CESGA

CENTRO DE SUPERCOMPUTACIÓN DE GALICIA

High Scalability Multipole Method. Solving Half Billion of Unknowns

J.C. Mouriño[1], A. Gómez[1], J.M. Taboada[2], L. Landesa[2], J.M. Bértolo[3], F. Obelleiro[3], J.L. Rodríguez[3]

[1] Supercomputing Center of Galicia (CESGA), Spain

[2] University of Extremadura, Spain

[3] University of Vigo, Spain

International Supercomputing Conference

Hamburg, June 23th 2009

Outline

- Introduction
- FMM-FFT algorithm
- Parallel implementation
- HEMCUVE. Challenge history
- Experimental Results
- Conclusions

Introduction

- Numerical solution of the integro-differential electromagnetic equations -> important in industry
- High frequency in large objects -> scalability limits
  - 0.78 m^2 at 3 GHz -> 15000 unknowns
  - real car at 79 GHz -> 400 million unknowns
- Great effort reducing computational time of MoM
- Fast Multipole Method (FMM) -> $O(N^{3/2})$ good scalability
- Multilevel FMM (MLFMA) -> O(N log N) poor scalability
- Modern HPC systems have thousands of cores
- FMM-FFT reduces complexity preserving scalability

Introduction

- **FMM**
  - The geometry is partitioned into M groups using an oct-tree decomposition algorithm
  - Parallelization:
    - k-space samples are independent from one another
    - The workload can be partitioned in the *k-space*
- **MLFMA**
  - The same algorithm can be recursively applied in a hierarchical multilevel oct-tree decomposition
  - Computational cost reduced to O(N log N)
  - Parallelization:
    - k-space samples are not independent due to interpolation/ anterpolation of fields across levels
    - Usually workload distributed by groups

FMM-FFT algorithm

- Full domain is divided in groups
- The translation stage in the FMM can be seen as a 3D circular convolution
- Use the FFT to speedup the translation stage
- Reduce complexity form quadratic to logarithmic
- The FMM-FFT combines a low complexity like MLFMA with good parallel scalability of the original FMM
- It constitutes a real alternative to the MLFMA to take advantage of modern supercomputers

Parallel implementation

- Three stage parallelization strategy
  - Far contributions: distribution of fields (*k-space* samples) among processors
  - Near contributions: distribution of oct-tree groups
  - Iterative solver: distribution of unknowns
- Optimal load balance and data locality, while minimizing memory footprint and communication requirements
- A single communication step is required at the end of the MVP

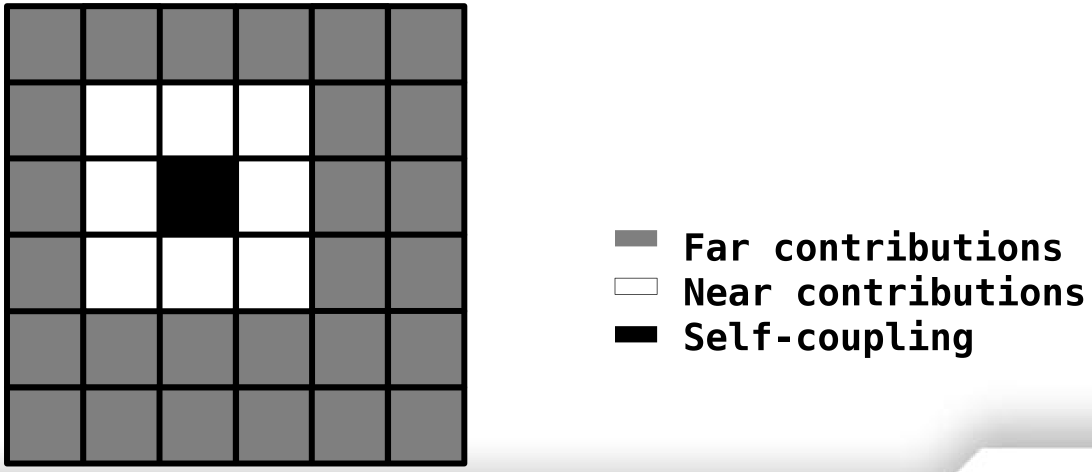

Nested FMM-FFT

- A refinement of the oct-tree decomposition is applied (one or more refinement steps)
- The far contributions are obtained at the coarse level of the oct-tree using a distributed FMM-FFT
- The near contributions are obtained at the fine level using a local shared memory FMM-FFT in each compute node
- **Dramatically reduction of memory requirements**

