CEA's Parallel industrial codes in electromagnetism

CEA/DAM - France

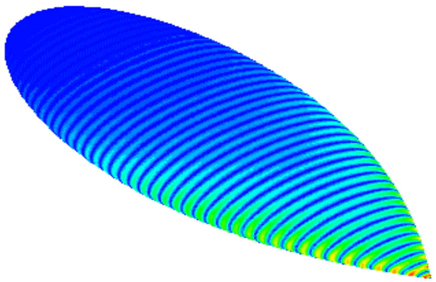

David GOUDIN

Michel MANDALLENA

Katherine MER-NKONGA

Jean Jacques PESQUE

Muriel SESQUES

Bruno STUPFEL

Who are we ?

CEA : French Atomic Energy Commission

CESTA : One center of CEA, located near Bordeaux, work on the electromagnetic behavior of complex target

Measurement devices

Simulation codes

- ARLAS
- ARLENE
- ODYSSEE

Terascale computer

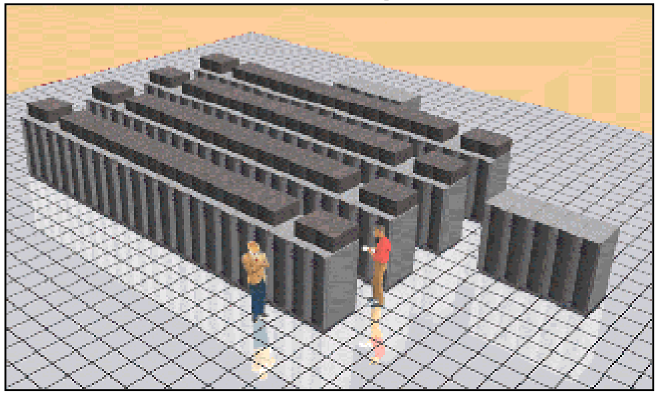

CEA's Parallel industrial codes in electromagnetism

The Simulation Program of CEA/DAM: reproduce the different stages of the functioning of a nuclear weapon through numerical simulation .

The project TERA : one of the 3 major systems of the Simulation Program

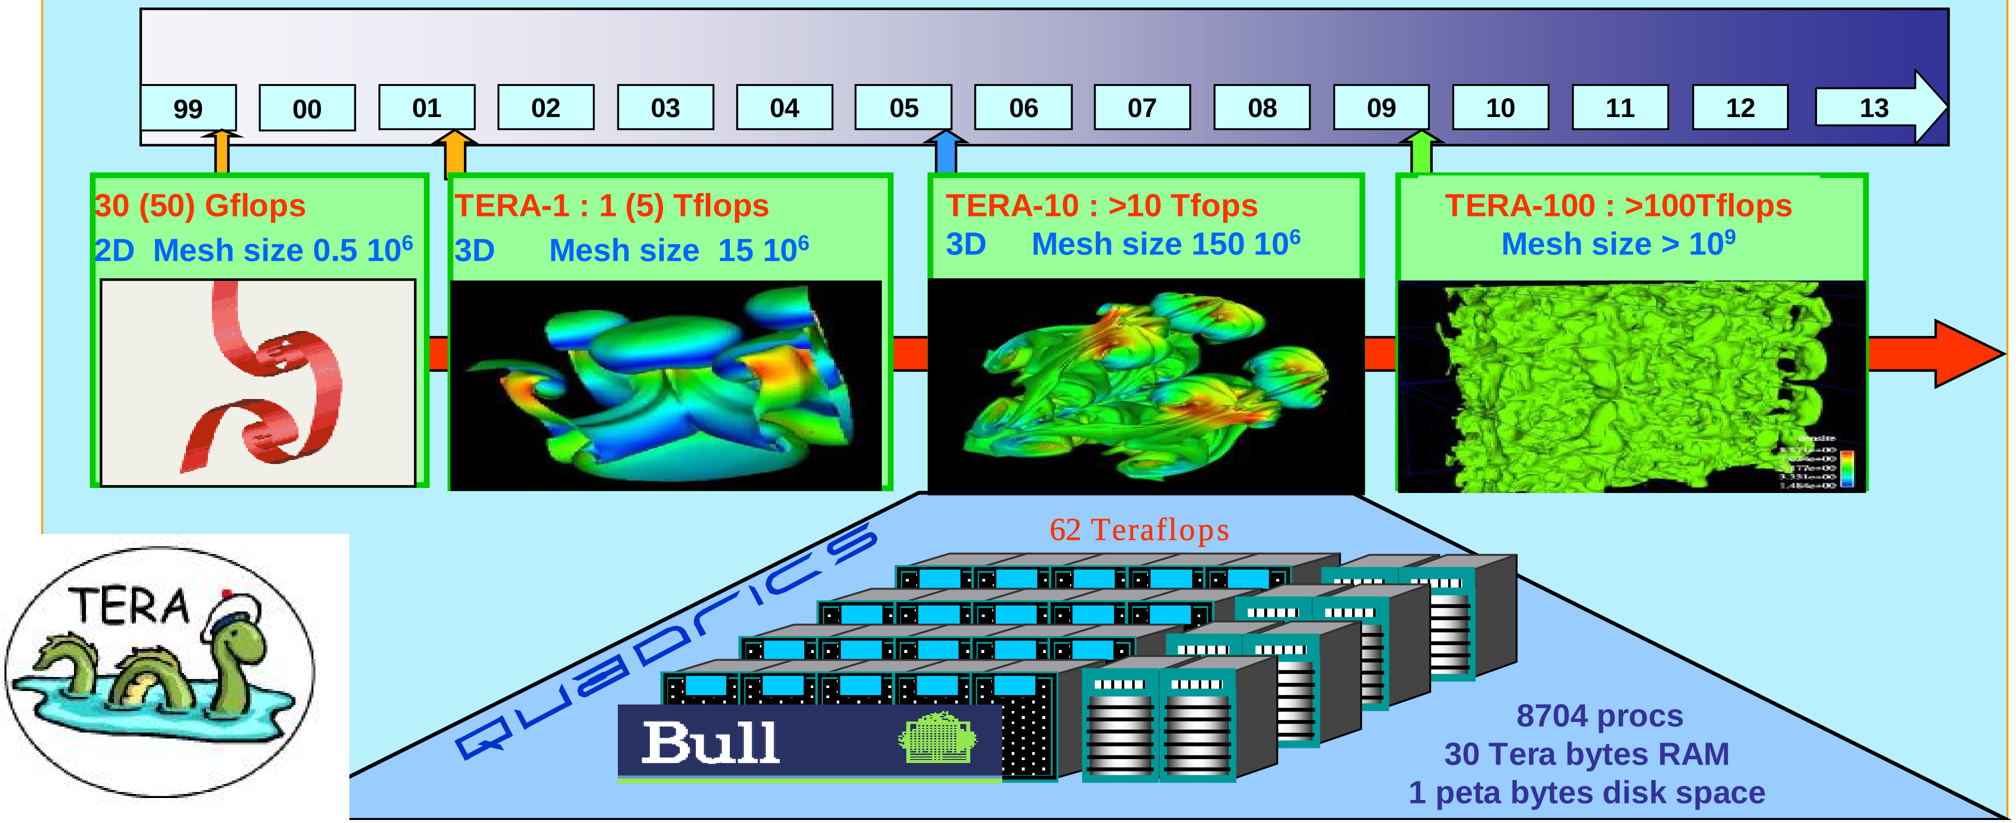

The supercomputer HP/Compaq TERA 1

- **5 Teraflops of peak performance**

- 2560 processors EV 68 (1Ghz)
  640 SMP nodes ES45 (4 processors/node)
- 2.5 Tera bytes of RAM
- 50 Tera bytes of disk space
- High-performance interconnection network (Quadrics)

⇨ **Benchmark Linpack : 3.98 Teraflops**

- 2005 - 2006 : the next level = 30 Teraflops

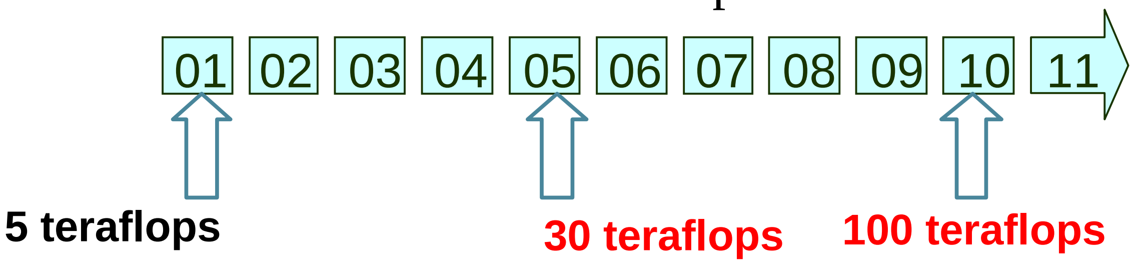

The supercomputer TERA-10

Main characteristics	
Installation	December 2005
SMP Nodes	544*16 proc.
Peak performance (Benchmark CEA)	> 60 Tflops (12,5 Tflops)
RAM	30 To
Disk space	1 Po - 56 nodes (54 OSS+2 MDS)
Consummation	< 2000 kW

