Lawrence Berkeley National Laboratory

Title
Efficient Scalable Parallel Higher Order Direct MoM-SIE Method With Hierarchically Semiseparable Structures for 3-D Scattering

Permalink
https://escholarship.org/uc/item/85k6f0t7

Authors
Manić, Ana B
Smull, Aaron P
Rouet, François-Henry
et al.

Peer reviewed

Efficient Scalable Parallel Higher Order Direct MoM-SIE Method With Hierarchically Semiseparable Structures for 3-D Scattering

Ana B. Manić, *Student Member, IEEE*, Aaron P. Smull, *Student Member, IEEE*, François-Henry Rouet, Xiaoye Sherry Li, and Branislav M. Notaroš, *Fellow, IEEE*

Abstract—A novel fast scalable parallel algorithm is proposed for the solution of large 3-D scattering problems based on: 1) the double (geometrical and current-approximation) higher order (DHO) method of moments (MoM) in the surface integral equation (SIE) formulation and 2) a direct solver for dense linear systems utilizing hierarchically semiseparable (HSS) structures. Namely, an HSS matrix representation is used for compression, factorization, and solution of the system matrix. In addition, a rank-revealing QR decomposition for memory compression is used, with a stopping criterion in terms of the relative rank tolerance value. A method for geometrical preprocessing of the scatterers based on the cobblestone distance sorting technique is employed in order to enhance the HSS algorithm accuracy and parallelization. Numerical examples show how the accuracy of the DHO HSS-MoM-SIE method is easily controllable by using the relative tolerance for the matrix compression. Moreover, the examples demonstrate low memory consumption, as well as much faster simulation time, when compared to the direct LU decomposition. The method enables dramatically faster monostatic scattering computations than iterative solvers and reduced number of unknowns when compared to low-order discretizations. Finally, great scalability of the algorithm is demonstrated on more than one thousand processes.

Index Terms—Curved parametric elements, direct solvers, fast solvers, hierarchically semiseparable (HSS) structures, higher order (HO) modeling, low-rank matrix approximation, method of moments (MoM), multilevel matrix compression, numerical algorithms, parallelization, polynomial basis functions, scalability, scattering, surface integral equation (SIE).

I. INTRODUCTION

RECENT trends in computational electromagnetics (CEM), in applications that involve calculating,

storing, and solving large and dense matrices, include applying fast, parallel (direct or iterative) solvers for the system of equations in conjunction with compressed storage of large matrices and their parts. Two general approaches emerge among fast methods attempting to reduce numerical and storage complexity: 1) the fast multipole method (FMM) [1] and 2) $\mathcal{H}$-matrices [2]–[4]. The idea behind both of them is to approximate the integral kernel by a degenerate kernel using so-called functional skeletons. In the case of multipole methods, these functions have to be known explicitly for each kernel, which means that the method and its behavior depend heavily on the physics behind the exact problem to be solved. On the other hand, in the case of algebraic methods, such as $\mathcal{H}$-matrices, matrix blocks are approximated by low-rank matrices.

$\mathcal{H}$-matrix algorithms were first introduced in [2]–[4], with their $\mathcal{H}^2$-matrix version being introduced in [5], and have been used efficiently with fast LU-based direct solvers or as preconditioners to fast iterative solvers. The $\mathcal{H}$-matrix methods are kernel independent so they are suitable for application to any type of integral equation (IE)-based formulation. In the CEM community, several applications of $\mathcal{H}$-type direct solvers to tackle surface IE (SIE) problems are presented in [6]–[10].

Matrix compression solvers rely heavily on a type of the low-rank matrix approximation method such as singular value decomposition [8], [11], rank-revealing QR (RRQR) decomposition [12], [13], or adaptive cross approximation (ACA), where ACA can be considered as rank-revealing LU (RRLU). ACA is well known and established method for fast matrix computation in CEM, introduced first to solve low-frequency (quasistatic) IE problems [14], and then combined with different matrix compression methods to solve systems of equations arising in high-frequency SIE methods [15]–[18].

In addition, semiseparable matrices, the ones that can easily be compressed and accurately approximated by their low-rank counterpart, and their application to Green's function integral kernels are discussed in [19]. In order to combine the beneficial features of semiseparable matrices and $\mathcal{H}$-matrix representation, hierarchically semiseparable (HSS) matrices were most recently introduced in [20], which exhibit great parallel performance for use in direct solvers when compared to their standard $\mathcal{H}$-matrix counterparts. The solution of problems in 2-D SIE method of moments (MoM) using an HSS compression algorithm is discussed in [21], where the

authors comment on the possible extension of their work to the 3-D case.

On the other hand, it is well known that by using higher order (HO) basis functions for current/field modeling in CEM, significant reductions in the number of unknowns, as well as faster system matrix computation/solution, can be achieved [22] when compared to the traditional low order (LO) modeling [23]. Tightly coupled with using HO basis functions is HO geometrical modeling [22], [24] and together they lay a foundation for double (geometrical and current-approximation) higher order (DHO) modeling. The DHO approach has been effectively used in both direct and iterative MoM-SIE solvers [25]–[29].

Besides developing fast algorithms to solve MoM-SIEs, the CEM researchers have intensely investigated parallelization of the fast algorithms coupled with direct solvers in order to speed up the simulations of electrically large electromagnetic structures. DHO MoM-SIE system matrix filling followed by a computationally expensive LU decomposition on a full matrix was implemented into parallel out-of-core hybrid GPU/CPU algorithm [28], while the performance of a similar method using HO basis functions was investigated on more than 4000 CPU cores on a distributed memory system [29]. Similarly, a parallel $\mathcal{H}$-LU direct solver using hybrid MPI-OpenMP that builds on the ability to combine both shared and distributed memory programming was used to analyze 3-D scattering problems with nearly 4 million unknowns [30]. Further, a parallel hierarchical ACA algorithm demonstrating an acceleration factor larger than 200 was presented in [31].

This paper proposes a novel fast scalable HO parallel algorithm for large and complex scattering, radiation, and propagation problems in CEM based on the DHO MoM-SIE modeling in the frequency domain (FD) [22], [24], [32], [33] in conjunction with a direct solver for dense linear systems using HSS matrices [34], namely, the DHO HSS-MoM-SIE method. We are developing asymptotically fast HO direct algorithms for MoM-SIE solutions which, in a nutshell, are an algebraic generalization to FMMs. In addition to being fast, they offer a promise of being memory- and communication-efficient and amenable to extreme-scale parallel computing. The main advantage of the HSS algorithm is in the linear-complexity ULV-type factorizations (whereas the conventional LU decomposition has cubic complexity). The HSS algorithm is shown to have excellent parallel scalability. Our work uses the recently developed new, state-of-the-art, algorithms for solving dense and sparse linear systems of equations based on the HSS algorithm [34]. The new HSS algorithm has been demonstrated to have a dramatic advantage in terms of time and space complexity (e.g., ∼70 times less memory for seismic imaging examples with matrix size $250\,000 \times 250\,000$) over the LU factorization algorithm, and to be extremely scalable. In addition, this paper employs an RRQR decomposition for the matrix (memory) compression. Its adaptive nature comes from the ability to use a stopping criterion, i.e., a relative rank tolerance value, which allows for the method to store only the low-rank approximation of the original matrix that satisfies a predefined accuracy. In order to take full advantage of the HSS algorithm, a method for geometrical preprocessing of scatterers

based on the cobblestone distance sorting technique [15] is utilized. MoM unknowns are divided among mesh groups so that unknowns having spatial locality (belonging to the same mesh group) also exhibit locality in the matrix system of equations. In IE methods, these spatially close unknowns have much more numerically significant interactions, whereas interactions between mesh groups are of smaller numerical rank [15]. Ultimately, this spatial localization of the numerical interactions is what is exploited to have such effective coupling with the HSS algorithm. The basic theory and preliminary results of the DHO HSS-MoM-SIE analysis are presented in a summary form in [35] and [36].

HSS construction is implemented in a multilevel fashion as described in [34], and essentially, its multilevel compression can be considered comparable to the one used in the multilevel ACA (MLACA) algorithm [17]. Furthermore, so-called multilevel "butterfly" algorithms [37]–[38], as well as the fast solver presented in [39], have a similar basis to the multilevel compression coupled with low-rank matrix representation. One of the main advantages the HSS algorithm has over methods like the MLACA algorithm is its strong parallel scalability. Furthermore, the possibility to adjust the numerical rank tolerance of the matrix or numerical rank tolerance of any of its subblocks gives more control over the method and its computational accuracy. Multilevel compressions utilize matrix rank reduction to the maximum extent and reduce the overall number of degrees of freedom (DoFs) describing the method, as well as the storage requirements, to their minimum. HSS's purely algebraic nature compresses the matrix of interest independently of the problem's geometry or electrical size. The direct ULV-type factorization also allows for a fast solution in the case of multiple excitation vectors (i.e., multiple right hand sides, a matrix of excitation vectors), whereas iterative solvers often require independent solving for each excitation. When coupled with the DHO approach to significantly reduce the number of unknowns in MoM-SIEs, the result is a fast parallel solver with excellent algorithmic complexity, strong parallel scalability, and controllable accuracy.