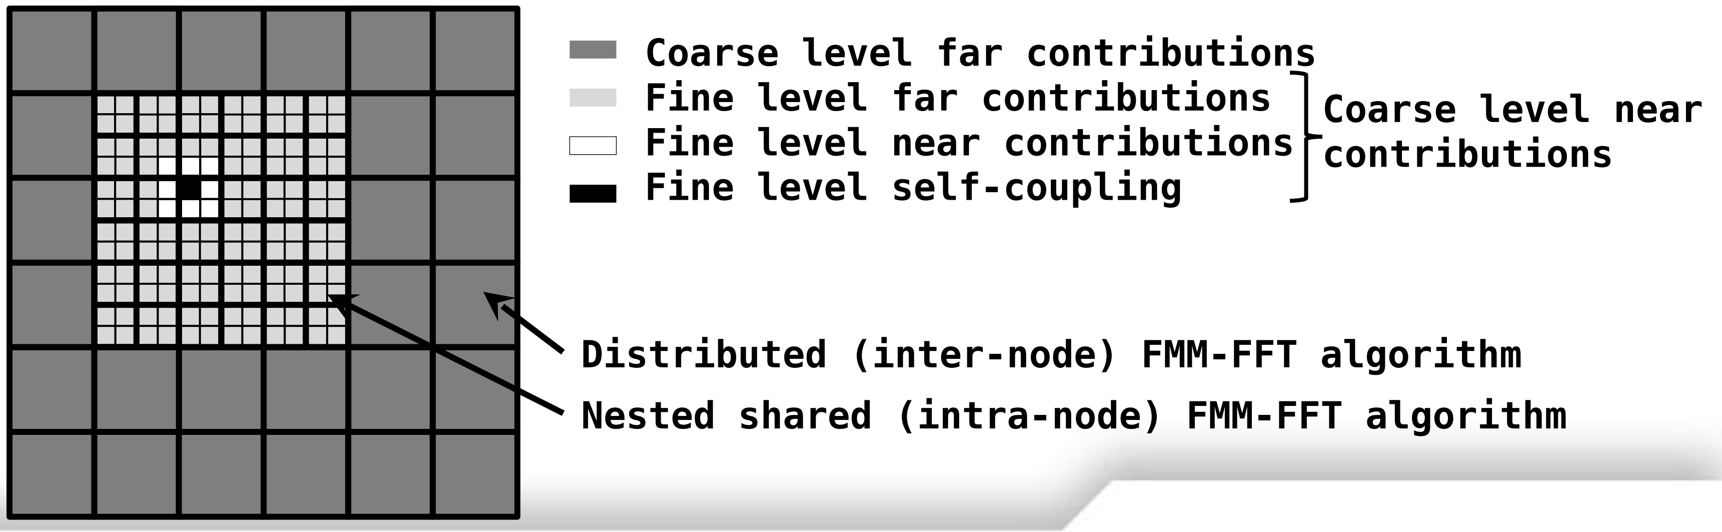

HEMCUVE. Challenge history

- Hybrid Electromagnetics: CESGA, University of Vigo and Extremadura
- C++ code for solving the previous algorithm
- In continuous development: increasing load balance, decreasing memory consumed, improving performance and scalability

