

MLFMA-FFT PARALLEL ALGORITHM FOR THE SOLUTION OF LARGE-SCALE PROBLEMS IN ELECTROMAGNETICS (INVITED PAPER)

J. M. Taboada

Department Tecnologías de los Computadores y de las
Comunicaciones, Escuela Politécnica
Universidad de Extremadura
Cáceres 10071, Spain

M. G. Araújo and J. M. Bértolo

Department Teoría do Sinal e Comunicacóns, E.T.S.E.
Telecomunicación
Universidade de Vigo
Vigo (Pontevedra) 36310, Spain

L. Landesa

Department Tecnologías de los Computadores y de las
Comunicaciones, Escuela Politécnica
Universidad de Extremadura
Cáceres 10071, Spain

F. Obelleiro and J. L. Rodriguez

Department Teoría do Sinal e Comunicacóns, E.T.S.E.
Telecomunicación
Universidade de Vigo
Vigo (Pontevedra) 36310, Spain

Abstract—An efficient hybrid MPI/OpenMP parallel implementation of an innovative approach that combines the Fast Fourier Transform (FFT) and Multilevel Fast Multipole Algorithm (MLFMA) has been successfully used to solve an electromagnetic problem involving 620 millions of unknowns. The MLFMA-FFT method can deal with extremely large problems due to its high scalability and its reduced computational complexity. The former is provided by the use of the

Corresponding author: M. G. Araujo (martaga@com.uvigo.es).

FFT in distributed calculations and the latter by the application of the MLFMA in shared computation.

1. INTRODUCTION

The constant development of the computer technology sets objectives more and more ambitious for the electromagnetic solvers. The methods not only must be fast, efficient and with low complexity but also must benefit from the present computational capabilities. The new supercomputers give the chance of solving large problems that were unattainable in the past or only tractable using inaccurate techniques. Consequently, as the capacity for solving large problems of hundreds of millions of unknowns grows, the electromagnetic numerical solvers and industrial necessities get closer. For all these reasons, in last years the well-known Method of Moments [1] made way for acceleration techniques as the Fast Multipole Method (FMM) [2] and its multilevel version, the MLFMA [3, 4]. In fact, the attention of many recent studies is concentrated on the improvement of the MLFMA parallelization over shared, distributed and mixed memory computers [5–15].

Another numerical technique that has gained interest is the FMM-Fast Fourier Transform (FMM-FFT) [16, 17]. The FMM-FFT preserves the natural parallel scaling propensity of the single-level FMM in the spectral (k -space) domain [18, 19]. Moreover, it requires less time than the FMM for computing the matrix-vector product (MVP) due to the translation stage acceleration given by the FFT. This method has demonstrated to be a good alternative to benefit from massively parallel distributed computers [19]. Recently, a hybrid Message Passing Interface (MPI)/OpenMP parallel implementation of a nested scheme of the FMM-FFT was proposed by the authors in [20–22]. In contrast with the conventional FMM-FFT, it shows a slightly worse parallel performance but in exchange for lower memory consumption.

Looking for the combination of the best characteristics of these previous works, a hybrid MPI/OpenMP parallel implementation approach is presented in this paper. The method uses the FFT for the distributed computation and MLFMA for the shared one, both implemented in an efficient hybrid MPI/OpenMP scheme. It will be shown that the proposed method, MLFMA-FFT, achieves an optimal performance in mixed memory supercomputers. A challenging numerical result with 620 millions of unknowns with useful applications in the automotive industry has been solved, which constitutes the

largest problem solved up to now.

The paper is organized as follows: Section 2 reviews the main aspects of the FMM-FFT, the Nested FMM-FFT and their parallel implementation. Section 3 outlines the MLFMA-FFT algorithm, some details about its parallelization and an assessment of the computational complexity. Section 4 presents the numerical examples, including a comparison between the method presented here and the parallel implementations developed by the authors in previous works, as well as the challenging problem mentioned above. Finally, the summary and conclusions are given in Section 5.

