

# Scaling Analysis of a Hierarchical Parallelization of Large Inverse Multiple-Scattering Solutions

Mert Hidayetoğlu, Carl Pearson, Izzat El Hajj, Weng Cho Chew, Levent Gürel, and Wen-Mei Hwu  
Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA  
[hidayet2,pearson,elhajj2,w-chew,lgurel,w-hwu]@illinois.edu

## ABSTRACT


We propose a hierarchical parallelization strategy to improve the scalability of inverse multiple-scattering solutions. The inverse solver parallelizes the independent forward solutions corresponding to different illuminations. For further scaling out on large numbers of computing nodes, each forward solver parallelizes the dense and large matrix-vector multiplications accelerated by the multilevel fast multipole algorithm. Numerical results on up to 1,024 of CPU nodes show that the former and latter parallelizations have 95% and 73% strong-scaling efficiencies, respectively.

## ACM Reference format:

Mert Hidayetoğlu, Carl Pearson, Izzat El Hajj, Weng Cho Chew, Levent Gürel, and Wen-Mei Hwu. 2017. Scaling Analysis of a Hierarchical Parallelization of Large Inverse Multiple-Scattering Solutions. In *Proceedings of Supercomputing, Denver, CO, November 2017 (SC17)*, 2 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

## 1 INTRODUCTION

Real-life inverse-scattering problems in imaging, remote sensing, nondestructive testing, and geophysical exploration often suffer from multiple-scattering phenomena, because the problem becomes mathematically nonlinear when higher-order scattering effects are involved. The distorted-Born iterative method (DBIM) follows a gradient-search approach with successive linearizations, where the derivatives are found semi-analytically with a series of forward-scattering problems involving the guessed object. Direct solution methods, such as LU decomposition and singular value decomposition, can have prohibitive computational complexity in terms of time and space since dense and large matrices are involved. A research aim is to marry inverse scattering with fast algorithms and parallel computing. The forward problems in DBIM are accelerated with modern low-complexity algorithms like the multilevel fast multipole algorithm (MLFMA) [2], and also a non-dependency among the forward solutions is exploited [1]. We recently took one step forward to tackle the world's largest and most challenging inverse-scattering problems with supercomputing. In our recent work, we parallelized the forward problems corresponding to different illuminations among up to 256 CPU and GPU nodes of NCSA Blue Waters supercomputer [3]. However, parallelizing only the illuminations does not give enough granularity to scale the solutions to larger numbers of nodes. Therefore, we proposed the parallelization of illuminations and MLFMA simultaneously [4]. This extended abstract provides a complete set of strong-scaling results of the proposed parallelization scheme.



**Figure 1: The proposed inverse-solution scheme parallelizes illuminations in one dimension and the corresponding forward solutions in another dimension of a computing-node array. Each dot represents a multi-core CPU and a pair represents a CPU node.**

## 2 IMPLEMENTATION

A nonlinear conjugate-gradient iterative scheme is implemented for inverse solutions. In each iteration, the negative-gradient direction and an approximate step size in the  $N$ -dimensional solution space are found semi-analytically through the distorted-Born approximation. As a result, three forward problems per illumination are solved in each inverse iteration. The DBIM solver tackles two-dimensional scalar Helmholtz-type scattering problems.

The forward problems corresponding to different illuminations are independent of each other, and therefore, can be solved in parallel. Figure 1 depicts the proposed configuration, where the illuminations are parallelized among eight rows of computing nodes. Each row employs an MLFMA solver, which is parallelized among 16 computing nodes. Each node employs a single MPI process and each process employs a set of shared-memory OpenMP threads to utilize the corresponding multi-core CPUs. A total of 128 nodes with 2,048 cores (each Blue Waters node has two CPUs with 8 floating-point cores) are employed in the figure. To organize the MPI communications among processes, local communicators are defined in each row and column of nodes. That is, each MPI process communicates with others through its row, column, and global communicators. The column and row communicators of a node with the global rank of 89 are highlighted in Figure 1.

For distributed-memory parallelization of the MLFMA solvers, the corresponding multilevel tree structures are partitioned among the row of processes. The BiCGS method is implemented for iterative forward solutions. In each iteration, the MLFMA operations are performed among the tree branches in parallel. When two branches belong to the same MPI process, the corresponding operations can be performed through the shared-memory OpenMP threads, however, the MPI communications are inevitable when two branches belong to different MPI processes. As a result, the MLFMA parallelization efficiency is spoiled.