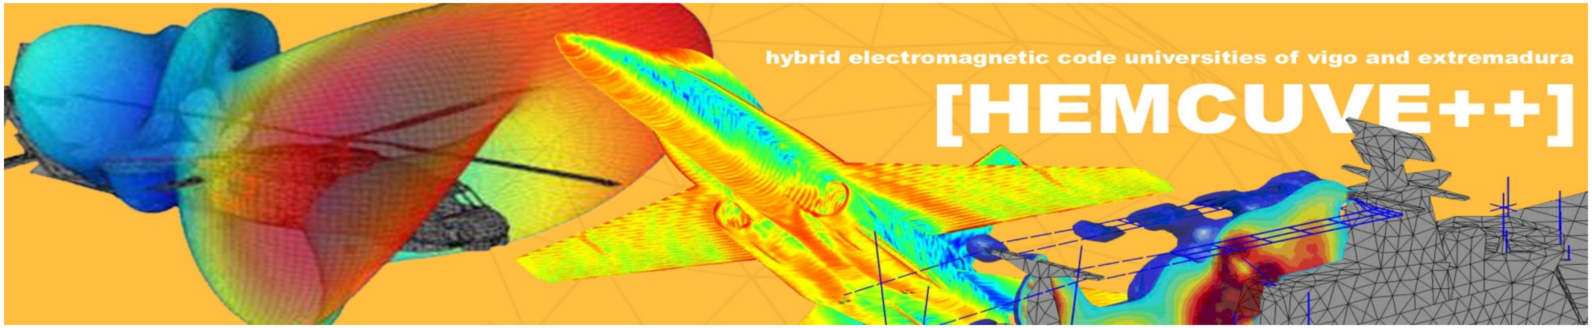

HEMCUVE. Challenge history

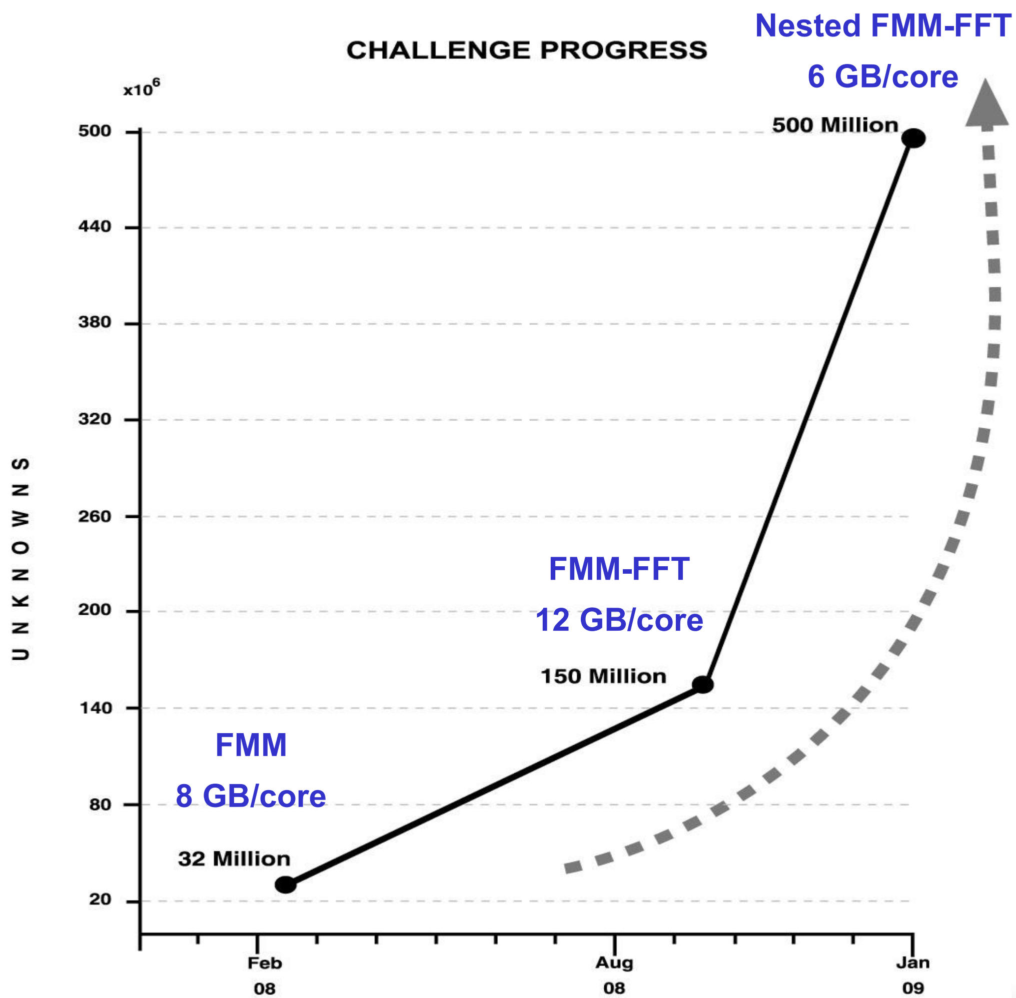

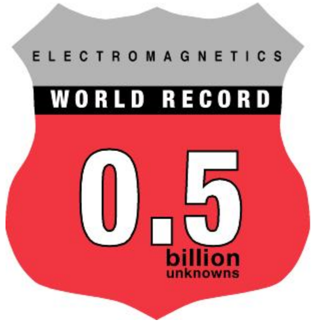

HEMCUVE. Challenge history

Scalability

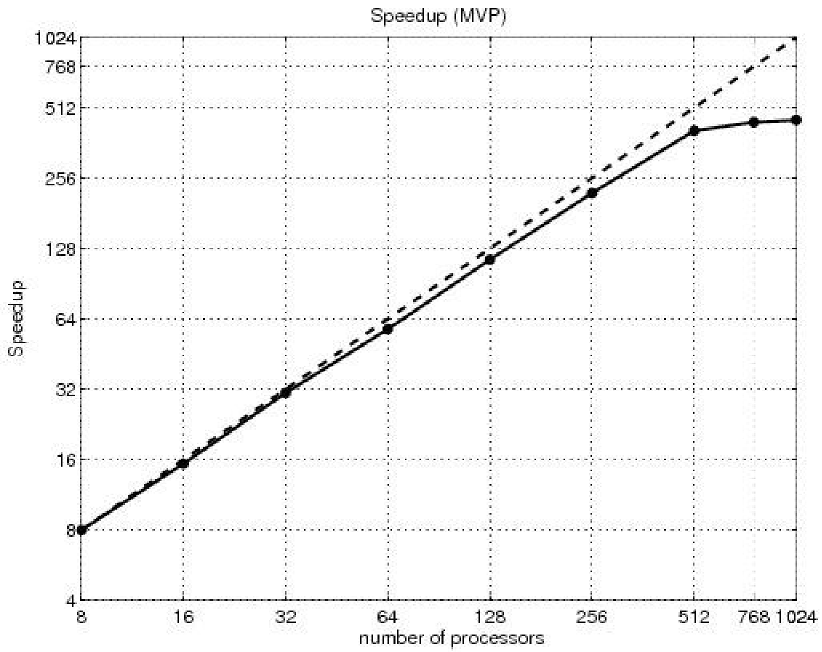

FMM

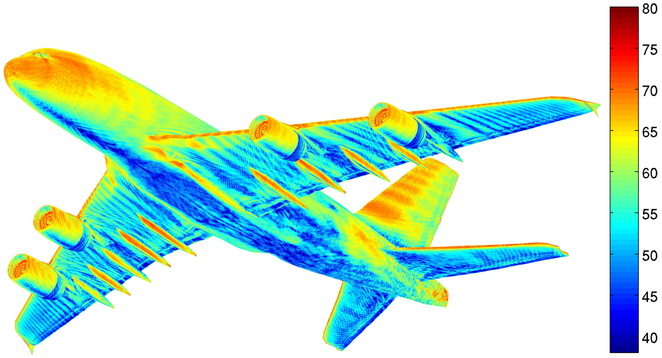

Induced currents in an Airbus A380 for an axial incidence of 1,2 GHz (30M)