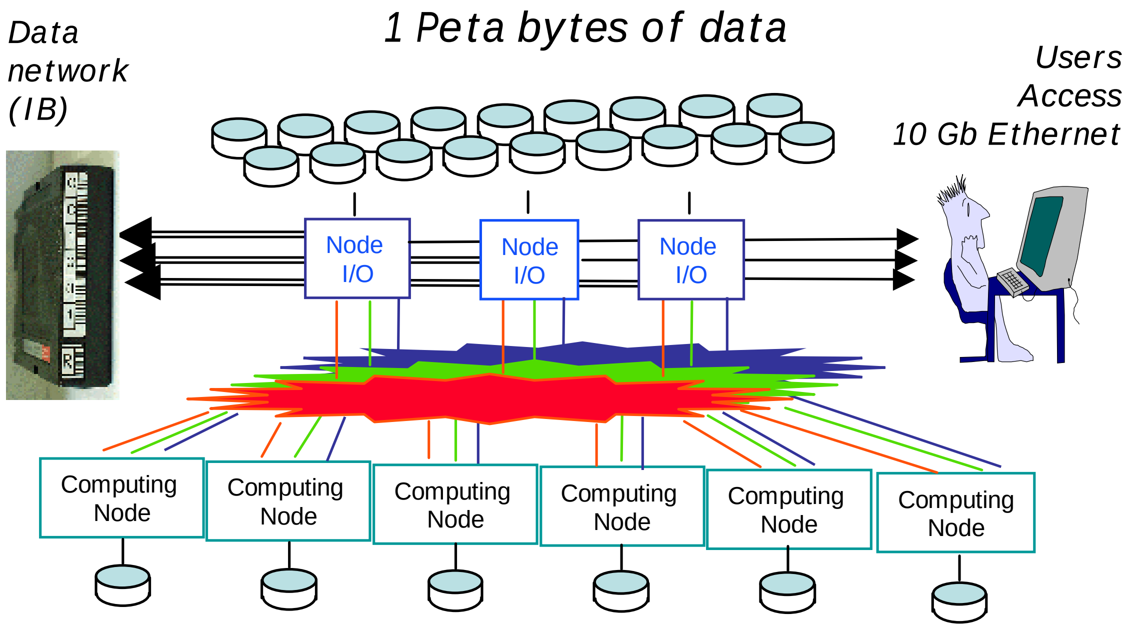

544 computing nodes

What is the problem ?

The simulation of the electromagnetic behavior of a complex target in frequency domain

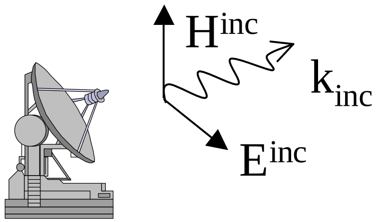

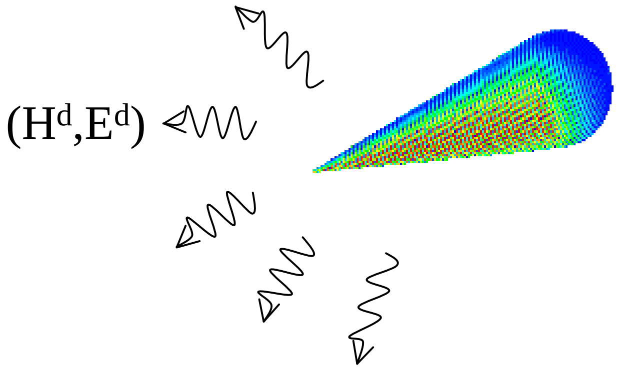

⇨ Numerical resolution of Maxwell's equations in the free space

- RCS (Radar Cross Section) computation
- Antenna computation

$$RCS = 10 \log_{10}(\sigma) \qquad \sigma = \lim_{r \to +\infty} 4\pi r^2 \frac{\left|\vec{E}^{d}(r)\right|^2}{\left|\vec{E}^{inc}(r)\right|^2}$$

3D Codes in production before 2004 :

ARLENE (since 1994)

ARLAS (since 2000)

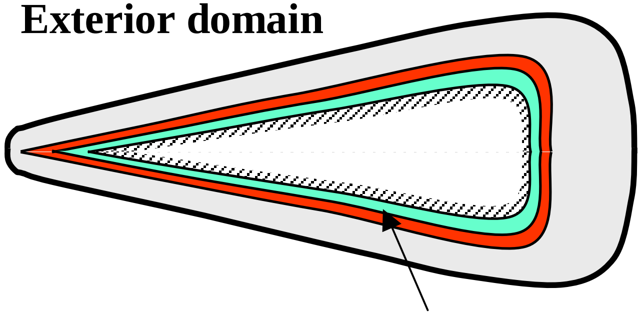

Based on the following numerical method

- Domain decomposition
- Strong coupling of numerical method between domains
  - -BIEM : Boundary Integral Equations Method
  - -PDE : Partial Derivative Equations Method
- Lead to solve a linear system A.x = b
- Solved by direct Method due to the characteristics of *A*

} Finite Elements Method 's discretization

EMILIO (industrial software tool, includes PaStiX, MPI version), developed in collaboration with ScAlApplix team (INRIA, Bordeaux)

CEA's Parallel industrial codes in electromagnetism

ARLENE

- **Fully BIEM**
- **Meshes at the surface**
  (interfaces between homogeneous isotropic medium)
  - number of unknowns reasonable
  - lead to a full matrix

$$\begin{pmatrix} A^s_{22} & A^s_{23} \\ A^s_{23} & A^s_{33} \end{pmatrix} \begin{pmatrix} M \\ J \end{pmatrix}$$

full part
sparse part

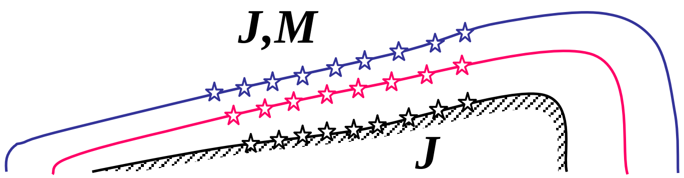

⇨ Parallelization : *very efficient*

ARLAS

- **Hybrid method**
  - free space problem by BIEM
  - interior problem by PDE
- **Hybrid meshes - on the outer boundary and in the volume**
  - big number of unknowns in the volume
  - lead to a matrix with full and sparse part

$$\begin{pmatrix} A^v_{11} & A^v_{12} & \mathbf{0} \\ A^v_{12} & A^v_{22} + A^s_{22} & A^s_{23} \\ \mathbf{0} & A^s_{23} & A^s_{33} \end{pmatrix} \begin{pmatrix} E \\ M \\ J \end{pmatrix}$$

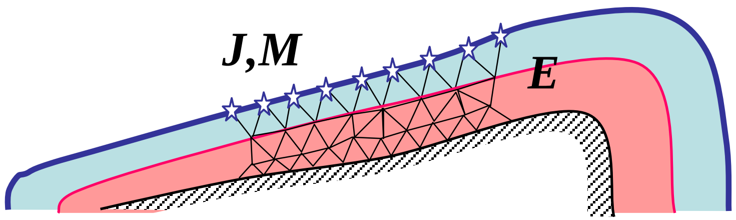

⇨ Parallelization : more difficult…

ARLENE

Solver

own parallel Cholesky-Crout solver :

the matrix is :

- symmetric
- complex
- non hermitian

ARLAS

Solver

Solved by Schur 's complement

⇨ by elimination of the electric field *E*

1. sparse matrix assembly : A^v_{11}

2. factorization by PaStiX, parallel sparse solver from EMILIO (INRIA-CEA)

3. computation of the Schur's complement : $(A^v_{11})^{-1} A^v_{12}$

4. dense matrix assembly $\begin{pmatrix} A^s_{22} & A^s_{22} \\ A^s_{22} & A^s_{22} \end{pmatrix}$

5. add the contribution to the dense matrix (3.)

$$\left(\begin{array}{c|c} -A^v_{12}(A^v_{11})^{-1} A^v_{12} + A^v_{22} + A^s_{22} & A^s_{23} \\ \hline A^s_{23} & A^s_{33} \end{array}\right)$$

6. factorization, resolution of the dense linear system

ARLENE - Performance tests on TERA 1

- **Cone sphere 2Ghz :** 104 000 unknowns
- **Cone sphere 5Ghz :** 248 000 unknowns
- **F117 :** 216 000 unknowns
- **NASA ALMOND** (a quarter object) - Workshop JINA '02 (18/11/02)
  - Conducting body : 237 000 unknowns (948 000 unknowns)

- **Sphere Cone** *2 GHZ :*
  - **104 000 unknowns**
  - matrix size 85.5 GBytes

Number of Procs	256	462	512	768	1024
CPU time	4781	2807	2338	1498	1414
Tflops	0.308	0.525	0.630	0.984	1.04
% / peak performance	60.15%	56.81%	61.5%	64.06%	50.8%

