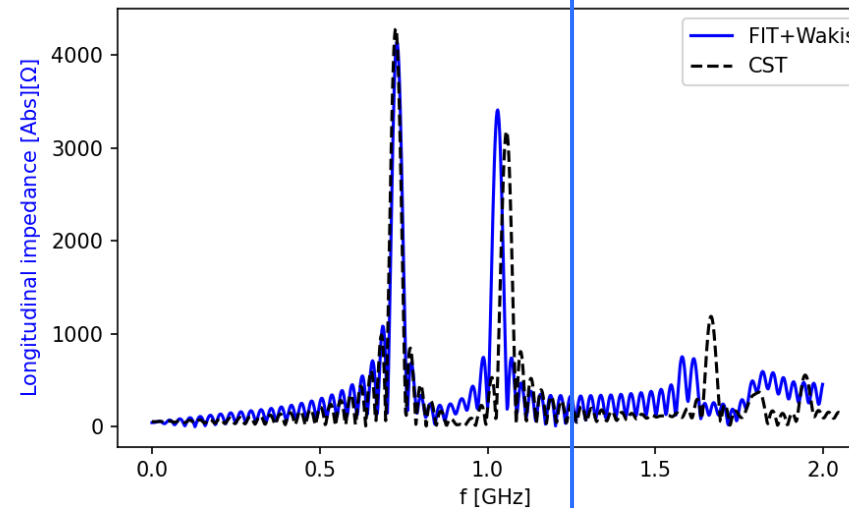
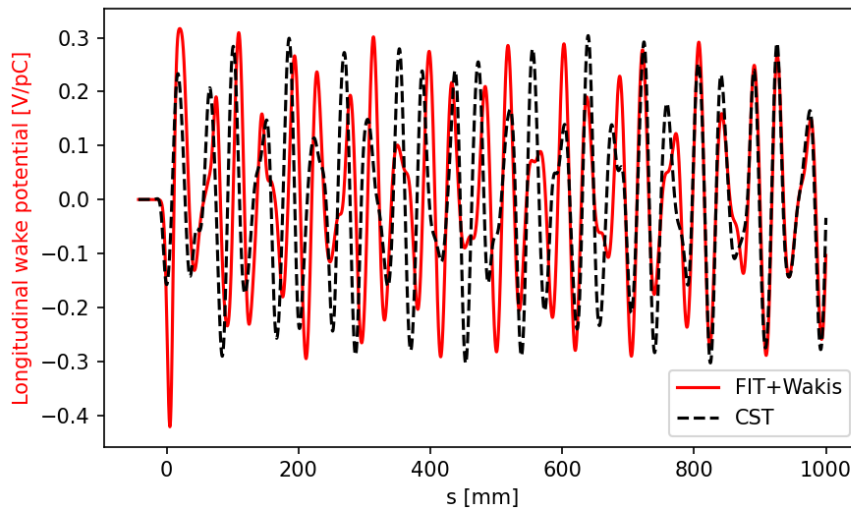
Just to try... I tested using this PML in a PEC cubic cavity with modes above cutoff

PEC



Cutoff frequency:
~1.25 GHz

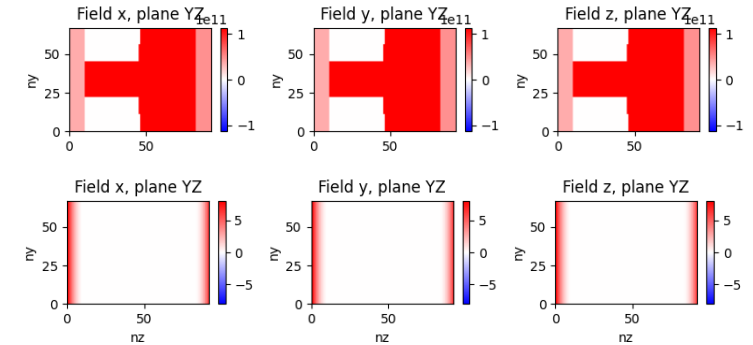
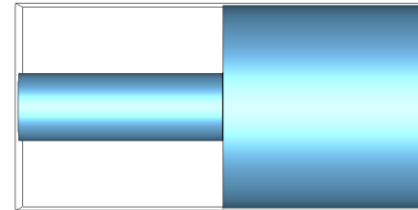
PML



