

MIT Open Access Articles

Compression of volume-surface integral equation matrices via Tucker decomposition for magnetic resonance applications

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Giannakopoulos, Ilias I, Guryev, Georgy D, Serralles, Jose EC, Georgakis, Ioannis P, Daniel, Luca et al. 2022. "Compression of volume-surface integral equation matrices via Tucker decomposition for magnetic resonance applications." IEEE Transactions on Antennas and Propagation, 70 (1).

As Published: 10.1109/TAP.2021.3090835

Publisher: Institute of Electrical and Electronics Engineers (IEEE)

Persistent URL: <https://hdl.handle.net/1721.1/143112>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



Compression of volume-surface integral equation matrices via Tucker decomposition for magnetic resonance applications

Ilias I. Giannakopoulos, *Member, IEEE*, Georgy D. Guryev, José E. C. Serrallés, *Member, IEEE*, Ioannis P. Georgakis, Luca Daniel, *Member, IEEE*, Jacob K. White, *Fellow, IEEE*, Riccardo Lattanzi, *Senior Member, IEEE*

Abstract—In this work, we propose a method for the compression of the coupling matrix in volume-surface integral equation (VSIE) formulations. VSIE methods are used for electromagnetic analysis in magnetic resonance imaging (MRI) applications, for which the coupling matrix models the interactions between the coil and the body. We showed that these effects can be represented as independent interactions between remote elements in 3D tensor formats, and subsequently decomposed with the Tucker model. Our method can work in tandem with the adaptive cross approximation technique to provide fast solutions of VSIE problems. We demonstrated that our compression approaches can enable the use of VSIE matrices of prohibitive memory requirements, by allowing the effective use of modern graphical processing units (GPUs) to accelerate the arising matrix-vector products. This is critical to enable numerical MRI simulations at clinical voxel resolutions in a feasible computation time. In this paper, we demonstrate that the VSIE matrix-vector products needed to calculate the electromagnetic field produced by an MRI coil inside a numerical body model with 1 mm^3 voxel resolution, could be performed in ~ 33 seconds in a GPU, after compressing the associated coupling matrix from $\sim 80 \text{ TB}$ to $\sim 43 \text{ MB}$.

Index terms— Cross approximation, Global Maxwell Tomography, graphical processing unit, magnetic resonance imaging, matrix-vector product, Tucker decomposition, volume-surface integral equation.

I. INTRODUCTION

MAGNETIC resonance (MR) imaging (MRI) provides high-resolution images of the interior anatomical and physiological structure of the human body, with exquisite soft-tissue contrast. The quality of MR images, as well as the achievable spatial and temporal resolution, depend on the available signal-to-noise ratio (SNR). SNR increases with the main magnetic field strength. This fact motivated the recent

development of 7 Tesla (T) clinical MR scanners and research-only scanners with field strengths as high as 11.7 T [1]. At ultra-high-field (UHF) MRI ($\geq 7 \text{ T}$), the radio frequency (RF) wavelength is short. This results in strong interactions between biological tissues and the electromagnetic (EM) field generated by the RF coils [2]–[5]. Such interactions could compromise image quality and patient safety. To address these issues, EM modeling is often used to predict and manipulate the EM field distribution during RF coil design.

Integral equation (IE) methods are suitable options for EM analysis in MRI. First, they do not suffer from grid dispersion errors [6], [7], in contrast with the finite-difference-time-domain (FDTD) and finite-element-methods (FEM), because the Green's functions in the underlying IE formulation act as an exact EM field propagator from a source to an observation point. Second, for the case of single-frequency problems, IE algorithms can be extensively customized with the use of numerical linear algebra techniques for fast and accurate simulations, tailored to specific applications [8]–[11].

For example, the MAgnetic-Resonance Integral Equation (MARIE) suite [11], [12] was developed to numerically compute the EM field distribution generated by RF coils in the human body during MRI. MARIE combines surface and volume integral equations (SIE, VIE), employing a triangular tessellation for the RF coils' conductors and a uniform voxelized grid discretization for the body models. RWG basis functions [13] and polynomial basis functions [12], [14] are used to compute the unknowns of the surface and volume IE, respectively. Matrix-vector products are accelerated using the fast Fourier transform (FFT).

The VIE computational engine of MARIE has been recently employed for the forward problem in Global Maxwell Tomography (GMT) [15], a technique that iteratively solves an ill-conditioned inverse problem to extract electrical properties from volumetric MR measurements. In the first experimental demonstration of GMT with a uniform phantom, constant incident fields were used for all iterations [15]. More recently, it was shown in simulation that GMT could accurately reconstruct brain electrical properties at 7 T using a tailored RF coil array [16]. However, in order to confirm this with in-vivo experiments, the currents on the coil conductors cannot be simply considered constant as in the initial experiment with a uniform phantom. Instead, the incident fields must be updated at each iteration of GMT to account for changes in the sample electrical properties distribution. Therefore, GMT must be

This work was supported by NIH R01 EB024536 and by NSF 1453675. It was performed under the rubric of the Center for Advanced Imaging Innovation and Research (CAI²R, www.cai2r.net), a NIBIB Biomedical Technology Resource Center (NIH P41 EB017183).

Ilias I. Giannakopoulos, Ioannis P. Georgakis and Riccardo Lattanzi are with Center for Advanced Imaging Innovation and Research (CAI²R), Department of Radiology, New York University Grossman School of Medicine, NY, USA.

Georgy D. Guryev, José E.C. Serrallés, Luca Daniel, and Jacob K. White are with the Research Laboratory of Electronics, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, USA.

Riccardo Lattanzi is also with the Bernard and Irene Schwartz Center for Biomedical Imaging, Department of Radiology, New York University Grossman School of Medicine, NY, USA and the Vilcek Institute of Graduate Biomedical Sciences, New York University Grossman School of Medicine, NY, USA.

implemented with a volume-surface IE (VSIE) framework, in which a coupling matrix represents the coil-body interactions in the IE system of equations [11]. Such approach requires a large amount of memory, which could prevent using clinically-relevant voxel resolutions and fine coil meshes.

The aim of this work is to use Tucker decomposition [17] to perform a column-wise compression of the VSIE coupling matrix, in order to limit the associated memory demand and enable the computation of the relevant matrix-vector products in GPUs. Our approach was motivated by previous work [10] on the reduction of the memory footprint of FFT-based VIE Green's function tensors and the acceleration of matrix-vector products in VIE using GPU. Tucker decomposition belongs to a larger family of tensor decompositions and have been used successfully in the past for matrix compression inside IE-based simulations for EM applications. Examples include EM simulations of realistic body models simulations for UHF MRI [10], [18] and capacitance extraction [19]–[21]. Other tensor decompositions could be used [22]–[24], but for the intrinsic 3D nature of the problem at hand, Tucker optimizes operations and memory complexity. In cases where the coil is placed far from the scatterer, the coupling matrix can be first compressed with a 2D cross approximation method [25]–[27] and then further compressed by applying our proposed technique to the resulting matrices. Towards this direction, we developed an algorithm based on the adaptive cross approximation (ACA) [28], [29] to efficiently perform our compression approach within the iterative loop of ACA.