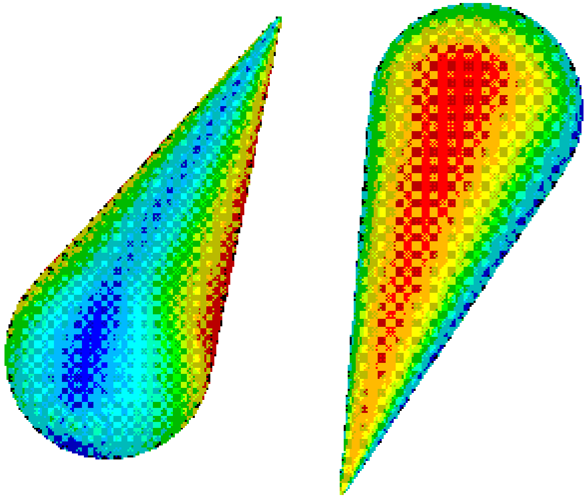

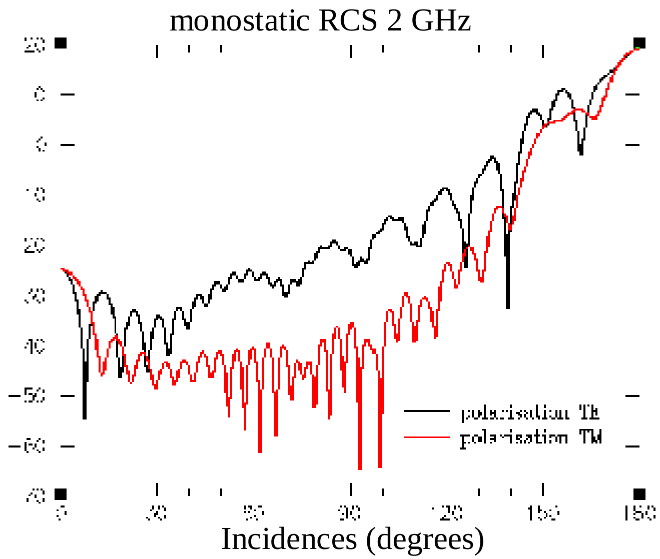

CEA's Parallel industrial codes in electromagnetism

- **Sphere Cone** *5 GHZ :*
  - **248 000 unknowns**
  - matrix size 492 GBytes

Number of Procs	1024	1536
CPU time	16 057	10 144
Tflops	1.27	2
% / peak performance	62%	65%

Electric currents for the 2 polarizations

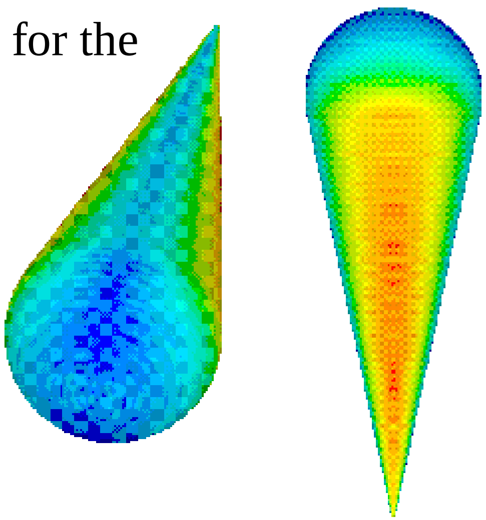

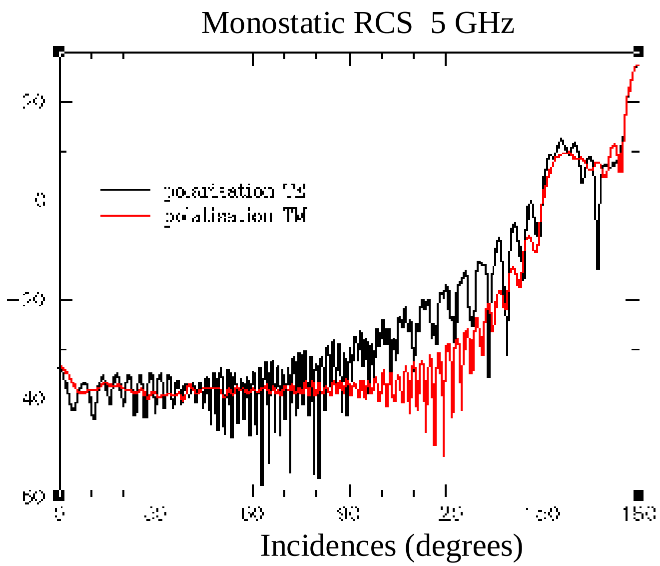

CEA's Parallel industrial codes in electromagnetism

- **F117** *500 Mhz*
  - •**216 425** unknowns
  - •matrix size : 375 Go

⇨ global CPU time on 1024 processors
10 684 seconds

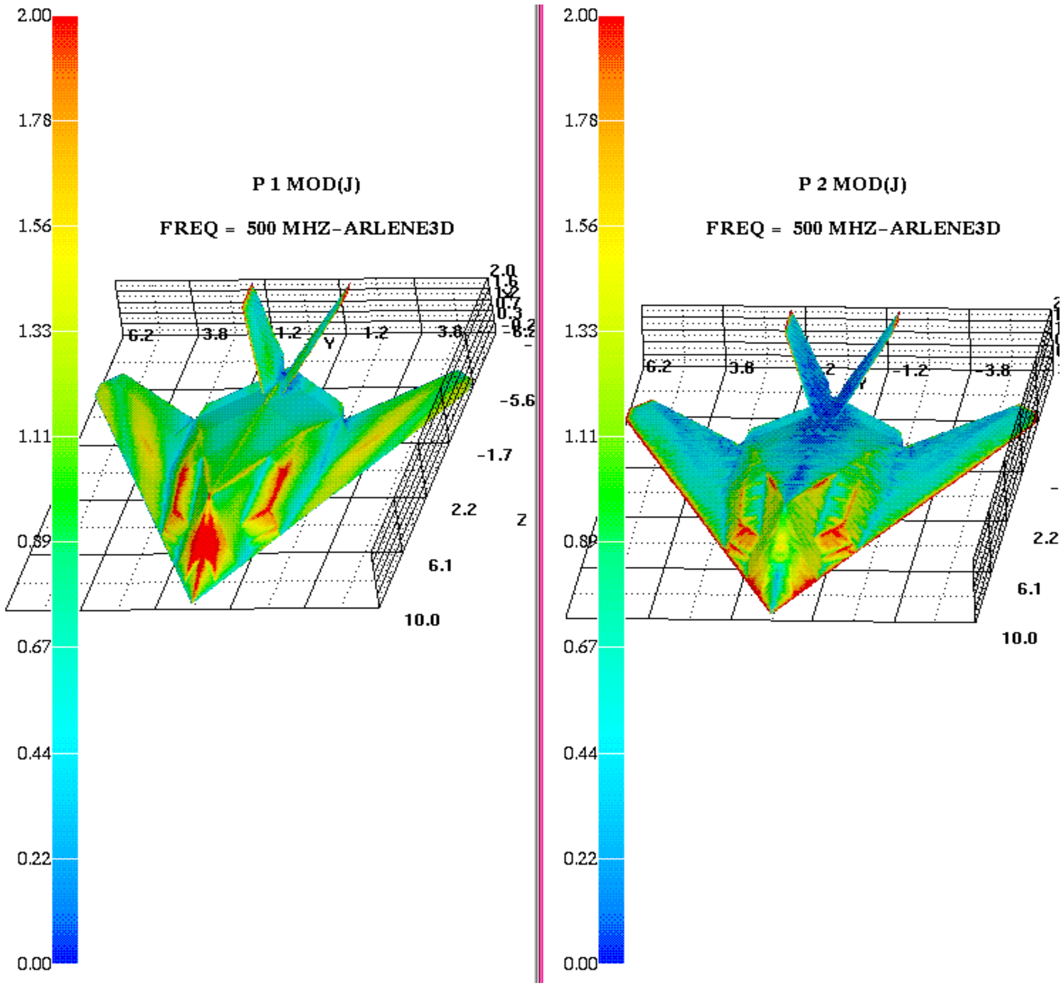

CEA’s Parallel industrial codes in electromagnetism

- **NASA Almond** *8 Ghz (Workshop JINA'02* **-** *case2* **-** *18 November 2002)*

  233 027 unknowns with 2 symmetries -> 932 108 unknowns
  - matrix size : 434 GBytes
  - 4 factorizations & 8 resolutions to compute the currents
  - global CPU time on 1536 processors : **36 083 seconds**

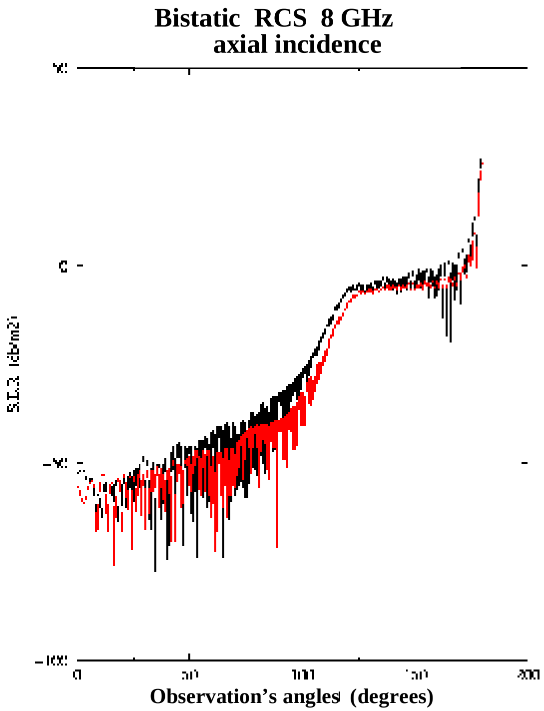

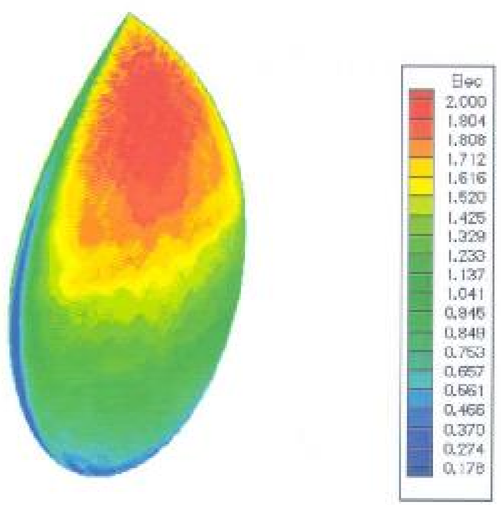

ARLAS - Performance tests on TERA 1

- **Monopole antenna :** 241 500 unknowns in the volume
  9 400 unknowns on the outer boundary

