

Hierarchical Parallelization of the Multilevel Fast Multipole Algorithm for the Efficient Solution of Large-Scale Scattering Problems[†]

Özgür Ergül^{1,2} and Levent Gürel^{1,2*}

¹Department of Electrical and Electronics Engineering

²Computational Electromagnetics Research Center (BiLCEM)

Bilkent University, TR-06800, Bilkent, Ankara, Turkey

E-mail: ergul@ee.bilkent.edu.tr, lgurel@bilkent.edu.tr

Introduction

The multilevel fast multipole algorithm (MLFMA) [1] provides fast and efficient solutions of scattering problems involving large objects with three-dimensional arbitrary geometries. However, accurate solutions of many real-life problems require discretizations with millions of unknowns, which may not be handled easily with sequential implementations of MLFMA. In order to achieve the solution of such large problems, it has been popular to increase the computational resources by parallelizing MLFMA on distributed memory architectures [2]. However, due to its complicated structure, parallelization of MLFMA is not trivial. Recently, we proposed a hierarchical parallelization strategy [3], which provides significantly higher efficiency than the previous parallelization approaches, especially as the number of processors increases. In this paper, we present the details of our algorithm and demonstrate its high efficiency on large sphere problems, involving as many as record-breaking 53 million unknowns.

Computational Properties of MLFMA

For perfectly-conducting objects, discretizations of surface integral equations, such as the combined-field integral equation (CFIE), lead to $N \times N$ dense matrix equations, where the matrix elements can be interpreted as the electromagnetic interactions of the discretization elements, i.e., basis and testing functions. For the discretization, we use Rao-Wilton-Glisson functions defined on planar triangles. The resulting matrix equations can be solved iteratively, where the required matrix-vector multiplications are performed efficiently by MLFMA as

$$\bar{\mathbf{Z}} \cdot \mathbf{x} = \bar{\mathbf{Z}}_{NF} \cdot \mathbf{x} + \bar{\mathbf{Z}}_{FF} \cdot \mathbf{x}. \quad (1)$$

In (1), the near-field interactions denoted by $\bar{\mathbf{Z}}_{NF}$ are calculated directly, while the far-field interactions denoted by $\bar{\mathbf{Z}}_{FF}$ are computed approximately via three stages, i.e., aggregation, translation, and disaggregation [1]. Without losing generality, we consider a smooth scatterer with an electrical dimension of kD , where $k = 2\pi/\lambda$ is the wavenumber. The scatterer is placed in a cubic box and the computational domain is recursively divided into subdomains until the box size is in the range from 0.15λ to 0.3λ . A multilevel tree structure with $L = \mathcal{O}(\log(kD)) = \mathcal{O}(\log N)$ levels is constructed by considering the nonempty boxes (clusters). Then, the far-field interactions are calculated in a cluster-by-cluster manner at different levels. At level l from 1 to L , the number of clusters can be approximated as $N_l \approx 4^{(1-l)}N_1$, where $N_1 = \mathcal{O}(N)$. For each cluster, radiated and incoming fields are sampled at $S_l = 2(T_l + 1)^2$ points, where T_l is the truncation number determined by the excess

[†]This work was supported by the Scientific and Technical Research Council of Turkey (TUBITAK) under Research Grants 105E172 and 107E136, by the Turkish Academy of Sciences in the framework of the Young Scientist Award Program (LG/TUBA-GEBIP/2002-1-12), and by contracts from ASELSAN and SSM. Computer time was provided in part by a generous allocation from Intel Corporation.

bandwidth formula. In general, the number of samples for a cluster is proportional to its size as measured by the wavelength; thus, $S_1 = \mathcal{O}(1)$ and $S_l \approx 4^{(l-1)}S_1$. Computational requirements of MLFMA for level l is proportional to $N_l S_l$. Since $N_l S_l \approx N_1 S_1 = \mathcal{O}(N)$, all levels of MLFMA have equal importance with $\mathcal{O}(N)$ complexity, leading to a total of $\mathcal{O}(N \log N)$ complexity.

