Parallel multigrid solver for time harmonic Maxwell equations

M. Chanaud[1] L. Giraud[1] D. Goudin[2] JJ. Pesqué[2] J. Roman[1]

[1]HiePACS project, INRIA Bordeaux Sud-Ouest, Talence, France
joint laboratory INRIA-CERFACS on HPC

[2]CEA/DAM/CESTA, Le Barp, France

Sparse Days 2011,
Sep. 6-7 2011, Toulouse, France

Context

Aim: compute large 3D stealthy objects Radar Cross Section (RCS).
We consider a concentric layer decomposition of a global domain Ω

- Innermost boundary (Σ_0) represent the studied object
- Ω_n outer boundary (Σ_n) represent the infinite domain

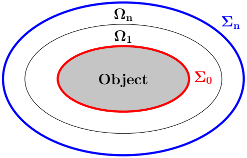

Must solve time harmonic Maxwell equations over all these domains

$$\nabla \times \nabla \times \vec{E} - k^2 \vec{E} = 0$$

This talk will focus on volume subdomain solutions.

Global domain resolution scheme

- Volume subdomains discretized by Nédélec's 3D first order finite elements
  $\longrightarrow$ sparse symmetric complex linear system. i
- Radiation condition applies on the outer surface Σ_n. Despres Integral Equations (EID) formulation
  $\longrightarrow$ dense complex linear system.

Each volume subdomain is coupled to its neighbors by a Robin boundary condition (transmission).

Global problem solved by iterative block Gauss-Seidel method:

- Direct method in volume subdomains
- Coupling boundary condition
- Iterative + Fast Multipole Method on outer surface subdomain (S_n)

Volume subdomains

Direct methods (LL^T, LDL^T factorization) are used on volume subdomains problems.
Performed in actual software suite by PaStiX parallel direct solver.

`http://pastix.gforge.inria.fr/`

Advantages & disadvantages

+ Solution accuracy
+ Predictable behavior (static mapping)
+ Multiple right-hand sides
+ Triangular solution complexity ($O(n^{4/3})$)
− Huge memory requirements
− Factorization complexity ($O(n^2)$)

Classical problems require tens of millions of unknowns per volume subdomain.

Full-Multigrid

Compute on multiple level of meshes.
Direct method on coarse mesh (small initial problem).
Solution interpolated on fine mesh and solved iteratively (smoother).

Advantages & disadvantages

- + Low memory consumption (*small* factorized matrix)
- + Linear multigrid complexity ($O(n)$)
- + Automatic uniform refinement
- – Less accurate than direct method

Nédélec's base element

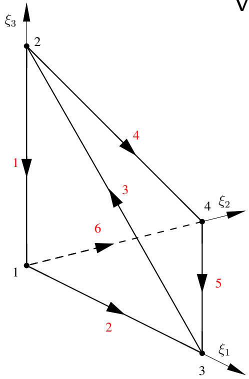

Unknowns are carried by edges rather than nodes.
Values are computed according to the following equation:

$$\alpha_i = \int_{\Gamma_i} \vec{E}(\xi_1, \ldots \xi_n) \cdot \vec{\tau}_i dl.$$

Where : Γ_i – the considered edge,
$\vec{\tau}_i$ – Γ_i unitary vector,
dl – displacement along Γ_i.

Multigrid origins

- Some iterative methods reduce High Frequency (HF) error modes
- Remaining error composed of Low Frequency (LF) error modes
- LF on fine grid becomes HF on coarse grid

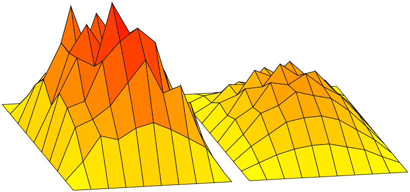

Figure: Error over a 10x10 regular mesh (before and after smoothing)

Multigrid V-cycle and full-multigrid

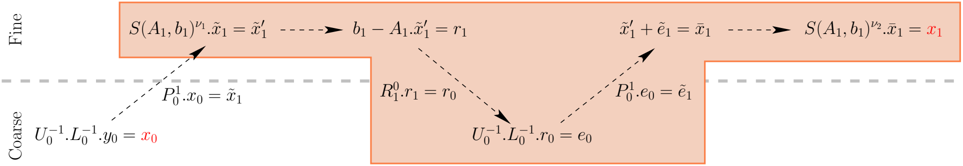

Multigrid V-Cycle

- Coarse grid factorization performed **once**
- Multiple coarse grid solution
- Residual restricted from fine to coarse mesh
- Solution/error prolongated from coarse to fine mesh

Mesh refinement and coarsening

Unstructured mesh refinement :

- uniform : easy coarsening
- adaptive : hard coarsening or storage needed

Two uniform refinements :

- edge refinement : split each edge
  - \+ good aspect ratio,
  - − large number of element (12x elements),
  - − coupling interface modified.