- **Cone sphere :** 1 055 000 unknowns in the volume
  14 600 unknowns on the outer boundary

- **Monopole antenna** *(Workshop JINA'02 - case3 - 18 November 2002)*

  241 500 unknowns in the volume
  9 400 unknowns on the outer boundary

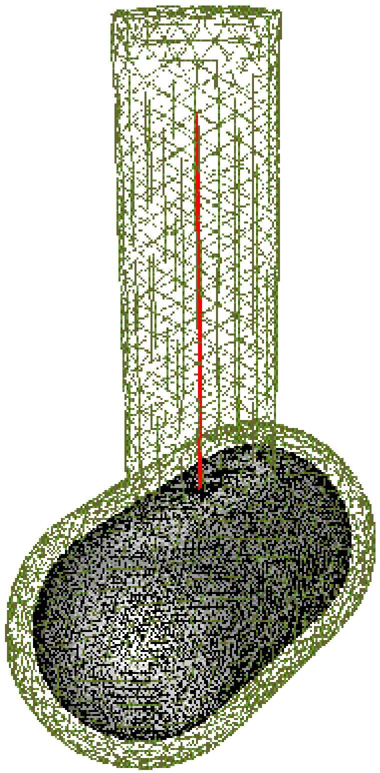

Number of Procs	factorization *sparse*	up&down resolutions	Schur's complement	factorization *full*
32	13.7s	0.3s	**1 467s**	28s

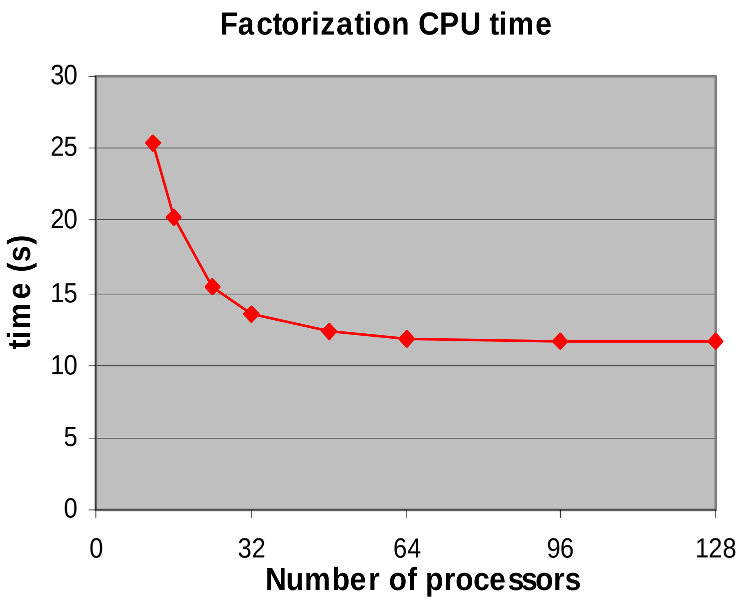

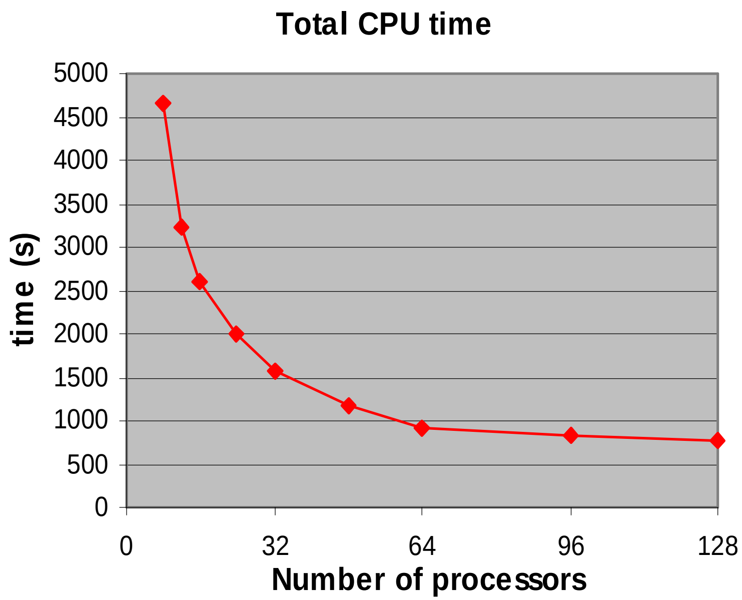

CEA's Parallel industrial codes in electromagnetism

- **Sphere cone :** 1 055 000 unknowns in the volume
  14 600 unknowns on the outer boundary

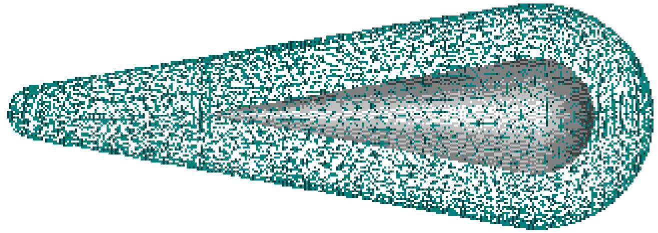

Mesh fragmented sight

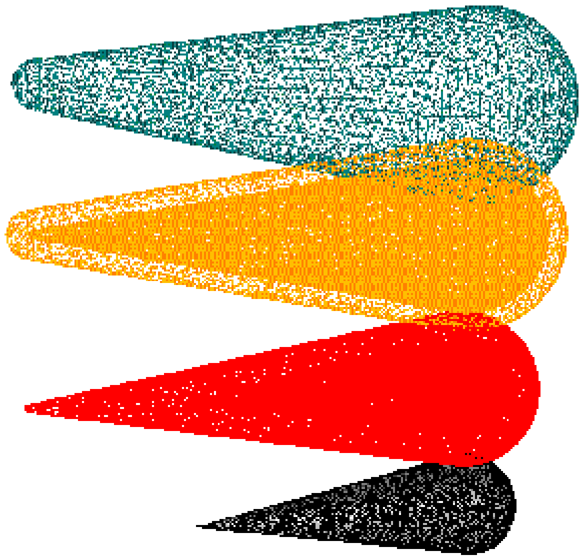

Number of Procs	factorization	up&down resolutions	Schur's complement
192	142.47s	1,9s	**14 298s**

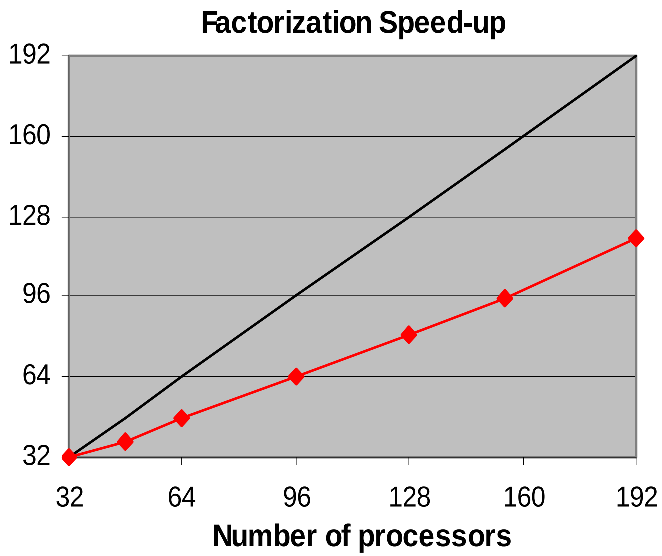

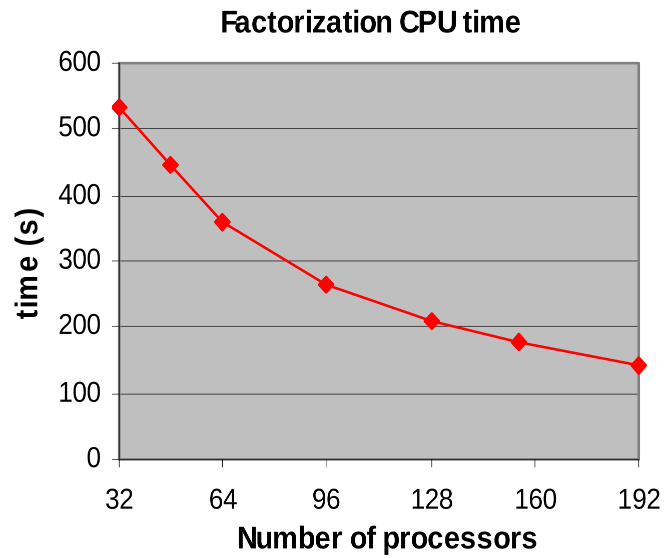

Conclusions

- **ARLENE :** - is limited to 450 000 unknowns on the outer boundary
  - reach : 65% of the performance peak until 1500 processors
  : 58% until 2048 processors