Specifically in terms of the comparison with the method in [6], as a major representative of $\mathcal{H}$-matrix applications in CEM, the efficiency of the $\mathcal{H}^2$-matrix solver [6] stems from the fact that it compresses only the parts of the system matrix that correspond to the interactions of the mesh groups located far from each other. The HSS algorithm inherits this behavior. Rather than calculating admissible blocks, the HSS algorithm simplifies the $\mathcal{H}^2$-matrix implementation by considering a single predefined block per tree node to be admissible. The $\mathcal{H}$-matrix type of solver has not been shown to exhibit scalable behavior when parallelized. Further, it is more suitable for use with iterative solvers, because finding the inverse matrix is computationally more demanding than in the case of the HSS algorithm.

Similarly, when specifically compared with the work in [40], as a major representative of the multiscale compressed block decomposition methods and the application of direct solvers in CEM, the matrix compression done in our HSS-based method is obtained by fully utilizing HSS structures, so the compressed system matrix has rather different form, storage

and compression layout. In addition, work presented in [40] utilizes an inverse of a system matrix to find an appropriate solution, while in this present work an *ULV* factorization is obtained. Also, this present method is parallel, while the method in [40] is implemented in a sequential fashion.

We find from numerical experiments that the DHO MoM-SIE HSS method, which inherits the aforementioned benefits of HSS matrix representation, is able to compress the involved system matrix while maintaining accuracy of results and results in smaller computational and memory requirements than either DHO modeling alone, or an HSS algorithm-based MoM in an LO modeling paradigm.

This paper is organized as follows. Section II gives an overview of the MoM-SIE methodology and the associated discretization using DHO modeling. In addition, geometrical preprocessing used to group surface quadrilaterals into mesh groups that achieves spatial-data locality in the system matrix is described. In Section III, the HSS algorithm, including matrix compression and the corresponding ULV factorization followed by the solution of the compressed matrix is discussed. The parallelization and communication between the processes in matrix filling and the HSS algorithm is outlined in Section IV. Section V provides numerical results and discussion, followed by the conclusions in Section VI.

II. DHO MoM-SIE Modeling of Metallic Scatterers

One of the most general and best established approaches to solving scattering CEM problems is the one based on the MoM in the SIE formulation and the FD [24]. Inherently, the MoM results in dense linear systems, so the compression and solver are applied to allow for fast and memory efficient execution.

A. Surface Integral Equation Formulation

The MoM-SIE methodology is applicable to the analysis of metallic and dielectric structures, where both electric and magnetic surface currents are introduced over boundary surfaces between homogeneous parts of the structure, and SIEs based on boundary conditions for both electric and magnetic field intensity vectors are solved with current densities as unknowns. This paper focuses on metallic structures only. However, extending this paper to include problems involving dielectrics is straightforward based on [22].

If a structure made of a perfect electric conductor (PEC) is excited by a time-harmonic electromagnetic field of electric field intensity $\mathbf{E}^{\text{inc}}$ at the angular frequency ω, then the scattered field $\mathbf{E}^{\text{scat}}$ can be expressed in terms of the surface electric currents of density $\mathbf{J}_s$ using the boundary condition for the tangential fields at the surface S of the structure as follows:

$$(\mathbf{E}^{\text{scat}}(\mathbf{J}_s) + \mathbf{E}^{\text{inc}})_{\text{tang}} = 0, \quad \mathbf{E}^{\text{scat}}(\mathbf{J}_s) = -j\omega\mathbf{A} - \nabla\Phi \quad (1)$$

$$\mathbf{A} = \mu \int_S \mathbf{J}_s g \, dS, \quad \Phi = \frac{j}{\omega\varepsilon} \int_S \nabla_s \cdot \mathbf{J}_s g \, dS \quad (2)$$

where $\mathbf{A}$ and Φ are the magnetic vector and electric scalar potentials, respectively, $g = e^{-j\omega\sqrt{\varepsilon\mu}R}/4\pi R$, is the Green's


Fig. 1. (a) Generalized curved parametric quadrilateral patch for DHO MoM-SIE modeling [22]. (b) Sketch of the orthogonality factor for the first several maximally orthogonalized hierarchical basis functions of the eighth order [41]. The darkness of each square represents the magnitude of the inner products between pairs of basis functions. The set of basis functions is almost orthonormal.

function for the unbounded homogeneous medium with parameters ε and μ, and R is the distance of the field point from the source point. Hence, (1) and (2) constitute an electric field IE for $\mathbf{J}_s$ as unknown quantity, which is discretized using the MoM.

B. Double Higher Order Modeling

The DHO modeling consists of meshing the geometry of the electromagnetic structure using DHO surface elements, which means that both geometry as well as the unknown variable (surface current) are discretized using HO functions. In specific, surface of the structure is modeled using generalized curved quadrilaterals of arbitrary geometrical orders K_u and K_v, shown in Fig. 1(a), and the current density $\mathbf{J}_s$ over quadrilaterals is approximated by means of hierarchical vector basis functions of arbitrarily high current-expansion orders N_u and N_v [22]

$$\mathbf{r}(u, v) = \sum_{k=0}^{K_u} \sum_{l=0}^{K_v} \mathbf{r}_{kl} L_k^{K_u}(u) L_l^{K_v}(v), \quad -1 \le u, \, v \le 1 \quad (3)$$

$$\mathbf{J}_s = \sum_{i=0}^{N_u} \sum_{j=0}^{N_v-1} \alpha_{ij}^{(u)} P_{ij}^{(u)}(u, v) \frac{\mathbf{a}_u}{\Im}$$

$$+ \sum_{i=0}^{N_u-1} \sum_{j=0}^{N_v} \alpha_{ij}^{(v)} P_{ij}^{(v)}(u, v) \frac{\mathbf{a}_v}{\Im} \quad (4)$$

arranged in a maximally orthogonalized fashion [41], [42] as illustrated in Fig. 1(b). Here, L represent Lagrange interpolation polynomials, $\mathbf{r}_{kl}$ are position vectors of interpolation nodes, P are divergence-conforming polynomial bases, $\Im = |\mathbf{a}_u \times \mathbf{a}_v|$ is the Jacobian of the covariant transformation, and $\mathbf{a}_u = \partial\mathbf{r}/\partial u$ and $\mathbf{a}_v = \partial\mathbf{r}/\partial v$ are unitary vectors along the parametric coordinates. The unknown current-distribution coefficients $\{\alpha\}$ in (4) are determined by solving the SIE in (1), employing Galerkin method. The DHO (geometrical and current) modeling enables the use of large curved patches, which can greatly reduce the number of unknowns for a given problem and enhance the accuracy and efficiency of the computation.

Fig. 2.   Illustration of geometrical preprocessing of objects based on the cobblestone distance sorting technique [15].

Fig. 3.   Four-level postordered HSS tree, along with the associated index sets at each level. Here, the full index set $\mathcal{I} = \{0, 1, \ldots, 7\}$.

C. Geometrical Preprocessing Based on Cobblestone Distance Sorting Technique

The geometrical grouping of the quadrilaterals in the object mesh based on their spatial locality, as shown in Fig. 2, is done by applying the distance sorting technique [15] as outlined as follows.

To divide a number of unknowns into N^g groups we first calculate how many unknowns N^k should be in each group so as to have an even division. A box is created surrounding all remaining mesh patches. We calculate the projection of all boxed patches' centers onto the vector forming the diagonal of the box. The patch with the smallest projection is defined as the zero point for our current mesh group. All unknowns are then sorted by their distance from the zero point, and the first N^k from this sort are assigned to that particular mesh group. These patches are then removed from the mesh and the entire process is repeated until there are no remaining unknowns to be sorted. The groupings generated from this process are illustrated in Fig. 2 for two different geometric shapes.

The distance sorting technique is adapted to more appropriately fit a hierarchical method like the HSS algorithm. Instead of choosing N^k subgroups, the cobblestone method is used to sort the initial mesh into just two subgroups. Afterwards, each of these mesh groups is treated as an independent mesh, and individually sorted into two subgroups among itself. This process is applied recursively in a binary fashion, until the number of mesh groups is equal to the number of leafs in the desired HSS tree (or the number of processes), as described in Section III. If there are multiple processes per leaf, each mesh group corresponding to a leaf is then sorted appropriately by itself among the remaining processes. This hierarchical division of the mesh matches exactly with the desired partition defined by the HSS tree (Section III).

This grouping technique ensures data locality in the dense system matrix, which is greatly beneficial for achieving the properties needed for the HSS algorithm. Specifically, the cobblestone distance sorting technique divides the mesh into N^g mesh groups. Each matrix subblock determined by the coordinates (i, j) stores the interactions between MoM unknowns belonging to the ith and jth mesh groups. The mesh size of each mesh group (i.e., the number of unknowns) is predetermined by the number of processes, geometry, and other simulation specifics and is given as an input to the geometrical preprocessor. The outcome of the preprocessor

is that MoM unknowns belonging to the same mesh group, along with having their spatial locality, exhibit data locality in the system matrix: self-interaction blocks are on the diagonal, while near interactions tend to be closer to the diagonal and far interactions are further away from the diagonal in the system matrix. It is well known that the numerical rank of the matrix block describing the interactions between two groups decreases with an increase in distance between the groups [15], [43]. This matrix property plays a significant role in matrix compression during the HSS construction step.