2. BACKGROUND: PARALLEL FMM-FFT AND NESTED FMM-FFT

According to present research interest, our recent efforts are headed for solving large scale problems taking advantage of massively parallel distributed computers that are available nowadays. The work was first focused on the FMM-FFT mainly due to its natural parallel scaling propensity in the spectral domain. Additionally, the use of the FFT to accelerate the translation stage provides a great reduction of the MVP CPU time with respect to the conventional FMM. An electromagnetic problem of more than 150 million unknowns has been solved in [19] using a hybrid parallel implementation of the method. We have used the MPI library for communications between distributed nodes and OpenMP standard for threads inside each shared-memory node. Regarding the parallelization strategy, it has been shown that the workload for the far interactions of the MVP can be optimally distributed by splitting the independent k -space samples among processors. For the near-field interactions the distribution is based on the spatial oct-tree partition established in the initial stage of the method application. A distribution by unknowns is used for the parallel iterative solver. Further details of this FMM-FFT implementation can be found in [19].

Later on, we have proposed the Nested FMM-FFT [22] to further reduce the memory consumption. From one or more refinement steps of the hierarchical oct-tree decomposition, the nested scheme is applied to the near field interactions in the MVP, which are obtained at the finest oct-tree level by using one or more local shared-memory FMM-FFT algorithms inside each computing node. The far-field interactions are still obtained at the coarsest level of the geometry partition using a distributed FMM-FFT algorithm. A problem with more than a half billion unknowns has been solved using the Nested FMM-FFT algorithm. The main aspects of the parallelization strategy considered for the algorithm in [22] are recalled here: (i) A mixed

distribution for the far-field interactions is applied, by groups at the finest level and by k -space samples at the coarsest one. (ii) A well-balanced load distribution by groups is considered for the near-field interactions at the finest level. (iii) A distributed scheme by equal number of unknowns per processor is used for the iterative solver (we have used GMRES [23]). The required communications between nodes do not involve latency periods or explicit synchronization because of the efficient management provided by the MPI library.

These computational challenges have been awarded with two prizes, namely the *PRACE Award 2009* and the *Itanium Innovation Alliance Award 2009* in computationally intensive applications.

3. THE PROPOSED METHOD: PARALLEL MLFMA-FFT

Because of its low computational cost, the MLFMA is a reference algorithm for the rigorous solution of large problems in computational electromagnetics. However, it is known that it suffers from bad parallel scaling on distributed memory computers. The difficulties faced in successfully parallelizing it in computer clusters have been widely discussed by many research groups [8, 10, 14, 15]. Considering the favorable MLFMA features in shared-memory computers, and taking into account the high scalability behavior of the FMM-FFT in distributed computers while maintaining a low numerical complexity, a proper combination of both techniques seems to be a good alternative to take a step further after the works of [19] and [22]. The MLFMA-FFT presented here combines both methods to get an optimal scheme from the point of view of modern mixed-memory supercomputers. The computational requirements of the resulting algorithm are significantly lower than the requirements of the previous approaches, as it will be shown in Section 4.

The proposed MLFMA-FFT method initially requires a multilevel refinement of the hierarchical oct-tree decomposition. The far-field interactions are still obtained at the coarsest level of the geometry partition using a global distributed FMM-FFT algorithm to accomplish them. The same parallelization strategy of the previous works based on the distribution by k -space samples has been considered. Regarding the near-field contributions, they are obtained at the finer levels of the oct-tree by using one or more local MLFMA algorithms inside each shared-memory computing node. Since in the MLFMA-FFT method the MLFMA is not applied among distributed nodes, it does not suffer from poor parallel scaling. This new approach provides a significant reduction of the computational complexity in the solution of extremely large problems with hundreds of millions of

unknowns. A practical example of the MLFMA-FFT applicability for this kind of large problems is detailed next in Section 4.

