Scalable Parallel DBIM Solutions of Inverse-Scattering Problems

Mert Hidayetoğlu, Carl Pearson, Levent Gürel, Wen-mei Hwu, and Weng Cho Chew

Department of Electrical and Computer Engineering
University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA
hidayet2@illinois.edu

***Abstract*— We report scalable solutions of inverse-scattering problems with the distorted Born iterative method (DBIM) on large number of computing nodes. Distributing forward solutions does not scale well when the number of illuminations is not greater than the number of computing nodes. As a remedy, we distribute both forward solutions and the corresponding forward solvers to improve granularity of DBIM solutions. This paper provides a set of solutions demonstrating good scaling of the proposed parallelization strategy up to 1,024 computing nodes, employing 16,394 processing cores in total.**

I. Introduction

Iterative solutions of nonlinear inverse-scattering problems with the distorted Born iterative method (DBIM) require several forward solutions for each object illuminations in each iteration. Distributing these illuminations among message-passing interface (MPI) processes, where each process handles a certain number of illuminations, was first proposed in [1]. In monochromatic object reconstructions, this strategy provides a good load balancing since the iterative forward solvers involve the same matrix system with different right-hand sides, often yielding similar number of iterations. This strategy is also implemented on 256 computing nodes involving graphics processing units (GPUs), where each forward solution is obtained on a single GPU [2]. However, distributing illuminations is not scalable because of the finite granularity of the partitioning strategy. To improve the granularity, we simultaneously distribute the illuminations and corresponding forward solutions among MPI processes, which can employ more number of nodes than the number of illuminations.

II. Effect of the Number of Illuminations

Consider a scenario where a numerical Shepp-Logan phantom [3] is placed into a reconstruction domain with an equal side length of 102.4λ, where λ is the wavelength in the background medium. The object is illuminated and the scattered field is collected by 1,024 transmitters and receivers, respectively. The solution domain is discretized with 1,048,576 pixels, whose wave properties are unknown. When all illuminations are involved into the numerical reconstruction, the corresponding right-hand side has 1,048,576 complex numbers, equal to the number of unknowns. However, practically, including all illuminations is not only costly, but also redundant since the finite precision of the numerical solutions spoils the conditioning of the nonlinear inverse problem.

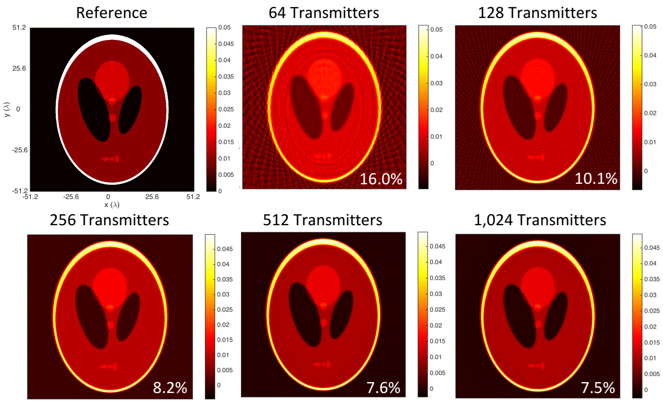

Fig. 1. Reconstructions of a numerical Shepp-Logan phantom with various number of illuminations. The relative object mismatches are noted in the lower-right corners. The reconstruction with 256 illuminations yields less than 10% mismatch, yielding a fair perception of the object details.

Fig. 1 shows five reconstructions of the numerical phantom with various numbers of illuminations. The DBIM solver performs 50 iterations without any regularization on the cost functional or intentional noise on the measured data. It can be seen that inclusion of all the 1,024 illuminations is not necessary for practical imaging purposes.

III. Parallelizing Illuminations

The DBIM solutions in Fig. 1 partitions the illuminations among computing nodes, where each node handles an equal number of illuminations. The forward solutions are obtained with the multilevel fast multipole algorithm (MLFMA), where each node employs a single solver running on multi-core processors with 16 floating-point cores. To employ the shared-memory cores, the MLFMA solvers are parallelized on 16 open multi-processing (OpenMP) threads, where each thread employs a single core. The forward solutions are obtained with BiCGStab iterative scheme.

TABLE I
Reconstruction Times with a Fixed Number of Nodes

Num. Illum.	**64**	**128**	**256**	**512**	**1,024**
Illum. per Node	**1**	**2**	**4**	**8**	**16**
Number of Nodes	64	64	64	64	64
Number of Cores	1,024	1,024	1,024	1,024	1,024
# MLFMA Solutions	9,600	19,200	38,400	76,800	153,600
Reconstruction (m)	6.25	12.42	25.06	49.93	99.12

978-1-5386-1732-8/17/$31.00 ©2017 IEEE

Table I shows the parallel reconstruction times with 64 nodes, whose results are shown in Fig. 1. The first row shows the number of illuminations and the second row shows the number of illuminations handled per node. For example, since there are only 64 nodes, 1,024 illuminations are serialized by handling 16 illuminations with a single node. The third and fourth rows show the number of nodes and cores, respectively. The fifth row shows the total number of MLFMA solutions in the inverse solution, where there are three solutions per illumination in each DBIM iteration. Each MLFMA solution involves approximately one million unknowns. The last row shows the wall times in minutes, from the beginning until the end of the executions (including setup and disk operations).