Other memory-friendly techniques are available for the fast implementation of matrix-vector products in VSIE simulations: the magnetic resonance Green's function (MRGF), the fast multipole method (FMM), the precorrected FFT method (pFFT), and the Multilevel Nonuniform Grid Algorithm (MNGA). The MRGF [11] is a model order reduction technique that can considerably accelerate the solutions of the VSIE system. However, the required computational time can be overwhelming when fine voxel resolutions and PWL basis functions are used. The FMM [30]–[32] has been extensively used for the compression of the Method of Moments (MoM) matrix appearing in IEs and could proven to be a good alternative for solving the VSIE system studied herein. Nevertheless, in the presented work, we are only interested in the compression of an off-diagonal block of the full MoM matrix (i.e., the coupling matrix), since the remaining blocks can be handled efficiently with other methods presented in [10], [12], [13]. The pFFT method [8] could be used to project the discretized coil's elements onto an extended VIE domain, where the Green's function tensors are compressible with the Tucker decomposition (pFFT+Tucker) [10]. However, this approach would be effective only when the coil is close to the scatterer, like for the close-fitting coil studied in section IV.B.1 [33]. In fact, in such situation the extended VIE domain would be larger than the original one by only a few voxels in each direction, allowing the matrix-vector products to fit in a GPU. As a result, the pFFT+Tucker approach could be more efficient than our proposed method for such geometries, although more complex to implement. Finally, the MNGA [34] aims to accelerate the solutions of the MoM system of

equations through interpolation of sampled fields generated by source distributions in spherical grids. However, this method works well for quasiplanar and elongated geometries, with dimensions larger than the operating wavelength, which is not the case in UHF MRI. While the previously described methods could be applied to VSIE simulations, since we are interested in problems discretized on a uniform 3D grid for UHF MRI frequencies, in this work we chose to explore approaches based on Tucker decomposition that allows us to exploit the low-rank properties encoded in the VSIE coupling matrix.

The remainder of this paper is organized as follows. In Section II, we summarize the relevant technical background and equations related to the standard VSIE method. We also show, as an example application, the use of the VSIE coupling matrix in the GMT inverse problem formulation. In addition, we outline the Tucker decomposition and the ACA method. In Section III, we introduce the Tucker-based assembly technique of the coupling matrix along with the novel algorithms for the matrix-vector products implementation. Moreover, we present a new algorithm for a memory friendly assembly of the coupling matrix, based on the combination of ACA with the Tucker decomposition. Section IV describes a series of numerical experiments aimed at investigating the multilinear rank and the memory requirements of the compressed matrix in different scenarios, along with time footprint of the matrix-vector product for various mesh discretizations. Section V discusses the results. Section VI summarizes the work and provides a few take home points. The following TABLE I lists the notations used in this work.

TABLE I
NOTATION

Notation	Description
a	Scalar
\mathbf{a}	Vector in \mathbb{C}^3
\mathbf{a}	Vector in \mathbb{C}^n
A	Matrix in $\mathbb{C}^{n_1 \times n_2}$
A^T	Transpose of matrix
A^*	Conjugate transpose of matrix
\mathcal{A}	Tensor in $\mathbb{C}^{n_1 \times n_2 \times n_3}$
\mathcal{A}	Struct containing Tucker compressed tensors
\mathcal{A}	Operator acting on vectors in \mathbb{C}^3
\times_i	n-mode products
i	Imaginary unit $i^2 = -1$

II. TECHNICAL BACKGROUND

A. Volume-Surface Integral Equation

1) *Coupled Linear System:* Let us consider an IE-based solver for the EM wave analysis in MRI applications. The body domain Ω can be modeled with the current-based VIE (JVIE), while the conducting surfaces (receive and transmit RF coil arrays, coil shields, and gradient coils) with the SIE. The

IEs are solved with the MoM [35]. The resulting system of equations can be written in the following block matrix form.

$$\begin{bmatrix} Z_{cc} & T_{bc}^T \\ T_{bc} & Z_{bb} \end{bmatrix} \begin{bmatrix} J_s \\ J_p \end{bmatrix} = \begin{bmatrix} V \\ 0 \end{bmatrix}. \quad (1)$$

Here, $Z_{cc} \in \mathbb{C}^{m \times m}$ is the Galerkin matrix that models interactions between the coil's discretization elements (triangles' common edges), with the aid of the free-space Green's function that appears in the SIE formulation. It associates the equivalent surface currents on the coil ($J_s \in \mathbb{C}^{m \times p}$), with the voltage excitation matrix $V \in \mathbb{C}^{m \times p}$. The coil conductors are modeled with triangular meshes, and the unknown surface equivalent currents are approximated with RWG basis functions [13]. m is the number of RWG, or non-boundary, triangle edges that appear on the mesh, whereas p is the number of excitation ports (i.e., number of coil's channels) of the conducting surfaces.

The matrix $Z_{bb} \in \mathbb{C}^{qn_v \times qn_v}$ is another Galerkin matrix which models the EM interactions of an external volumetric EM field, produced by the coil, with the body. Specifically, the matrix relates the polarization currents in Ω to an external incident EM field produced by the conducting surfaces. n_v is the number of voxels and q the number of basis functions per voxel. Differently than Z_{cc} , Z_{bb} requires a large amount of memory, even for coarse resolutions. To handle that, Ω can be discretized over a voxelized uniform grid, giving Z_{bb} a three-level Block-Toeplitz with Toeplitz Blocks (BTTB) structure. As a result, only the defining columns of the BTTB matrix need to be stored and the matrix-vector product can be accelerated using the FFT, as in [12], [36]–[42]. The unknown polarization currents ($J_p \in \mathbb{C}^{qn_v \times p}$) can be discretized with polynomial basis functions, either piecewise constant [12] (PWC, 3 unknowns per voxel) or piecewise linear [14] (PWL, 12 unknowns per voxel), and a single-voxel support.

