Itanium®-based Supercomputer Solves Largest Electromagnetic Simulations

> "We can now simulate massive, real electromagnetic problems which were impossible before."
> – Fernando Obelleiro, University of Vigo

Large electromagnetic simulations are instrumental for improving the design of ships, aircraft and cars as well as systems that emit radiation. For example, designers of onboard radiation systems must evaluate the electromagnetic compatibility (EMC) of complex conducting structures on ships, aircraft and satellites. Designers of radar systems must evaluate electromagnetic interference (EMI) and electromagnetic radiation hazards (EMR) that could be dangerous to personnel or electronic sensors. Defense systems manufacturers need to determine the radar cross-section (RCS) of conducting structures.

Electromagnetic simulations in high-frequency domains have been nearly impossible due to limits in software and computer scalability. For example, analyzing the electromagnetic RCS of new anti-collision vehicle systems in the assigned frequency (79 GHz) would require solving a problem of 200 to 400 million unknowns, using exact electromagnetic solvers such as Method of Moments (MoM).

For that reason, scientists are striving to develop fast, efficient algorithms to reduce the computational cost of electromagnetic simulations and thereby solve highly complex real-world problems. Two important advances in this area have been the fast multipole method (FMM) algorithm and the appearance of supercomputers with relatively inexpensive clusters amassing very large memory and computational resources. These concurrent advances in algorithms and hardware have motivated the development of faster, more efficient and accurate numerical solutions in computational electromagnetics.

Researchers at the Supercomputing Center of Galicia (CESGA), the University of Vigo and the University of Extremadura in Spain have applied a Fast Fourier Transform extension of the conventional FMM, known as FMM-FFT, which dramatically lowered the complexity of FFM while preserving its parallel scalability. This scalability, combined with performance improvements, can make very effective use of large, parallel, high-performance supercomputers.

The team validated this approach by solving a challenging problem with more than 500 million unknowns using 1,024 Intel® Itanium® processor cores and with 6 TB of memory, which constituted the largest problem solved in computational electromagnetics at that time. The solution demonstrated that EMC can be solved computationally, even for big problems and at high frequencies, within sectors such as aerospace, automotive, naval, biomedicine and electronic warfare.

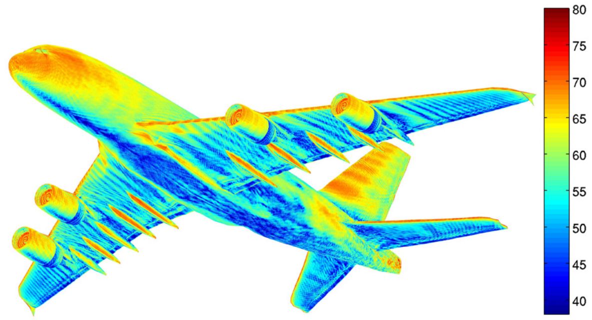

Scaling to hundreds of millions of unknowns makes electromagnetic simulation viable in industries such as aerospace, automotive, naval, biomedicine and electronic warfare.

Challenges

The biggest challenge to high scalability in computational electromagnetics was finding the optimal balance of four criteria: workload balancing, data locality, memory footprint and communication requirements. Using the FMM-FFT, the team could distribute the workload equally among processors while keeping good data locality, by applying a k-space parallelization strategy. This efficient distribution relied on the fact that, in FMM-FFT, each sample in k-space is completely independent of the others.

Taking these considerations into account, the team devised a three-stage parallelization for the FMM-FFT that led to optimal load balancing and data locality, while minimizing memory footprint and communication requirements:

- Distribution of far-fields among processors to account for the far-field interactions
- Distribution of oct-tree groups for near-field interactions
- Distribution of unknowns for the iterative solver

The parallelization of the far-field contributions of the matrix-vector-product (MVP) required partitioning the data structures in the spectral domain. This meant that, during the setup, each node has only to evaluate and store its assigned directions of the k-space. The subsequent MVP operations perform independently, without communication among processors, making the translation stage (one of the most critical stages for the scalability) clean and efficient.

The near-field interactions were parallelized through standard domain decomposition. The selection of a proper load-balancing algorithm was crucial for scalability: the distribution of equal number of groups in each node does not work for general problems, because all groups have different numbers of basis functions and nearby groups.