Scalability

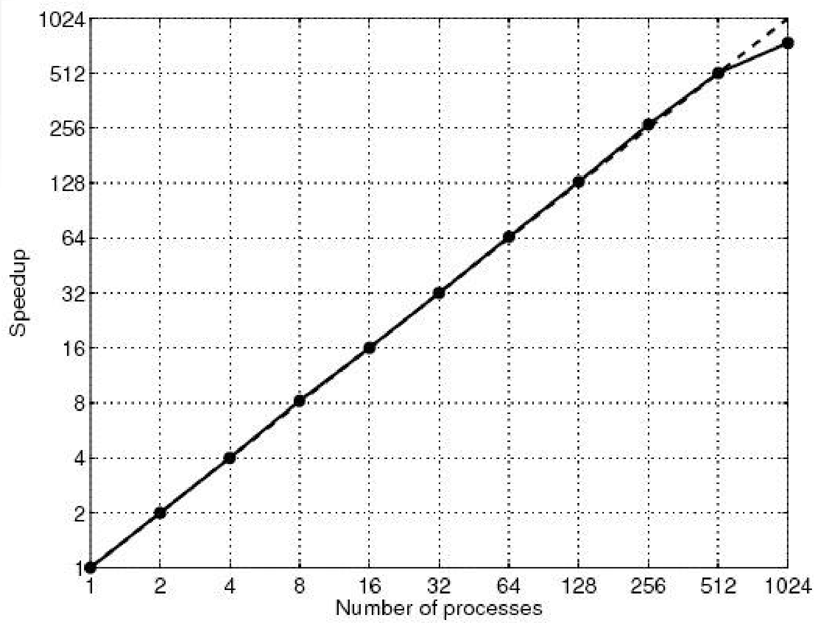

FMM-FFT

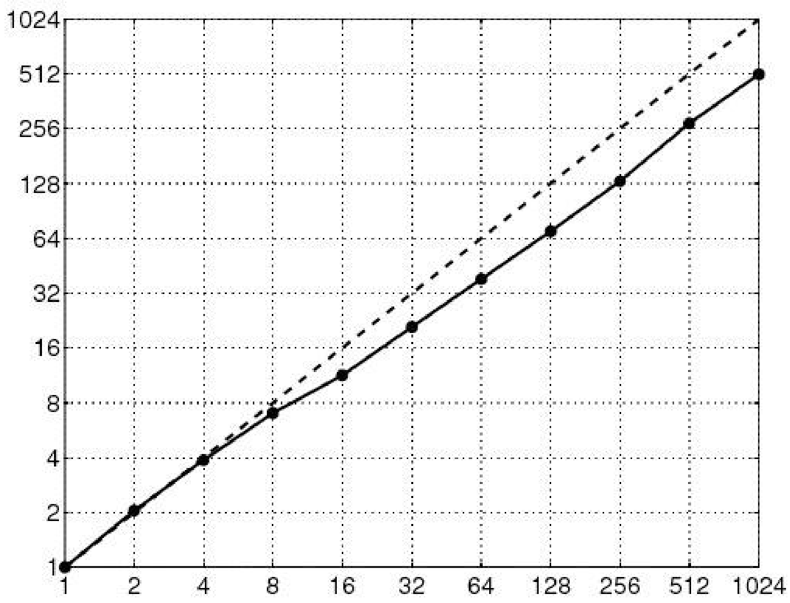

Nested FMM-FFT

Using a 10M unknowns problem

Experimental Results

FinisTerrae

More than: **16,000 GFLOPS** **2,580 CPUs** **19,640 GB Memory**

SuSE Linux Enterprise Server 10

Experimental Results

CESGA

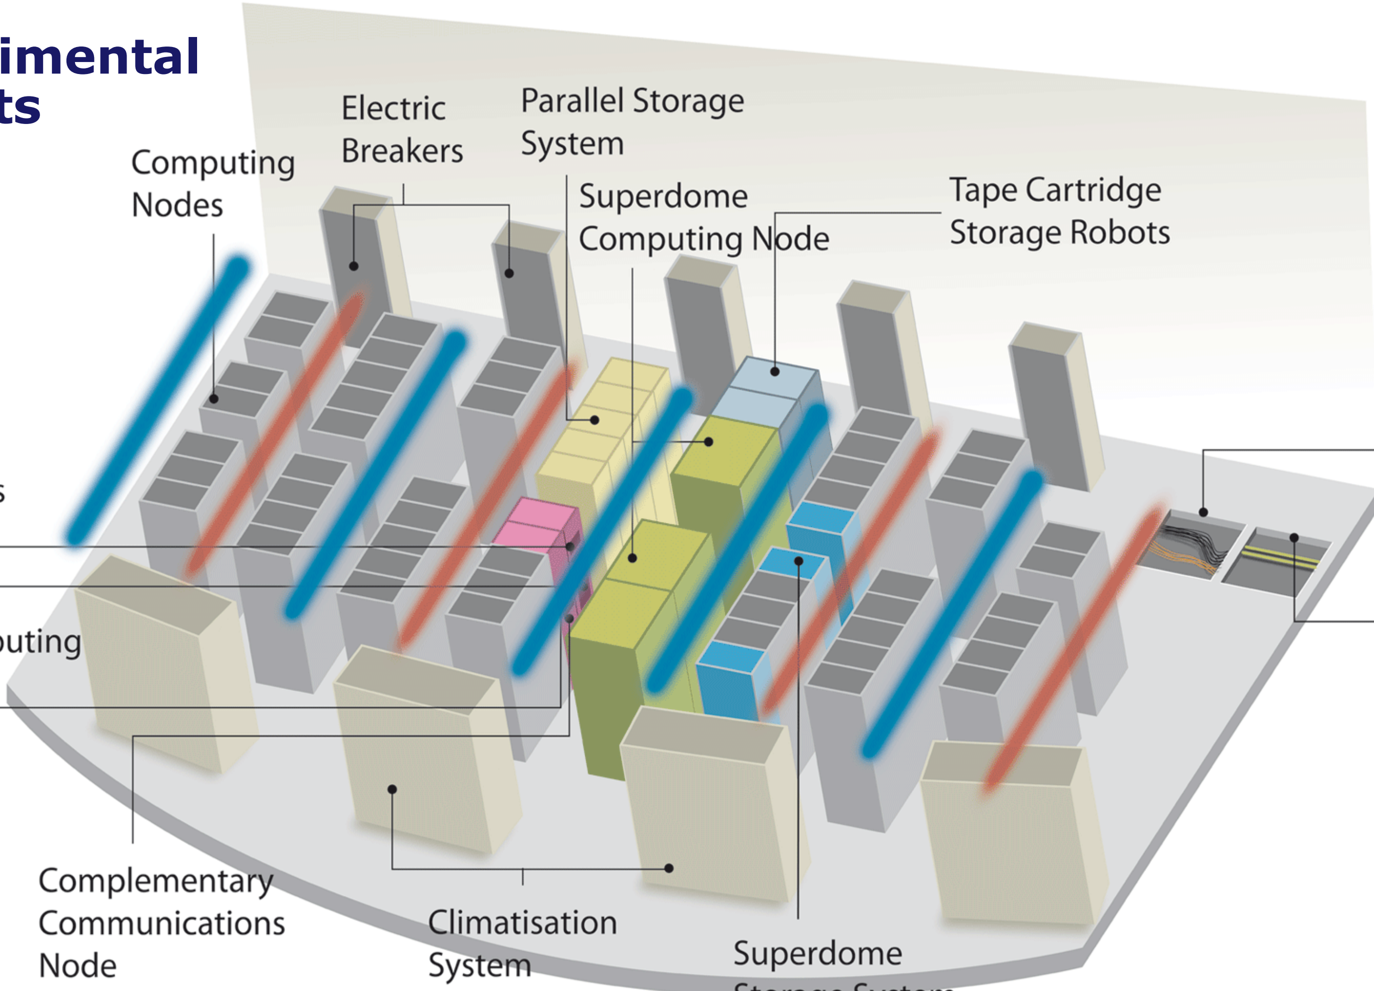

Technical Specs.

Surface Area: 140 m²

Weight: 35.000 Kg

Storage:

2.200.000 GB on tape

390.000 GB on disk

Memory:

19.670 GB

2.528 Processing Cores

142 nodes, each with 16 cores & 128 GB memory

1 node with 128 cores & 1.024 GB memory

1 node with 128 cores & 384 GB memory

Node Interconnect INFINIBAND
4x DDR at 20 Gbps

85 Km of interconnect cable

Open Software: Linux, Lustre, Globus...