The presence of conductive tissue near the coil conductors perturbs J_s from their values in free-space. In fact, the voltage excitations at the coil's ports create incident EM fields that scatter from the dielectric body back to the coil conductors, changing their current distribution. The coupling matrix $T_{bc} \in \mathbb{C}^{qn_v \times m}$ is used to account for this effect, by modeling the coupling interactions between the dyadic Green's function [43] of the SIE and the VIE formulations. Specifically, in equation (1), T_{bc} maps electric surface equivalent currents to electric fields through the \mathcal{N} Green's function operator of VIE:

$$\mathcal{N}(\mathbf{s}) \triangleq \nabla \times \nabla \times \int_{\Omega} g(\mathbf{r} - \mathbf{r}') \mathbf{s}(\mathbf{r}') d^3 \mathbf{r}'. \quad (2)$$

g is the free-space Green's function, or fundamental Helmholtz solution, and it is equal to

$$g(\mathbf{r} - \mathbf{r}') = \frac{e^{-ik_0|\mathbf{r} - \mathbf{r}'|}}{4\pi|\mathbf{r} - \mathbf{r}'|}, \quad (3)$$

where k_0 is the free-space wavenumber, \mathbf{r} the source point, and \mathbf{r}' the observation point. Each element of the VSIE coupling matrix is a 5D integral formed from the inner product between the discretized \mathcal{N} operator applied on a VIE basis function, and an RWG basis function.

2) *VSIE implementation of GMT*: GMT estimates tissue electrical properties from MR measurements by solving an inverse problem [15]. In GMT, the cost function compares actual measurements against simulated measurements of the relative b_1^+ fields generated by multiple sources (e.g., multipoint transmit coils) inside a sample and iteratively updates the estimate of the sample's electrical properties. GMT was initially demonstrated using the JVIE formulation for the solutions of the forward problem, therefore ignoring the effect of the dielectric sample on the incident fields. However, these interactions must be taken into account for accurate in-vivo experiments with close-fitting RF coils. In other words, the GMT framework must be ported from a VIE to a VSIE formulation, in which the incident fields are not constant but calculated at each GMT iteration as

$$\begin{aligned} E^{\text{inc}}(\epsilon_r) &= T_{bc} J_s(\epsilon_r), \\ H^{\text{inc}}(\epsilon_r) &= Z_{bc} J_s(\epsilon_r). \end{aligned} \quad (4)$$

where E^{inc} and H^{inc} are the discretized incident electric and magnetic fields respectively. ϵ_r is the complex permittivity. Z_{bc} maps the equivalent surface electric currents to the magnetic fields with the aid of the \mathcal{K} operator:

$$\mathcal{K}(\mathbf{s}) \triangleq \nabla \times \int_{\Omega} g(\mathbf{r} - \mathbf{r}') \mathbf{s}(\mathbf{r}') d^3 \mathbf{r}'. \quad (5)$$

In addition, in the new implementation, the gradient of the GMT's cost function will require to solve a Hermitian adjoint system of equations that includes multiplications with the conjugate transpose of Z_{bc} .

Matrix-vector products involving the coupling matrix are typically performed without storing the full matrix, due to its intractably large size. In the case of iterative inverse problem solutions, such as in GMT, this approach could considerably increase the computation time, because it requires the re-assembly of the full matrix at each iteration. In the next sections, we propose a compression algorithm that reduces the computational burden, by enabling one to assemble the full coupling matrix only once and then perform just the matrix-vector multiplications in each of GMT's iterations.

B. Numerical Linear Algebra Methods

1) *Tucker Decomposition*: A 3D tensor $\mathcal{A} \in \mathbb{C}^{n_1 \times n_2 \times n_3}$ can be decomposed with the Tucker model [17] in the following form:

$$\begin{aligned} \mathcal{A} &= \mathcal{G} \times_1 U^1 \times_2 U^2 \times_3 U^3, \text{ or} \\ \mathcal{A}_{ijk} &= \sum_{\chi=1}^{r_1} \sum_{\psi=1}^{r_2} \sum_{\zeta=1}^{r_3} \mathcal{G}_{\chi\psi\zeta} U_{i\chi}^1 U_{j\psi}^2 U_{k\zeta}^3. \end{aligned} \quad (6)$$

Here $U^\gamma \in \mathbb{C}^{n_\gamma \times r_\gamma}$, $\gamma = 1, 2, 3$, are unitary matrices, dubbed as Tucker factors, while $\mathcal{G} \in \mathbb{C}^{r_1 \times r_2 \times r_3}$ is the Tucker core. The dimensions of the Tucker core indicate the multilinear (or Tucker) ranks of \mathcal{A} . The symbols \times_γ are called n -mode products and perform a convolution over the \times_γ axis, for example, the \times_1 product performs the following operation:

$$\mathcal{P} = \mathcal{G} \times_1 U^1 \Leftrightarrow \mathcal{P}_{i\psi\zeta} = \sum_{\chi=1}^{r_1} \mathcal{G}_{\chi\psi\zeta} U_{i\chi}^1. \quad (7)$$

Here, $\mathcal{P} \in \mathbb{C}^{n_1 \times r_2 \times r_3}$. The expansion of \mathcal{A} in equation (6) can be truncated to a desired tolerance, and return an approximation of \mathcal{A} . A visual representation of Tucker decomposition can be seen in Fig. 1.




Fig. 1. Visual representation of Tucker decomposition.

To compute the above-mentioned Tucker components, one has to choose a suitable compression algorithm. The higher order singular value decomposition (HOSVD) is an orthogonal Tucker decomposition, widely used because it has a proven upper error bound [44]. Moreover, the algorithm is based entirely on singular value decomposition (SVD), which provides a robust and stable approximation of the initial tensor. Note that SVD requires the assembly of the initial array, which could be challenging for large tensors. In such cases, one could implement the Tucker decomposition using a 3D cross approximation algorithm [45].

2) *Cross Approximation*: A matrix $A \in \mathbb{C}^{n_1 \times n_2}$ can be approximated with the so-called 2D cross approximation method [25], [26] as follows:

$$A \approx UV^*. \quad (8)$$

Here, $U \in \mathbb{C}^{n_1 \times r_c}$ and $V \in \mathbb{C}^{n_2 \times r_c}$. r_c represents the column rank of matrix A . Cross approximation algorithms construct the decomposition of A by using only some rows and columns of it, differently than SVD, which depends on the availability of the full matrix. Several algorithms have been developed over the previous decades for the implementation of cross approximation, with two being the most used ones: the ACA [28], [29], and the maximum volume-based cross algorithm [27], [46]. The latter requires the implementation of LU, QR and SVD for its efficient implementation. Therefore, the memory demand of the algorithm increases drastically for large tall matrices, such as the coupling matrices T_{bc} and Z_{bc} in the case fine voxel resolutions. On the other hand, the memory demand of ACA is only dictated by the size of matrices U and V .

III. TUCKER-BASED COMPRESSION ALGORITHM

In VSIE, the columns of the coupling matrix describe interactions between coil and body basis functions through the Green's functions of equations (2) and (5); therefore, they represent well-separated geometrical blocks. Due to the 3D nature of the problem, the key idea of our proposed compression algorithm is to reshape these columns as tensors and approximate them with the low multilinear Tucker model. This

compression strategy enabled us to develop a new method to efficiently perform the matrix-vector product and an extension to ACA, which are described later in this section.