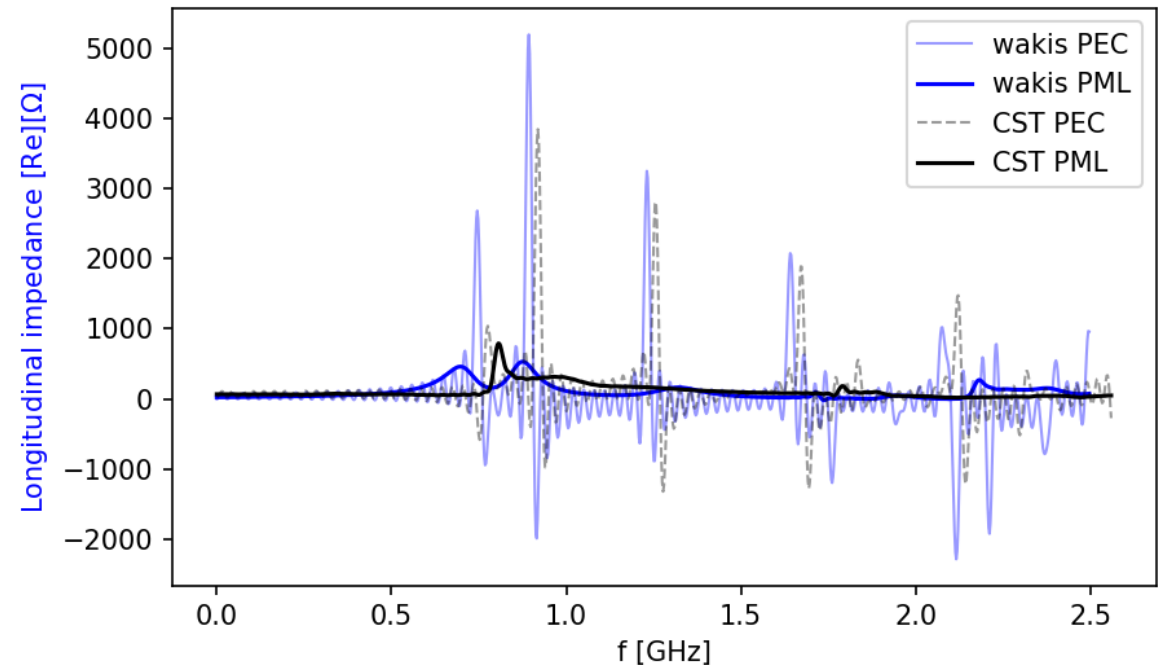
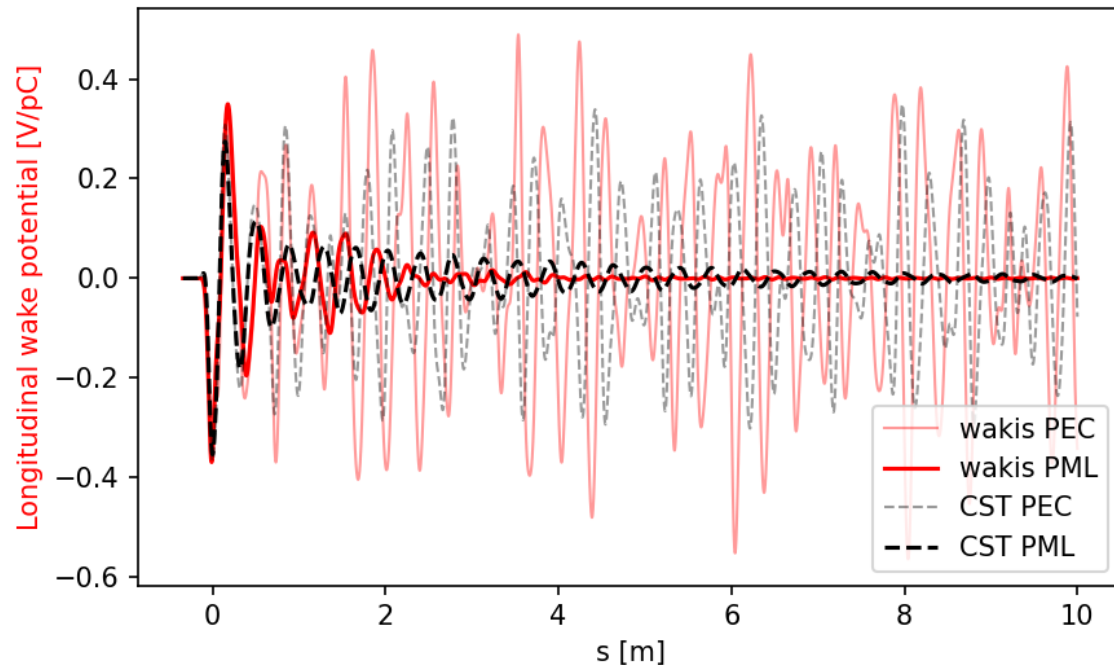
PML on Impedance simulation (III)

Worst case for PML: **PEC step-out transition**

- Same mesh as CST
- 10 layer PML



Benchmark with CST Wakefield Solver



Surprisingly good results



Conclusions

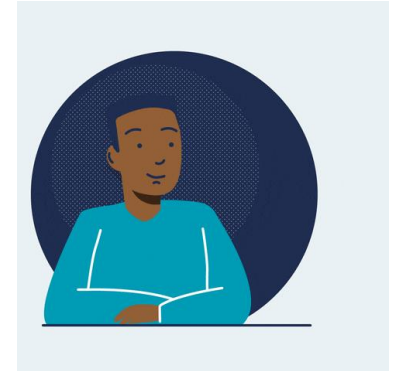
✓ Progress on the code:

- First users, documentation, work on Gitub workflow to keep a stable version
- GPU acceleration successful, speedup x120 vs CPU for limit case 20 million cells
 - New bottleneck: memory alloc -> **parallelize?**
- Low beta simulations:
 - Potential to understand space charge problem, create model to remove it (?)
 - Remove injection perturbation by enlarging domain only during excitation (?) only needed for beta=1
- PML:
 - Mathematical derivation and 1st implementation done
 - Under reflection test -> 70% reflection
 - **Seems to work already for impedance simulations?**

