Fast and accurate solutions of extremely large integral-equation problems discretised with tens of millions of unknowns

L. Gürel and Ö. Ergül

The solution of extremely large scattering problems that are formulated by integral equations and discretised with tens of millions of unknowns is reported. Accurate and efficient solutions are performed by employing a parallel implementation of the multilevel fast multipole algorithm. The effectiveness of the implementation is demonstrated on a sphere problem containing more than 33 million unknowns, which is the largest integral-equation problem ever solved to our knowledge.

Introduction: For numerical solutions of scattering problems in electromagnetics, integral-equation formulations provide accurate results when they are discretised appropriately by using small elements with respect to wavelength [1]. Simultaneous discretisations of the scatterer and the integral equations lead to dense matrix equations, which can be solved iteratively using efficient acceleration methods, such as the multilevel fast multipole algorithm (MLFMA) [2]. However, accurate solutions of many real-life problems require discretisations with millions of elements, which result in matrix equations with millions of unknowns. For the solutions of these large-scale problems, MLFMA must be parallelised, but this is not trivial owing to the complicated structure of the algorithm [3–5]. In this Letter, we report an implementation of the parallel MLFMA that is able to solve problems discretised with tens of millions of unknowns. Specifically, we present the results of a scattering problem involving a sphere of radius 96λ, where an accurate solution requires a discretisation with more than 33 million unknowns. To the best of our knowledge, this is the solution of the largest integral-equation problem reported up to now.

Parallel implementation of MLFMA: MLFMA performs the matrix-vector multiplications related to an $N \times N$ dense matrix equation in $\mathcal{O}(NL)$ time using $\mathcal{O}(NL)$ memory, where $L = O(\log N)$ is the number of levels of the tree structure, which is constructed by recursively dividing the computational domain into sub-domains (clusters). MLFMA calculates the far-field interactions between the radiating (basis) and receiving (testing) elements in a group-by-group manner consisting of three stages: aggregation, translation and disaggregation. For each matrix-vector multiplication required by the iterative solver, these stages are performed on the tree structure in a multilevel manner. There are also $O(N)$ near-field interactions that are calculated directly and stored in the memory to be used multiple times.

In the aggregation step, radiation patterns of the clusters are computed from the bottom of the tree structure to the top. Before the iterations, radiation patterns of the basis functions are calculated and stored in the memory. Owing to the nature of the Helmholtz equation, sampling rates of the radiation patterns depend on the sizes of the clusters. Using the excess-bandwidth formula and considering the worst-case scenario [6], we determine the number of samples for each level according to the desired accuracy. The samples are chosen uniformly in the ϕ direction while they are selected as the Gauss-Legendre points in the θ direction. During the aggregation process, sampling rates of the consecutive levels are matched by employing a local Lagrange interpolation algorithm with enhanced accuracy [7].

For the parallelisation of the aggregation process, we choose a level of distribution (LoD) to divide the clusters among the processors. Using a load-balancing algorithm, the levels below the LoD are distributed among the processors by assigning each cluster and its parent cluster to the same processor. In this way, aggregation operations can be performed independently in each processor from the bottom of the tree structure up to the LoD without any communication [8]. In the higher levels above the LoD, however, radiation patterns are distributed among the processors, instead of the clusters [5]. This is required in order to improve the load balancing since the higher levels include fewer clusters with densely-sampled fields. Then, an all-to-all communication is required at the LoD to switch between the two strategies applied in the lower and higher levels of the tree structure. We also note that one-to-one communications are required in the higher levels, where the fields are distributed among the processors and the interpolations in a processor require samples that are stored in other processors [5].

In the translation stage of MLFMA, radiated fields of the clusters are converted into incoming fields for other clusters. Translations are performed between pairs of clusters when the clusters are far from each other while their parent clusters are electromagnetically close to each other. Using a one-box-buffer scheme, there are $O(1)$ translation operations for each cluster in any level. Translation functions to perform these operations are calculated and stored in memory before the iterations. Using regularly-spaced cubic clusters, we significantly reduce the number of different translation functions required for each level [9]. In addition, calculation of the translation operators is accelerated by using local interpolation techniques that are optimised according to the desired level of accuracy [10]. Each translation operator is an infinite summation that must be truncated [2], where the truncation number is also determined by the excess-bandwidth formula [6].