3.1. Parallelization Issues

For the parallel implementation of the MLFMA-FFT, we have selected a hybrid parallel programming combining MPI with the OpenMP standard, which fits perfectly with mixed-memory computer systems. Regarding the treatment of the far-field interactions at the coarsest level of the oct-tree, an optimal parallel performance is achieved applying a distribution of the work by k -space samples and then obtaining the far coupling by means of a distributed FMM-FFT algorithm. This is a common issue with the parallel implementations adopted in [19] for the FMM-FFT and in [22] for the nested version. The main differences between the methods are given by the procedure used for dealing with the near couplings. In the MLFMA-FFT method, this task is accomplished by the MLFMA. The partition of the work is based on a distribution by oct-tree groups at the coarsest level whose partial contributions to the MVP are then computed by several MLFMA algorithms operating strictly inside each shared-memory computing node.

A simplified 2D problem is depicted in Figure 1 in order to illustrate the work-load balancing distribution in the MLFMA-FFT algorithm. Figure 1(a) represents the complete problem at the coarsest

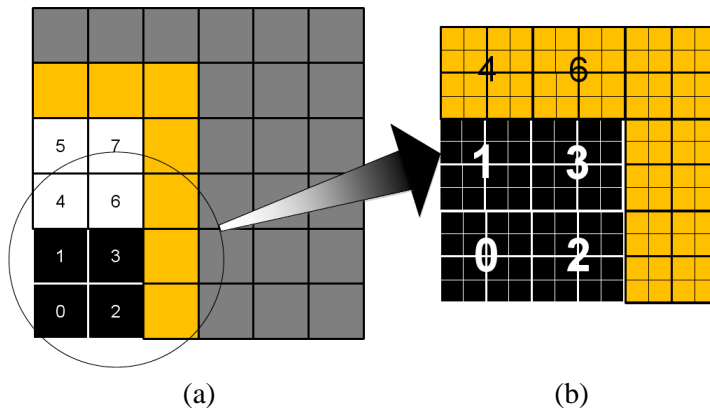


Figure 1. Work-load balancing among processors in the MLFMA-FFT: (a) Complete problem at the coarsest level. (b) Intranode problem (in node 1) comprising the finest level to one before the coarsest level.

level. The black boxes with labels from 0 to 3 identify the coarsest level observation groups that are assigned exclusively to one node (e.g., the first node), in order to obtain their near contribution to the MVP. This contribution consists of the radiation incoming from their interaction source groups (groups 0 to 3 plus their respective adjacent groups in the example of Figure 1, as it is illustrated in Figure 1(b)). The white boxes (4 to 7) are the observation groups assigned to the next node and so on, taking into account that a work-load distribution algorithm is used to look for a well-balanced scheme. At this point, the near partial contribution of a given node to the MVP is computed strictly inside the node by means of the MLFMA algorithm. The MLFMA is applied at the finer levels of the oct-tree (from the finest to one level before the coarsest) as it is shown in Figure 1(b). This partial contribution does not overlap with any other, because of the exclusive distribution of observation groups among nodes. To avoid internode communications at the finer levels, the information relative to the source nearby groups that are required by the MLFMA must be also kept in the node. This is illustrated in Figure 1(b): apart from the observation groups 0 to 3 and their respective children at finer levels (represented in black color), the computing node must also have the data relative to the adjacent groups and their children at lower levels (represented here in light gray color — yellow in the color version of the paper). The storage of these border groups data causes a slight memory footprint replicating some data among “nearby” nodes, but it is worth mentioning that it happens in exchange for data locality and the elimination of communication needs at the MLFMA stage. Besides, using the hierarchical oct-tree ordering of [24] minimizes the number of border source groups that must be stored in each node.

In order to make clear the MLFMA-FFT operation the Figure 2 has been included. Throughout this and the following sections, the parameters K , N , M and n referring to the number of k -space samples, the total number of unknowns of the problem, the number of non-empty groups and the number of nodes, respectively, will be used. The subscripts c and f are added to indicate coarsest and finest levels, respectively. It can be observed in the figure that the MLFMA algorithms work in levels comprised between the finest and the preceding to the coarsest one. As it was detailed before, the near coupling contributions to the MVP at the coarsest level are computed by the MLFMA algorithm inside each node. Since the MLFMA is a well-known algorithm, these calculations are not presented in detail in the figure, which is focused on showing how the MLFMA-FFT tackles the far-field interactions. After the required interpolation of the outcoming fields to the coarsest level, we can see in Figure 2