FINISTERRAE

Experimental Results

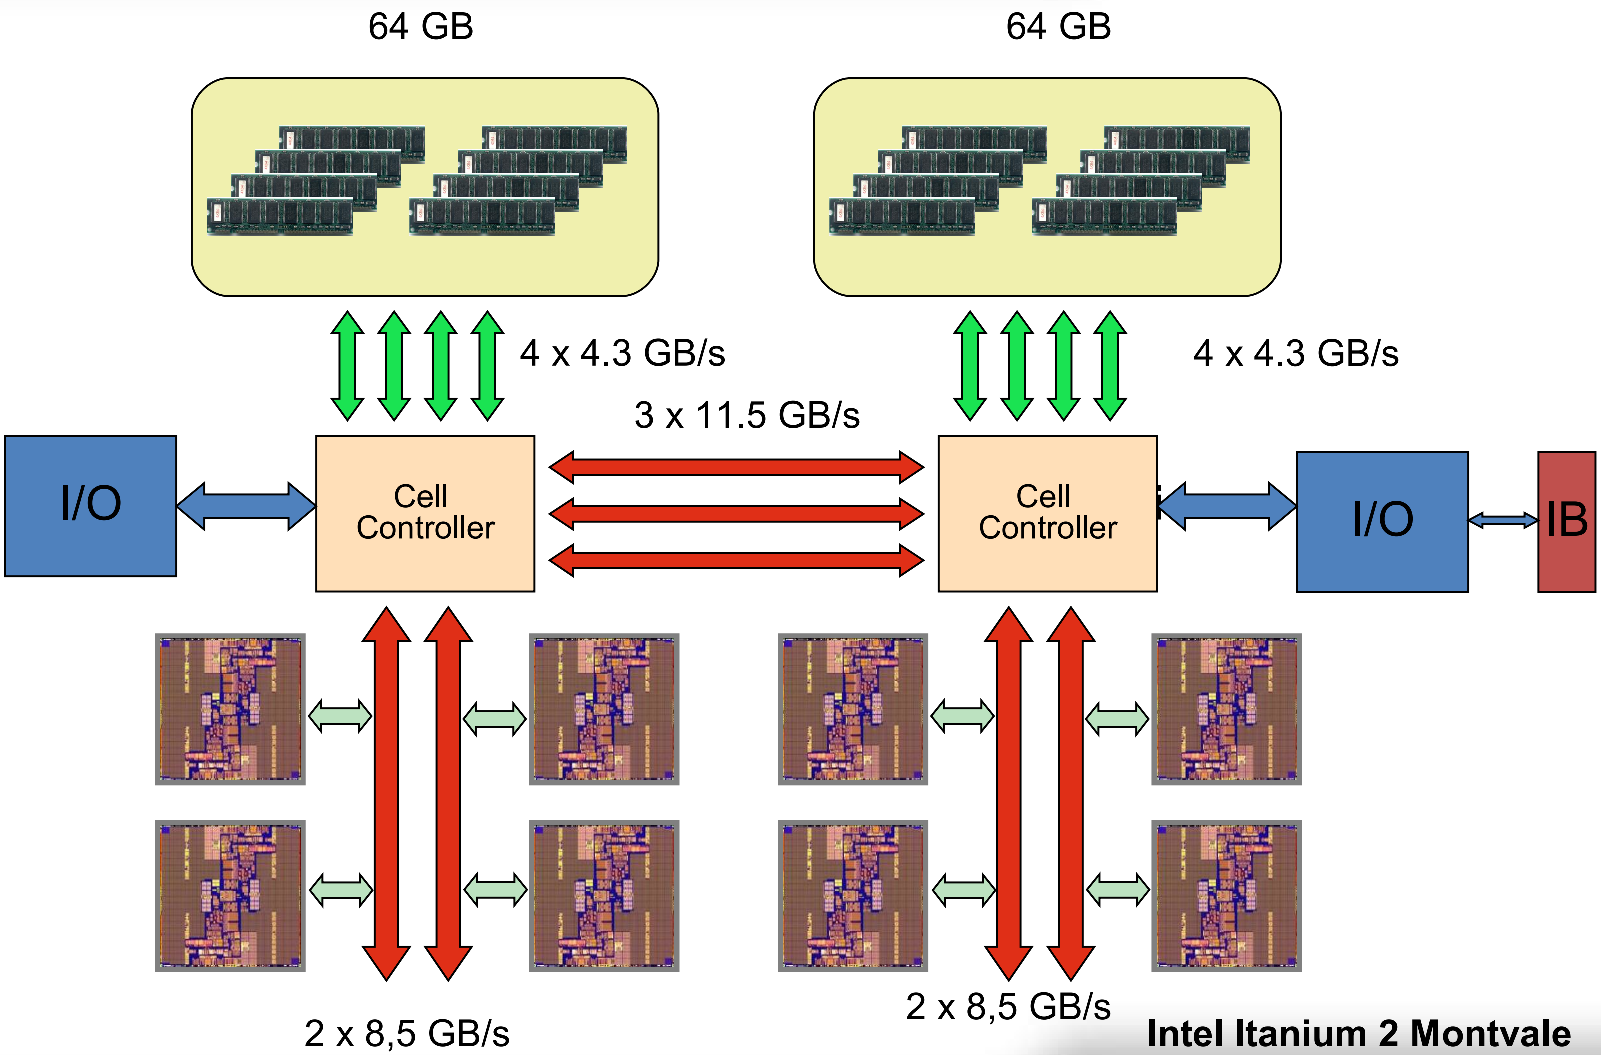

Intel Itanium 2 Montvale

Experimental Results

- Intel C++ Compiler version 11.0.069
- Intel MPI version 3.2.0.011
  - Limitation of 2 GB per message
- Intel Cluster MKL version 10.0.2.018

The Challenge problem:

- RCS of a PEC sphere with 728.36λ diameter
- 64 HP rx7640 nodes (1,024 cores) and 6TB of memory
- 10 GMRES iterations
- Residual error below $5\cdot10^{-5}$
- More than 30 execution hours

Experimental Results

Sphere diameter	728.36λ
Frequency	300 MHz
Number of Unknowns	500,159,232
Nodes / cores per node	64 / 16
Min / max memory per node	89.2GB / 99.9GB
Total memory	6 TB
Number of iterations	10
Setup / solution time	5h / 26 h

Experimental Results

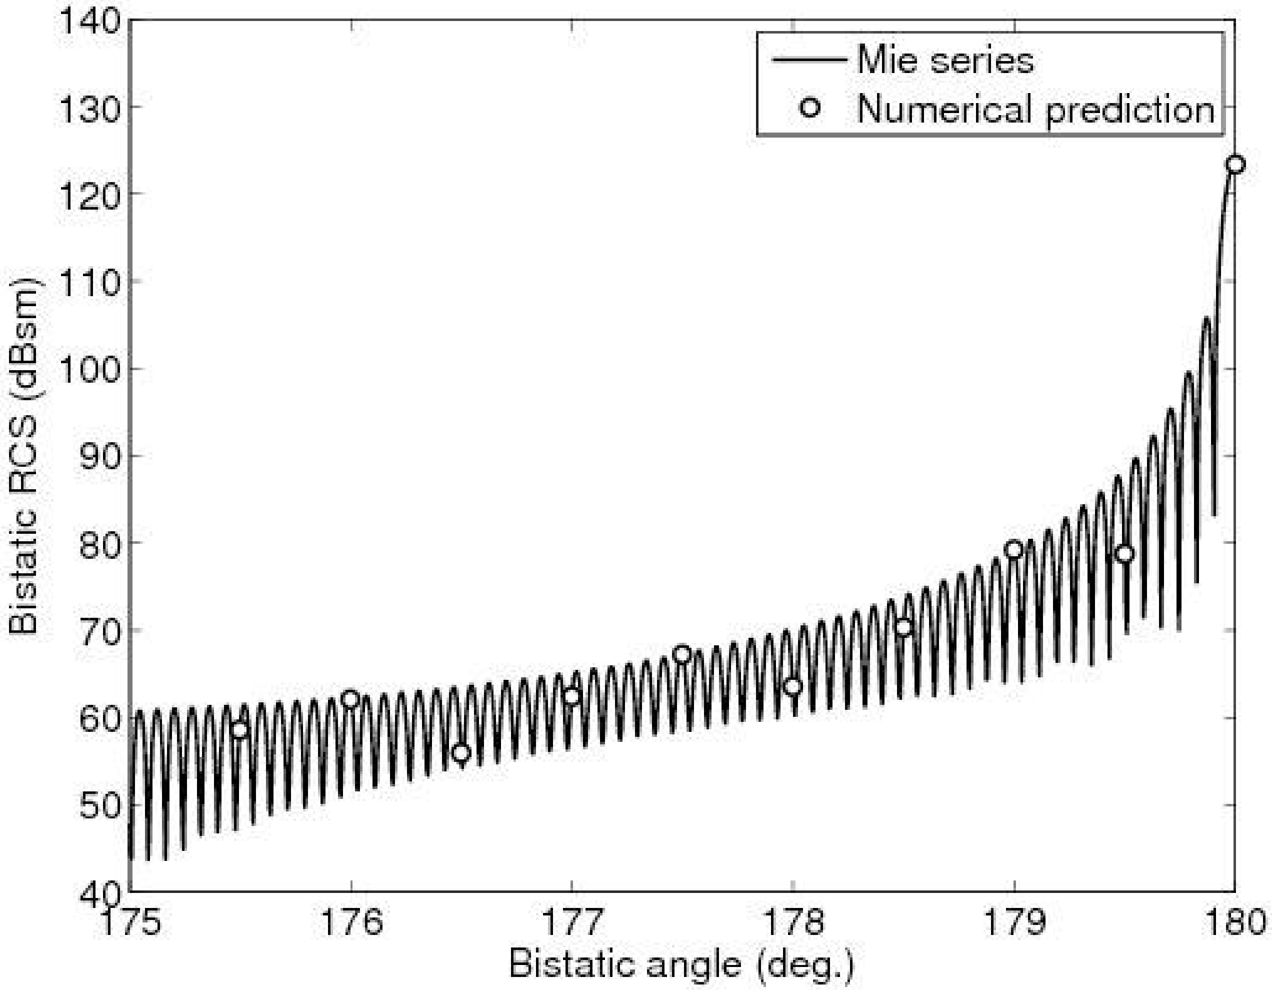

Experimental Results

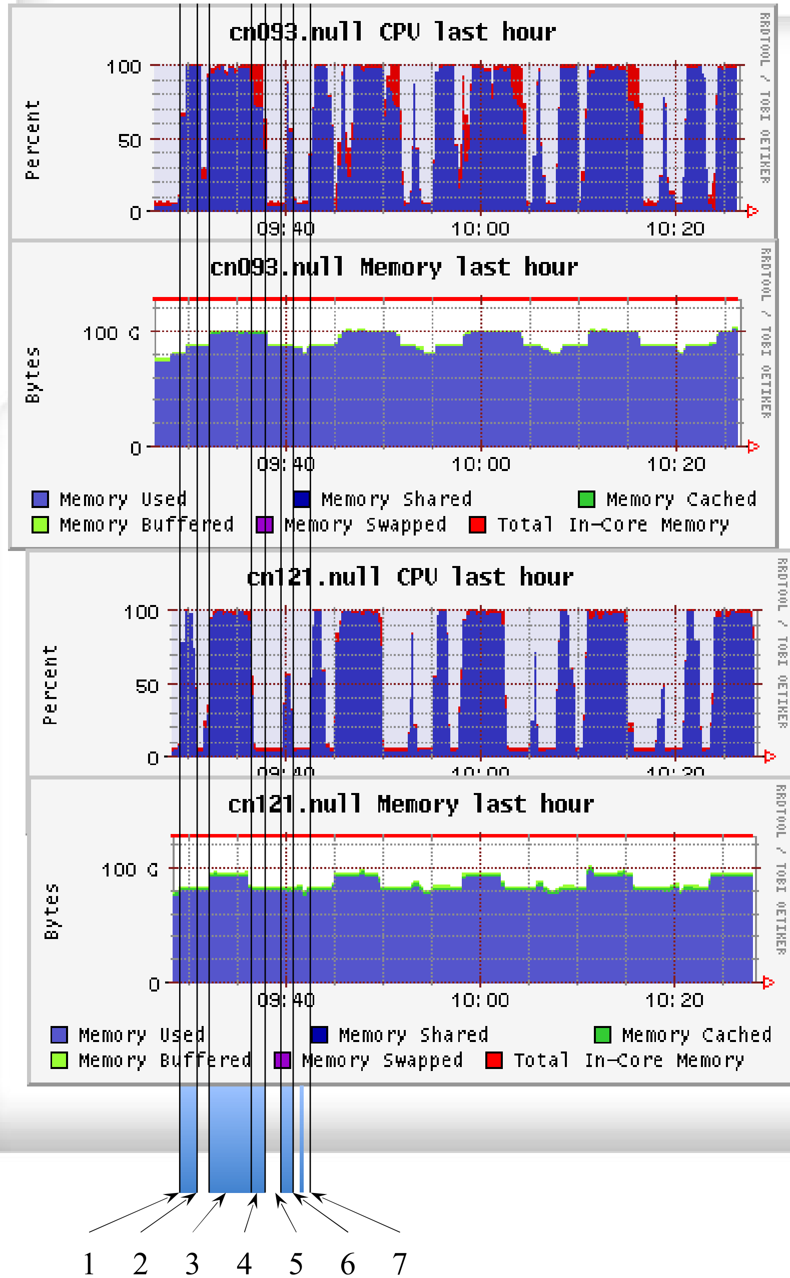

1: Aggregation, interpolation and near-field coupling
2: Communication Alltoallw
3: 3D-FFT translation
4: Restructuration of incoming fields
5: Communication Alltoallw
6: Disaggregation
7: Communication and GMRES (no OMP parallelization)