A. Matrix Assembly

Each of the m columns of the coupling matrices Z_{bc} and T_{bc} can be seen as the concatenation of q vectors, where each vector represents the component-wise interaction between one RWG element on the coil and the basis functions of all the voxels in the body domain. For PWC, $q = 3$, whereas for PWL, $q = 12$. Since these vectors model interactions between remote discretization elements, they have low-rank properties [47]. To exploit the low-rank, each column of the coupling matrix can be reshaped as q 3D tensors $\mathcal{Z}^{kj} \in \mathbb{C}^{n_1 \times n_2 \times n_3}$, $k = 1 : q$, $n_v = n_1 \times n_2 \times n_3$, which are compressible with the Tucker decomposition [48]. A graphical description of the algorithm is shown in Fig. 2 for Z_{bc} and PWC basis functions.




Fig. 2. Visual representation of the Tucker-based algorithm for the compression of the Z_{bc} matrix, in the case of PWC basis functions. Each vector can be reshaped into a 3D tensor that is then compressed via Tucker decomposition.

If the coupling matrix is first approximated with the ACA as UV^* , then our approach can still be used to compress the r_c columns of U . In fact, cross approximation is a well-conditioned operation, therefore the Tucker ranks of the reshaped columns of U will be similar to the ones of the reshaped columns of the coupling matrix. The V matrix here is usually much smaller than U and does not require any additional compression.

TABLE II shows the total memory footprint associated with the assembly of the coupling matrix: Full assembly, assembly with ACA, and assembly with our proposed method by compressing either the columns of the coupling matrix (Tucker) or the columns of U (ACA+Tucker). The memory required after compressing the coupling matrix with Tucker is $n_v / (r^3 + 3nr)$ times smaller than the memory required by the full matrix, where n and r refer to the tensor's linear dimension and Tucker rank, respectively. If our Tucker-based compression method is instead applied after the ACA assembly, then the total compression improves by a factor of $\sim m/r_c$, given that r_c is small. TABLE II also shows the computational complexity of the assembly operations. The multiplicative constant factor c_1 , which is present in all cases, represents the cost to compute the elements of the coupling matrix and is usually large. In fact, each element requires a 5D integration, whose computational cost depends on the number

of both surface and volume quadrature integration points. As a result, the assembly of the matrix is extremely inefficient and should be implemented in parallel for multiple voxel-RWG basis function interactions.

Note that in certain cases, for example when the coil is close to the body, ACA may not achieve a meaningful compression and would not be advantageous to combine it with Tucker decomposition. In such cases, the preferable approach would be to divide the coupling matrix in q blocks of size $n_v \times m$, assembled them in parallel, and then compress their tensor components with a Tucker-based method like HOSVD. Alternatively, if the coupling matrix is sufficiently small, one could assemble it in its full form and then apply Tucker directly to compress it.

TABLE II
COMPLEXITY FOR CONSTRUCTING THE COUPLING MATRIX

Assembly Method	Operations	Memory
Full	$\mathcal{O}(c_1 q n_v m)$	$q n_v m$
ACA	$\mathcal{O}(c_1 r_c^2 (q n_v + m))$	$q n_v r_c + m r_c$
Tucker	Full + $\mathcal{O}(r q n_v m)$	$q(r^3 + 3nr)m$
ACA+Tucker	ACA + $\mathcal{O}(r q n_v r_c)$	$q(r^3 + 3nr)r_c + m r_c$

B. Matrix-Vector Product

Decompressing the coupling matrix to compute the matrix-vector product $\mathbf{y} = Z_{bc}\mathbf{x}$, like in equations (4), may not be possible due to computer or GPU memory limitations. To address this, we propose a novel approach to efficiently compute the matrix-vector product without fully decompressing the coupling matrix. We initiate \mathbf{y} as a vector of zeros. Inside a loop that cycles over the RWG basis functions, we decompress the q tensors of a column $j \in [1, m]$, reshape them as vectors, and concatenate them to form the j -th column of the original Z_{bc} matrix. The vector-scalar product between the j -th column and \mathbf{x}_j is then computed, and the result increments the elements of \mathbf{y} . The same algorithm can be followed for the matrix-matrix product $Y = Z_{bc}X$.

The conjugate transpose matrix-vector product $\mathbf{y} = Z_{bc}^* \mathbf{x}$ is required for the computation of the gradient of the cost function in the VSIE-based GMT implementation. This case is slightly different than the standard matrix-vector product: inside the loop cycling through the RWG functions, a row-column vector product must be computed between the conjugate transpose of the decompressed j -th column of Z_{bc} and \mathbf{x} , which yields the scalar \mathbf{y}_j . The algorithm remains instead the same for the conjugate transpose matrix-matrix products. Both algorithms (for p -columned matrices X and Y) are summarized in the pseudocode below:

Algorithm 1: $Y = Z_{bc}X$

```

1: for  $k=1:q$  do
2:    $Y^k = 0$ 
3: for  $j=1:m$  do

```

```

4:   for  $k=1:q$  do
5:     Decompress  $Z^{kj}$ 
6:      $Y^k += Z^{kj}(:, \mathbf{x}_j)$ 
7:  $Y = \begin{bmatrix} Y^1 \\ \vdots \\ Y^q \end{bmatrix}$ 

```

Algorithm 2: $Y = Z_{bc}^* X$

```

1:  $Y = 0$ 
2: for  $j=1:m$  do
3:   for  $k=1:q$  do
4:     Decompress  $Z^{kj}$ 
5:      $Y_j = \begin{bmatrix} Z^{1j}(:, \mathbf{x}_j)^* \\ \vdots \\ Z^{qj}(:, \mathbf{x}_j)^* \end{bmatrix} X$ 

```

In Algorithm 1, $X \in \mathbb{C}^{m \times p}$, $Y \in \mathbb{C}^{q n_v \times p}$ (vice-versa for Algorithm 2), and $Z^{kj}(:, \cdot)$ is the reshaped column vector of the tensor component Z^{kj} . The algorithms remain the same if Z_{bc} is compressed with ACA first ($Z_{bc} = UV^*$). One has to replace Z_{bc} with U , m with r_c , and assign $X = V^* X$ for Algorithm 1, and $Y = VY$ for Algorithm 2. Both the standard and the conjugate transpose matrix-vector products have the same complexity, shown in TABLE III, for the full, ACA, Tucker, and ACA+Tucker compressed cases. The full matrix-vector product is faster than the Tucker-compressed approach by a factor of $(r + p)/p$, which depends on the number of columns of X and Y and the Tucker rank. ACA can be faster than the full case for small values of r_c . Although the approach based on Tucker decomposition is slower because it requires additional flops compared to the other methods, it is more likely to fit in GPUs, thanks to its small memory footprint.