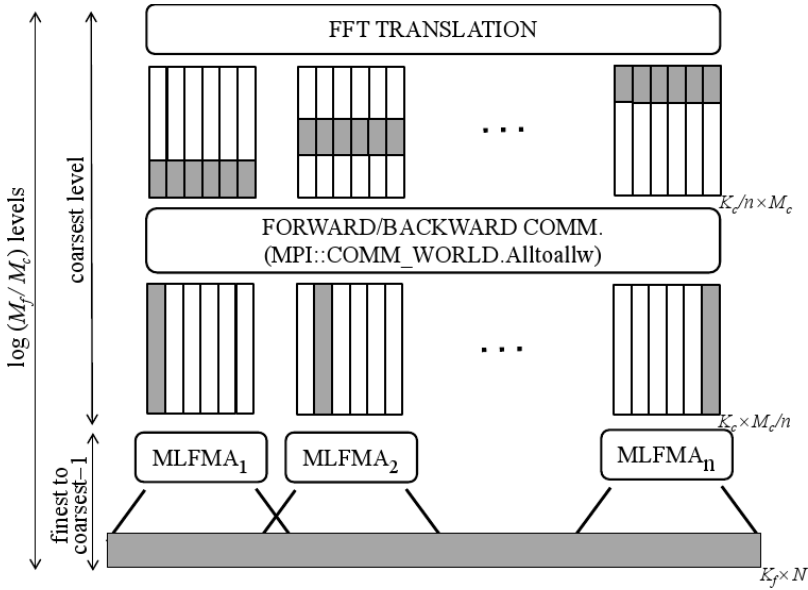


Figure 2. MLFMA-FFT graphic scheme.

that each node has the complete set of directions, K_c , for its assigned M_c/n observation groups. At this point, an all to all communication is performed in order to obtain the partial K_c/n samples assigned to each node for all the M_c groups at the coarsest level (distribution by fields). These inter-node communications are efficiently carried out during the MVP in a single step by using the asymmetric MPI **Alltoallw** operation. The **Alltoallw** high-level command makes possible to accomplish all the required communications without latency periods or explicit synchronization because of the efficient management provided by the MPI library. With regard to the iterative solver, the work is distributed by equal number of unknowns per processor.

Regarding the parallel efficiency of the method, given that the way of carrying out the distributed computation in the MLFMA-FFT and the Nested FMM-FFT algorithms is a common point, a similar scalability performance of both methods is expected. The parallel speed-up curve can be checked in [22].

3.2. Computational Complexity

To analyze the computational cost of the MLFMA-FFT method we need to assess both the cost of the FFT translations carried out at the coarsest level and the cost of the MLFMA algorithms applied intranode

involving the rest of levels. The relations $K_c \propto N/M_c$ and $Q_c \approx M_c^{3/2}$ are taken into account, where Q_c is the total number of groups at the coarsest level (including non-empty and empty groups). Under these conditions, the FFT-based translations cost at the coarsest level can be written as $\mathcal{O}(K_c Q_c \log Q_c) \approx \mathcal{O}(N \sqrt{M_c} \log M_c)$ [22].

Regarding the MLFMA, it is known that an efficient implementation achieves a $\mathcal{O}(N)$ computational cost at each level and requires approximately $\log(N)$ levels, then leading to a global cost of $\mathcal{O}(N \log N)$. In the method proposed in this work, $\mathcal{O}(M_c)$ MLFMA algorithms of $\mathcal{O}(N/M_c)$ unknowns are applied to account for the near interactions throughout the multilevel oct-tree. Given that the number of required levels will be approximately $\log(M_f/M_c)$, a computational cost $\mathcal{O}(N/M_c \log(M_f/M_c) M_c) = \mathcal{O}(N \log(M_f/M_c))$ is obtained. If the condition $M_f \propto N$ of small finest level groups is respected in order to guarantee the MLFMA efficiency, this cost can be expressed as $\mathcal{O}(N \log(N/M_c))$.