III. HSS Theory

This section contains an overview of HSS structures, their construction, factorization, and solution, defined and explained in more detail in [34].

A. HSS Structures

The HSS form of a general $N \times N$ matrix A relies on something called an HSS tree, denoted by $\mathcal{T}$, which defines a hierarchical partitioning of the set of indices of the matrix, $\mathcal{I} = \{1, 2, \ldots, N\}$. A binary cluster tree is a binary tree such that every node i has associated with it an index subset, $t_i \subset \mathcal{I}$. We denote the left and right child nodes of a particular node as $c1$ and $c2$, respectively. An HSS tree is a full binary cluster tree such that any nonleaf node i's children $c1$ and $c2$ have the property that $t_{c1} \cap t_{c2} = \emptyset$ and $t_{c1} \cup t_{c2} = t_i$, i.e., a node's children define an even finer partition of that node's index set. Consequently, $A_{|t_i \times t_j}$ denotes the submatrix of A formed by the columns with indices in t_i and the rows with indices in t_j.

We say that an HSS tree $\mathcal{T}$ is postordered if the nodes in the tree are enumerated in such a way that for every nonleaf node $i \in T$, its children $c1$ and $c2$ satisfy $c1 < c2 < i$. Following this ordering scheme, a full L-level postordered HSS tree will consist of $2k - 1$ (where $k = 2^{L-1}$) ordered nodes, in which the root node is always numbered as $2k - 1$. The index set associated with the root node is $t_{2k-1} = \mathcal{I}$, the full index set. As we traverse down the HSS tree starting from the root, the partitioning of the index set described by the nodes on each level becomes finer and finer. We also number the levels of the tree with the leaf level being 1, and the level number increasing as we get closer to the root. A visualization of a four-level HSS tree and corresponding index partitioning is given in Fig. 3.

Fig. 4. Illustration of a corresponding HSS form. Diagonal blocks of the matrix are stored in full, while off-diagonal blocks are stored as low-rank approximations.

With the structure of an HSS tree firmly in place, the HSS compressed representation $\hat{A}$ of a matrix A can be defined from the bottom up by defining a set of so-called generator matrices, D_i, U_i, R_i, B_i, and V_i, at the ith node of the HSS tree (we discuss HSS construction and corresponding computation of these generator matrices in Section III-B). At the lowest (leaf) level, a matrix D_i corresponds to the block of the original matrix on the diagonal at the finest partition, $A_{|t_i \times t_i}$. With the unknown ordering scheme described in Section II, these diagonal blocks consist primarily of the interactions between basis functions which are spatially close to each other in the MoM-SIE mesh. Because of the strong $1/R$ spatial dependence of the electrodynamic Green's function, blocks close to the main diagonal of the system matrix are numerically much more significant (higher rank) than the off diagonal blocks [15], [43]. This numerical structure of the system matrix is what allows us to exploit the compressibility of the system matrix during HSS construction. Thus, the HSS algorithm is perfectly suited for its application in many CEM problems.

At the higher (nonleaf) levels, a subset of the generator matrices of a node i are block-wise defined in terms of generators of i's children, $c1$ and $c2$, as so

$$D_i = \hat{A}_{|t_i \times t_i} = \begin{bmatrix} D_{c1} & U_{c1}B_{c1}V_{c2}^H \\ U_{c2}B_{c2}V_{c1}^H & D_{c2} \end{bmatrix} \quad (5)$$

$$U_i = \begin{bmatrix} U_{c1}R_{c1} \\ U_{c2}R_{c2} \end{bmatrix}, \quad V_i = \begin{bmatrix} V_{c1}W_{c1} \\ V_{c2}W_{c2} \end{bmatrix} \quad (6)$$

where the superscript "H" denotes the Hermitian transpose. The remaining generator matrices, R_i, B_i, and W_i are stored at each level. With this structure in place, at the root node $2k-1$, we have that $D_{2k-1} = \hat{A}$, the full HSS compressed representation of A. Fig. 4 depicts a block example of the 8×8 (four-level) HSS representation of a matrix defined on the corresponding HSS tree given in Fig. 3. Leaf level D matrices are calculated and stored in a fully dense manner, while other matrices will be calculated and saved in a compressed form obtained by the RRQR decomposition.

The maximum numerical rank r for a given rank tolerance τ of all compressed blocks is called the HSS rank of matrix A. We say that a matrix A has the low-rank property and can be efficiently compressed and solved using the HSS algorithm if r is small comparing to the matrix size.

B. HSS Construction

HSS construction is done in two stages: first a row compression is applied, followed by a column compression. We utilize RRQR decompositions [13] for matrix compression at all stages. The RRQR decomposition is a modified Graham–Schmidt algorithm designed to terminate after a certain rank tolerance is reached.

The RRQR algorithm is as follows. We may compress and approximate a general $M \times N$ matrix $A \approx QT$ with a rank tolerance τ by iterating through the columns (or the first M columns if $M < N$) of the matrix and doing the following at each iteration i. We will use the notation a_i to designate the ith column of our current matrix A.

1) Find the column of the current matrix A with the largest two-norm. We call this the jth column, a_j.
2) Swap a_i and a_j.
3) Set the ith diagonal element of T to be the two-norm of a_i, $t_{ii} = \|a_i\|_2$. If t_{ii}/t_{11} is below the selected tolerance τ, we terminate the algorithm, and we define the rank r of A to be i.
4) Calculate the ith column of Q as $q_i = a_i/t_{ii}$.
5) Calculate the remainder of the ith row of T as $t_i^H = q_i^H [a_{i+1}, a_{i+2}, \ldots, a_N]$.
6) Update the remaining columns of A as

$$[a_{i+1}, a_{i+2}, \ldots, a_N] = [a_{i+1}, a_{i+2}, \ldots, a_N] - q_i t_i^H.$$

The matrix A is now compressed and stored as a low-rank approximation given by the product of an $M \times r$ orthogonal matrix Q and an $r \times N$ upper trapezoidal matrix T.

Row compression is applied in an upward sweep along the HSS tree, beginning at the leaf level, where at each node i, a local row block of the global system matrix $A_{|t_i \times \mathcal{I}}$ is filled in full as explained in Section II. To this row block, an initial RRQR decomposition is applied to the portion which does not lie on the diagonal block, $A_{|t_i \times \mathcal{I}\backslash t_i}$ (the notation $\mathcal{I}\backslash t_i$ refers to the set of all indices which are in $\mathcal{I}$, but not in t_i). We approximate this block via an RRQR decomposition as

$$A_{|t_i \times (\mathcal{I}\backslash t_i)} \approx U_i A_{|\hat{t}_i \times (\mathcal{I}\backslash t_i)} \quad (7)$$

which defines the column block HSS generators U_i at the leaf level nodes. We use the notation $\hat{t}_i$ to express that the new row index set is no longer a part of the global index set, but rather a set of indices corresponding to the new row compressed form of that matrix block.

At every nonleaf level i, row compression is done by inheriting submatrices of i's children's row compressed blocks, $A_{|t_{c1} \times \mathcal{I}\backslash t_i}$ and $A_{|t_{c2} \times \mathcal{I}\backslash t_i}$. Again these correspond to the portions of the compressed matrices which do not lie on the diagonal block at the next level up the HSS tree. These two forwarded blocks are concatenated vertically and then

compressed again as

$$\begin{pmatrix} A_{|\hat{t}_{c1}\times\ (\mathcal{I}\backslash t_i)} \\ A_{|\hat{t}_{c2}\times\ (\mathcal{I}\backslash t_i)} \end{pmatrix} \approx \begin{pmatrix} R_{c1} \\ R_{c2} \end{pmatrix} A_{|\hat{t}_i\times\ (\mathcal{I}\backslash t_i)}. \tag{8}$$

The remaining portions of i's children's corresponding row compressed matrices are left as they were

$$A_{|t_{c1}\times t_{c2}} \approx U_{c1}A_{|\hat{t}_{c1}\times t_{c2}}, \quad A_{|t_{c2}\times t_{c1}} \approx U_{c2}A_{|\hat{t}_{c2}\times t_{c1}}. \tag{9}$$

Recursively applying this algorithm upward along the HSS tree until the root node is reached defines all of the necessary row generators U_i and R_i. It also defines all row-compressed matrices of the form $A_{|\hat{t}_i\times t_j}$ where i and j are sibling nodes (i.e., nodes which are both children of the same node).

Column compression is applied in a similar fashion via an upward sweep along the HSS tree. We define at every leaf node i a new row index set $\bar{t}_i$, which consists of all the rows in all the row-compressed matrices for which there exist entries in the columns indexed by t_i. Formally, this may be expressed as

$$\bar{t}_i = \bigcup_{l=1}^{L-1} \hat{t}_{\text{sib}(\text{ances}(i,l))} \tag{10}$$

where $sib(i)$ refers to the sibling node of i and $ances(i,l)$ refers to i's ancestor node at the lth level of the HSS tree [note that $ances(i,1) = i$].