TABLE III
MATRIX-VECTOR PRODUCT COMPLEXITY

Matrix Form	Operations Complexity
Full	$\mathcal{O}(q n_v m p)$
ACA	$\mathcal{O}(q n_v r_c p) + \mathcal{O}(r_c m p)$
Tucker	$\mathcal{O}(r q n_v m) + \mathcal{O}(q n_v m p)$
ACA+Tucker	$\mathcal{O}(r q n_v r_c) + \mathcal{O}(q n_v r_c p) + \mathcal{O}(r_c m p)$

C. Tucker-based ACA

If the coupling matrix is first compressed with ACA, the previous methods for matrix assembly and matrix-vector product could still be applied to the matrix U of the cross approximation. However, for the case of realistic body models discretized with fine voxel resolutions, the traditional implementation of ACA (a detailed description can be found in [49]) might fail due to random access memory (RAM) overflow because of the size of U (see section IV.B.2). To address this, we propose an extension of ACA in which the matrix U is assembled directly in a compressed form, based on our

proposed Tucker decomposition technique. The algorithm is summarized in pseudocode below:

Algorithm 3: ACA of $Z \in \mathbb{C}^{m_1 \times m_2}$. Assembly with compressed U

```

1:  $\epsilon = \text{tolerance}, i = 1, \mathbf{s}_1 = 0$ 
2:  $\mathcal{U} = [], V = []$ 
3: for  $k = 1 : \min(m_1, m_2)$  do
4:    $\mathbf{r} \leftarrow Z_{:,i}$ 
5:   if  $k > 1$  then
6:      $[f_1, f_2, f_3, p] \leftarrow i$ 
7:     for  $l = 1 : \text{size}(\mathcal{U}, 2)$  do
8:        $[\mathcal{G}, U^1, U^2, U^3] \leftarrow U_{pl}$ 
9:        $\mathbf{t}_l = \mathcal{G} \times_1 U_{f_1,:}^1 \times_2 U_{f_2,:}^2 \times_3 U_{f_3,:}^3$ 
10:       $\mathbf{r} += -\mathbf{t}V^*$ 
11:     $j \leftarrow \text{index of max element of } \mathbf{r}$ 
12:     $\mathbf{y} \leftarrow (\mathbf{r}/\mathbf{r}_j)^*$ 
13:     $\mathbf{x} \leftarrow Z_{:,j}$ 
14:    if  $k > 1$  then
15:       $\mathbf{x} += -\text{Alg1}(\mathcal{U}, V_{j,:}^*)$ 
16:     $\mathbf{s}_{k+1} \leftarrow \mathbf{s}_k + (\|\mathbf{x}\| \|\mathbf{y}\|)^2$ 
17:    if  $k > 1$  then
18:       $\mathbf{s}_{k+1} += 2 \sum [\text{Re}\{(\text{Alg2}(\mathcal{U}, \mathbf{x})) \odot (V^* \mathbf{y})\}]$ 
19:    Reshape  $\mathbf{x}$  to  $q$   $\mathcal{X}^q$  tensors.
20:     $\mathcal{U} = [\mathcal{U} \text{ HOSVD}(\mathcal{X}^1, 3\epsilon) \dots \text{HOSVD}(\mathcal{X}^q, 3\epsilon)]$ 
21:     $V = [V \ \mathbf{y}]$ 
22:    if  $\|\mathbf{x}\| \|\mathbf{y}\| \leq \epsilon \sqrt{\mathbf{s}_{k+1}}$  then break
23:     $\mathbf{x} = |\mathbf{x}|, \ \mathbf{x}_i = 0$ 
24:     $i \leftarrow \text{index of max element of } \mathbf{x}$ 

```

Here \mathcal{U} is a struct of size $q \times r_c$ ($q = 3$ for PWC, $q = 12$ for PWL, r_c is the rank of Z) than contains tensors. Each time a new column of U is computed, it is reshaped to q tensors, which are then compressed with a truncated HOSVD of tolerance 3ϵ (line 19,20). The HOSVD tolerance has to be higher than the ACA tolerance since the irrelevant numerical digits ($< 1e-3$) appearing in U are incompressible. We found that a 3 times higher tolerance is a good choice for our numerical examples. To perform matrix- and conjugate transpose matrix-vector products with the compressed U we followed Algorithms 1 (line 15) and 2 (line 18). Finally, when a row of U is requested in Algorithm 3, we first calculate the voxel and basis function component corresponding to that row (line 6) and then decompress, using equation (6), only the required elements from the Tucker compressed components of U (line 8,9). The proposed algorithm avoids RAM overflowing, but it is slower than the traditional ACA due to the multiple tensor decompositions. Nevertheless, it could always be accelerated via GPU, since its memory demand is as low as for one column of the coupling matrix.

IV. NUMERICAL EXPERIMENTS

A. Tucker Rank Behavior

In this section, we study the low-Tucker rank properties of the Z_{bc} coupling matrix. We considered multiple geometrical scenarios and altered the distance between the conductive

surface (coil) and the VIE domain, the operating frequency and the conductive surface's discretization. The tensor components of the columns of the coupling matrix were compressed with the HOSVD algorithm and a tolerance of $1e-8$, which yielded a relative error similar to the tolerance for all cases, due to the robustness of the SVD itself. Such error can be considered negligible, because the tolerance of the iterative solver used in FFT-based VIE systems is usually orders of magnitude higher.

1) *Tucker Rank vs. Distance:* It is well established that the Green's function integro-differential operators between well-separated geometries present low-rank properties [47]. Here we studied the relation between the low multilinear rank of the compressed coupling matrix Z_{bc} and the distance between the body's domain and the coil. We set the frequency to 298.06 MHz, the operating frequency of 7 Tesla MRI. We modeled a single perfectly electric conducting (PEC) loop coil of radius $\rho = 0.50$ m and discretized it with 125 triangular elements. The coil was centered in $(0, d, 0)$, where d was varied as 0.55, 0.6, ..., 1 m. The domain was a cuboid with edge length of 1 m, centered at $(0, 0, 0)$ and discretized with voxels of 1 cm isotropic resolution and PWC basis functions (Fig. 3). As a result, the tensor's size was $101 \times 101 \times 101$ and the memory required by the fully assembled Z_{bc} was 5848 MBs.




Fig. 3. Loop-cubic domain geometry. The loop coil was shifted on the \hat{y} direction, for 10 discrete distances between 0.55 to 1 m from the center of the cube.