Hierarchical Parallelization of MLFMA

Parallelization of MLFMA is not trivial due to its complicated multilevel tree structure. Simple parallelization strategies based on distributing the clusters among the processors usually fail to provide efficient solutions, especially when the number of processors is large [4]. A hybrid parallelization strategy was suggested by applying two different partitioning schemes in the lower and higher levels of the tree structure [2],[4]. Using the hybrid strategy, clusters in the lower levels are still distributed among the processors, while the clusters in the higher levels are shared by all processors and the samples of the fields are distributed. Although this approach provides higher efficiency compared to the simple parallelization strategy, its performance may not be sufficient. Recently, we developed a hierarchical parallelization strategy based on partitioning both the clusters and the samples of the fields at each level [3]. This strategy is well suited for the multilevel structure of MLFMA and provides higher efficiency than the simple and hybrid parallelization approaches. In this paper, we present the details of our algorithm.

Partitioning of the Tree Structure: Using the hierarchical parallelization strategy, we distribute both the clusters and the samples of the radiated and incoming fields among the processors. We carefully choose the numbers of partitions by considering the numbers of clusters and the samples at each level separately. As an example, for a parallelization among $p = 2^i$ processors, where $i \geq (L - 1)$, clusters and samples are distributed into $2^{(1-l)}p$ and $2^{(l-1)}$ partitions, respectively, for $l = 1, 2, \dots, L$. Then, the number of clusters assigned to each processor is proportional to $(N_1/p)2^{(1-l)}$, which decreases by the factor of 2 from a level to the next upper level. On the other hand, the number of samples per processor is proportional to $S_1 2^{(l-1)}$. The samples are partitioned only along the θ direction as described in [2], and we keep the number of samples in the ϕ direction proportional to $(T_1 + 1)2^{(l-1)}$ per processor. Then, the number of samples in the θ direction is approximately constant for the entire tree structure, which is an important advantage of using the hierarchical parallelization strategy. In the hybrid parallelization approach, partitioning the samples among p processors may lead to a poor load-balance, especially when p is comparable to $(T_1 + 1)$, i.e., the total number of samples in the θ direction.

Aggregation/Disaggregation Stages: During the aggregation (disaggregation) stage of MLFMA, radiated (incoming) fields are calculated at the centers of the clusters from the bottom (top) of the tree structure to the highest (lowest) level. Fig. 1 presents an example for the partitioning of the clusters and the samples of the fields at some levels l and $(l + 1)$, when MLFMA is parallelized among 16 processors. At level l , 4×4 partitioning is used, i.e., the number partitions for both the clusters (horizontal direction) and the samples of the fields (vertical direction) are 4. By performing the aggregation operations involving local interpolations and exponential shifts, radiated fields at the centers of the clusters in level $(l + 1)$ are calculated. Since the samples are partitioned, interpolations require one-to-one communications between the processors, as detailed in [2]. For the specific partitioning scheme shown in Fig. 1, communications are performed within four separate groups including the processors in the same columns, i.e., (1,5,9,13), (2,6,10,14), (3,7,11,15), and (4,8,12,16). Using the hierarchical partitioning strategy, distribution of the samples




Figure 1: Partitioning of the clusters and the samples of the fields using the hierarchical parallelization strategy.

into large numbers of partitions is avoided. Therefore, communications are required mostly between the processors located next to each other. For example, processor 6 in Fig. 1 communicates mainly with processors 2 and 10, but not with processor 14, unless the interpolation order is extremely large. Following the aggregation operations, the partitioning is modified from 4×4 to 8×2 as depicted in Fig. 1 via one-to-one communications. Data is exchanged between the processors that are paired as follows: (1,2), (3,4), (5,6), (7,8), (9,10), (11,12), (13,14), and (15,16). In the disaggregation stage, all operations described above are performed in a reverse manner; following the data exchanges between the processors, incoming fields are shifted and antinterpolated, and the resulting data is communicated within the four groups.