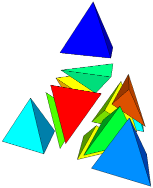

- cell refinement : one new vertex per elt.
  - \+ coupling interfaces preserved,
  - \+ slow problem size increase (4x elements),
  - − aspect ratio (flat elements).

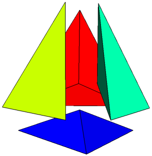

Jacobi smoother

Jacobi decomposition

Consider the following componentwise decomposition of A :

$$A = D + L + U.$$

Where : D – diagonal matrix
L – lower triangular
U – upper triangular

The Jacobi iteration is defined by

$$x^{k+1} = x^k + D^{-1}(b - Ax^k)$$

Damped Jacobi iteration

$$x^{k+1} = (1-\omega)x^k + \omega(x^k + D^{-1}(b - Ax^k))$$
$$x^{k+1} = x^k + \omega(D^{-1}(b - Ax^k))$$

Matrix-free Jacobi

From the damped Jacobi equation

$$x^{k+1} = x^k + \omega(D^{-1}(b - Ax^k)),$$

one Jacobi iteration consists "only" in a matrix-vector product.

Given the assembly equality

$$A = \sum_{elements} A_e,$$

matrix-vector multiplication can be performed without global matrix assembly

- $\Rightarrow$ Low memory consumption
- $\Rightarrow$ Only vectors are stored
- $\Rightarrow$ Every element matrix computed multiple times ("cheap")

Parallel aspects

- First approximation computed by a parallel direct method
- Unknowns renumbered by nested dissection algorithm and halo approximate minimum degree (PT-Scotch software)

  `http://scotch.gforge.inria.fr/`
- Static mapping computed by PaStiX

→ Use this data distribution for the iterative solver.

+ No need to distribute data accross processors before each direct solve
+ Load balance preserved by uniform refinement
− Edge distribution ≠ cell distribution

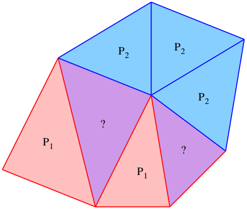

Multigrid solver – Software design

- Elementary matrices computed “on the fly”
- EM code acts as a service provider (direct solver driven)
- Multigrid library in C
- EM code developed in Fortran

→ Need a coupling code

Multigrid solver – Software design

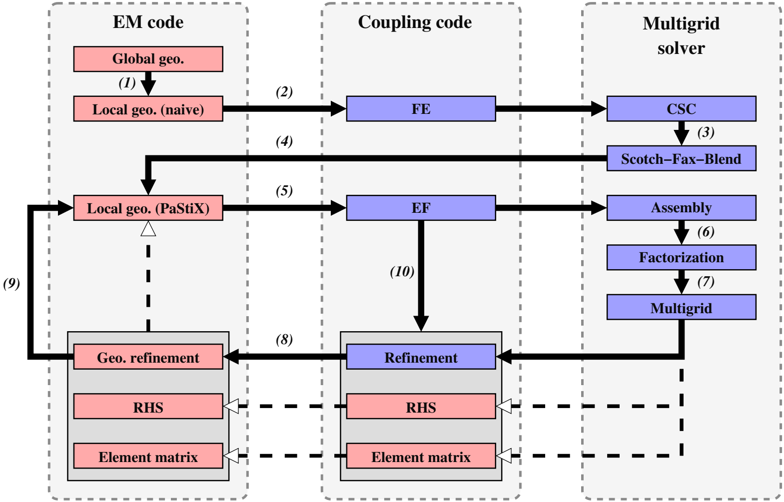

Validation – Variable material properties

Sphere test case

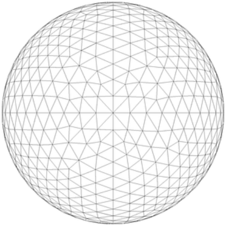

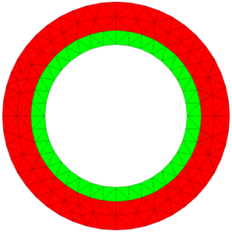

- Constant mesh and multigrid cycles (FMG 10 V(2,2))
- Increase k by modifying coating material properties (green material in the graph above)

Validation – Variable material properties

Table: DDL per wavelength, sphere test case

freq.	lev.0	lev.1	lev.2	ϵ	μ	n	k
low	12	17	23	2	1	1.41	4.44
low	6	9	11	4	2	2.82	8.89
low	3	4	6	8	4	5.64	17.77
low	12	17	23	1+1i	1+1i	1.41	4.44
low	6	9	11	2+2i	2+2i	2.82	8.89
low	3	4	6	4+4i	4+4i	5.64	17.77
unks.	14 625	57 909	231 045				
edges	11mm	8mm	6mm				