Fig. 4 illustrates the reduction of the maximum rank (maximum of all Tucker ranks for all components) of the coupling matrix (right axis), and the total memory of the compressed matrix using the algorithm described in Section IV (left axis). It is evident that the greater the distance between the domain and the coil, the lower the multilinear ranks and the required memory. The compression factor varied between ~ 50 and 190, depending on the distance.

2) *Tucker Rank vs. Frequency:* The work presented in [47] showed that the rank of the integral operators for 3D problems increases linearly with respect to the operating frequency. This was confirmed in [10], for the BTTB defining tensors of the FFT-based VIE systems (discretized integral operators). These tensors were columns of the corresponding Galerkin




Fig. 4. Memory footprint (left) and maximum rank (rank) of the compressed Z_{bc} matrix, for different distances between the loop and the cubic domain.

MoM matrices and modeled the interactions between one voxel's basis function and the whole domain via the \mathcal{N} or \mathcal{K} operators. In the present study, the tensors are columns of the coupling matrix and model the interactions between one RWG basis function and the whole body domain via the same operators. Since in both cases the interactions between separated geometry blocks are modeled, one can expect a similar behavior for the Tucker ranks.

To confirm this, we performed a frequency sweep (300, 600, ..., 2700MHz) for the setup in Fig. 3. The coil was discretized with 125 elements, whereas the voxel's isotropic resolution was set to $\lambda/20$, with λ being the wavelength. We repeated the calculations for three positions of the coil ($d = 0.55, 0.65$, and 0.8 m). The memory footprint (left) of the compressed matrix, along with the maximum rank (right), are shown in the dual axis chart of Fig. 5. The memory footprint increased linearly with frequency, whereas the maximum rank grew at different rates for the three investigated cases. This is expected because the maximum rank represents the worst-case scenario among all ranks, whereas the memory footprint summarizes the overall effect of all ranks.

3) *Tucker Rank vs. Surface Mesh*: Let us consider a fixed mesh for the domain and refine only the surface mesh of the coil. As the coil mesh is refined, the edges of the triangular elements become smaller and the number of columns of the coupling matrix increases, with each column representing more remote element interactions between an edge of the triangular mesh and the voxelized domain. As a result, we should expect a drop in the Tucker ranks. To verify this, we used the same domain of the previous sections, and a PEC equilateral triangle with centroid at $(0, 0.55\text{m}, 0)$ and one vertex at $(0, 0.55\text{m}, 0.5\text{m})$. The triangle was discretized with 10 different meshes, starting from triangular element's edge of 0.5m and reducing it by a factor of $\sqrt{2}$, which resulted in 4, 6, 11, 30, 48, 102, 184, 358, 727, and 1480 elements. Fig. 6 reports the maximum rank as a function of the length of the triangular element's edge, confirming that the rank is smaller




Fig. 5. Memory footprint (left) and maximum rank (rank) of the compressed Z_{bc} matrix, for different operating frequencies. Results are shown for three different distances between the loop and the domain.

when the PEC triangle's mesh is finer.




Fig. 6. Maximum rank of the compressed Z_{bc} matrix, for various PEC triangle's meshes. The rank drops as we refine the mesh.

B. Application to VSIE-based MRI Simulations

Here we aim to validate the performance of the proposed algorithms for the assembly of the coupling matrix Z_{bc} , and the matrix-vector implementation for two VSIE-based MRI applications. Both numerical experiments were implemented in Matlab, except for the matrix assembly part which was written in C++. For the GPU computations, we used an NVIDIA Quadro Volta GV100 32GB HBM2 PCIe. For the CPU computations, in the first experiment we used a server with CentOS 6.9 operating system and an Intel(R) Xeon(R) CPU E5-2699 v3 at 2.30GHz, while for the second experiment we used a server with Ubuntu 18.04.5 LTS operating system and an Intel(R) Xeon(R) Gold 6248 CPU at 2.50GHz. We parallelized on 12 workers where needed.

1) *Head Coil Experiments*: We first demonstrated the proposed compression method for an 8-ports close-fitting head coil, previously designed for GMT [16], which we loaded with the “Billie” realistic head model from the virtual family population [50] (Fig. 7). The operating frequency was set to 298 MHz.




Fig. 7. Coil-head geometry. The RF coil (discretized with 2380 triangular element edges) was loaded with the voxelized realistic human head model “Billie” (discretized with voxels of 1 mm isotropic resolution).

The VSIE-based implementation of GMT requires performing operations on the coupling matrix Z_{bc} and its conjugate transpose. We analyzed the memory footprint reduction for the compressed coupling matrix and measured the computation time for both the matrix- and conjugate transpose matrix-vector products using the algorithms presented in section III. The coil was discretized with both a coarse (516 RWG) and a fine (2380 RWG) mesh resolution. For the VIE domain enclosing the head, we tested three different voxel resolutions, namely 5, 2, and 1 mm³, which resulted in $34 \times 38 \times 45$, $84 \times 96 \times 116$, and $168 \times 188 \times 222$ voxels, respectively. Both PWC (3 unknowns per voxel) and PWL (12 unknown per voxel) VIE basis functions were considered.

We used a tolerance of $1e-8$ for HOSVD, which would ensure accurate estimation of electrical properties in an actual GMT experiment. Since the coil closely fits the head, ACA (or SVD-based methods in general) are expected to provide negligible compression with a tolerance of $1e-8$. We confirmed this for the case of PWC basis functions, 5 mm³ voxel resolution, and fine coil discretization, for which, in fact, we found that 2238 of the 2380 singular values would be needed to accurately represent Z_{bc} , compressing the matrix from 6.18 GB to 6.07 GB. Consequently, for the head coil experiments we did not use the Tucker-based ACA algorithm, but instead we compressed the columns of the coupling matrix only with the HOSVD-based method.

a) *Memory Compression*: The memory footprint for the assembly of the coupling matrix Z_{bc} is shown in TABLE IV. The memory required to assemble the full matrix was considerably larger than for the HOSVD-compressed matrix. For example, for PWC basis functions, voxel resolution of

1 mm³, and fine coil mesh, the required memory in the full matrix case was > 740 GBs, whereas the compressed matrix required only 2.6 GBs. Note that in the challenging case of PWL basis functions, 1 mm³ voxel resolution, and fine coil mesh, it was not feasible to apply our compression method. In fact, the memory requirements even for just one of the q blocks of the matrix (see Section IV. A.) were prohibitively large for our server. While we could have still implemented the HOSVD compression by further dividing the matrix in smaller blocks, that would have required ~ 1 month of computations. An alternative method for such costly cases is mentioned in the discussion and will be pursued in future work.