In the lower levels below the LoD, some of the translations can be performed in each processor without any communication, while the rest are related to the clusters that are assigned to different processors so that communications to complete these translations are inevitable. We carefully organise the required data transfers by matching the processors appropriately using a communication map. For the upper levels above the LoD, all translations can be performed without any communication; this is another advantage of distributing fields instead of clusters [5]. After the translations, the disaggregation stage is performed as the inverse of the aggregation process. Incoming fields are calculated for each cluster from the top of the tree structure to the lowest level. The incoming field to a cluster is a combination of the incoming field to its parent cluster and the incoming fields due to the translations. We use transpose interpolation to accurately match the different sampling rates of the successive levels [7]. At the end of the disaggregation, a numerical integration is performed for each testing function in the lowest level to complete the matrix-vector multiplications related to the far-field interactions. Finally, matrix-vector multiplications related to near-field interactions are performed directly. For high efficiency, it is crucial to distribute the near-field interactions among the processors according to a load-balancing strategy, which usually leads to different partitioning schemes for the near-field and far-field interactions [8].

Table 1: MLFMA solution of a sphere problem with 33 791 232 unknowns

Geometry size (diameter)	192λ
Number of processors	16
Number of levels	9
Smallest cluster size	0.19λ
Total number of clusters	5 904 951
Number of clusters in lowest level	4 344 205
Number of near-field interactions	3 732 101 432
Truncation numbers (2 digits of accuracy)	6 to 546
Number of iterations (BiCGStab and 10^{-3} residual error)	21
Setup time (minutes)	177
Solution time (minutes)	265
Time for matrix-vector multiplication (s)	370
Memory for translation functions (GB)	2
Memory for radiation/receiving patterns (GB)	56
Memory for nearfield interactions (GB)	28
Memory for aggregation/disaggregation arrays (GB)	79

Results: To demonstrate the efficiency and accuracy of our implementation, we present the results of a sphere problem with radius 96λ. The discretisation of the problem with a mesh size of $\lambda/10$ leads to 33 791 232 unknowns when Rao-Wilton-Glisson [11] functions are employed as the basis and testing functions on triangular domains. The scattering problem is formulated with the combined-field integral equation [1] and iteratively solved by a biconjugate-gradient-stabilised (BiCGStab) algorithm. The solution is performed on a cluster of quad-core Intel Xeon 5355 processors connected via an Infiniband network and the results are summarised in Table 1, where we list the clustering information, processing times and memory usage. Using a block-diagonal preconditioner, only 21 iterations are

required to reduce the residual error below 10^{-3}. Parallelising the solution into 16 processes, the iterative solution is completed in 265 min. Finally, the bistatic radar cross-section (RCS) values are shown in Fig. 1, where the computed values sampled at 0.1° are in agreement with the analytical curve obtained by a Mie-series solution. In the Figure, 180° corresponds to the forward-scattering direction and the root-mean-square error [3] of the RCS is only 0.915 dB in the 170–180° range.