Again, RRQR decompositions are performed hierarchically to blocks of the row-compressed matrices to generate the final, fully compressed matrices. The column compression of the local matrix at the leaf level is given by

$$A_{|\bar{t}_i\times t_i} \approx A_{|\bar{t}_i\times \tilde{t}_i}V_i^H. \tag{11}$$

Blocks of these leaf level matrices are again inherited up the tree, just as in the row compression step given by (9). For a nonleaf node i, we define the horizontal compression in terms of submatrices of i's children's compressed matrices. We concatenate blocks of the column compressed matrices from each child, and perform the last series of RRQR decompositions which define the generator matrices W_{c1} and W_{c2}

$$\begin{pmatrix} A_{|\bar{t}_i\times\tilde{t}_{c1}} & A_{|\bar{t}_i\times\tilde{t}_{c2}} \end{pmatrix} = A_{|\bar{t}_i\times\tilde{t}_i} \begin{pmatrix} W_{c1} \\ W_{c2} \end{pmatrix}^H. \tag{12}$$

The portions of the compressed matrices from the children which are not forwarded become the final generator matrices, B_i for the child nodes

$$B_{c1} = A_{|\hat{t}_{c1}\times\tilde{t}_{c2}}, \quad B_{c2} = A_{|\hat{t}_{c2}\times\tilde{t}_{c1}}. \tag{13}$$

Construction of the matrix A in its HSS form and all generators corresponding to the nonroot nodes are obtained by the completion of the column compression stage. Note that the compression of the right-hand side is done in the same fashion as the compression of the system matrix.

C. HSS ULV Factorization and Solution

Once the HSS matrix form has been constructed, and all relevant generator matrices have been calculated, an HSS *ULV* factorization is systematically applied to the HSS compressed form of the matrix A in (5) in order to find orthogonal (U, V)

and triangular (L) matrices. The ULV factorization is done beginning at the lowest nonleaf level node i by multiplying the local matrix D_i given in (5) by specially constructed orthogonal matrices Q and P, calculated from i's children's generators

$$\begin{bmatrix} Q_{c1}^H & 0 \\ 0 & Q_{c2}^H \end{bmatrix} \begin{bmatrix} D_{c1} & U_{c1}B_{c1}V_{c2}^H \\ U_{c2}B_{c2}V_{c1}^H & D_{c2} \end{bmatrix} \begin{bmatrix} P_{c1}^H & 0 \\ 0 & P_{c2}^H \end{bmatrix}. \tag{14}$$

The Q matrices are formed by standard Gram–Schmidt QL factorizations of the children's column block generator matrices, U_{c1} and U_{c2}, in order to introduce zeros on their off-diagonal row blocks

$$U_{c1} = Q_{c1}\begin{bmatrix} 0 \\ \tilde{U}_{c1} \end{bmatrix}, \quad U_{c2} = Q_{c2}\begin{bmatrix} 0 \\ \tilde{U}_{c2} \end{bmatrix}. \tag{15}$$

Now, $\tilde{U}_{c1}$ and $\tilde{U}_{c2}$ are lower square triangular matrices of size r_{c1} and r_{c2}, respectively, with r_{ck} being the rank of the matrix U_{ck}. The obtained Q matrices are used to further define $\hat{D}_{ck} = Q_{ck}^H D_{ck}$, $k = 1, 2$, which are partitioned as

$$\hat{D}_{ck} = \begin{bmatrix} \hat{D}_{ck;1,1} & \hat{D}_{ck;1,2} \\ \hat{D}_{ck;2,1} & \hat{D}_{ck;2,2} \end{bmatrix}, \quad k = 1, 2 \tag{16}$$

so that $\hat{D}_{ck;2,2}$ is a square matrix of size equal to r_{ck}. Finally, the P matrices are defined by the following LQ factorization:

$$(\hat{D}_{ck;1,1} \quad \hat{D}_{ck;1,2}) = (\tilde{D}_{ck;1,1} \quad 0)P_{ck}. \tag{17}$$

Following the steps given in (14)–(17) and applying the orthogonality properties of the involved P and Q matrices, (14) may be expressed in a new, more convenient form:

$$\begin{bmatrix} \begin{bmatrix} \tilde{D}_{c1;1,1} & 0 \\ \tilde{D}_{c1;2,1} & \tilde{D}_{c1;2,2} \end{bmatrix} & \begin{bmatrix} 0 \\ \tilde{U}_{c1}B_{c1}\begin{bmatrix} \tilde{V}_{c2;1}^H & \tilde{V}_{c2;2}^H \end{bmatrix} \end{bmatrix} \\ \begin{bmatrix} 0 \\ \tilde{U}_{c2}B_{c2}\begin{bmatrix} \tilde{V}_{c1;1}^H & \tilde{V}_{c1;2}^H \end{bmatrix} \end{bmatrix} & \begin{bmatrix} \tilde{D}_{c2;1,1} & 0 \\ \tilde{D}_{c2;2,1} & \tilde{D}_{c2;2,2} \end{bmatrix} \end{bmatrix}. \tag{18}$$

It is important to note that all nonzero off-diagonal blocks in (18) are of small dimensions relative to the full matrix. At this stage, it is easy to redefine the generators at the next level up node using only a small part of its children's factorized generators

$$D_i = \begin{bmatrix} \tilde{D}_{c1;2,2} & \tilde{U}_{c1}B_{c1}\tilde{V}_{c2;2}^H \\ \tilde{U}_{c2}B_{c2}\tilde{V}_{c1;2}^H & \tilde{D}_{c2;2,2} \end{bmatrix}$$

$$U_i = \begin{bmatrix} \tilde{U}_{c1}R_{c1} \\ \tilde{U}_{c2}R_{c2} \end{bmatrix}, \quad V_i = \begin{bmatrix} \tilde{V}_{c1;2}W_{c1} \\ \tilde{V}_{c2;2}W_{c2} \end{bmatrix}. \tag{19}$$

Note that, for example, a square matrix D_i is still saved in a compressed fashion as before, but its new dimension is only $r_{c1} + r_{c2}$, while immediately after the HSS construction stage, it was the size of $t_{c1} \cup t_{c2}$.

The algorithm described in (14)–(19) is then performed in an upward sweep along the HSS tree. When the root node is reached, a dense LU factorization with partial pivoting is performed on a square matrix of dimensions far smaller than those of the starting matrix (19) along with the corresponding HSS compressed form of the excitation matrix, defines the matrix system at the ith node—these unknowns do not correspond to the final solution, but a partial factorization. Once this
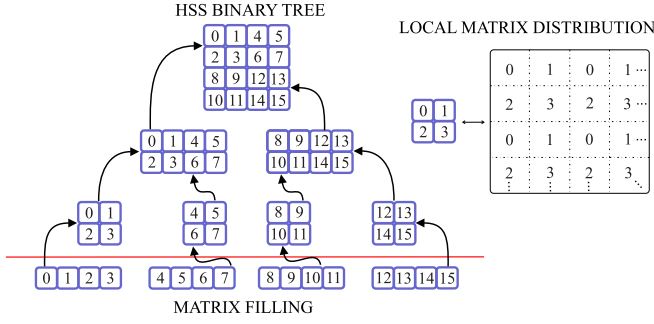
Fig. 5. Illustration of a process context change throughout matrix filling and HSS construction/solution on a level 3 HSS tree using 16 processes. Matrix filling is done by partitioning the matrix into row chunks, while all matrix operations during HSS construction/solution are done utilizing 2-D block cyclic distributions. The process grid arrangements in the binary tree represent the block cyclic distribution patterns for the local matrices at each tree node during HSS construction and factorization.
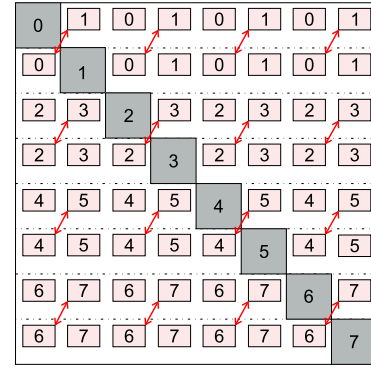

Fig. 6. Illustration of pairwise communication, visualized on eight processes (one per node of the HSS tree). Immediately after the first row compression stage, the processes have exchanged matrix blocks according to the red arrows. Note that when, for example, processes 6 and 7 merge into one context, the data is in the proper block cyclic distribution for the next level of row compression.

is solved on the parent node, it is straightforward to perform elimination on (18) and calculate the solutions corresponding to the children nodes. Solving the matrix is thus done in a downward sweep along the HSS tree—the final solution is obtained once the leaf level is reached.

IV. PARALLELIZATION STRATEGY

The parallelization strategy of the HSS-MoM-SIE method is adapted from [34] for construction, factorization, and solution, and from [44] for MoM-SIE matrix filling. The communication methods between the processes are designed to fit with the communication layer of the state-of-the-art dense linear algebra libraries ScaLAPACK [45] and basic linear algebra communication subprograms (BLACS) [46]. Both libraries are utilized extensively for the execution of algebraic computations on dense matrix blocks.