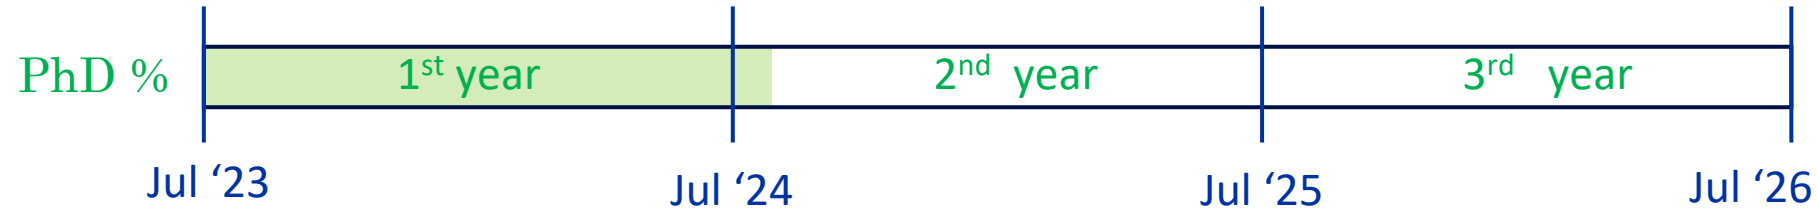
Conclusions & Future

✓ Progress on the code:

- First users, documentation, work on Github workflow to keep a stable version
- GPU acceleration successful, speedup x120 vs CPU for limit case 20 million cells
 - New bottleneck: memory alloc -> **parallelize?**
- Low beta simulations:
 - Potential to understand space charge problem, create model to remove it (?)
 - Remove injection perturbation by enlarging domain only during excitation (?) only needed for beta=1
- PML:
 - Mathematical derivation and 1st implementation done
 - Under reflection test -> 70% reflection
 - **Seems to work already for impedance simulations?**
- Still many things to do:
 - Good conductors, staircased geometry, grid refinement, moving window, frequency dependent materials, frequency domain monitors, TESTS **Where to focus?** IPAC in October



Wake solver milestones:



meetings [#17](#), [#18](#), [#19](#), [CEI](#)

0. Wake potential and impedance

1. FIT Maxwell Equations in 3D

2. PEC, Periodic, PMC boundaries

3. Embedded Boundaries: geometry import from **.stl** files

4. Beam injection J_z

5. Materials: ϵ, μ, σ

In progress

x. Documentation

6. on GPU & memory opt.

7. Low beta & SC

8. Open boundaries & PML

9. Numerical tests vs analytic

To Do

10. Leontovich condition for good conductors >100 S/m

11. Staircased geometry: PBA/grid refinement

12. Moving window solver (?)

13. Frequency-dependent materials

14. ...

Thank you 😊 !!!