The adequate selection of M_c must be discussed at this point. On the one hand, the number of coarsest groups must be large enough (small-sized groups) to guarantee that the MLFMA algorithms can be run independently inside each node avoiding internode communications. Besides, this makes easier to obtain a well-balanced work-load distribution (that must be obtained regarding the number of unknowns contained in the groups assigned to each node), since this distribution is performed at the coarsest level of the oct-tree. On the other hand, the selection of a small M_c number maximizes the volume of interactions that are computed by means of the MLFMA, and thereby using the shared memory configuration, which constitutes the most efficient alternative. Taking into account all these considerations, the number of MLFMA algorithms must be chosen proportional to the number of available nodes, n . Therefore, with $M_c \propto n$ and given that n depends only on the available computational resources and not on the problem size, the asymptotic relation $M_c \approx \mathcal{O}(1)$ can be established.

Therefore, observing the two main terms involved in the computational cost of the MLFMA-FFT detailed above, it can be concluded that the global cost is dominated by the MLFMA cost and it tends to $\mathcal{O}(N \log N)$.

Regarding the memory consumption, the estimation corresponding to the FFT coarsest level translations is $\mathcal{O}(Q_c) \approx \mathcal{O}(M_c \sqrt{M_c})$. For the MLFMA algorithms the memory requirements are given by a $\mathcal{O}(N \log(M_f/M_c))$ term. According to the previous assumptions for M_f and M_c , the prevailing memory requirements of the MLFMA-FFT method also tend to $\mathcal{O}(N \log N)$.

In order to illustrate the cost estimations detailed above, several




Figure 3. MLFMA-FFT cpu time requirement for a sphere with different number of unknowns.




Figure 4. MLFMA-FFT memory consumption for a sphere with different number of unknowns.

executions of a sphere with different number of unknowns have been carried out using the MLFMA-FFT algorithm. Figures 3 and 4 show the computational cost and the memory consumption, respectively. Both of them fit the curve $\mathcal{O}(N \log N)$ according to the given estimations.

4. NUMERICAL RESULTS

The examples of the following Sections 4.1 and 4.2 have been addressed with an Electric Field Integral Equation (EFIE) based Method of Moments formulation, in which the well-known Rao-Wilton-Glisson (RWG) basis functions [25] have been applied both in the discretization of the geometry and the Galerkin's testing procedure. No preconditioning was used in the simulations and the iterative solver employed was GMRES.

4.1. Evolution of the Parallel Implementations

In order to make a comparison between the FMM-FFT, the Nested FMM-FFT and the MLFMA-FFT parallel algorithms, the problem of a sphere of 150 millions of unknowns has been solved by means of the three methods using the same GMRES iterative solver parameters (then obtaining an identical residual error below 10^{-2}).

The calculations with the FMM-FFT were performed using the HPC supercomputer Finis Terrae, installed in the Supercomputing Center of Galicia (CESGA). Finis Terrae consists of 142 cc-NUMA HP Integrity rx7640 with 8 dual core Intel Itanium 2 Montvale processors at 1.6 GHz with 18 MB L3 cache and 128 GB per node. The solution for this sphere employed 64 nodes involving a total of 1,024 processors. The nodes are interconnected through a high efficiency Infiniband network (4xDDR), and the operating system is Linux SLES 10. For this result, we have used the Intel C++ Compiler version 10.1.012, Intel MPI version 3.1.038 for communications and Intel Cluster MKL version 10.0.2.018 for matrix/vector linear algebra operations. The Nested FMM-FFT and MLFMA-FFT simulations were carried out using the LUSITANIA supercomputer, installed in the Centro Extremeño de Investigación, Innovación Tecnológica y Supercomputación (CénitS). Lusitania is made up of 2 HP Integrity SuperDome SX2000 nodes with 64 dual core Itanium2 Montvale processors at 1.6 GHz (18 MB cache). Only 64 processors per node were used for the sphere simulations. The operating system is Linux SLES 10 and we have used the Intel C++ Compiler version 11.0.069, Intel MPI version 3.2.0.011 for communications and Intel Cluster MKL version 10.0.2.018

for matrix/vector linear algebra operations.