MoM-SIE matrix filling is done as described in Section II, where the parallel matrix filling builds on top of the geometrical processing as modified from [44]. If the total number of processes running the simulation is N_{procs}, then the number of mesh groups used to partition the mesh during preprocessing is $N^g = \sqrt{N_{procs}}$. The processes form a block partitioning of the matrix, where each matrix block may be described by its coordinates (p, q), corresponding to the interactions between the pth and qth mesh groups.

Matrix filling is done at the leaf level, where the 2-D process grid is divided into k (the number of leaf nodes) contexts, where each context defines a process subgrid of size $N^g / k \times N^g$, as illustrated in Fig. 5 for the case of 16 processes. Each of these subgrids consists of multiple processes operating blockwise on a row chunk of the full matrix. A set of processes operating on a matrix block together like this is known as a context. The context corresponding to the leaf-level node i is used to calculate the chunk of the MoM-SIE system matrix $A_{|t_i \times \mathcal{I}}$ as well as the corresponding right-hand side vector set.

For optimal communication and computation, dense matrix blocks shared by multiple processes on a context are stored in a distributed 2-D block cyclic fashion in accordance with the ScaLAPACK library routines [45]. Any overlap in computation (i.e., elements whose unknowns may be split between processes), may be avoided by effective use of the BLACS communication routines.

Matrix distribution and process grouping at each level of the HSS algorithm follows the basic HSS tree structure. The processes are utilized in such a way that they are combined and split into hierarchical groupings (known as contexts) as the algorithm traverses up and down the tree. At each node of the HSS tree, the process grid (and hence the block cyclic distribution) is redefined in terms of its child nodes' contexts. It inherits the contexts from both nodes, and their process grid is concatenated in order to form a new process grid. This is done hierarchically in a way that preserves as close to a square process grid as possible, as seen in Fig. 5.

During the construction and factorization stages, intercontext communication can largely be avoided, relying largely on only intracontext communication. An example of this communication process can be visualized as in Fig. 6. The communication in HSS construction is visualized here as performed with eight processes, each corresponding to one row block of the matrix at the row level; immediately following the leaf level row compression, blocks of the row compressed matrices $A_{|\hat{t}_i \times \ (\mathcal{I} \setminus t_i)}$ are pairwise exchanged between pairs of processes as follows: $0 \leftrightarrow 1$, $2 \leftrightarrow 3$, $4 \leftrightarrow 5$, and $6 \leftrightarrow 7$, as the pairs of processes are merged into new contexts ($\{0, 1\}$, $\{2, 3\}$, $\{4, 5\}$, and $\{6, 7\}$). This sets up the matrix in the proper block cyclic distribution for the next level of row compression. At every level during the row compression, to maintain the proper distribution during context merging, all that is required are similar pairwise exchanges, which may all be done in parallel.

More details behind the communication at each step in the parallel HSS solver can be found in [34]. An example of intracontext communication throughout the computation phase is the following. In the RRQR algorithm, when finding the norm of each column, the norm of each local column is determined and then the BLACS function is used for fast summing in a column-wise fashion for norm calculations on a 2-D process grid.

V. NUMERICAL RESULTS AND DISCUSSION

This section provides numerical results obtained by the DHO HSS-MoM-SIE analysis. All simulations were run on the

TABLE I
SIMULATION PARAMETERS FOR THE RESULTS GIVEN IN FIG. 7

RRQR rel. tolerance	Maximal rank	Time [s]	Storage [GB]	ζ [dB] ($\phi = 0$ cut)	ζ [dB] ($\theta = 90°$ cut)
$\tau = 1 \cdot 10^{-2}$	311	6.55	0.04	-11.823	-10.412
$\tau = 5 \cdot 10^{-3}$	351	6.90	0.04	-16.700	-17.189
$\tau = 2 \cdot 10^{-3}$	399	6.94	0.05	-21.279	-20.970
$\tau = 1 \cdot 10^{-3}$	436	7.08	0.06	-22.118	-21.936
$\tau = 5 \cdot 10^{-4}$	468	7.24	0.07	-22.002	-22.064
$\tau = 2 \cdot 10^{-4}$	510	7.83	0.08	-22.097	-22.079

TACC Stampede supercomputer, made accessible through The extreme science and engineering discovery environment [47]. Stampede has 6400 compute nodes where each node contains two Xeon Intel 8-Core 64-bit E5-processors (16 cores on each node). The core frequency is 2.7 GHz and supports eight floating-point operations per clock period with a peak performance of 21.6 GFLOPS/core or 346 GFLOPS/node. Each node contains 32 GB of memory (2 GB/core). Nodes are interconnected with Mellanox FDR InfiniBand technology in a two-level fat-tree topology [48]. The first example is a spherical scatterer for comparison with the exact Mie series solutions. This example is well suited for the demonstration of the effectiveness of using HO modeling compared to LO modeling, as well as a strong parallel scalability. The second example is a NASA almond scatterer to show that this method performs well for electrically large problems, as well as objects with complex shapes.

A. Example 1: Spherical Scatterer

As the first example of our application as well as a validation of the HSS-MoM-SIE method, we consider the scattering analysis of a PEC sphere of diameter $d = 4.666\lambda_0$, where λ_0 is the free-space wavelength. First, consider an HO mesh of the scatterer, of which the allowed size for a single patch is to be less than or equal to a wavelength. After the meshing procedure, the patch size s is approximately $0.9\lambda_0 \leq s \leq 0.95\lambda_0$ for all 96 geometrically second-order ($K_u = K_v = 2$) curvilinear quadrilateral patches. The adopted current approximation orders are $N_u = N_v = 4$, which results in a total of $N = 3072$ MoM-SIE unknowns. Fig. 7 shows the normalized bistatic radar cross section (RCS) σ_{3D}/λ_0^2 as a function of the scattered angle, in two characteristic plane cuts. The exciting plane wave is incident from the direction defined by $(\theta_{inc}, \phi_{inc}) = (90°, 0°)$ in spherical coordinates. For the set of results given in Fig. 7, the number of levels in the full postordered HSS tree is chosen to be 5, which results in 16 leaf nodes, while the number of processes used in the parallel simulation is 64. In the same figure, we observe the convergence of results to the analytical Mie solution with the decrease of τ, the user selected RRQR relative tolerance (used in the matrix compressions during the construction step).

In addition, Table I provides information on the average error, maximal rank, memory consumption, and total simulation time (including matrix filling, HSS construction, factorization, and solution times) for different simulations given
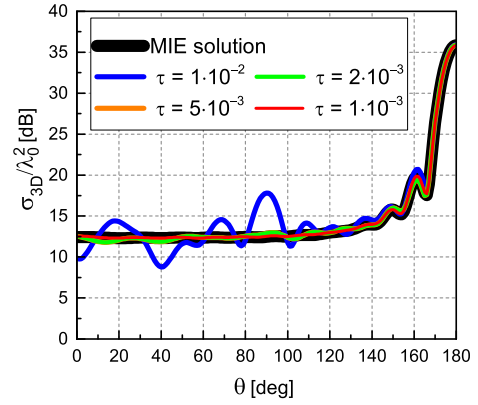

Fig. 7. Normalized bistatic RCS of a spherical PEC scatterer computed by the HSS-MoM-SIE method using preprocessed mesh given in the Fig. 2 and by the Mie's series. (a) $\phi = 0°$ cut. (b) $\theta = 90°$ cut.

in Fig. 7. The average relative error is obtained as the average of the absolute error between the normalized bistatic RCS calculated by the numerical method and by the analytical Mie's series, respectively. The averaging is done by taking into account the error in a number (N^{dir}) of directions describing the bistatic RCS plane

$$\zeta = 10\log\frac{\left(\sum_{i=0}^{N^{dir}} \left|\sigma_i^{num} - \sigma_i^{MIE}\right|\right)}{\left(\sum_{i=0}^{N^{dir}} \left|\sigma_i^{MIE}\right|\right)}. \quad (20)$$

Based on the convergence of different graphs given in Fig. 7, as well as the average errors given in Table I, it can be concluded that the accuracy of the results is easy to control by the relative tolerance used in the RRQR. In addition, by inspecting the results given in Fig. 7 and Table I, as well as the results of the wide range of performed simulations (different tree levels and RRQR relative tolerance) using the same HO model, scattering results for the PEC sphere that can readily be considered as accurate are obtained in simulations with the maximal rank $r \geq 490$. Maximal rank in the HSS algorithm corresponds to the maximal number of independent DoFs needed for an accurate numerical simulation [49]. The number of DoFs needed to accurately model a scatterer should

Fig. 8. Error of the normalized bistatic RCS of the PEC scatterer computed by the HSS-MoM-SIE method with respect to the exact Mie solution. (a) $\phi = 0°$ cut. (b) $\theta = 90°$ cut. We compare LO modeling with HO models. L denotes the number of levels in the HSS tree.

| Model type	RRQR rel. tolerance	Maximal rank	ζ [dB] ($\phi = 0$ cut)	ζ [dB] ($\theta = 90°$ cut)
High Order	$\tau = 1 \cdot 10^{-3}$	436	-22.118	-21.936
Low Order	$\tau = 1 \cdot 10^{-2}$	477	-6.901	-7.767
	$\tau = 5 \cdot 10^{-3}$	538	-14.152	-14.485
	$\tau = 2 \cdot 10^{-3}$	620	-23.619	-27.039

Fig. 9. Performance and scalability of the HSS-MoM-SIE method applied to the simulation of an LO PEC sphere model with the number of levels in the HSS tree equal to 5.

depend solely on the scatterer properties and not on the used discretization [37], [43], [49].