TABLE IV
MEMORY REQUIREMENTS (GBs) OF Z_{bc}

Voxel Res.	Assembly	PWC-coarse	PWC-fine	PWL-coarse	PWL-fine
5 mm ³	Full	1.34	6.18	5.36	24.74
	HOSVD	0.20	0.88	0.86	3.77
2 mm ³	Full	21.57	99.52	86.30	398.09
	HOSVD	0.40	1.74	1.79	7.49
1 mm ³	Full	161.73	745.99	646.95	2983.99
	HOSVD	0.63	2.62	2.85	N/A

Fig. 8 shows that the compression factor, defined as the memory of the full matrix over the memory of the compressed one, decreased as the voxel resolutions (h) of the VIE domain’s grid became coarser. The behavior of the compression factor was similar for PWC or PWL basis functions, either with fine or coarse mesh discretization. This confirms the excellent stability of our compression method.




Fig. 8. Compression factor of the compressed matrix Z_{bc} . Results are shown for all investigated head and coil discretizations.

Fig. 9 shows the maximum Tucker rank, obtained with HOSVD, for all tensor components of the coupling matrix. The rank decreased slower than the compression factor (Fig.

8) for coarser discretizations of the VIE domain. For example, when PWC basis functions and fine coil resolution were used (PWC-fine), the maximum rank decreased by only 1.5 times (from 42 to 28) when the isotropic voxel size increased from 1 to 5 mm³, which corresponds to a 5 times smaller grid in all directions. For all cases, the maximum rank was smaller for finer coil meshes, which is in agreement with the results shown in section V.A.3.

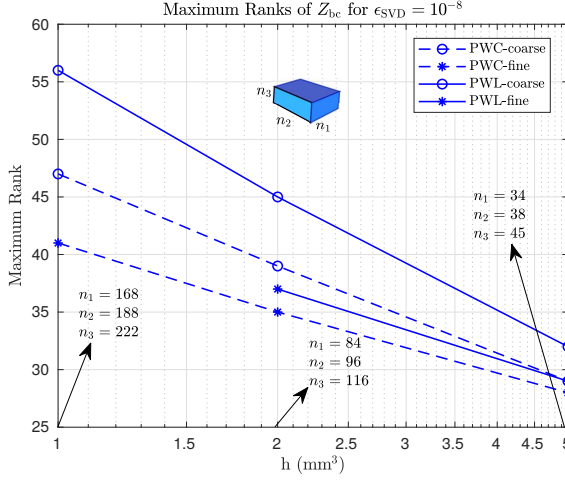


Fig. 9. Maximum rank as a function of voxel resolution of the VIE domain. The maximum rank was calculated among all Tucker ranks of the decomposed tensors, which were obtained with HOSVD. Results are shown for all investigated head and coil discretizations. For each voxel resolution, the size of the corresponding VIE discretization grid is indicated.

b) Computation Time: TABLE V reports the computation time for the assembly of the full and the HOSVD-compressed coupling matrix (rounded to the nearest higher second). For PWC, we used one quadrature integration point per triangle and voxel, while for PWL, two for each triangle and eight for each voxel. For the low resolution matrices, the assembly time for the compressed matrix was larger than the one for the full matrix by a factor $\leq q$, since in the HOSVD case the compressed matrix was assembled as q sequential compressed matrix blocks. For the larger matrices, our server could not perform the assembly of the full matrix, due to the prohibitively large memory requirements, but it was able to assemble the compressed matrix using our proposed method.

TABLE VI and VII summarize the computation times for the matrix- and conjugate transpose matrix-vector products. Compared to the full form case (Full-CPU), the compressed matrix-vector product requires additional operations for the decompression of the tensors. While Algorithms 1 and 2 can reduce the memory requirements of the matrix-vector products, the time footprint varies based on how these algorithms are implemented. In particular, the nested loops over the m RWG functions can be either parallelized on a CPU, if the RAM can support multiple tensor decompositions in parallel (HOSVD-CPU), or performed sequentially using a GPU (HOSVD-GPU). In our tests, we multiplied Z_{bc} with $X \in \mathbb{C}^{m \times 8}$ (TABLE VI) and Z_{bc}^* with $\Phi \in \mathbb{C}^{qn_v \times 8}$ (TABLE VII), where both X and Φ were random matrices to keep the

TABLE V
TIME FOOTPRINT (HH:MM:SS) OF Z_{bc} ASSEMBLY

Voxel Res.	Assembly	PWC-coarse	PWC-fine	PWL-coarse	PWL-fine
5 mm ³	Full	00:00:12	00:00:45	00:03:31	00:14:17
	HOSVD	00:00:34	00:02:33	00:32:27	02:21:21
2 mm ³	Full	00:02:48	N/A	01:09:49	N/A
	HOSVD	00:06:16	00:33:28	07:15:45	30:32:42
1 mm ³	Full	N/A	N/A	N/A	N/A
	HOSVD	00:27:59	02:31:52	52:17:11	N/A

results general. The eight columns of X could correspond, for example, to the currents associated with the eight channels of the coil in Fig. 7.

TABLE VI
TIME FOOTPRINT (HH:MM:SS) OF $Y = Z_{bc}X$

Voxel Res.	Form	PWC-coarse	PWC-fine	PWL-coarse	PWL-fine
5 mm ³	Full-CPU	00:00:01	00:00:02	00:00:02	00:00:06
	HOSVD-CPU	00:00:03	00:00:07	00:00:07	00:00:24
	HOSVD-GPU	00:00:03	00:00:11	00:00:10	00:00:44
2 mm ³	Full-CPU	00:00:06	N/A	00:00:25	N/A
	HOSVD-CPU	00:00:25	00:01:46	00:01:37	00:06:24
	HOSVD-GPU	00:00:04	00:00:14	00:00:13	00:00:56
1 mm ³	Full-CPU	N/A	N/A	N/A	N/A
	HOSVD-CPU	00:02:54	00:11:25	00:11:30	N/A
	HOSVD-GPU	00:00:13	00:00:53	00:00:52	N/A

TABLE VII
TIME FOOTPRINT (HH:MM:SS) OF $\Psi = Z_{bc}^* \Phi$

Voxel Res.	Form	PWC-coarse	PWC-fine	PWL-coarse	PWL-fine
5 mm ³	Full-CPU	00:00:01	00:00:02	00:00:02	00:00:06
	HOSVD-CPU	00:00:02	00:00:04	00:00:04	00:00:26
	HOSVD-GPU	00:00:02	00:00:07	00:00:05	00:00:21
2 mm ³	Full-CPU	00:00:06	N/A	00:00:22	N/A
	HOSVD-CPU	00:00:13	00:00:45	00:00:49	00:03:25
	HOSVD-GPU	00:00:03	00:00:11	00:00:10	00:00:42
1 mm ³	Full-CPU	N/A	N/A	N/A	N/A
	HOSVD-CPU	00:01:20	00:05:19	00:05:23	N/A
	HOSVD-GPU	00:00:11	00:00:45	00:00:41	N/A

For 5 mm³ isotropic voxel resolution, the Full-CPU matrix-vector product was the fastest for all cases, because the coupling matrix is small. For 2 and 1 mm³ voxel resolution, the HOSVD-GPU implementation was the fastest. Note that the Full-CPU case could not be performed for high voxel and coil mesh resolutions, due to the excessive memory requirements. The HOSVD-CPU was slower than HOSVD-GPU, except for the 5 mm³ voxel resolution.