Table I shows that the number of MLFMA solutions increases linearly with the number of illuminations. As a result, with a fixed number of nodes, the reconstruction times increase proportional to the number of illuminations. Table II shows reconstruction times with full distribution of illuminations among nodes such that each node handles the forward problems of a single illumination. For example, the solutions with 64 and 1,024 illuminations employ 64 and 1,024 nodes, yielding 1,024 and 6,394 cores in total, respectively. With this parallelization strategy, the number of nodes increases with the number of illuminations as well and, due to the good scaling of parallelization, the reconstruction times do not increase significantly. If we consider 64-node solutions in Table I as a baseline, Table II shows that the reconstruction with 1,024 illuminations can be efficiently partitioned among 1,024 nodes with 93% efficiency, yielding 14.88 times speedup.

TABLE II
SCALING WITH SINGLE NODE PER MLFMA SOLVER

Num. Illum.	**64**	**128**	**256**	**512**	**1,024**
Illum. per Node	**1**	**1**	**1**	**1**	**1**
Number of Nodes	64	128	256	512	1,024
Number of Cores	1,024	2,048	4,096	8,192	16,384
Reconstruction (m)	6.25	6.39	6.58	6.63	6.66
Speedup	1.00	1.94	3.80	7.53	14.88
Efficiency (%)	100	97	95	94	93

IV. PARALLELIZING ILLUMINATIONS AND MLFMA

The reconstructions with small number of illuminations under-utilize the computational resources, i.e., there is not enough (or not necessary) number of illuminations to spread the reconstructions among large number of nodes. As a remedy, we parallelize the forward solvers over distributed-memory MPI processes simultaneously with illuminations. Since MLFMA is implemented for solving two-dimensional scattering problems from inhomogeneous dielectric profiles, a corresponding volume integral equation is solved, and hence, the MLFMA tree structure is well-balanced, i.e., there is no empty cluster. Therefore a simple partitioning of the MLFMA tree structure provides a good load balancing up to 16 processes. The BiCGStab solvers also run in parallel on the processes for reducing MPI communications during the iterative forward solutions.

TABLE III
SCALING WITH SIMULTANEOUS PARALLELIZATION

Num. Illum.	**64**	**128**	**256**	**512**	**1,024**
Illum. per Node	**1**	**1**	**1**	**1**	**1**
Nodes per MLFMA	**16**	**8**	**4**	**2**	**1**
Number of Nodes	1,024	1,024	1,024	1,024	1,024
Number of Cores	16,384	16,384	16,384	16,384	16,384
Reconstruction (m)	0.58	0.97	1.78	3.41	6.66
Speedup	10.78	12.80	14.07	14.48	14.88
Efficiency (%)	67	80	88	91	93

Table III shows the reconstructions with simultaneous parallelization of illuminations and MLFMA among 1,024 computing nodes. For example, in order to use all the nodes with 64 illuminations, MLFMA employs 16 nodes, where each MLFMA solver handles scattering problems corresponding to a single illumination, yielding 1,024 nodes. As shown in the table, the reconstruction with 64 illuminations is 10.76 times faster with 1,024 nodes with respect to 64 nodes, yielding parallelization efficiency of 67%. This shows that parallelizing MLFMA solutions is less efficient than parallelizing illuminations. This is mainly due to the intense MPI communications performed in each MLFMA multiplication in contrast to pleasingly parallel nature of illuminations. Similarly, the reconstructions with 128, 256, and 512 illuminations are 12.80, 14.07, and 14.88 times faster, respectively, on 1,024 nodes rather than on 64 nodes. Since the largest reconstruction can use all the 1,024 nodes with parallelizing illuminations only, a simultaneous parallelization is not invoked.

V. CONCLUSIONS

A set of DBIM solutions demonstrating their scalable parallelization is presented. A simple parallelization partitions the object illuminations among computing nodes, where each node handles an equal number of illuminations. This provides good efficiency but poor granularity. A simultaneous parallelization employs more number of nodes by parallelizing forward solutions among distributed-memory processes as well, which provides better granularity but worse efficiency. The results demonstrate that the parallel DBIM implementation scales out well up to 1,024 nodes with 67% to 93% efficiency, depending on the number of object illuminations, i.e., problem size.

ACKNOWLEDGMENT

This work was supported by NVIDIA GPU Center of Excellence and NCSA Petascale Application Improvement Discovery Program grants, NSF grant EECS-1609195, and UIUC COE Strategic Research Initiative Grant.

REFERENCES

[1] A. J. Hesford and W. C. Chew, "A frequency-domain formulation of the Fréchet derivative to exploit the inherent parallelism of the distorted Born iterative method," *Wave Random Complex*, vol. 16, no. 4, pp. 495–508, Feb. 2007.

[2] M. Hidayetoğlu, C. Pearson, W. C. Chew, L. Gürel, and W.-M. Hwu, "Large inverse-scattering solutions with DBIM on GPU-enabled supercomputers," *Applied and Computational Electromagnetics Symp. (ACES 2017)*, Florence, Italy, Mar. 2017.

[3] L. A. Shepp and B. F. Logan, "The Fourier reconstruction of a head section," *IEEE Trans. Nucl. Sci.*, vol. 21, no. 3, pp. 21–43, June 1974.