Further analysis that compares low- and high-order modeling results shows the advantage of HO modeling when capturing the real rank of the scattering problem. In particular, the adopted LO model of the same PEC scatterer consists of 7776 geometrically first-order ($K_u = K_v = 1$) quadrilaterals with the maximal size of $0.12\lambda_0$. The adopted current approximation orders are $N_u = N_v = 1$, resulting in a total of $N = 15552$ unknowns. After the initial discretization, when compared to the HO model, the number of unknowns in the LO model is more than 5 times larger. However, due to the physical properties of the scatterer, the maximal numerical rank of the compressed HSS matrix should be approximately the same for both models.

On the other hand, Fig. 8 shows the relative error metric $\sigma_i^{num}/\sigma_i^{MIE}$ in the RCS for three choices of LO simulations and one HO simulation, confirming that in order to achieve similar accuracy, an LO simulation needs about one-and-a-half times the numerical rank in the compressed system. This is confirmed by the information in Table II that contains the average error for the four simulations given in Fig. 8.

Further, note that the true LO modeling based on Rao-Wilton-Glisson (RWG) functions [1] defined over flat triangular patches uses from 500 to 600 unknowns per square wavelength [17], [18], [48], which leads to a truly LO model of a PEC sphere with around 35 000 to 40 000 unknowns, making the number of unknowns more than 10 times larger than in the adopted HO model.

Next, we test the scalability performance of the HSS-MoM-SIE method—in the same example. Because the HO model with only 96 quadrilaterals is not well fit for simulations on hundreds of processes, the scalability testing of the method is done on an LO model of the sphere. All simulations used in the test are set up for the tolerance $\tau = 5 \cdot 10^{-4}$ and level 5 full postordered HSS tree. Note that the number of leaves in the tree is 16, which, for the particular example, corresponds to the minimal number of processes that can be used in a parallel run of the HSS-MoM-SIE code. The runtime of the parallel code on 16 processes is thus adopted to be the baseline computational time used in the speed-up calculations. To measure scalability, we run the same model on 64, 256, and 1024 processes, and observe an excellent scalability performance in Fig. 9.

B. Example 2: NASA Almond

In the second example, the HSS-SIE-MoM code is applied to analyze scattering from a NASA almond [50], an established benchmarking structure for monostatic RCS computations.
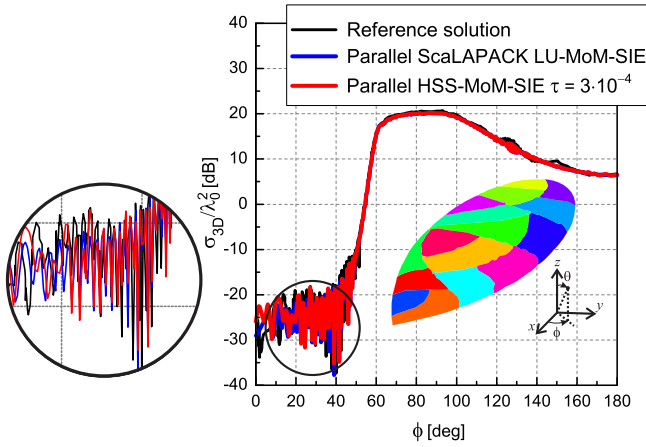
Fig. 10. Normalized monostatic RCS of a PEC NASA almond computed at 50 GHz by the full-storage direct ScaLAPACK LU-MoM-SIE and HSS-MoM-SIE methods. Note that the full LU factorization solution and HSS-MoM solutions (both from same geometrical model) show good agreement.

In specific, we consider a PEC almond of the maximal size $42\lambda_0$ at a frequency of 50 GHz, with the overall surface of the scatterer being equal to $1111\lambda_0^2$. The constructed HO model of the almond uses a total of $16\,384$ curvilinear quadrilateral elements with $K_u = K_v = 2$ and the current approximation in different directions on different patches ranging from $N_{u/v} = 1$ to $N_{u/v} = 3$ depending on the electrical dimensions of the quadrilateral element. The final number of unknowns (that would be even smaller if larger patches and even higher $N_{u/v}$ were used) is $149\,756$. A similar scatterer is analyzed in [40] using the model with around $450\,000$ unknowns. In addition, applying, for comparison, the true lower order quadrilateral modeling to the almond scatterer requires $524\,288$ unknowns defined over $262\,144$ patches with $K_u = K_v = 1$ and $N_{u/v} = 1$ on all the patches.

Fig. 10 shows the monostatic scattering computations of the DHO model of the PEC almond obtained by the HSS-MoM-SIE method and validated by the full-storage direct solver using ScaLAPACK LU decomposition [44] simulation of the same model, as well as against the LO modeling results [40]. The normalized monostatic RCS is calculated for 361 different directions, in the $z = 0$ plane, with the polarization of the incident electric field along the z-axis.

HSS construction for this example is done using a relative tolerance of $\tau = 3\cdot10^{-4}$ on a five-level HSS tree. The relative rank tolerance τ was selected heuristically in such a way as to satisfy the needs for the electrical size of the problem. In general, scatterers which are electrically large will require a smaller rank tolerance than electrically small scatterers to maintain the same level of solution accuracy, so the second example will require a smaller τ. The maximal rank of the compressed matrix in this example comes out to be 3926. The simulation is run in parallel on 256 processes, with matrix calculation and HSS construction times being 846 and 2596 s, respectively. Further, the HSS factorization time is 41.95 s, while the total backsubstitution time for all 361 excitation vectors is 5.36 s. ScaLAPACK LU decomposition time on

the same number of processes is 8565 s, which in terms of the performance can be compared to total time of the HSS construction and factorization: 2638 s. The compressed matrix storage in the DHO HSS-MoM-SIE simulation amounts to 20.78 GB, while the full matrix storage for the same model would require 180 GB of memory. In addition, the LO quadrilateral model described above would require 2.2 TB. Hence, we observe great advantages of the DHO modeling coupled with the direct solver and compression of the MoM-SIE matrix in the analysis of electrically large objects with multiple excitations (right-hand side values).

VI. CONCLUSION

This paper has proposed a novel fast scalable parallel algorithm and solver for large scattering problems based on (geometrical and current-approximation) DHO MoM in the SIE formulation and the FD in conjunction with a direct solver for dense linear systems with HSS structures, namely, with a HSS matrix representation for compression, factorization, and solution of the system matrix. In addition, an RRQR decomposition for memory compression has been used, with a stopping criterion in terms of the relative tolerance value, allowing for the method to store only the low-rank approximation of the original matrix that satisfies predefined accuracy. In order to enhance the HSS construction and parallelization, a method for geometrical preprocessing of the scatterers based on the cobblestone distance sorting technique has been employed, such that the MoM unknowns belonging to the same mesh group and thus having spatial locality also exhibit the data locality in the matrix system of equations.

Numerical examples have shown how the accuracy of the DHO HSS-MoM-SIE method is easily controllable by using the relative tolerance for the matrix compressions. Moreover, the examples have demonstrated low memory consumption, as well as much faster simulation times, when compared to direct LU decompositions. Finally, great scalability of the algorithm has been demonstrated on more than thousand processes.

Overall, the DHO HSS-MoM-SIE method and its future extensions and advancements are asymptotically faster direct algorithms for IE solutions that are memory and communication-efficient and amenable to extreme-scale parallel computing. They also are purely algebraic and kernel-independent and enable dramatically faster monostatic scattering and other multiple-excitation computations than iterative solvers. The combination of the DHO and HSS approaches, along with the distance sorting technique and parallel method, fully utilizes problem's properties to reduce the solution complexity (both computation time and storage) by minimizing the problem's DoFs based on both its physical and algebraic properties. The presented coupling and implementation of the two approaches is a natural way to advance the MoM-SIE modeling, since compressing an LO matrix may be considered less efficient. For example, based on the presented numerical results, it can be concluded that the prime problem size reduction is obtained by combining the DHO and HSS approaches together.

This paper has focused on metallic scatterers. Nonetheless, due to the truly algebraic nature of the method, its extension to any electromagnetic system solved by the MoM-SIE analysis, or any other IE method, is straightforward.

REFERENCES

[1] W. C. Chew, J.-M. Jin, E. Michielssen, and J. Song, *Fast and Efficient Algorithms in Computational Electromagnetics*. Boston, MA, USA: Artech House, 2001.

[2] W. Hackbusch, "A sparse matrix arithmetic based on H-matrices. Part I: Introduction to H-matrices," *Computing*, vol. 62, no. 2, pp. 89–108, 1999.