All these factors had to be taken into account to obtain a well-balanced distribution in which the processors evaluate near-field interactions without any external communication.

At the end of the MVP calculation, the computed far and near interactions are summed in each node. A single communication step is then required, in an almost negligible amount of time, to sum up all the partial contributions computed by each node.

The development process

HEMCUVE (Hybrid ElectroMagnetic Code University of Vigo and Extremadura) is a C++ development of the above-mentioned algorithm. It has been in continuous evolution to improve the load balancing, decrease the amount of memory consumed, and achieve a better performance and scalability. The objective: to produce a software solution that could solve a real electromagnetism problem of one-half-billion unknowns.

The initial parallel version used a single-level FMM with high inherent scalability because it minimized communications and facilitated load balancing. The first parallelization of the algorithm was implemented using a distributed memory scheme.

Benefits of Itanium

The well-known ability of the Intel® Itanium® processor to manage huge amounts of memory becomes crucial for this kind of leading-edge application. The rigorous integral-equation-based solvers in computational electromagnetics consume massive amounts of RAM, in this case, a total of 6 TB of RAM with close to 100 GB of RAM per computing node. The large cache memory in the Intel Itanium processors was also a key factor in the successful completion of the calculation.

Challenges	Results
Tradeoffs among work-load balancing, data locality, memory footprint and communication requirements.	Nested FMM-FFT strategy provides efficient distribution of samples among massively parallel processing nodes.
Integral-equation-based solvers in computational electromagnetics consume massive amounts of RAM.	64-bit Intel Itanium processor manages huge amounts of shared memory – 6 TB in this example.

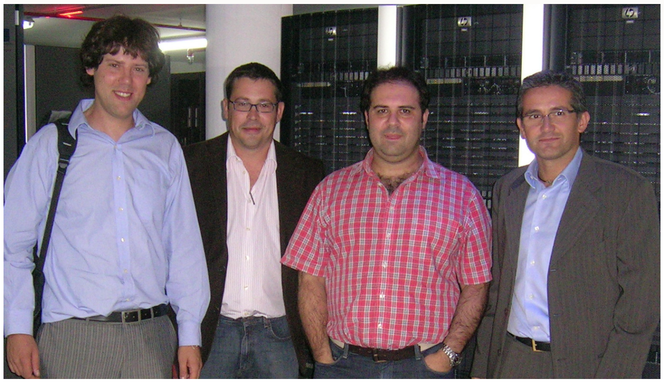

Some members of the HEMCUVE team (Dr. Taboada, Dr. Rodríguez, Dr. Mouriño and Prof. Obelleiro).

> "What we have solved today with Itanium would have been science fiction just a few years ago. Now it is a reality!"
>
> – Fernando Obelleiro, University of Vigo

After these first experiences, technicians from CESGA joined the project, granting access to its supercomputers and providing their expertise in HPC applications. The team ported the initial version of the application to the Itanium platform at the end of 2007, and was able to take advantage of its code support for large scalability. To decrease the communication requirements among multiple cores, the team introduced a hybrid parallel algorithm, using Message Passing Interface (MPI) protocol between distributed nodes and the Open Multiprocessing (OpenMP) API within each shared-memory node.

When the first big problem (about 32 million unknowns) was executed using this code in February 2008, the team detected some unbalancing and bottlenecks that limited scalability. A new version solved those problems and introduced the new FMM-FFT method. With this new version, the team executed a simulation with 150 million unknowns in August 2008, which was the biggest simulation in the world at that time.