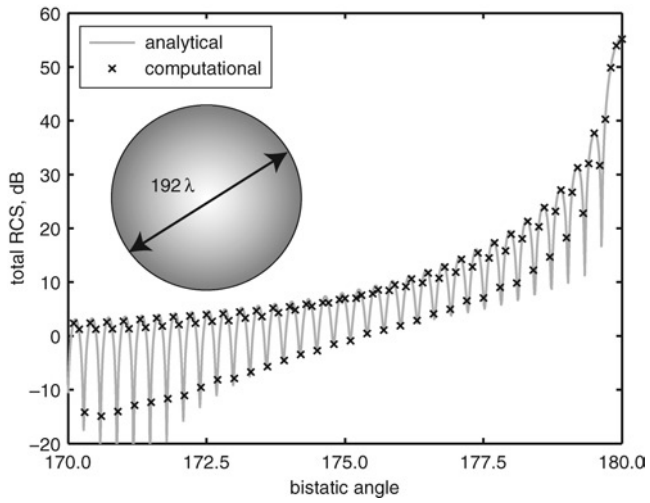

Fig. 1 *Bistatic RCS of sphere of radius 96λ*

Computational values obtained by solution of 33 791 232-unknown problem are in agreement with analytical curve obtained by Mie-series solution

Conclusions: We have presented the integral-equation solution of a scattering problem involving a sphere of radius 96λ discretised with 33 791 232 unknowns, which corresponds to the solution of a dense matrix equation with more than 10^{15} nonzero elements. This is the largest integral-equation problem reported to date. By employing an efficient implementation of the parallel MLFMA, it becomes possible to solve such large-scale problems on relatively inexpensive computational platforms.

Acknowledgments: This work was supported by the Scientific and Technical Research Council of Turkey (TUBITAK) under Research Grant 105E172, 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.

© The Institution of Engineering and Technology 2007
5 March 2007
Electronics Letters online no: 20070639
doi: 10.1049/el:20070639

L. Gürel (*Computational Electromagnetics Research Center (BiLCEM) and Department of Electrical and Electronics Engineering, Bilkent University, Bilkent, Ankara TR-06800, Turkey*)

E-mail: lgurel@bilkent.edu.tr

Ö. Ergül (*Department of Electrical and Electronics Engineering, Bilkent University, Bilkent, Ankara TR-06800, Turkey*)

References

1 Poggio, A.J., and Miller, E.K.: 'Computer techniques for electromagnetics' (Permagon Press, Oxford, UK, 1973, chap. 4)
2 Song, J., Lu, C.-C., and Chew, W.C.: 'Multilevel fast multipole algorithm for electromagnetic scattering by large complex objects', *IEEE Trans. Antennas Propag.*, 1997, **45**, (10), pp. 1488–1493
3 Hastriter, M.L.: 'A study of MLFMA for large-scale scattering problems', PhD thesis, University of Illinois at Urbana-Champaign, 2003
4 Sylvand, G.: 'Performance of a parallel implementation of the FMM for electromagnetics applications', *Int. J. Numer. Meth. Fluids*, 2003, **43**, pp. 865–879
5 Velamparambil, S., and Chew, W.C.: 'Analysis and performance of a distributed memory multilevel fast multipole algorithm', *IEEE Trans. Antennas Propag.*, 2005, **53**, (8), pp. 2719–2727
6 Chew, W.C., Jin, J.-M., Michielssen, E., and Song, J.: 'Fast and efficient algorithms in computational electromagnetics' (Artech House, Boston MA, USA, 2001)
7 Ergül, Ö., and Gürel, L.: 'Enhancing the accuracy of the interpolations and anterpolations in MLFMA', *IEEE Antennas Wirel. Propag. Lett.*, 2006, **5**, pp. 467–470
8 Ergül, Ö., and Gürel, L.: 'Efficient parallelization of multilevel fast multipole algorithm'. Proc. European Conference on Antennas and Propagation (EuCAP), 2006, 350094
9 Velamparambil, S., Chew, W.C., and Song, J.: '10 million unknowns: Is it that big?', *IEEE Antennas Propag. Mag.*, 2003, **45**, (2), pp. 43–58
10 Ergül, Ö., and Gürel, L.: 'Optimal interpolation of translation operator in multilevel fast multipole algorithm', *IEEE Trans. Antennas Propag.*, 2006, **54**, (12), pp. 3822–3826
11 Rao, S.M., Wilton, D.R., and Glisson, A.W.: 'Electromagnetic scattering by surfaces of arbitrary shape', *IEEE Trans. Antennas Propag.*, 1982, **AP-30**, (3), pp. 409–418