The main configuration aspects of each method and their corresponding time and memory requirements are detailed in Table 1. It can be observed that a low amount of memory is required in the nested and the MLFMA-FFT implementations in contrast with the FMM-FFT consumption, more than 10 times higher. As it was also expected, the MLFMA-FFT takes up a solution time almost 5 times lower than the time required by the Nested FMM-FFT with a comparable configuration. It has been also included in Table 1 the total CPU time, which takes into account the total number of processors employed. The results are in accordance with the expected method behavior described in previous sections and point out that the MLFMA-FFT is the most efficient alternative to take advantage of the computational resources with mixed (shared/distributed) architectures.

Table 1. Evolution of methods: analysis of a 150,039,552 unknowns sphere.

	FMM-FFT	Nested FMM-FFT		MLFMA-FFT
Groups dimensions: nest/coarsest level	2λ	0.25λ/8λ		
Multipole terms	20	5/62	5/7/12/20/35/62	
Num. total/non-empty groups: finest level coarsest level	8, 000, 000/185, 453	4, 065, 356, 736/10, 937, 250 125, 000/11, 637		
Supercomputer	Finis Terrae (CESGA)	LUSITANIA (C�nitS)		
num. nodes	64	2		
processors per node	16	64		
total processors	1024	128		
Min./max. peak mem. (in node)	76.3/84.7 GB	227.4/228.8 GB	243.7/245.3 GB	
Total memory	5.4 TB	456.2 GB	489.0 GB	
GMRES: num. iter./restart	11/10			
Setup/solution time	1.1 / 4.98 h	1.35 / 25.22 h	1.51 / 5.25 h	
Total CPU time: setup/solution time × total processors	1126.4/5099.5 h	172.8/3228.2 h	193.3/672.0 h	

4.2. Practical Numerical Example

The 79 GHz radar cross section (RCS) analysis of a car (CITROËN C3) using the MLFMA-FFT has been carried out. The 77–81 GHz frequency band has been designed for the automotive collision warning future Short Range Radars (SRR). For this reason, modeling the electromagnetic behavior of a car at 79 GHz is of great interest for the automotive industry. The large size of the required analysis has made difficult to obtain suitable results up to now. Instead of resorting to asymptotic approaches with reduced accuracy, a reliable result can be achieved by means of the MLFMA-FFT method.

This example has been performed employing the LUSITANIA supercomputer described above. A total of 1.6 TB of RAM and 2 HP Integrity SuperDome SX2000 nodes with 128 processors have been used. The model of the car is made up of 620 million unknowns. Both configuration and solution related data are gathered in Table 2. The GMRES(50) residual error obtained after 5 iterations was below $5 \cdot 10^{-2}$.

The 79 GHz bistatic RCS result is shown in Figure 5. A front incidence ($\theta = 90^\circ, \phi = 270^\circ$) has been considered. Due to the rapid fluctuation of the RCS pattern with changing aspect angle, a window of 2° has been selected to calculate the median value of the RCS in the backward direction, which results 0.34 dBsm [26]. This RCS analysis is complemented by the surface current density distribution

Table 2. Technical data for the solution of a 620 millions of unknowns problem.

Frequency	79 GHz
Number of unknowns	620, 739, 632
Groups dimensions finest/coarsest level	$0.2\lambda/12.8\lambda$
Number of levels	7
Multipole terms	4/7/11/18/29/52/95
Number of total/non-empty groups finest level coarsest level	16, 841, 845, 020/31, 201, 960 66, 600/7, 848
Num. of nodes/processors per node	2/128
Total processors	256
Min./max. peak memory in node	816/821 GB
Total memory	1.6 TB
Num. of iterations/GMRES restart	5/50
Setup/solution time	3.65/43.2 h




Figure 5. Bistatic RCS of a 620 million unknowns car at 79 GHz. Elevation angle $\theta = 90^\circ$.