Experimental Results

Citröen C3 at 24.125 GHz (radar frequency) 40M unknowns

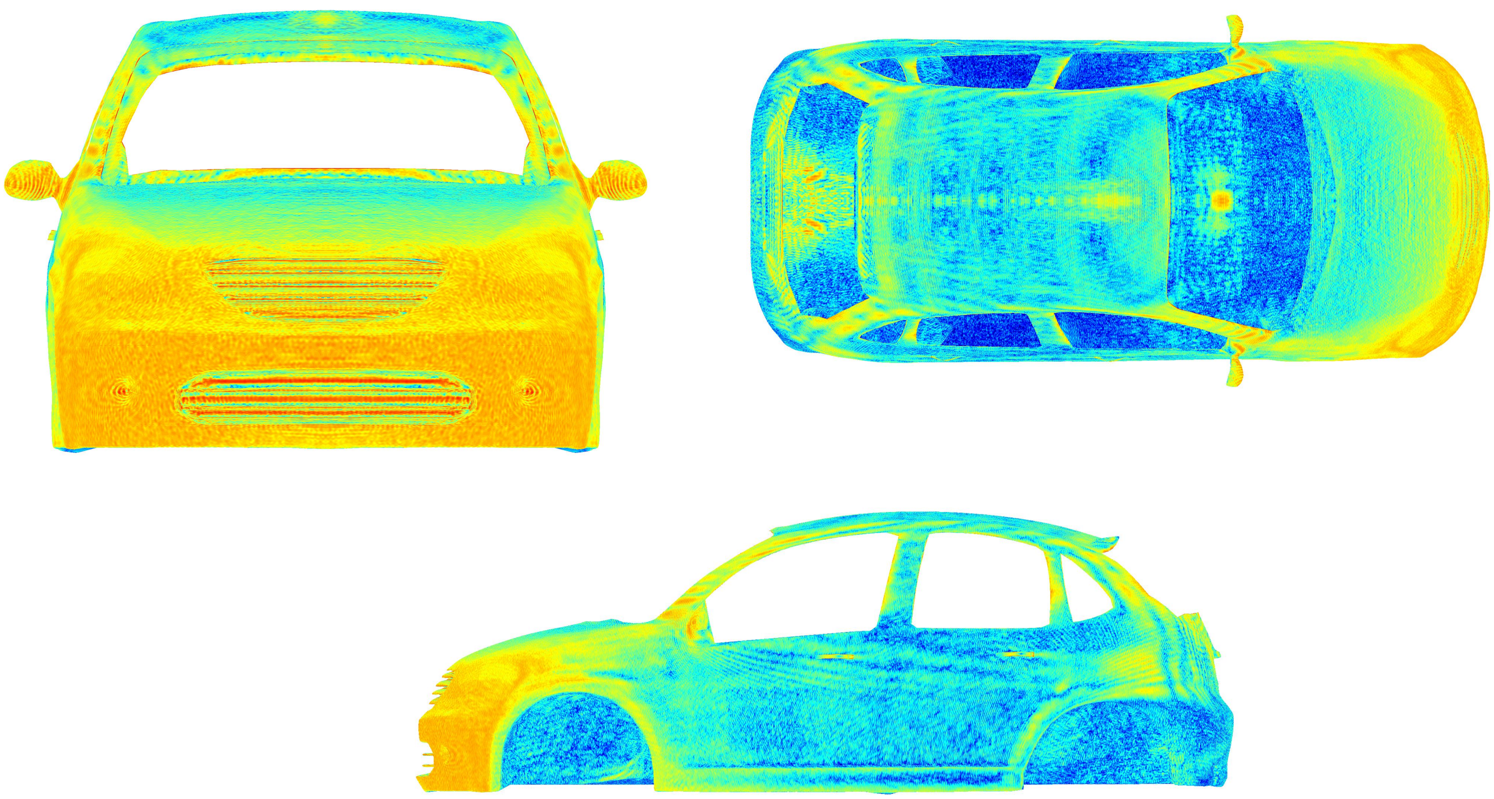

Experimental Results

Citröen C3 at 79 GHz (anti-collision frequency) < 400M

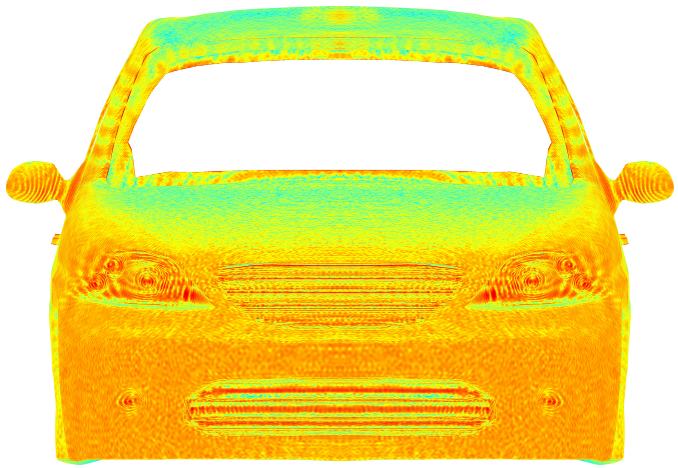

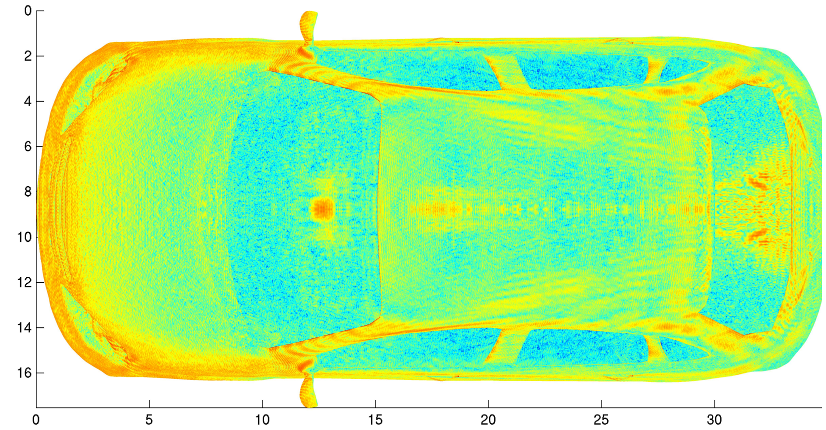

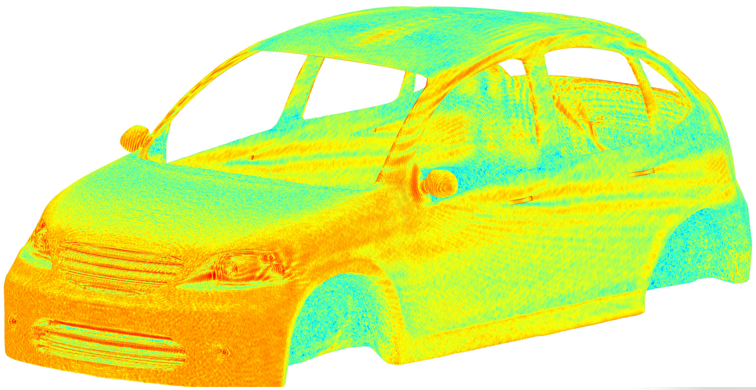

Conclusions

- The FMM-FFT is a real alternative for the analysis of very large general electromagnetic problems
- Very good scalability behavior up to 1,024 parallel processors
- Analysis of challenging problems up to 0.5 billion unknowns
- Suitable for calculating real structures at real frequencies