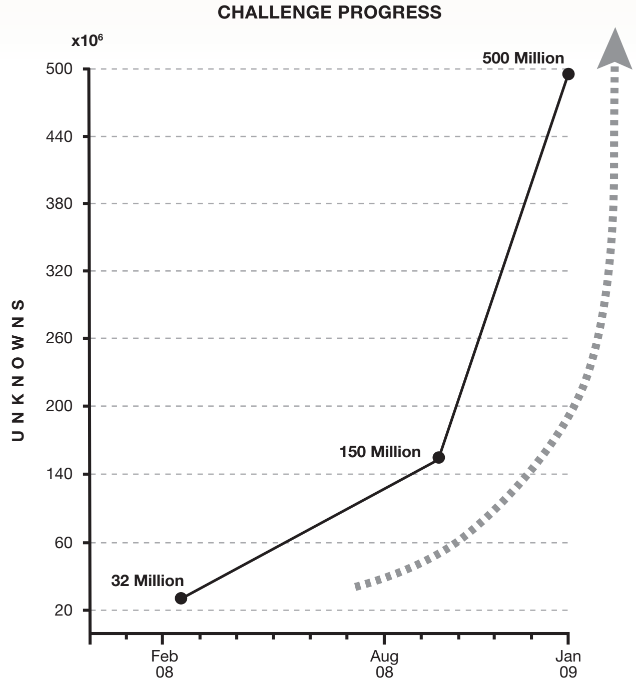

Each release of the HUMCUVE algorithm improved the load balancing, decreased the amount of memory consumed, and achieved a better performance and scalability.

As the amount of memory became a critical issue – e.g., in problems involving several tens or hundreds of million unknowns – the team further reduced memory usage by applying a nested FMM-FFT scheme based on a multilevel oct-tree partition of the geometry. The far-field interactions are obtained at the coarse level of this partition. However, the near-field interactions are evaluated using a local FMM-FFT solver at the finest level of the oct-tree, thus providing a significant reduction of memory usage. Nevertheless, the reduction in the case of the aggregation matrices is obtained at the expense of additional communications inside the MVP.

Results

The sheer size of the target electromagnetic simulation required the capacity of the Finis Terrae computer at CESGA. A supercomputer with one of the best memory-to-processor ratios in the world, the Finis Terrae is based on a sophisticated HP Cluster Platform 6000* driven by HP Integrity* servers with Intel Itanium processors. CESGA works closely with research centers throughout Europe to help physicists, chemists, engineers, and earth and life scientists apply supercomputer speed and capacity to highly complex research and analysis.

> "We have used efficiently more than 6 TB of RAM memory for a single problem. It is amazing!"
>
> – José Manuel Taboada, University of Extremadura, Spain

Using Finis Terrae, the team analyzed the challenging problem consisting of the bistatic RCS of a PEC sphere with 728.36 λ (wavelength) diameter and 500,159,232 unknowns, the world record at the time. Running on a total of 1,024 processors with 6 TB of memory, the setup time was about five hours and the iterative solution took less than 26 hours.

This problem has a known analytic solution (being used commonly by the community to test software) and the numerical solution agreed well with the analytical model. As a result, the team proved that electromagnetic computation can solve even the biggest problems at the highest frequencies.

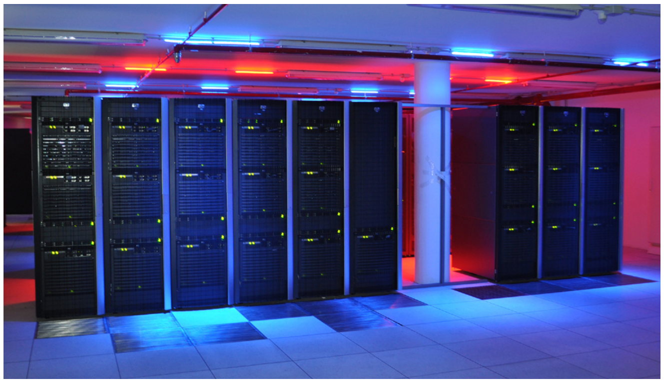

Using the Finis Terrae supercomputer, the authors analyzed a problem with 500,159,232 million unknowns running on a total of 1,024 processors with 6 TB of memory.

System Configuration

Hardware: HPC supercomputer Finis Terrae, installed in the Supercomputing Centre of Galicia (CESGA) – this solution employed 64 HP rx7640* nodes with a total of 1,024 processors and 6 TB of memory.

Operating system: SUSE SLES 10.*

Application: custom algorithm developed using Intel® C++ Compiler version 11.0.069, and Intel® MPI version 3.2.0.011 for inter-node communications. For matrix/vector linear algebra operations, the Intel® Cluster MKL version 10.0.2.018 was used.

Printed in USA 1009/JJ/OCG/XX/PDF Please Recycle