[3] W. Hackbusch and B. N. Khoromskij, "A sparse H-matrix arithmetic. Part II: Application to multi-dimensional problems," *Computing*, vol. 64, no. 1, pp. 21–47, Jan. 2000.

[4] W. Hackbusch, L. Grasedyck, and S. Börm, "An introduction to hierarchical matrices," *Math. Bohem.*, vol. 127, no. 2, pp. 229–241, 2002.

[5] S. Börm, "H^2-matrices—Multilevel methods for the approximation of integral operators," *Comput. Vis. Sci.*, vol. 7, no. 3, pp. 173–181, Oct. 2004.

[6] W. Chai and D. Jiao, "An H^2-matrix-based integral-equation solver of reduced complexity and controlled accuracy for solving electrodynamic problems," *IEEE Trans. Antennas Propag.*, vol. 57, no. 10, pp. 3147–3159, Oct. 2009.

[7] W. Chai and D. Jiao, "Dense matrix inversion of linear complexity for integral-equation-based large-scale 3-D capacitance extraction," *IEEE Trans. Microw. Theory Techn.*, vol. 59, no. 10, pp. 2404–2421, Oct. 2011.

[8] T. Wan, Z. N. Jiang, and Y. J. Sheng, "Hierarchical matrix techniques based on matrix decomposition algorithm for the fast analysis of planar layered structures," *IEEE Trans. Antennas Propag.*, vol. 59, no. 11, pp. 4132–4141, Nov. 2011.

[9] W. Chai and D. Jiao, "Linear-complexity direct and iterative integral equation solvers accelerated by a new rank-minimized H^2-representation for large-scale 3-D interconnect extraction," *IEEE Trans. Microw. Theory Techn.*, vol. 61, no. 8, pp. 2792–2805, Aug. 2013.

[10] W. Chai and D. Jiao, "Fast H-matrix-based direct integral equation solver with reduced computational cost for large-scale interconnect extraction," *IEEE Trans. Compon. Packag. Manuf. Technol.*, vol. 3, no. 2, pp. 289–298, Feb. 2013.

[11] A. Heldring, J. M. Rius, J. M. Tamayo, J. Parron, and E. Ubeda, "Fast direct solution of method of moments linear system," *IEEE Trans. Antennas Propag.*, vol. 55, no. 11, pp. 3220–3228, Nov. 2007.

[12] G. H. Golub and C. Van Loan, *Matrix Computations*, 3rd ed. Baltimore, MD, USA: The Johns Hopkins Univ. Press, 1996.

[13] M. Gu and S. C. Eisenstat, "Efficient algorithms for computing a strong rank-revealing QR factorization," *SIAM J. Sci. Comput.*, vol. 17, no. 4, pp. 848–869, Jul. 1996.

[14] K. Zhao, M. N. Vouvakis, and J.-F. Lee, "The adaptive cross approximation algorithm for accelerated method of moments computations of EMC problems," *IEEE Trans. Electromagn. Compat.*, vol. 47, no. 4, pp. 763–773, Nov. 2005.

[15] J. Shaeffer, "Direct solve of electrically large integral equations for problem sizes to 1 M unknowns," *IEEE Trans. Antennas Propag.*, vol. 56, no. 8, pp. 2306–2313, Aug. 2008.

[16] J. Laviada, F. Las-Heras, M. R. Pino, and R. Mittra, "Solution of electrically large problems with multilevel characteristic basis functions," *IEEE Trans. Antennas Propag.*, vol. 57, no. 10, pp. 3189–3198, Oct. 2009.

[17] J. M. Tamayo, A. Heldring, and J. M. Rius, "Multilevel adaptive cross approximation (MLACA)," *IEEE Trans. Antennas Propag.*, vol. 59, no. 12, pp. 4600–4608, Dec. 2011.

[18] A. Heldring, J. M. Tamayo, C. Simon, E. Ubeda, and J. M. Rius, "Sparsified adaptive cross approximation algorithm for accelerated method of moments computations," *IEEE Trans. Antennas Propag.*, vol. 61, no. 1, pp. 240–246, Jan. 2013.

[19] F. Gesztesy and K. A. Makarov, "(Modified) fredholm determinants for operators with matrix-valued semi-separable integral kernels revisited," *Integral Equ. Operator Theory*, vol. 47, no. 4, pp. 457–497, Dec. 2003.

[20] S. Chandrasekaran *et al.*, "Some fast algorithms for sequentially semiseparable representations," *SIAM J. Sci. Comput.*, vol. 27, no. 2, pp. 341–364, 2015.

[21] E. Corona, P.-G. Martinsson, and D. Zorin, "An O(N) direct solver for integral equations on the plane," *Appl. Comput. Harmon. Anal.*, vol. 38, no. 2, pp. 284–317, Mar. 2015.

[22] M. Djordjevic and B. M. Notaros, "Double higher order method of moments for surface integral equation modeling of metallic and dielectric antennas and scatterers," *IEEE Trans. Antennas Propag.*, vol. 52, no. 8, pp. 2118–2129, Aug. 2004.

[23] S. M. Rao, D. R. Wilton, and A. W. Glisson, "Electromagnetic scattering by surfaces of arbitrary shape," *IEEE Trans. Antennas Propag.*, vol. 30, no. 3, pp. 409–418, May 1982.

[24] B. M. Kolundzija and A. R. Djordjevic, *Electromagnetic Modelling of Composite Metallic and Dielectric Structures*. Boston, MA, USA: Artech House, 2002.

[25] K. C. Donepudi, J.-M. Jin, and W. C. Chew, "A higher order multilevel fast multipole algorithm for scattering from mixed conducting/dielectric bodies," *IEEE Trans. Antennas Propag.*, vol. 51, no. 10, pp. 2814–2821, Oct. 2003.

[26] O. Borries, P. Meincke, E. Jørgensen, and P. C. Hansen, "Multilevel fast multipole method for higher order discretizations," *IEEE Trans. Antennas Propag.*, vol. 62, no. 9, pp. 4695–4705, Sep. 2014.

[27] L. P. Zha, Y. Q. Hu, and T. Su, "Efficient surface integral equation using hierarchical vector bases for complex EM scattering problems," *IEEE Trans. Antennas Propag.*, vol. 60, no. 2, pp. 952–957, Feb. 2012.

[28] X. Mu, H.-X. Zhou, K. Chen, and W. Hong, "Higher order method of moments with a parallel out-of-core LU solver on GPU/CPU platform," *IEEE Trans. Antennas Propag.*, vol. 62, no. 11, pp. 5634–5646, Nov. 2014.

[29] Y. Zhang, Z. Lin, X. Zhao, and T. K. Sarkar, "Performance of a massively parallel higher-order method of moments code using thousands of CPUs and its applications," *IEEE Trans. Antennas Propag.*, vol. 62, no. 12, pp. 6317–6324, Dec. 2014.

[30] H. Guo, J. Hu, and Z. Nie, "An MPI-OpenMP hybrid parallel H-LU direct solver for electromagnetic integral equations," *Int. J. Antennas Propag.*, vol. 2015, Dec. 2015, Art. no. 615743.

[31] A. Schroder, H.-D. Bruns, and C. Schuster, "Fast evaluation of electromagnetic fields using a parallelized adaptive cross approximation," *IEEE Trans. Antennas Propag.*, vol. 62, no. 5, pp. 2818–2822, May 2014.

[32] M. Djordjević and B. M. Notaroš, "Higher-order hierarchical basis functions with improved orthogonality properties for moment-method modeling of metallic and dielectric microwave structures," *Microw. Opt. Technol. Lett.*, vol. 37, no. 2, pp. 83–88, Apr. 2003.

[33] B. M. Notaros, B. D. Popovic, J. P. Weem, R. A. Brown, and Z. Popovic, "Efficient large-domain MoM solutions to electrically large practical EM problems," *IEEE Trans. Microw. Theory Techn.*, vol. 49, no. 1, pp. 151–159, Jan. 2001.

[34] S. Wang, X. S. Li, J. Xia, Y. Situ, and M. V. de Hoop, "Efficient scalable algorithms for solving dense linear systems with hierarchically semiseparable structures," *SIAM J. Sci. Comput.*, vol. 35, no. 6, pp. C519–C544, 2013.

[35] A. B. Manić, F.-H. Rouet, X. S. Li, and B. M. Notaroš, "Efficient EM scattering analysis based on MoM, HSS direct solver, and RRQR decomposition," in *Proc. IEEE Int. Symp. Antennas Propag.*, Vancouver, BC, Canada, Jul. 2015, pp. 1660–1661.

[36] B. M. Notaroš, A. B. Manić, X. S. Li, and F.-H. Rouet, "Controlling the accuracy of double higher order surface integral equation modeling by relative tolerance for matrix compression," in *Proc. Appl. Comput. Electromagn. Soc. Conf. (ACES)*, Honolulu, HI, USA, Mar. 2016, pp. 165–166.

[37] E. Michielssen and A. Boag, "A multilevel matrix decomposition algorithm for analyzing scattering from large structures," *IEEE Trans. Antennas Propag.*, vol. 44, no. 8, pp. 1086–1093, Aug. 1996.