Figure 6. Detail of the current density over the car surface at 24 GHz.

corresponding to a frequency of 24 GHz shown enlarged in Figure 6. This kind of analysis can provide useful information for the design initial stages of the automotive short range radar systems which are still under development.

5. CONCLUSION

The MLFMA-FFT algorithm presented in this work takes advantage of the high scalability behavior of the FMM-FFT for the distributed

computation and of the algorithmical efficiency of the MLFMA for shared-memory computers. Interesting analysis can be carried out by means of this technique due to its ability for profiting from present supercomputing resources. In fact, the analysis of the largest problem solved up to now with more than 620 million unknowns presented in this work has demonstrated the potential of the MLFMA-FFT algorithm for solving extremely large problems with practical or industrial interest. The electromagnetic behavior of a car has been studied at 79 GHz providing results with direct applications in the future radar sensor design for the automotive industry.

ACKNOWLEDGMENT

This work was supported by the Spanish Government (projects TEC2008-06714-C02-01, TEC2008-06714-C02-02, and CONSOLIDER-INGENIO2010 CSD2008-00068), and by Xunta de Galicia (project INCITE08PXIB322250PR). The authors would like to thank the Supercomputing Center of Galicia (CESGA) for their support in the use of Finis Terrae supercomputer and the Centro Extremeño de Investigación, Innovación Tecnológica y Supercomputación (CénitS) for their support in the use of LUSITANIA supercomputer, which were used to run the simulations.

REFERENCES

1. Harrington, R. F., *Field Computation by Moment Method*, IEEE Press, New York, 1993.
2. Coifman, R., V. Rokhlin, and S. Wanzura, "The fast multipole method for the wave equation: A pedestrian prescription," *IEEE Antennas Propag. Mag.*, Vol. 35, No. 3, 7–12, Jun. 1993.
3. Song, J. M. and W. C. Chew, "Multilevel fast multipole algorithm for solving combined field integral equations of electromagnetic scattering," *Microw. Opt. Tech. Lett.*, Vol. 10, 14–19, Sep. 1995.
4. Song, J. M., C. C. Lu, W. C. Chew, and S. W. Lee, "Fast Illinois solver code (FISC)," *IEEE Antennas Propag. Mag.*, Vol. 40, No. 3, 27–34, Jun. 1998.
5. Velampambil, S., J. M. Song, W. C. Chew, and K. Gallivan, "Scaleme: A portable scaleable multipole engine for electromagnetic and acoustic integral equation solvers," *IEEE Antennas Propag. Soc. Int. Symp.*, Vol. 3, 1774–1777, Jun. 1998.
6. Velampambil, S., J. E. Schutt-Aine, J. G. Nickel, J. M. Song, and W. C. Chew, "Solving large scale electromagnetic problems