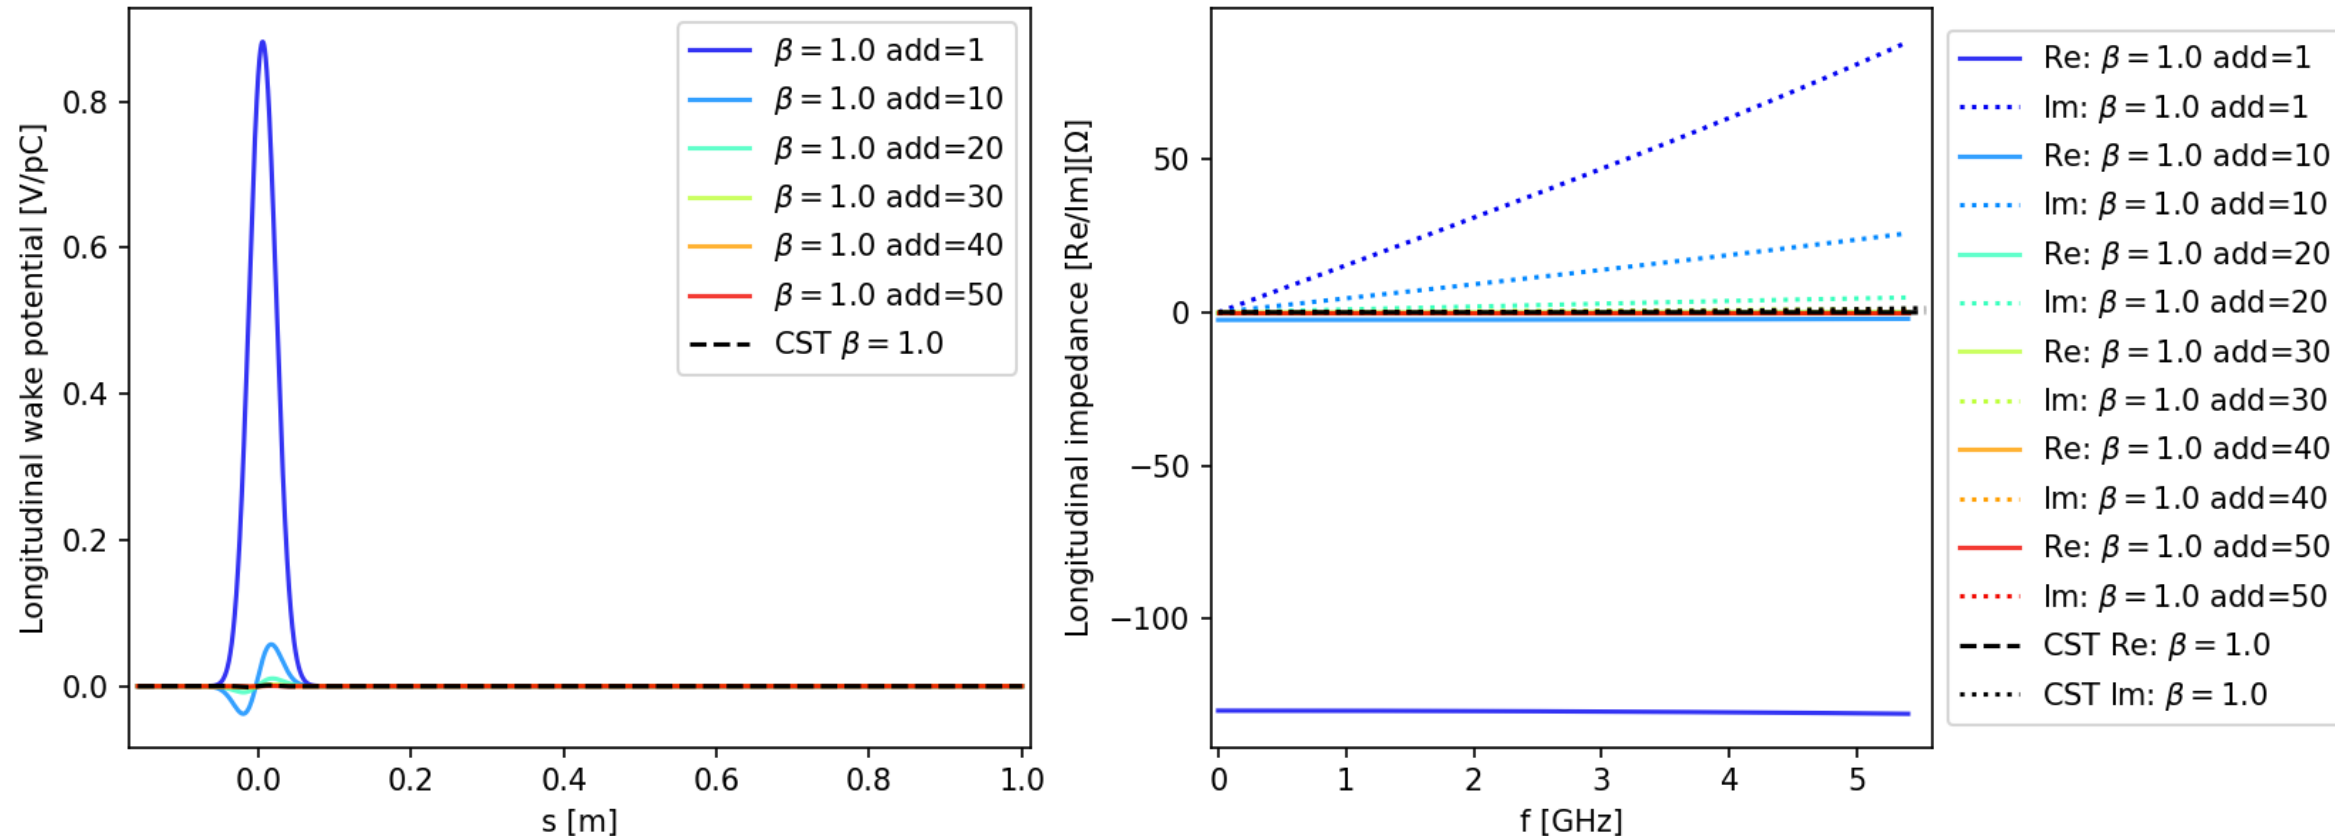


Electromagnetic and Wake Solver Development
meeting #20

Elena de la Fuente García (BE-ABP-CEI)