Translation Stage: In MLFMA, translations are performed between the aggregation and disaggregation stages to translate the radiated fields into incoming fields. Since the clusters are partitioned, some of the translations are between the clusters that are located in different processors. Therefore, one-to-one communications are required between the processors. For level l in Fig. 1, communications are performed within four separate groups including the processors in the same rows, i.e., (1,2,3,4), (5,6,7,8), (9,10,11,12), and (13,14,15,16).

Results

Fig. 2(a) presents the parallelization efficiency for MLFMA solutions of a scattering problem involving a sphere of radius 60λ discretized with 13,278,096 unknowns. The problem is solved on a cluster of Intel Xeon 7300 processors (2.93 GHz) connected via an Infiniband network, as the number of processors changes from 4 to 128. In addition to the hierarchical parallelization strategy, we also use the simple and hybrid parallelization strategies. The parallelization efficiency is defined as

$$\epsilon_p = 4T_4 / (pT_p), \quad (2)$$

where T_p is the processing time required for the solution with p processors. All parallelization strategies are optimized with load-balancing algorithms. Fig. 2(a) shows that the overall efficiency, including the setup and iterative solution parts, is increased significantly by using the hierarchical parallelization strategy compared to the simple and hybrid strategies. For 128 processors, the hierarchical parallelization provides 65% efficiency (corresponding to 21-fold speedup with respect to the 4-processor solution), which is significantly higher than 26% and 50% efficiencies provided by the simple and hybrid parallelization strategies, respectively.




Figure 2: (a) Parallelization efficiency for the solution of a scattering problem involving a sphere of radius 60λ discretized with 13,278,096 unknowns. (b) Bistatic RCS (in dB) of a sphere of radius 120λ discretized with 53,112,384 unknowns from 170° to 180° , where 180° corresponds to the forward-scattering direction.

Using the hierarchical parallelization strategy, we are able to solve very large problems discretized with tens of millions of unknowns. As an example, we consider the solution of a scattering problem involving a sphere of radius 120λ discretized with 53,112,384 unknowns. This is the largest integral-equation problem solved to date. Scattering problem is formulated with CFIE and solved with 2-digits of accuracy using a 9-level MLFMA parallelized into 16 and 32 processors. Overall time including the setup and the iterative solution (21 BiCGStab iterations for 0.001 residual error) parts is 547 minutes and 289 minutes, respectively, with 16 and 32 processors. The speed up obtained by increasing the number of processors from 16 to 32 is 1.89, corresponding to 95% efficiency. To present the accuracy of the solutions, Fig. 2(b) depicts the normalized bistatic radar cross section ($\text{RCS}/\pi a^2$, where a is the radius of the sphere in meters) values in decibels (dB). We observe that the computational values are in agreement with the analytical values obtained by a Mie-series solution.

Conclusion

We present the details of a hierarchical parallelization strategy, which provides higher efficiency than the previous parallelization approaches for MLFMA. We demonstrate the effectiveness of our algorithm on sphere problems involving large numbers of unknowns, such as a sphere of radius 120λ discretized with 53 million unknowns.

References

- [1] J. Song, C.-C. Lu, and W. C. Chew, "Multilevel fast multipole algorithm for electromagnetic scattering by large complex objects," *IEEE Trans. Antennas Propag.*, vol. 45, no. 10, pp. 1488–1493, Oct. 1997.
- [2] S. Velampambal and W. C. Chew, "Analysis and performance of a distributed memory multilevel fast multipole algorithm," *IEEE Trans. Antennas Propag.*, vol. 53, no. 8, pp. 2719–2727, Aug. 2005.
- [3] Ö. Ergül and L. Gürel, "Hierarchical parallelisation strategy for multilevel fast multipole algorithm in computational electromagnetics," *Electronics Lett.*, vol. 44, no. 1, pp. 3–5, Jan. 2008.
- [4] Ö. Ergül and L. Gürel, "Efficient parallelization of multilevel fast multipole algorithm," in *Proc. European Conference on Antennas and Propagation (EuCAP)*, no. 350094, 2006.