[38] H. Guo, J. Hu, and E. Michielssen, "On MLMDA/butterfly compressibility of inverse integral operators," *IEEE Antennas Wireless Propag. Lett.*, vol. 12, pp. 31–34, 2013.

[39] Y. Brick, V. Lomakin, and A. Boag, "Fast direct solver for essentially convex scatterers using multilevel non-uniform grids," *IEEE Trans. Antennas Propag.*, vol. 62, no. 8, pp. 4314–4324, Aug. 2014.

[40] A. Heldring, J. M. Rius, J. M. Tamayo, J. Parrón, and E. Úbeda, "Multiscale compressed block decomposition for fast direct solution of method of moments linear system," *IEEE Trans. Antennas Propag.*, vol. 59, no. 2, pp. 526–536, Feb. 2011.

[41] M. M. Kostic and B. M. Kolundzija, "Maximally orthogonalized higher order bases over generalized wires, quadrilaterals, and hexahedra," *IEEE Trans. Antennas Propag.*, vol. 61, no. 6, pp. 3135–3148, Jun. 2013.

[42] I. Bogaert and F. P. Andriulli, "Maximally orthogonal high-order basis functions have a well-conditioned Gram matrix," *IEEE Trans. Antennas Propag.*, vol. 62, no. 8, pp. 4096–4104, Aug. 2014.

[43] A. Heldring, J. M. Tamayo, and J. M. Rius, "On the degrees of freedom in the interaction between sets of elementary scatterers," presented at the 3rd Eur. Conf. Antennas Propag. (Eucap), Berlin, Germany, Mar. 2009.

[44] A. B. Manić, E. Chobanyan, M. M. Ilic, and B. M. Notaros, "Parallelization of double higher order FEM and MoM techniques," in *Proc. IEEE Antennas Propag. Soc. Int. Symp.*, Memphis, TN, USA, Mar. 2014, pp. 2511–2514.

[45] *ScaLAPACK-Scalable Linear Algebra PACKage*, accessed on Mar. 7, 2017. [Online]. Available: http://www.netlib.org/scalapack

[46] *BLACS*, accessed on Mar. 7, 2017. [Online]. Available: http://www.netlib.org/blacs

[47] J. Towns *et al.*, "XSEDE: Accelerating scientific discovery," *Comput. Sci. Eng.*, vol. 16, no. 5, pp. 62–74, Sep./Oct. 2014, doi: 10.1109/MCSE.2014.80.

[48] *Stampede User Guide*, accessed on Mar. 7, 2017. [Online]. Available: https://portal.tacc.utexas.edu/user-guides/stampede#overview

[49] O. M. Bucci and G. Franceschetti, "On the degrees of freedom of scattered fields," *IEEE Trans. Antennas Propag.*, vol. 37, no. 7, pp. 918–926, Jul. 1989.

[50] A. C. Woo, H. T. G. Wang, M. J. Schuh, and M. L. Sanders, "EM programmer's notebook-benchmark radar targets for the validation of computational electromagnetics programs," *IEEE Antennas Propag. Mag.*, vol. 35, no. 1, pp. 84–89, Feb. 1993.

Ana B. Manić (S'12) was born in Belgrade, Serbia, in 1986. She received the B.S. degree from the Department of Electronics, School of Electrical Engineering, Belgrade University, Belgrade, in 2009, and the Ph.D. degree from the Department of Electrical and Computer Engineering, Colorado State University, Fort Collins, CO USA, in 2015.

She was involved in different applications of the method of moments in computational electromagnetics such as method of moments coupling with the finite element method, diakoptics domain decomposition method, fast direct solvers, and accurate and efficient calculation of singular and hypersingular integrals arising in the Galerkin-type method of moments system of equations.

Aaron P. Smull (S'15) was born in Santa Rosa, CA, in 1993. He received the B.S. degree from Colorado State University, Fort Collins, CO USA, in 2015, where he is currently pursuing the M.S. degree in electrical engineering.

His current research interests include the development of numerical algorithms for high-frequency electromagnetic scattering analysis and the design of novel electromagnetic and optical devices based on exotic material parameters.

François-Henry Rouet was born in La Roche-sur-Yon, France, in 1985. He received the M.Sc. degree in applied mathematics and the Ph.D. degree in computer science from the Institut National Polytechnique de Toulouse, Toulouse, France, in 2009 and 2012, respectively.

He was a Post-Doctoral Research Fellow with the Scalable Solvers Group, Lawrence Berkeley National Laboratory, Berkeley, CA, USA, from 2012 to 2015. He is currently a Research Scientist with Livermore Software Technology Corporation, Livermore, CA, USA, where he is involved in multiphysics finite element code LS-DYNA. He also collaborates on MUMPS and STRUMPACK linear algebra projects. His current research interests include parallel computing, sparse matrix computations, and graph algorithms.

Xiaoye Sherry Li received the Ph.D. degree in computer science from the University of California, Berkeley, CA, USA, in 1996.

She was involved in diverse problems in high performance scientific computations, including parallel computing, sparse matrix computations, high precision arithmetic, and combinatorial scientific computing. She is currently a Senior Scientist with the Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley. She is also the Lead Developer of SuperLU, a widely-used sparse direct solver, and has contributed to the development of several other mathematical libraries, including ARPREC, LAPACK, PDSLin, STRUMPACK, and XBLAS. She has co-authored over 95 publications, and contributed to several book chapters. She has collaborated with many domain scientists to deploy the advanced mathematical software in their application codes, including those from accelerator engineering, chemical science, earth science, plasma fusion energy science, and materials science.

Dr. Li is a SIAM Fellow and an ACM Senior Member. She has served on the editorial boards of the *SIAM Journal of Scientific Computing* and the *ACM Transactions on Mathematical Software*, as well as on many program committees of the scientific conferences.

Branislav M. Notaroš (M'00–SM'03–F'16) was born in Zrenjanin, Yugoslavia, in 1965. He received the Dipl.Ing. (B.S.), M.S., and Ph.D. degrees in electrical engineering from the University of Belgrade, Belgrade, Yugoslavia, in 1988, 1992, and 1995, respectively.

From 1996 to 1999, he was an Assistant Professor with the School of Electrical Engineering, University of Belgrade. From 1998 to1999, he was a Visiting Scholar at the University of Colorado at Boulder, Boulder, CO, USA. He was an Assistant Professor from 1999 to 2004 and an Associate Professor from 2004 to 2006 with the Department of Electrical and Computer Engineering, University of Massachusetts Dartmouth, Dartmouth, MA, USA. From 2006 to 2012, he was an Associate Professor with the Department of Electrical and Computer Engineering, Colorado State University, Fort Collins, CO, USA, where he is currently a Professor and a University Distinguished Teaching Scholar, as well as the Director of the Electromagnetics Laboratory. He has authored more than 200 journal and conference papers, and three workbooks in electromagnetics and in fundamentals of electrical engineering (basic circuits and fields). He has authored textbooks *Electromagnetics* (Prentice Hall, 2010), *MATLAB-Based Electromagnetics* (Prentice Hall, 2013), and *Conceptual Electromagnetics* (CRC Press, 2017). His current research interests include computational electromagnetics, higher order numerical methods, antennas, scattering, microwaves, metamaterials, the characterization of snow and rain, surface and radar precipitation measurements, RF design for MRI at ultrahigh magnetic fields, and electromagnetics education.

Dr. Notaroš is a fellow of the Applied Computational Electromagnetics Society (ACES). He was a recipient of the 2005 IEEE MTT-S Microwave Prize, Best Paper Award for the IEEE TRANSACTIONS ON MICROWAVE THEORY AND TECHNIQUES, the 1999 IEE Marconi Premium, Best Paper Award for the IEE Proceedings on Microwaves, Antennas and Propagation, the 1999 URSI Young Scientist Award, the 2005 UMass Dartmouth Scholar of the Year Award, the 2004 UMass Dartmouth College of Engineering Dean's Recognition Award, the 1992 Belgrade Chamber of Industry and Commerce Best M.S. Thesis Award, the 2009, 2010, 2011, and 2014 Colorado State University Electrical and Computer Engineering Excellence in Teaching Awards, the 2010 Colorado State University College of Engineering George T. Abell Outstanding Teaching and Service Faculty Award, the 2012 Colorado State University System Board of Governors Excellence in Undergraduate Teaching Award, the 2014 Colorado State University Provost's N. Preston Davis Award for Instructional Innovation, the 2012 IEEE Region 5 Outstanding Engineering Educator Award, the 2014 Carnegie Foundation for the Advancement of Teaching Colorado Professor of the Year Award, the 2015 American Society for Engineering Education ECE Distinguished Educator Award, and the 2015 IEEE Undergraduate Teaching Award. He served as the General Chair for the 11th International Workshop on Finite Elements for Microwave Engineering—FEM2012, 2012, Estes Park, CO, USA, and a Guest Editor of the Special Issue on Finite Elements for Microwave Engineering, Electromagnetics, Vol. 34, Issue 3–4, 2014. He has served on the Board of Directors of the ACES since 2016, and as the Chair of the Technical Committee for the Commission B of the U.S. National Committee, International Union of Radio Science since 2014.