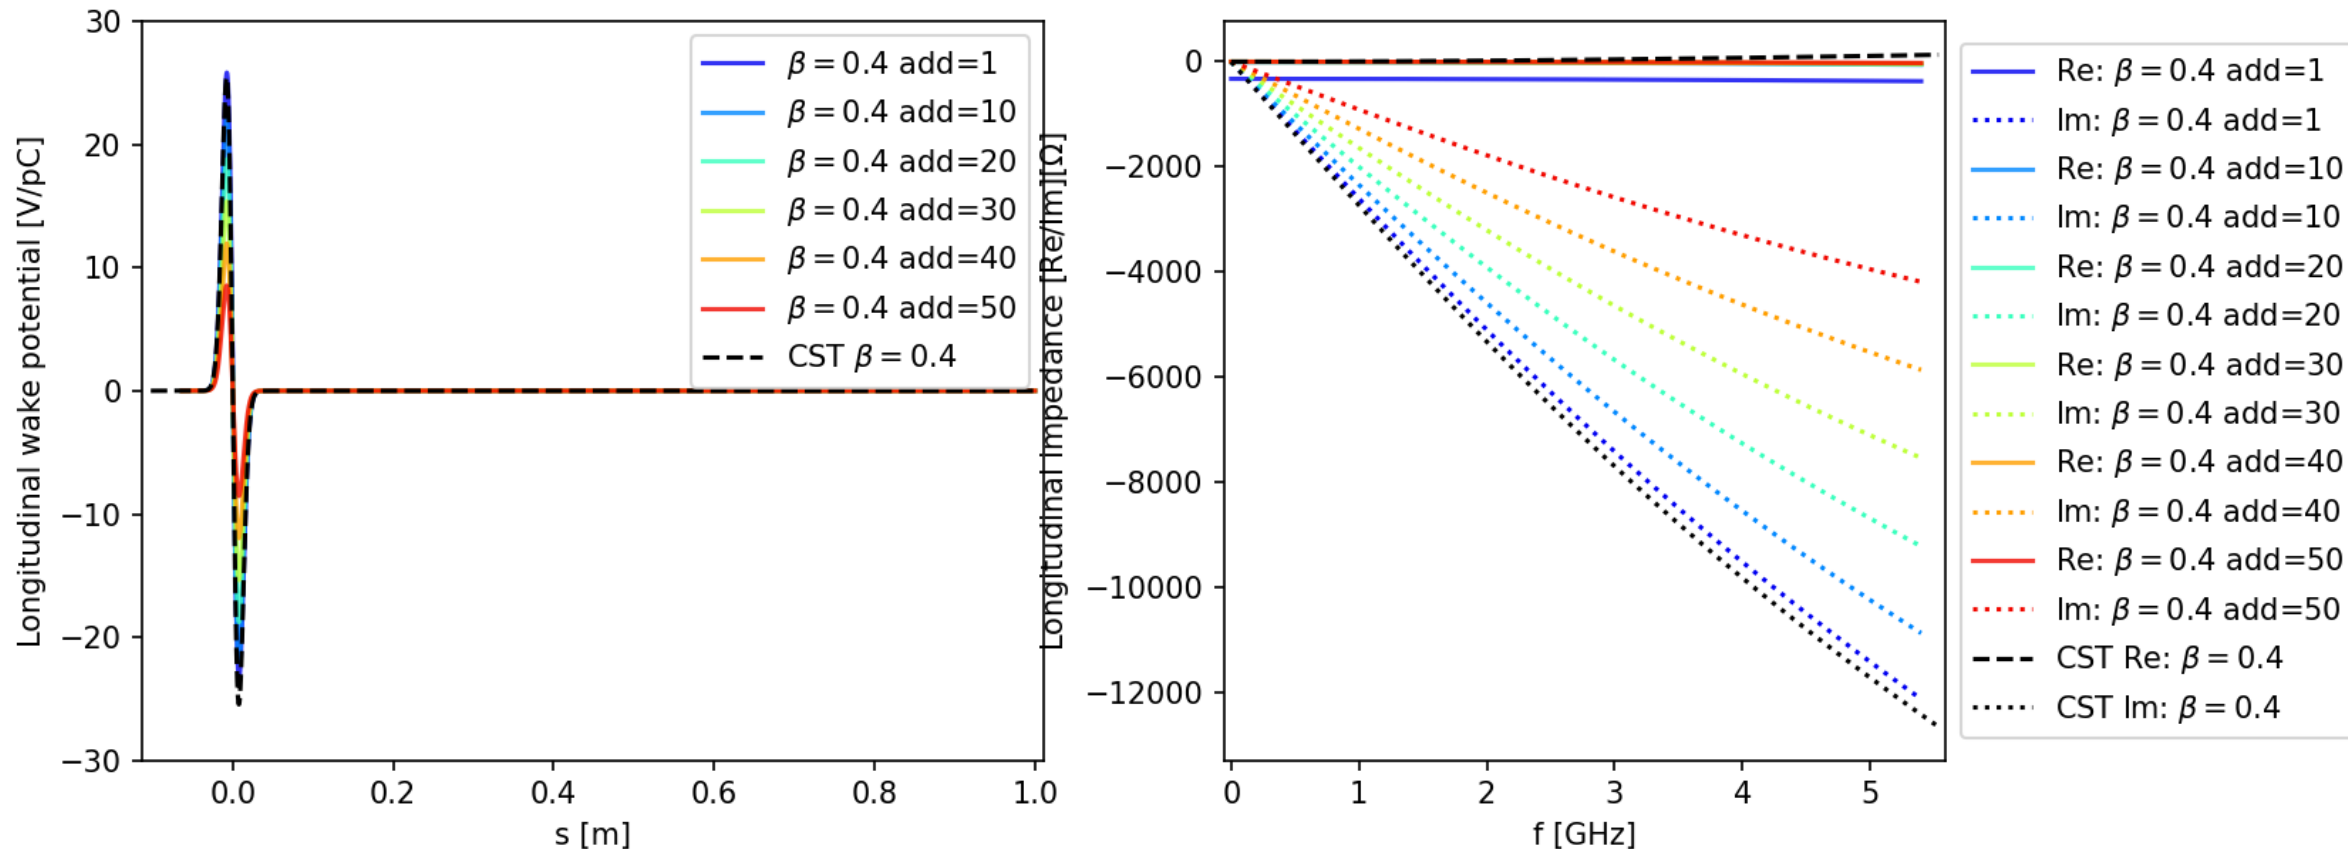
Backup: beta=1, different add_space

Benchmark with CST Wakefield Solver