- **ARLAS :** - is limited to 3 million unknowns in the volume
  25 000 unknowns on the outer boundary

But : for our future applications it's not enough
we'll need more than :
- 30 million unknowns in the volume
- 500 000 unknowns on the outer boundary

Solution ? A new 3D code : ODYSSEE (since 2004)

Based on :

* a domain decomposition method (DDM) partitioned into concentric sub-domains
* a Robin-type transmission condition on the interfaces
* on the outer boundary, the radiation condition is taken into account by a new Boundary Integral formulation called EID
* we solve this problem with 1 or 2 iterative method's level
  - Gauss-Seidel for the global problem

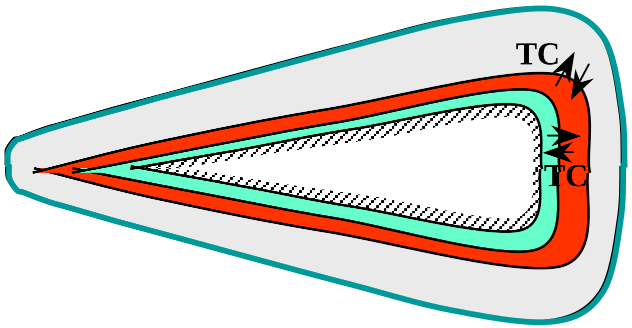

Inside each sub-domain

- , iterative solver (// conjugate gradient)
- the PaStiX direct solver

for the free space problem

- , a // Fast Multipole Method
- a direct Scalapack solver

What about our constraints ?

- 3 dimensional targets
- Complex materials :
  - inhomogeneous medium layer
  - isotropic or/and anisotropic medium layer
  - conductor or high index medium layer taken into account by a Leontovich impedance condition
  - wire, conducting patches, …
- We need a very good accuracy : low level of RCS
- We want to reach medium and high frequencies
  ⇒ body's size about 20 to 100 λ
- We want a parallel industrial code which is able to run on the super computer we have

Numerical issues

- **Unbounded domain** : the domain can be cut by any kind of ABC
  - ⇨ Absorbing Boundary Condition = not accurate enough for
  - ⇨ our applications

    The radiation condition is taken into account by a Boundary Integral Equation

    þ leads to a linear system with a **full matrix**
- **Discretization**
  - ⇨ Mesh's size = f(λ) þ the higher the frequency, the bigger the problem's size is

ODYSSEE needs ➡

- Numerical methods
- High Performance Computing

ODYSSEE - Key ideas

- Domain Decomposition Method
  the computational domain is partitioned into concentric sub-domains
- On the outer boundary, the radiation condition is taken into account by a new BIEM called EID
- A Multi Level Fast Multipole Algorithm has been applied to this IE in order to reduce the CPU time

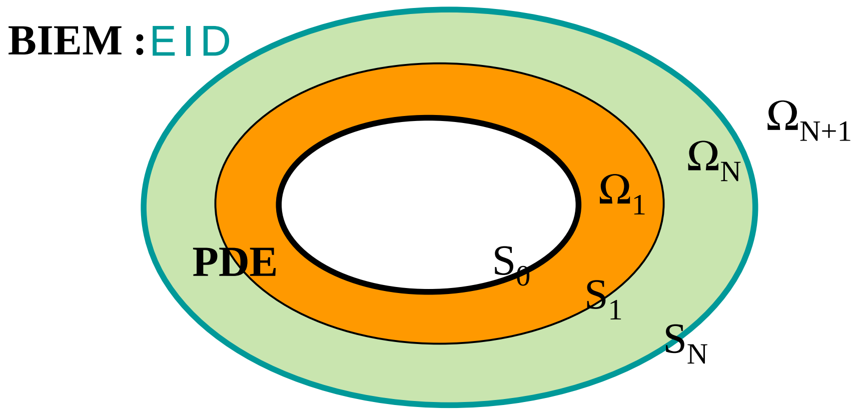

Domain Decomposition Method

B. Stupfel - B. Desprès (1999)

Feature : DDM partitioned into concentric sub-domains

$$\Omega = \bigcup_{i=1,N+1} \Omega_i, \quad S = \bigcup_{i=0,N} S_i, \quad S_{i-1} \cup S_i = \partial\Omega_i$$

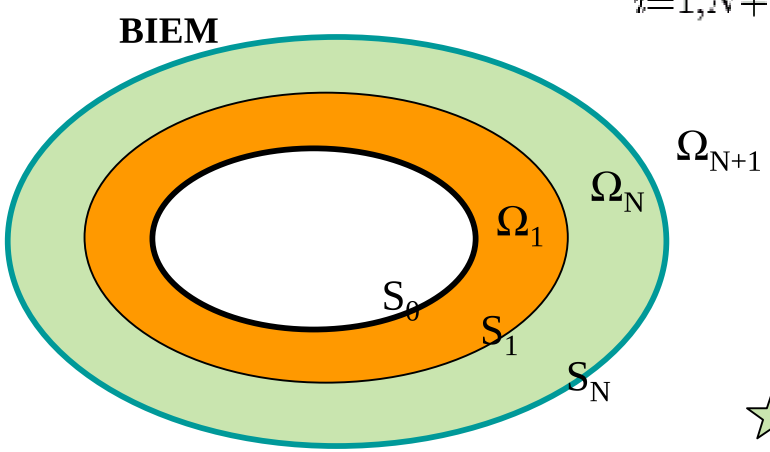

- Leontovitch impedance condition on S_0
- Transmission condition on S_i

$$T^{\pm}\mathbf{E} = \pm\mathbf{n} \wedge \left(\mu^{-1}\nabla \wedge \mathbf{E}\right) - ik_0\mathbf{E}_{tg}$$

- Coupling with the outer problem :
  TC <==> Impedance condition on S_N with Z=1
- We used an iterative resolution to solve the global problem

The BIEM - EID

B. Desprès (1996)

- Minimization of a positive quadratic functional associated to incoming and outgoing electromagnetic waves, with a set of linear constraints
- F. Collino & B. Desprès (2000) : link with the classical BIEM
- Very accurate method but leads to solve 2 coupled linear systems with 4 unknowns by edge
  - 2 linear combinations of the electric and magnetic currents
  - 2 Lagrangian unknowns (because of the constraints)
- Very efficient with an impedance condition Z=1
  - memory space
  - the 2 linear system become uncoupled

Z=1 ⇒ exactly the case we are interested in

The MLFMA applied to EID

K. Mer-Nkonga (2001)

- Potential of interaction

$$\phi(\mathbf{x}) = \sum_{|\mathbf{x}'|<R} \frac{e^{ik|\mathbf{x}-\mathbf{x}'|}}{4\pi|\mathbf{x}-\mathbf{x}'|}\rho(\mathbf{x}') = \sum_{|\mathbf{x}'|<R} G(\mathbf{x},\mathbf{x}')\rho(\mathbf{x}')$$

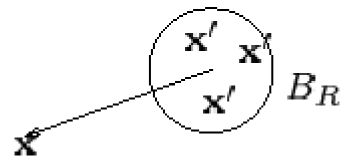

- Separation of interaction between far and near sources + approximation for the interaction with the far sources :

$$\phi = \phi^{near} + \phi^{far} \simeq \phi^{near} + \tilde{\phi}^{far}$$

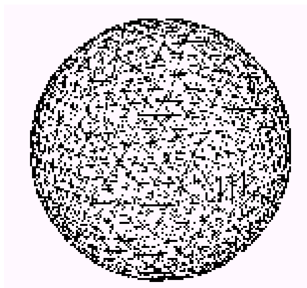

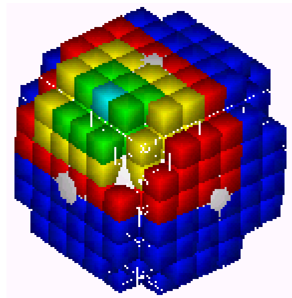

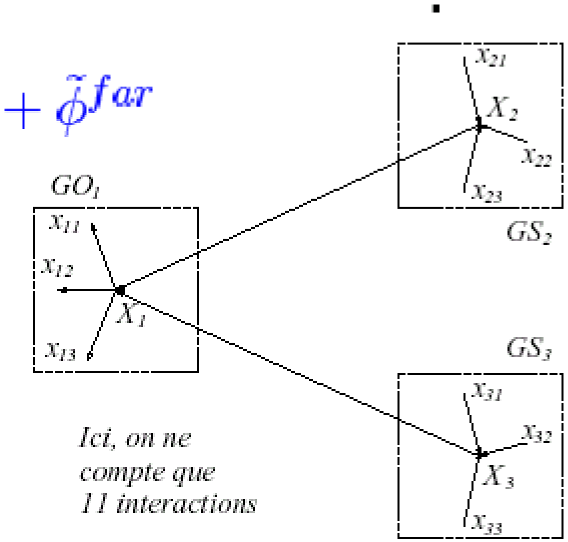

- In the EID context :
  - we have to deal with 2 operators

$\tilde{G} = \mathcal{R}e(G)$ et $\tilde{G} = \mathcal{I}m(G)$: a linear combination

Multi scale parallel algorithm

The global problem is solved by a Gauss-Seidel method ⇨ Sequential step

⇨ The parallelization is at the level of the resolution of each sub-domain

Nominal strategy

For each inner problem use very efficient parallel sparse direct methods : EMILIO (with solver PaStiX) on NPr processors

For the outer problem use iterative solver with a parallel FMM on NPu processors

NPr >> NPu

⇨ We use this unbalanced number of processors to implement a multi level parallel algorithm

⇨ Decomposition in pool of processors, each of them is in charge of one right hand side (useful for the 2 polarizations and multi incidences) .