- using a linux cluster and parallel MLFMA,” *IEEE Antennas Propag. Soc. Int. Symp.*, Vol. 1, 636–639, 1999.
7. Velamparambil, S., W. C. Chew, and M. L. Hastriter, “Scalable electromagnetic scattering computations,” *IEEE Antennas Propag. Soc. Int. Symp.*, Vol. 3, 176–179, 2002.
 8. Velamparambil, S., W. C. Chew, and J. M. Song, “10 Million unknowns: Is it that big?” *IEEE Antennas Propagat. Mag.*, Vol. 45, No. 3, 43–58, Apr. 2003.
 9. Sylvand, G., “Performance of a parallel implementation of the FMM for electromagnetics applications,” *Int. J. Numer. Meth. Fluids*, Vol. 43, 865–879, 2003.
 10. Velamparambil, S. and W. C. Chew, “Analysis and performance of a distributed memory multilevel fast multipole algorithm,” *IEEE Trans. Antennas Propag.*, Vol. 53, No. 8, 2719–2727, 2005.
 11. Gürel, L. and Ö. Ergül, “Fast and accurate solutions of extremely large integral-equation problems discretised with tens of millions of unknowns,” *Electronics Letters*, Vol. 43, No. 9, 499–500, Apr. 2007.
 12. Ergül, Ö. and L. Gürel, “Parallel-MLFMA solution of CFIE discretized with tens of millions of unknowns,” *Proc. 2nd Europ. Conf. Antennas Propag. (EuCAP)*, Nov. 2007.
 13. Pan, X.-M. and X. X.-Q. Sheng, “A sophisticated parallel MLFMA for scattering by extremely large targets,” *IEEE Antennas Propagat. Mag.*, Vol. 50, 129–138, Jun. 2008.
 14. Ergül, Ö. and L. Gürel, “Efficient parallelization of the multilevel fast multipole algorithm for the solution of large-scale scattering problems,” *IEEE Trans. Antennas Propag.*, Vol. 56, No. 8, 2335–2345, Aug. 2008.
 15. Ergül, Ö. and L. Gürel, “A hierarchical partitioning strategy for an efficient parallelization of the multilevel fast multipole algorithm,” *IEEE Trans. Antennas Propag.*, Vol. 57, No. 6, 1740–1750, Jun. 2009.
 16. Wagner, R., J. M. Song, and W. C. Chew, “Montecarlo simulation of electromagnetic scattering from two-dimensional random rough surfaces,” *IEEE Trans. Antennas Propag.*, Vol. 45, No. 2, 235–245, Feb. 1997.
 17. Waltz, C., K. Sertel, M. A. Carr, B. C. Usner, and J. L. Volakis, “Massively parallel fast multipole method solutions of large electromagnetic scattering problems,” *IEEE Trans. Antennas Propag.*, Vol. 55, No. 6, 1810–1816, Jun. 2007.
 18. Landesa, L., J. M. Taboada, F. Obelleiro, J. L. Rodriguez,

- C. Mouriño, and A. Gomez, "Solution of very large integral-equation problems with single-level FMM," *Microw. Opt. Technol. Lett.*, Vol. 51, No. 6, 20–28, Dec. 2009.
19. Taboada, J. M., L. Landesa, F. Obelleiro, J. L. Rodriguez, J. M. Bertolo, M. G. Araujo, J. C. Mouriño, and A. Gomez, "High scalability FMM-FFT electromagnetic solver for supercomputer systems," *IEEE Antennas Propagat. Mag.*, Vol. 51, No. 6, 20–28, Dec. 2009.
 20. Taboada, J. M., L. Landesa, J. M. Bertolo, F. Obelleiro, J. L. Rodriguez, C. Mouriño, and A. Gomez, "High scalability multipole method for the analysis of hundreds of millions of unknowns," *Proc. IEEE Europ. Conf. Antennas Propag. (EuCAP)*, Berlin, Germany, Mar. 2009.
 21. Landesa, L., J. M. Taboada, J. L. Rodriguez, F. Obelleiro, J. M. Bertolo, J. C. Mouriño, and A. Gomez, "Analysis of 0.5 billion unknowns using a parallel FMMFFT solver," *Proc. IEEE Antennas Propag. Soc. Int. Symp.*, Charleston (South Carolina), USA, Jun. 2009.
 22. Araújo, M. G., J. M. Taboada, F. Obelleiro, J. M. Bertolo, L. Landesa, J. Rivero, and J. L. Rodríguez, "Supercomputer aware approach for the solution of challenging electromagnetic problems," *Progress In Electromagnetics Research*, Vol. 101, 241–256, 2010.
 23. Saad, Y. and M. Schultz, "GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems," *SIAMJ. Sci. Statist. Comput.*, Vol. 7, 856–869, 1986.
 24. Gumerov, N. A., R. Duraiswami, and E. A. Borovikov, "Data structures, optimal choice of parameters and complexity results for generalized multilevel fast multipole methods in d dimensions," UM Computer Science Technical Report CS-TR-4458 UMIACS, University of Maryland, 2003.
 25. Rao, S. M., D. R. Wilton, and A. W. Glisson, "Electromagnetic scattering by surfaces of arbitrary shape," *IEEE Trans. Antennas Propagat.*, Vol. 30, 409–418, May 1982.
 26. Knott, E. F., J. F. Shaeffer, and M. T. Tuley, *Radar Cross Section*, Artech House, Inc., Norwood, MA, 1985.