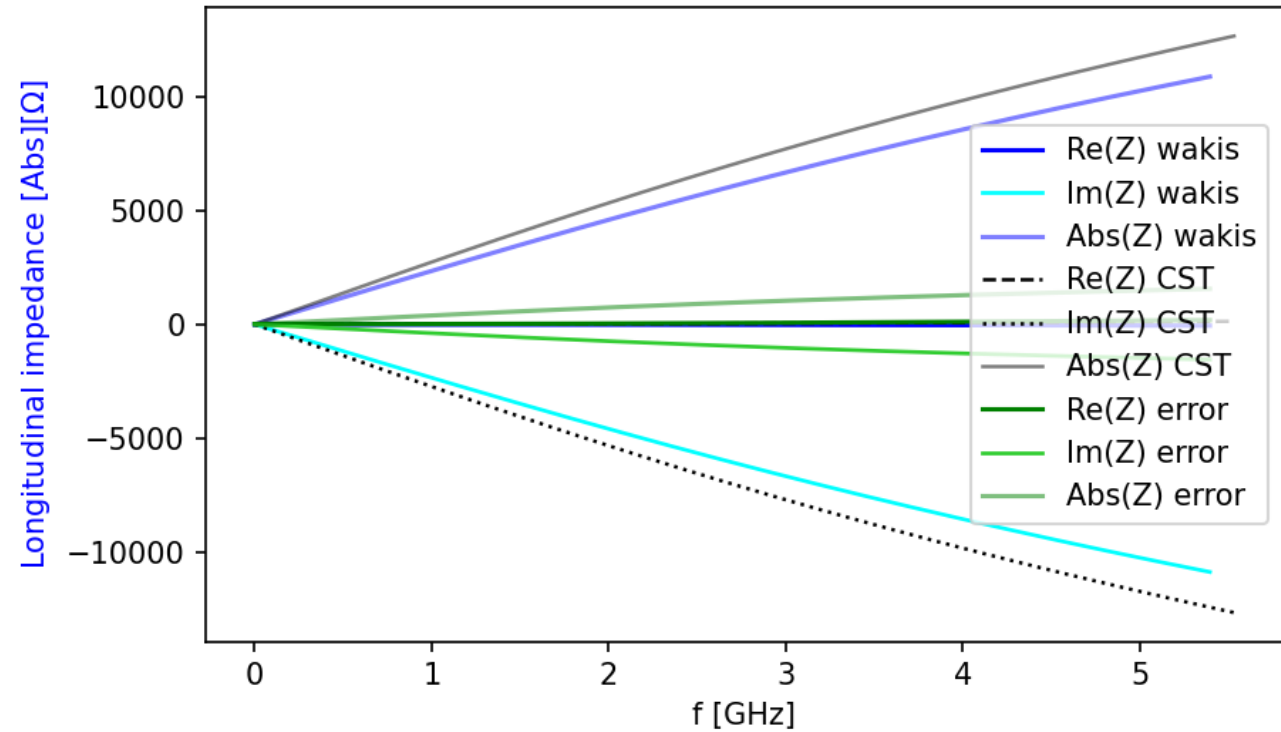
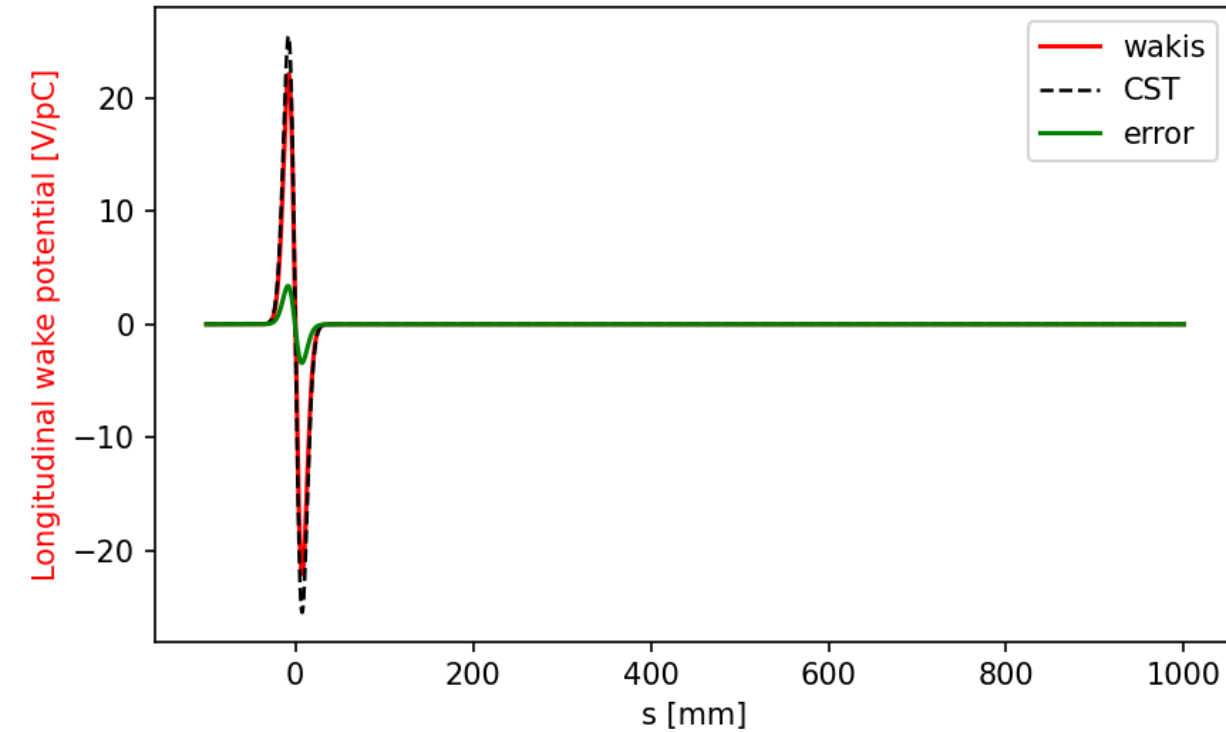
Backup: beta=0.4, different add_space