CEA’s Parallel industrial codes in electromagnetism

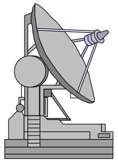

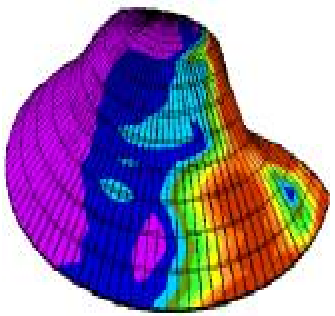

Why EMILIO ?

- 1997 : Needs of the CESTA of an industrial software tool that would be robust and versatile to solve very large sparse linear systems (structural mechanics, electromagnetism).

EMILIO is developed in collaboration with the ScAlApplix team (INRIA, Bordeaux).

- EMILIO gives efficient solution to the problem of the parallel resolution of sparse linear systems by direct methods, thanks to the high performance direct solver PaStiX (developed by ScAlApplix)

- Organized in two parts :
    - a sequential pre-processing phase with a global strategy,
    - a parallel phase.

EMILIO, PaStiX, Scotch ?

- General description
- Ordering and block symbolic factorization
- Parallel factorization algorithm
- Block partitioning and mapping
- Work in progress

EMILIO : General description

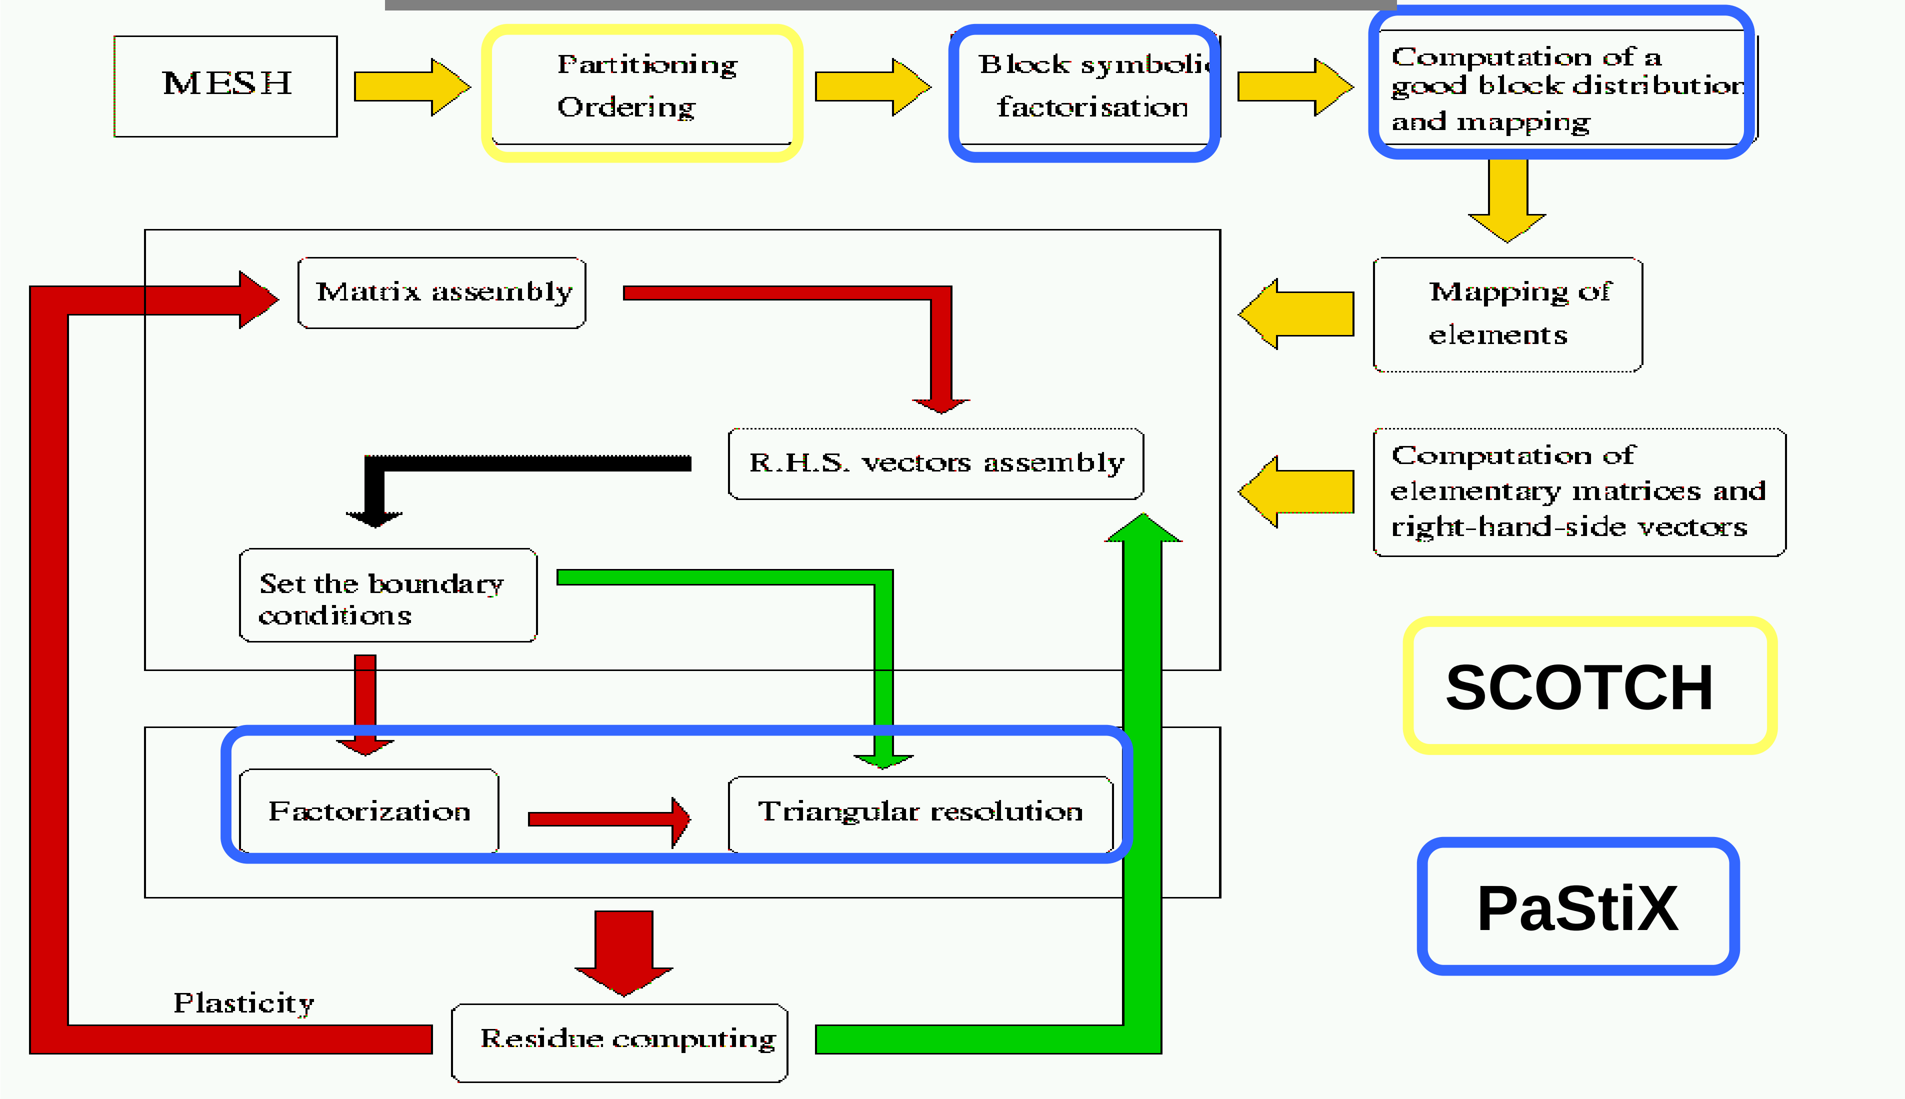

Ordering (Scotch)

- Collaboration ScAlApplix (INRIA) - P. Amestoy (ENSEEIHT / IRIT-Toulouse).
- Tight coupling

→ Size threshold to shift from Nested Dissection to Halo Approximate Minimum Degree

→ Partition P of the graph vertices is obtained by merging the partition of separators and the supernodes computed by block amalgamation over the columns of the subgraphs ordered by HAMD.