Figure 2: Reconstruction of a large Shepp-Logan phantom with a maximum contrast of 0.05 involving 1,024 transmitters and receivers. After about 50 iterations, most of the image details can be perceived very well. This, indeed, is the largest DBIM solution achieved in near-real time.

### 3 SCALING RESULTS

To demonstrate the scaling of the parallel inverse solver, a large Shepp-Logan numerical phantom is reconstructed in a solution domain with  $102.4\lambda$  side-length, where  $\lambda$  is the wavelength of the illuminating plane waves. The discretization with  $0.1\lambda$  yields 1,048,576 pixels whose wave properties are unknown. For the reconstruction shown in Figure 2, 1,024 transmitters and receivers are arranged around the object, 50 DBIM iterations are performed, and 38,400 forward-scattering problems are solved, each involving  $1,048,576 \times 1,048,576$  dense interaction matrices to be inverted. A total of approximately two million MLFMA multiplications are performed for the BiCGS solutions of forward problems.

Two strong scaling experiments are carried out since we parallelize the inverse solver over two dimensions, i.e., illuminations and MLFMA, one in each dimension. The multi-core CPU nodes, with 16 cores each, are used in the Blue Waters supercomputer. In the first experiment, as we increase the number of nodes, we distribute 1,024 illuminations across the nodes in a load-balanced manner, while holding constant the number of nodes that solves the forward problems. The scaling results are shown in Figure 3, where a run on 64 nodes with 1,024 cores is taken as a baseline. The high parallel efficiency reflects the independent handling of different illuminations. When scaling from 1,024 to 16,384 cores along this axis of parallelism, 94.7% efficiency is achieved.

In the second experiment, as we increase the number of nodes, we distribute the MLFMA operations across the additional nodes while maintaining 16 illuminations per MLFMA node-group. The scaling results are shown in Figure 4. Parallelization of MLFMA across multiple nodes incurs communication among the corresponding MPI processes, so the parallel efficiency in this problem decomposition is not as high as that of parallelizing illuminations only. Scaling from 1,024 to 16,384 cores yields a parallel efficiency of 72.9%.

### 4 CONCLUSIONS

A hierarchical parallelization of DBIM provides scalable multiple-scattering solutions on large numbers of CPU nodes. An inverse




Figure 3: Parallelizing illuminations scales almost perfectly (95% efficiency on 1,024 CPU nodes) due to the independent nature of the corresponding forward solutions.




Figure 4: Scaling efficiency of parallel MLFMA is hampered by intense MPI communications performed in each matrix-vector multiplication for the iterative forward solutions.

problem involving a large Shepp-Logan phantom is solved on up to 1,024 CPU nodes of the Blue Waters supercomputer in order to demonstrate the strong-scaling efficiency of the proposed parallelization scheme. The results show that parallelizing illuminations has almost perfect scaling efficiency of 95% because of the independent nature of forward-scattering solutions, however, parallelization of MLFMA has 73% efficiency due to MPI communications in MLFMA multiplications. Nevertheless, the proposed strategy improves granularity and allows spreading DBIM solutions on large numbers of nodes.

### ACKNOWLEDGMENTS

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 Initiatives grant.

### REFERENCES

- [1] Andrew J. Hesford and Weng Cho Chew. 2006. A frequency-domain formulation of the Fréchet derivative to exploit the inherent parallelism of the distorted Born iterative method. *Waves in Random and Complex Media* 16 (2006), 495–508.
- [2] Andrew J. Hesford and Weng Cho Chew. 2010. Fast inverse scattering solutions using the distorted Born iterative method and the multilevel fast multipole algorithm. *The Journal of the Acoustical Society of America* 128, 2 (2010), 679–690.
- [3] Mert Hidayetoğlu, Carl Pearson, Weng Cho Chew, Levent Gürel, and Wen-Mei Hwu. 2017. Large inverse-scattering solutions with DBIM on GPU-enabled supercomputers. In *Proceedings of ACES Symposium*. Florence, Italy.
- [4] Mert Hidayetoğlu, Carl Pearson, Levent Gürel, Wen mei Hwu, and Weng Cho Chew. 2017. Scalable parallel DBIM solutions inverse-scattering problems. In *Proceedings of CEM'17 Workshop*. Barcelona, Spain, 65–66.