Benchmark with CST Wakefield Solver



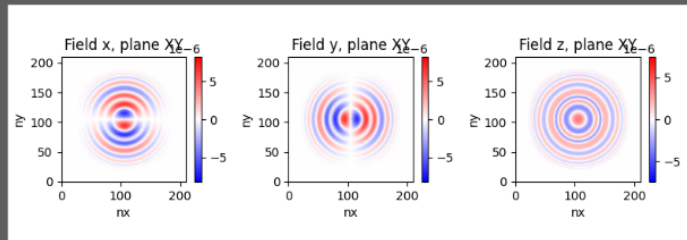
Backup: beta=0.4, error add_space=10

Benchmark with CST Wakefield Solver

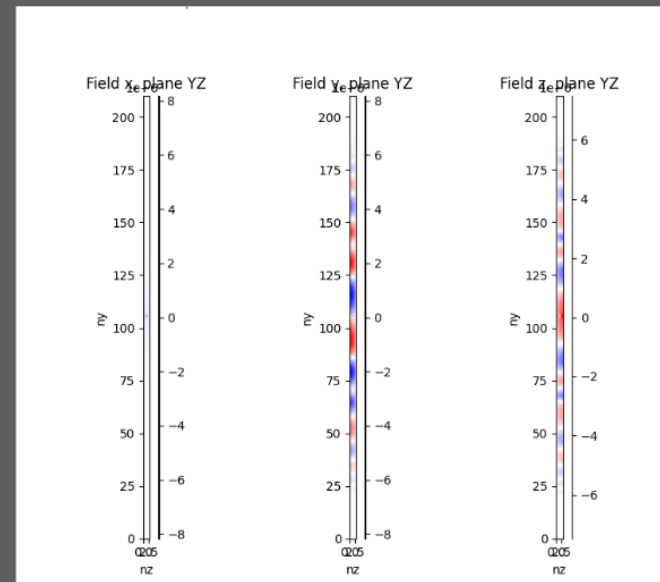


Backup: issues 3D to 2D sim

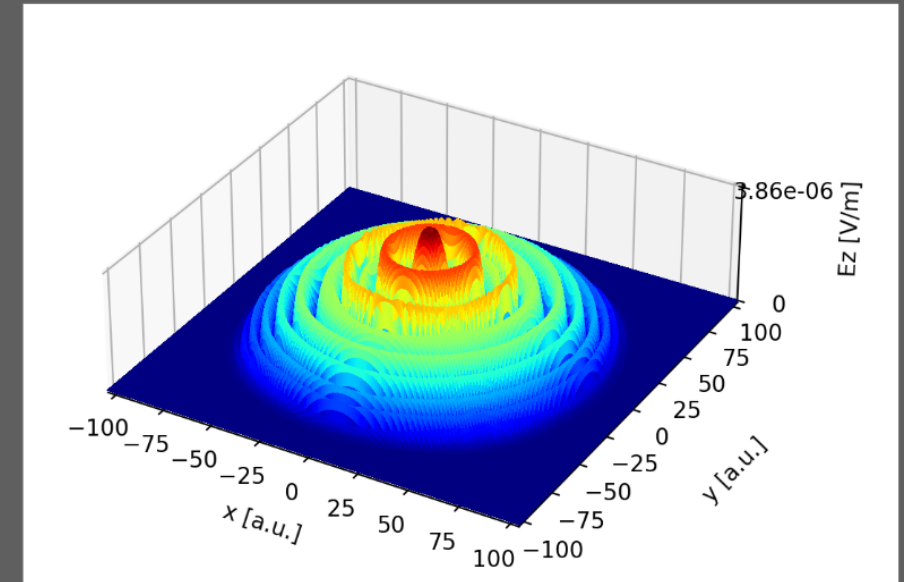
plane XY



plane YZ



3D view



Berenguer PML update equations

Not working, all fields = 0

$$\begin{aligned}
 E_y^{n+1}(i, j+1/2) &= e^{-\sigma_x(i) \Delta t/\epsilon_0} E_y^n(i, j+1/2) - \frac{(1 - e^{-\sigma_x(i) \Delta t/\epsilon_0})}{\sigma_x(i) \Delta x} \\
 &\times [H_{zx}^{n+1/2}(i+1/2, j+1/2) + H_{zy}^{n+1/2}(i+1/2, j+1/2) \\
 &- H_{zx}^{n+1/2}(i-1/2, j+1/2) - H_{zy}^{n+1/2}(i-1/2, j+1/2)] \quad (37)
 \end{aligned}$$

$$\begin{aligned}
 H_{zx}^{n+1/2}(i+1/2, j+1/2) &= e^{-\sigma_x^*(i+1/2) \Delta t/\mu_0} H_{zx}^{n-1/2}(i+1/2, j+1/2) \\
 &- \frac{(1 - e^{-\sigma_x^*(i+1/2) \Delta t/\mu_0})}{\sigma_x^*(i+1/2) \Delta x} \\
 &\times [E_y^n(i+1, j+1/2) - E_y^n(i, j+1/2)], \quad (38)
 \end{aligned}$$