- Without *Halo*

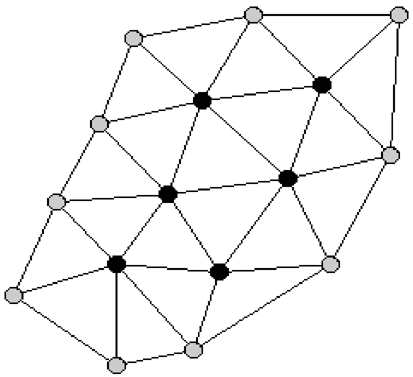

- With *Halo*

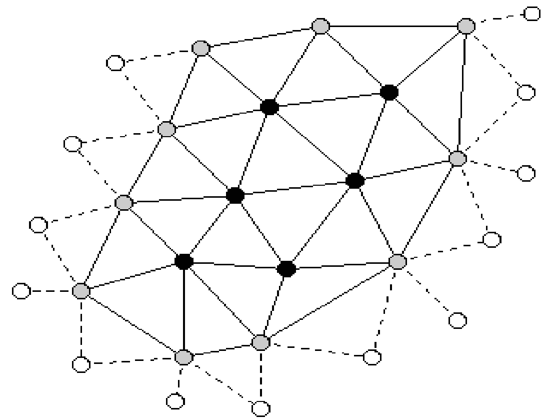

Symbolic Block factorization (PaStiX)

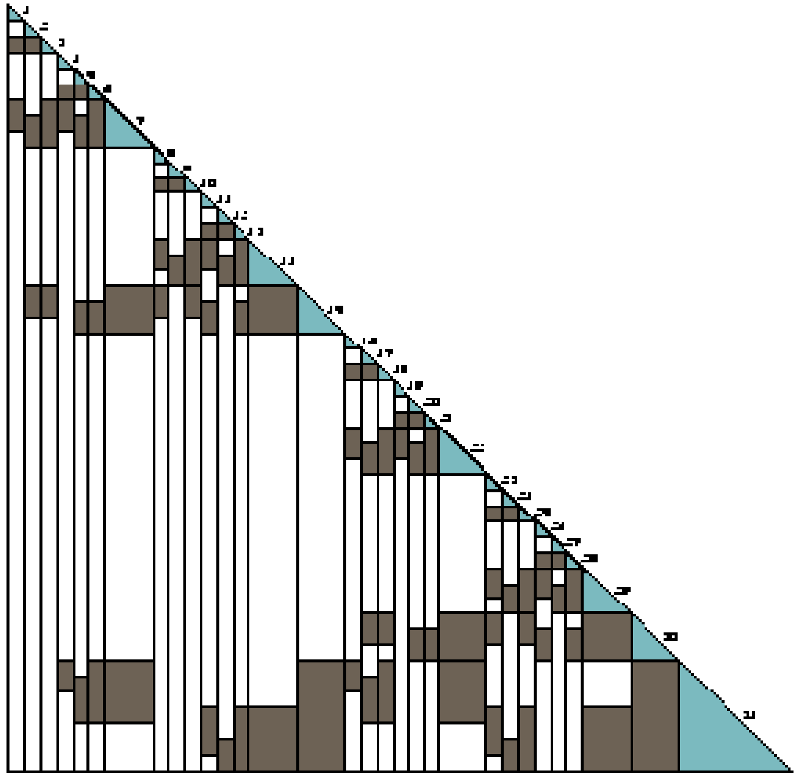

- Linear time and space complexities
- Data block structures
  ⇒ only a few pointers
  ⇒ use of BLAS3 primitives for numerical computations

Parallel factorization algorithm (PaStiX)

- Factorization without pivoting

→ Static regulation scheme

- The algorithm we deal with is a parallel supernodal version of sparse $L.D.L^t$, $L.L^t$, $L.U$ with 1D/2D block distributions

- Block or column block computations
  - → block algorithms are highly cache-friendly
  - → BLAS3 computations

- Processors communicate using aggregated update blocks only

- Homogeneous & Heterogeneous architectures with predictable performance (SMP)

CEA's Parallel industrial codes in electromagnetism

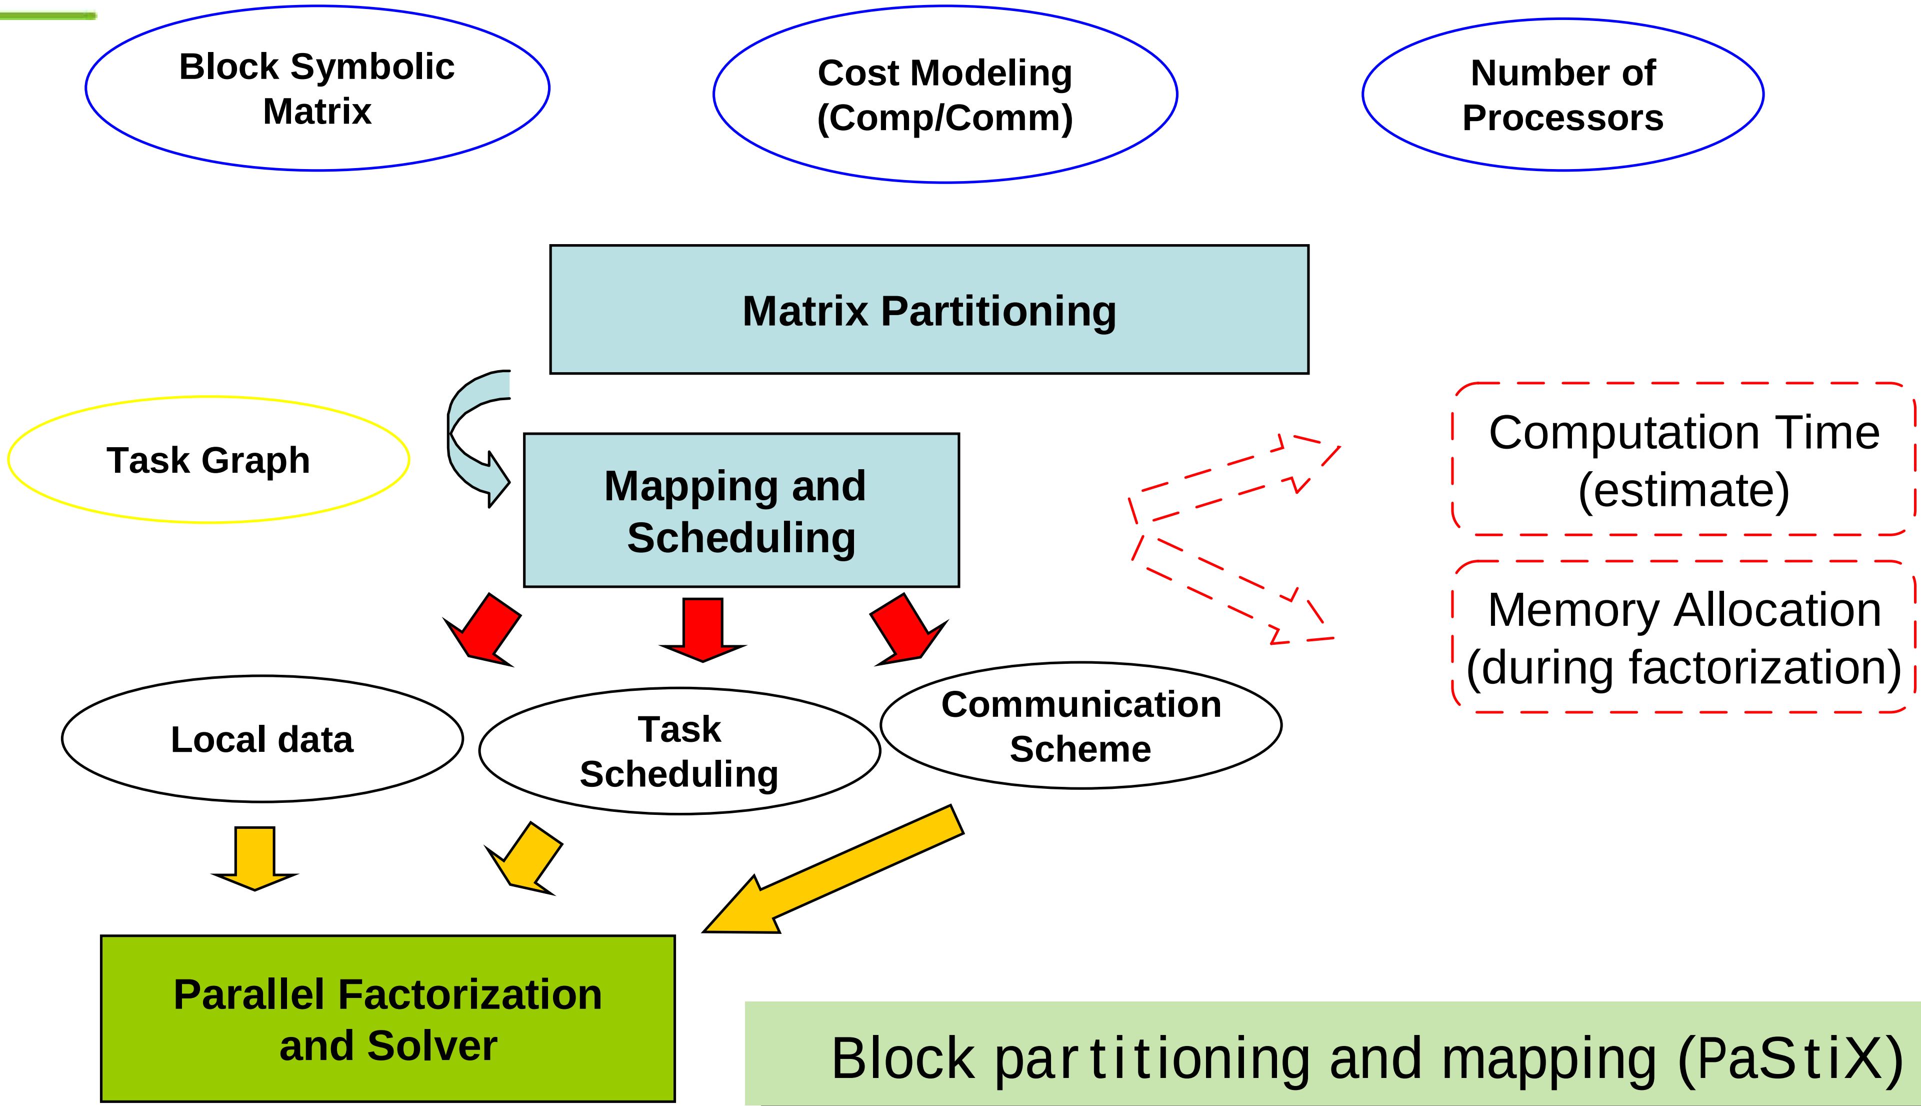

Block partitioning and mapping (PaStiX)