Validation – Variable material properties

Multigrid solver :

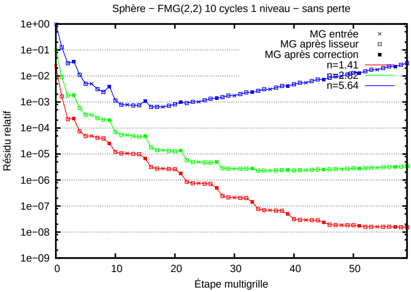

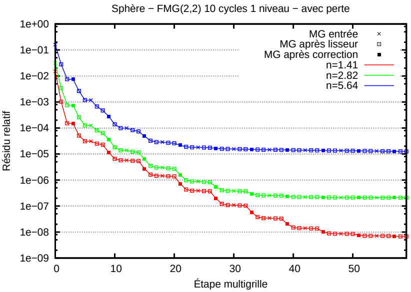

Figure: MG convergence for lossless material (left) and absorbing material (right)

Validation – Variable material properties

Multigrid preconditioned GMRES :

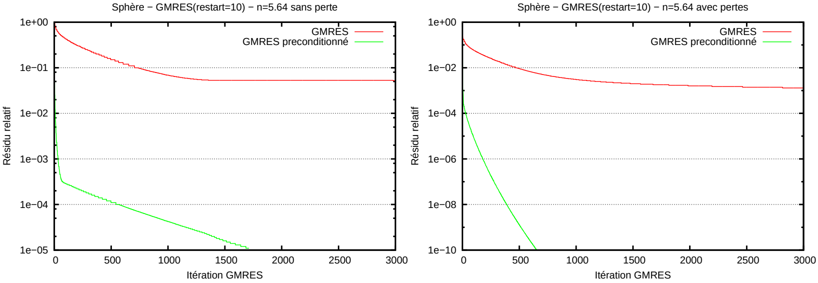

Figure: GMRES convergence for lossless material (left) and absorbing material (right)

Validation – Variable frequency

Haltere test case

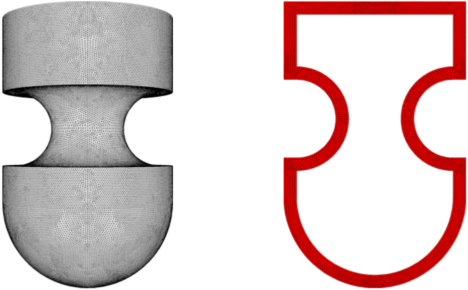

- Constant mesh and multigrid cycles (FMG 10 V(2,2))
- Increase k by modifying the incident wave frequency

Validation – Variable frequency

Table: DDL per wavelength, `haltere` test case (fixed n=1.41).

freq.	lev.0	lev.1	lev.2	lev.3	lev.4	lev.5	n	k
1 GHz	20	29	38	49	62	78	1.41	29.6
2 GHz	10	14	19	24	31	39	1.41	59.2
3 GHz	7	10	13	16	21	26	1.41	88.8
5 GHz	4	6	8	10	12	16	1.41	148.1
unks.	1.2M	5.2M	21.1M	84.6M	338.5M	1.3B		
edges	10mm	7mm	5mm	4mm	3mm	2mm		

Validation – Variable frequency

1 level multigrid (relative residual) :

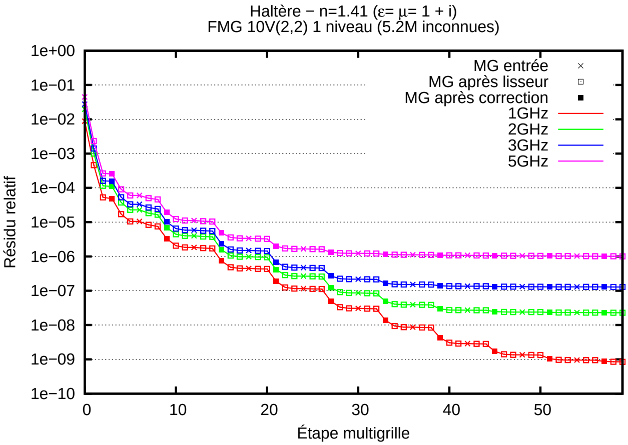

Validation – Variable frequency

2 level multigrid (relative residual):