Future Work:

- new improved version: billion unknowns using 2,000 processors

Acknowledgements

- Spanish government
  - TEC2005-07355-C02-01
  - TEC2005-07355-C02-02
  - TEC2008-06714-C02-01
  - TEC2008-06714-C02-01
  - CONSOLIDER-INGENIO2010 CSD2008-00068
- Junta de Extremadura
  - 3PR05A002
- Xunta de Galicia
  - PGIDIT05TIC32201PR
  - INCITE08PXIB322250PR
- CESGA for their support in the use of Finis Terrae Supercomputer

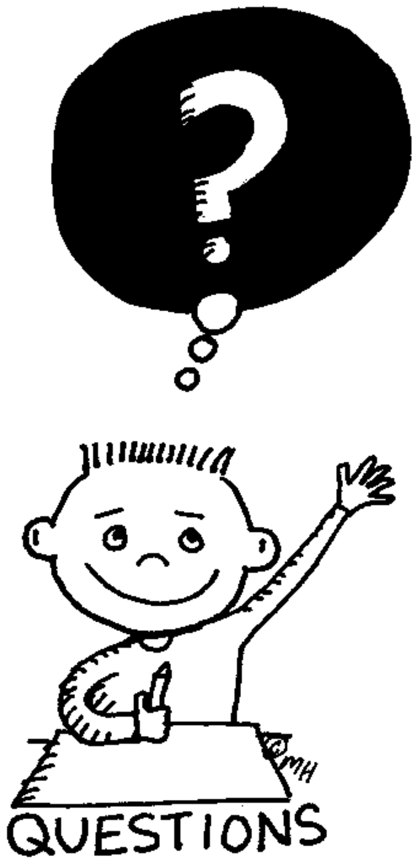

THANK YOU!

QUESTIONS?

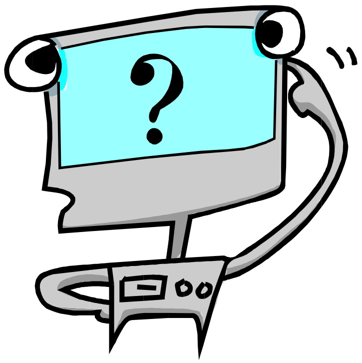

CONTACT:

Carlos Mouriño

jmourino@cesga.es

http://www.cesga.es