Current work on Scotch

- On-going work on PT-Scotch (Parallel version)

Current work for PaStiX : Hybrid iterative-direct block solver

- Starts from the acknowledgement that it is difficult to build a generic and robust pre-conditioner
  - Large scale 3D problems
  - High performance computing
- Derives direct solver techniques to compute incomplete factorization based preconditioner
- What's new ? : (dense) **block** formulation
- Incomplete block symbolic factorization:
  - Remove blocks with algebraic criteria
  - Sensitive to structure, not in a first time, to numerical value
- Provides incomplete LDLt, Cholesky, LU (with static pivoting for symmetric pattern)

CEA's Parallel industrial codes in electromagnetism

- Mesh of a Sphere Cone :
  - * **10 millions** unknowns in 6 sub-domains (in the volume)
  - * **220 000** edges for EID (on the outer boundary)

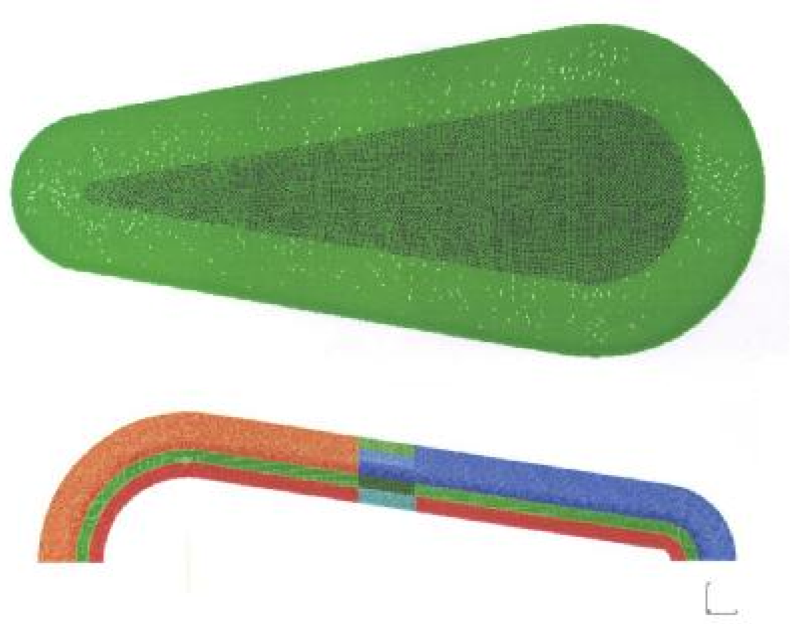

- x 10 millions volumic unknowns
- x 220 000 egdes (surfacic coupling)
- x 6 sub-domains

- x **7 global iterations**
- x **64 Processors**
- x **CPU Time : 6h15'**
- x **max. memory / proc : 2.6 Gbytes**

CEA's Parallel industrial codes in electromagnetism

- Field E Picture (PhD F. Vivodtzev)

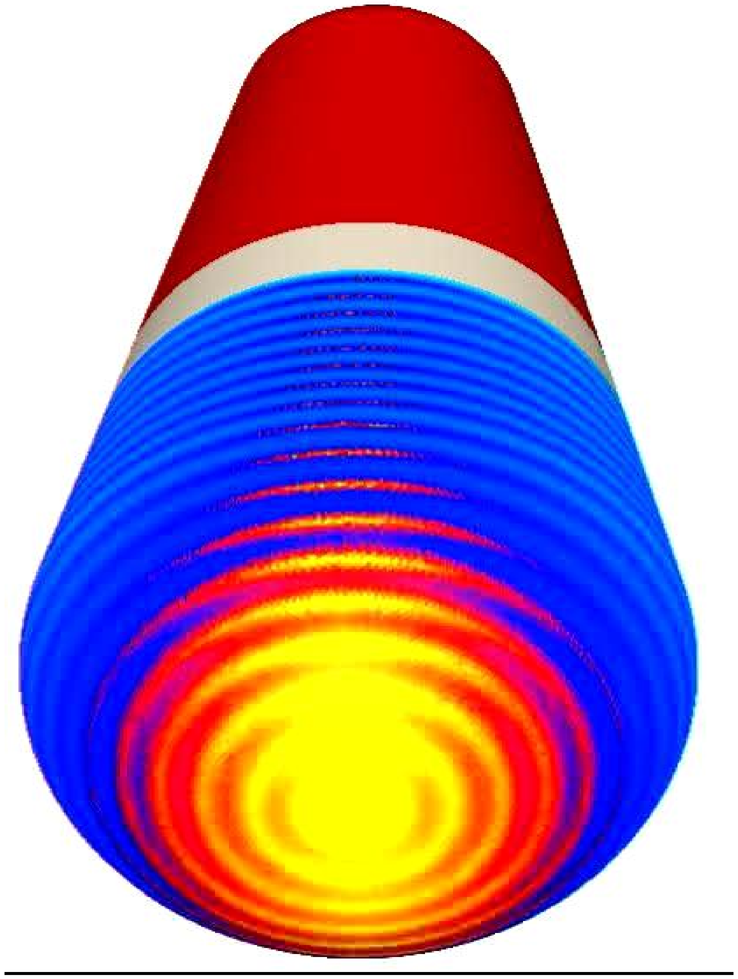

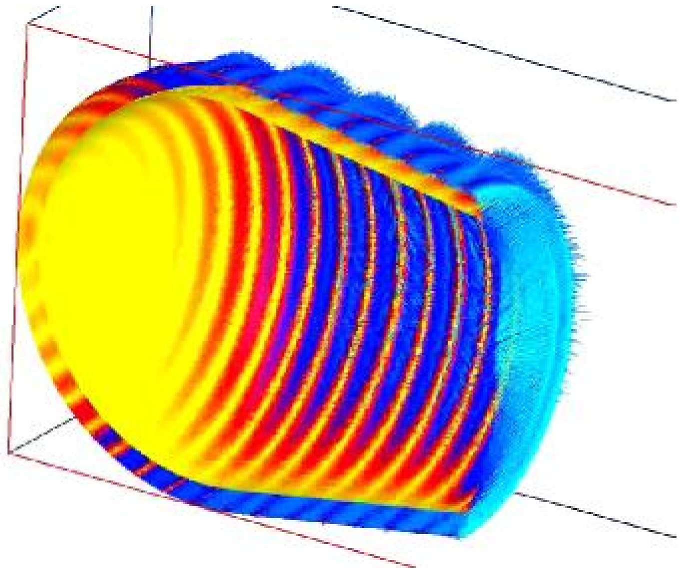

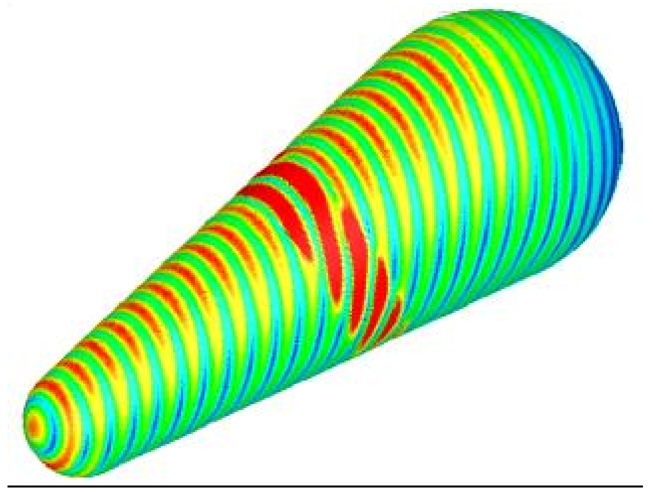

Conclusions for Parallel codes in electromagnetism

- We have developed a 3D code which is able to take into account all the constraints we have

- We did with success the validation of all the physics we have put in it by comparison with :
  - other codes we have (full BIEM, 2D axis symmetric)
  - measurements

The next step will be to reach (for the end of this year) :

- 30 millions unknowns : total amount of DOF (Degrees Of Freedom) in the volume
- 1 million unknowns : on the outer boundary

A recent result of the solver PaStiX :

10 millions unknowns, 64 procs Power 4 on 2 SMP nodes, 64 giga bytes of RAM per node,

NNZL : 6,7 thousand millions, 43 Tera operations needed to factorize,

CPU time : 400 seconds to factorize

Thank you

Merci

Grazie Mille

Danke Schön

Gracias

} for your attention