$$\begin{aligned}
 e^{n+1.5} &= \exp(-\tilde{D}_\epsilon^{-1} \tilde{D}_\kappa \Delta t) e^{n+0.5} \\
 &+ (1 - \exp(-\tilde{D}_\epsilon^{-1} \tilde{D}_\kappa \Delta t)) \tilde{D}_\kappa^{-1} D_A^{-1} C D_s D_\mu^{-1} b^{n+1} \\
 &- (1 - \exp(-\tilde{D}_\epsilon^{-1} \tilde{D}_\kappa \Delta t)) \tilde{D}_\kappa^{-1} j^{n+1} \\
 h^{n+1} &= h^n - \Delta t \tilde{D}_s D_\mu^{-1} D_A^{-1} C e^{n+0.5}
 \end{aligned}$$

```

def one_step_etd(self):
    '''[TODO] Not working
    ...

    if self.step_0:
        self.set_ghosts_to_0()
        self.step_0 = False

    #cleanup
    del self.itDaiDepsDstC

    #pre-compute
    a, b = 1.0, self.sigma.toarray()
    isigma = np.divide(a, b, out=np.zeros_like(b), where=b!=0)

    self.iDsigma = diags(isigma, shape=(3*self.N, 3*self.N), dtype=float)
    self.Dexp = diags(np.exp(-self.ieps.toarray()*self.sigma.toarray()*self.dt),
                      shape=(3*self.N, 3*self.N), dtype=float)
    self.oneMinusDexp = diags(1.0-np.exp(-self.ieps.toarray()*self.sigma.toarray()*self.dt),
                              shape=(3*self.N, 3*self.N), dtype=float)
    self.itDaiDsigmaDstC = self.itDa * self.iDsigma * self.Ds * self.C.transpose()

    del a, b, isigma
    self.attrcleanup()

    self.H.fromarray(self.H.toarray() -
                    self.dt*self.tDsiDmuiDaC*self.E.toarray()
                    )

    self.E.fromarray(self.Dexp*self.E.toarray() +
                    (self.oneMinusDexp)*self.itDaiDsigmaDstC*self.H.toarray() -
                    (self.oneMinusDexp)*self.iDsigma*self.J.toarray()
                    )

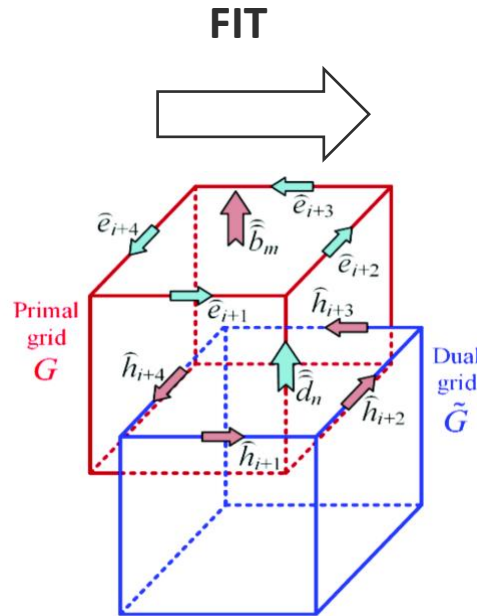
    self.J.fromarray(self.Dsigma*self.E.toarray())

    #update ABC
    if self.activate_abc:
        self.update_abc()
    
```

FIT theory: Grid Maxwell Equations

$$\left\{ \begin{aligned} \oint_{\partial A} \mathbf{E} \cdot d\mathbf{s} &= - \iint_A \frac{\partial \mathbf{B}}{\partial t} \cdot d\mathbf{A} \\ \oint_{\partial A} \mathbf{H} \cdot d\mathbf{s} &= - \iint_A \left(\frac{\partial \mathbf{D}}{\partial t} + \mathbf{J} \right) \cdot d\mathbf{A} \\ \oiint_{\partial V} \mathbf{B} \cdot d\mathbf{A} &= 0 \\ \oiint_{\partial V} \mathbf{D} \cdot d\mathbf{A} &= \iiint_V \rho \, dV \end{aligned} \right.$$

$$\underline{\underline{\mathbf{D}}} = \underline{\underline{\epsilon}} \mathbf{E}, \quad \underline{\underline{\mathbf{B}}} = \underline{\underline{\mu}} \mathbf{H}, \quad \underline{\underline{\mathbf{J}}} = \underline{\underline{\sigma}} \mathbf{E} + \rho \mathbf{v}$$



Grid Maxwell Equations

$$\mathbf{C} \mathbf{D}_s \mathbf{e} = - \mathbf{D}_A \frac{\partial \mathbf{b}}{\partial t}$$

$$\tilde{\mathbf{C}} \tilde{\mathbf{D}}_s \mathbf{h} = \tilde{\mathbf{D}}_A \left(\frac{\partial \mathbf{d}}{\partial t} + \mathbf{j} \right)$$

$$\mathbf{S} \mathbf{D}_A \mathbf{b} = \mathbf{0}$$

$$\tilde{\mathbf{S}} \tilde{\mathbf{D}}_A \left(\frac{\partial \mathbf{d}}{\partial t} + \mathbf{j} \right) = \mathbf{0}$$

$$\mathbf{d} = \tilde{\mathbf{D}}_\epsilon \mathbf{e}, \quad \mathbf{b} = \mathbf{D}_\mu \mathbf{h}, \quad \mathbf{j} = \tilde{\mathbf{D}}_\sigma \mathbf{e} + \mathbf{D}_\rho \mathbf{v}$$

$$\epsilon = \left(\epsilon_r + \frac{\sigma}{j\omega} \right) \epsilon_0$$

- Operators
- Spatial matrixes
- Material properties

What we care about:
Update equations

$$\mathbf{h}^{n+1} = \mathbf{h}^n - \Delta t \tilde{\mathbf{D}}_s \mathbf{D}_\mu^{-1} \mathbf{D}_A^{-1} \mathbf{C} \mathbf{e}^{n+0.5}$$

$$\mathbf{e}^{n+1.5} = \mathbf{e}^{n+0.5} + \Delta t \mathbf{D}_s \tilde{\mathbf{D}}_\epsilon^{-1} \tilde{\mathbf{D}}_A^{-1} \tilde{\mathbf{C}} \mathbf{h}^n$$

$$- \tilde{\mathbf{D}}_\epsilon^{-1} \mathbf{j}_{beam}^n - \tilde{\mathbf{D}}_\epsilon^{-1} \tilde{\mathbf{D}}_\sigma \mathbf{e}^{n+0.5}$$

We need to build all these matrices and then apply these equations every timestep !