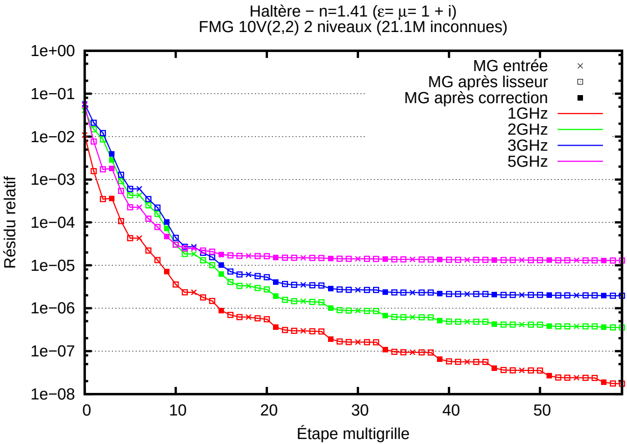

Validation – Strong scaling

- Haltere test case
- Fixed coarse mesh
- Variables number of cores
- Run on TERA100 nodes (4 Intel Xeon processors, 32 cores/node)
- Use only 2 cores per node (memory constraints)

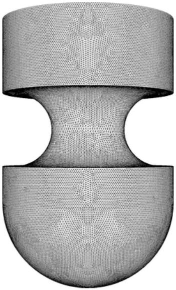

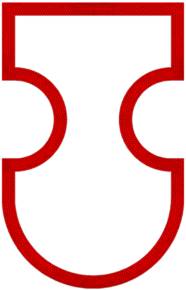

Validation – Strong scaling

Table: fine 84.6M – coarse 1.2 M – FMG(2,2) 3 levels

Cores	Init.	CSC	Assemb.	Facto.	Solve	10V
16	123.98	0.48	0.42	79.62	0.55	2285.58
32	118.67	0.31	0.23	82.72	0.30	1369.26
64	103.10	0.15	0.12	68.18	0.19	710.57
128	102.69	0.08	0.06	60.94	0.11	399.23
256	130.32	0.03	0.03	75.27	0.21	194.42
512	170.97	0.01	0.01	78.81	0.06	100.17

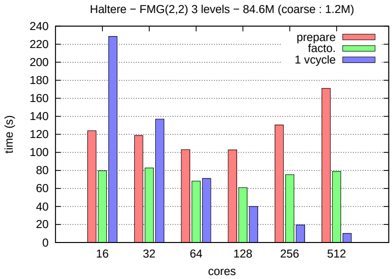

Validation – Strong scaling

Table: fine 84.6M – coarse 5.2 M – FMG(2,2) 2 levels

Cores	Init.	CSC	Assemb.	Facto.	Solve	10V
128	229.34	0.36	0.24	103.04	0.26	354.23
256	286.71	0.10	0.12	95.40	0.20	174.91
512	369.69	0.05	0.06	98.56	0.19	85.89
1024	480.22	0.02	0.03	98.96	0.26	65.88
2048*	1331.51	0.01	0.02	107.95	0.27	47.29

(*) 4 cores per nodes instead of 2 (16384 reserved cores).

Validation – Strong scaling

Table: fine 338.5M – coarse 21.1 M – FMG(2,2) 2 levels

Cores	Init.	CSC	Assemb.	Facto.	Solve	10V
256	662.22	0.43	0.52	117.96	0.64	628.74
512	881.20	0.20	0.23	125.56	0.56	345.00
1024	1196.74	0.13	0.14	151.56	0.56	213.49

Validation – Strong scaling

Table: fine 1.3B – coarse 21.1 M – FMG(2,2) 3 levels

Cores	Init.	CSC	Assemb.	Facto.	Solve	10V	(20V)
1024	1160.28	0.12	0.14	147.73	0.55	857.69	(1715.38)

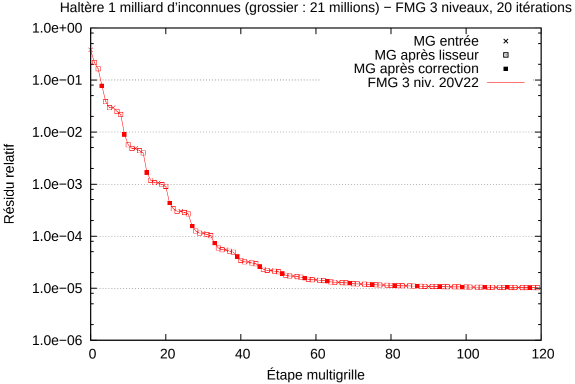

Conclusion

Conclusion :

- Parallel multigrid solver fully functionnal and scalable
- Can be used as a solver for simple problems (absorbing materials)
- Can be used as a preconditioner for complex problems (lossless materials)

Perspectives :

- Lossless materials : apply a complex shift to the preconditioner
- Modify refinement to ensure tetrahedron aspect ratio
- Adaptive distribution depending on levels