2) *Body Coil Experiments*: For the second MRI experiment, we simulated the volume bodycoil of a commercial 3T MRI scanner [51], [52] and we loaded it with “Billie”, from the virtual family population [50] (Fig. 10). The frequency was set to 123 MHz, corresponding to 3 Tesla MRI. The coil has 32 legs, a radius of 35.5 cm, length of 45 cm, and is centered at (0, 0, 0). We also modeled the system conductive shield, which has a radius of 37.2 cm, a length of 1.5 m and is centered at (0, 0, 0). The distance between the coil and the cuboid VIE domain enclosing “Billie” was 15.5 cm and 24.5 cm in the x and y , respectively. In contrast with the previous case where the coil tightly fitted the head, here the coil is remote enough to allow a good compression of the coupling matrix Z_{bc} with ACA. For this experiment, we used PWC and PWL basis functions and three voxel resolutions (5, 2, and 1 mm³), which corresponded to $81 \times 44 \times 108$, $205 \times 109 \times 270$, and $409 \times 219 \times 541$ voxels for the VIE domain. For the coil and the shield we used 9450 RWG basis functions. Two quadrature integration points were used for each triangle and eight for each voxel, both for PWC and PWL basis functions.




Fig. 10. Coil-body geometry. The RF coil and shield (discretized with 9450 triangular element edges) was loaded with part of the voxelized realistic human body model “Billie” (discretized with voxels of 2 mm isotropic resolution).

a) *Matrix Assembly*: TABLE VIII, summarizes the memory requirements and the assembly time for the coupling matrix Z_{bc} . The ACA tolerance was set to $1e-3$ to achieve good compression. The ACA rank of Z_{bc} was 250 for the 5 mm³ cases and 287 for the 2 and 1 mm³ cases. The maximum Tucker rank of U was between 15 and 18 for all cases. In the 5 mm³ case, our results show that ACA could offer an excellent compression of the coupling matrix and the assembly

could be rapidly performed in CPU. For 2 mm³, ACA’s memory requirements were large and ACA was outperformed in speed by our proposed Algorithm 3 (Tucker-based ACA), for which the low memory footprint allowed using a GPU. For 1 mm³ resolution, the standard ACA algorithm could not be performed even on a server equipped with hundreds of GB’s of RAM, due to overwhelming memory requirements. On the other hand, our proposed ACA extension in Algorithm 3 kept the memory demand small, enabling for fast matrix assembly in GPU. Note that the full matrix assembly was only possible for 5 mm³ voxel resolution and PWC basis functions.

TABLE VIII
MEMORY DEMAND (GB) AND
TIME FOOTPRINT (HH:MM:SS) OF Z_{bc} ASSEMBLY

Voxel Res.	Form	PWC		PWL	
		memory	time	memory	time
5 mm ³	Full-CPU	162.60	00:15:49	650.42	N/A
	ACA-CPU	4.33	00:01:53	17.24	00:06:39
	Algorithm 3-GPU	0.036	00:03:26	0.036	00:09:58
2 mm ³	Full-CPU	2548	N/A	10195	N/A
	ACA-CPU	77.44	00:37:16	309.65	04:30:36
	Algorithm 3-GPU	0.041	00:26:57	0.042	01:28:53
1 mm ³	Full-CPU	20471	N/A	81884	N/A
	ACA-CPU	621.75	N/A	2486	N/A
	Algorithm 3-GPU	0.041	03:35:36	0.042	15:20:41

For the coarser case of 5 mm³ voxel resolution and PWC basis functions, the time footprint of Algorithm 3 for CPU (not shown in the table) was 00:17:50, which is ~ 5 times slower than for the GPU execution.

b) *Matrix-Vector Product Performance*: The time footprints for the matrix-vector product between the compressed coupling matrix Z_{bc} and a random vector $\mathbf{x} \in \mathbb{C}^{m \times 1}$ are shown in TABLE IX. ACA-CPU corresponds to performing the product $U(V^* \mathbf{x})$ in CPU. For ACA+HOSVD-GPU, Z_{bc} was compressed with Algorithm 3, and the matrix-vector product was performed with Algorithm 1 in GPU. For 5 mm³ voxel resolution, the efficiency is similar for both approaches. For 2 mm³ voxel resolution, ACA+HOSVD-GPU outperformed ACA by a factor of 3, because, due to its low memory demand, it could be executed on a GPU, whereas ACA could not.

The relative error of \mathbf{y} obtained with ACA+HOSVD-GPU relative to Full-CPU (ground truth) is shown on the right axis of Fig. 11 for the case of 5 mm³ voxel resolution and PWC basis functions. The plot shows how the error changes as a function of the tolerance ($1e-3, \dots, 1e-8$) used for ACA. In particular, the relative error remained approximately an order of magnitude higher than ACA’s tolerance. Fig. 11 also shows plots for the ACA rank and the maximum Tucker rank of Z_{bc} (values on the left axis). Both ranks increased as the tolerance of ACA was decreased. We expect similar results for the other

TABLE IX
TIME FOOTPRINT (HH:MM:SS) OF $\mathbf{y} = \mathbf{Z}_{bc}\mathbf{x}$

Voxel Res.	Form	PWC	PWL
5 mm ³	Full-CPU	00:00:11	N/A
	ACA-CPU	00:00:01	00:00:01
	ACA+HOSVD-GPU	00:00:01	00:00:02
2 mm ³	Full-CPU	N/A	N/A
	ACA-CPU	00:00:05	00:00:18
	ACA+HOSVD-GPU	00:00:02	00:00:06
1 mm ³	Full-CPU	N/A	N/A
	ACA-CPU	N/A	N/A
	ACA+HOSVD-GPU	00:00:09	00:00:33

cases, but we were unable to assemble the full coupling matrix, due to its vast memory footprint.




Fig. 11. (left axis) ACA rank and maximum Tucker rank (obtained with HOSVD) of \mathbf{Z}_{bc} . (right axis) Error in $\mathbf{y} = \mathbf{Z}_{bc}\mathbf{x}$ calculated with ACA+HOSVD-GPU relative to Full-CPU.

V. DISCUSSION

We showed that our new Tucker-based algorithms could effectively compress the coupling VSIE matrices in the case of fine voxel resolutions. Thanks to the achieved compression factors, the matrix-vector products can be performed in GPU's, yielding faster execution. The proposed approach will be key for a VSIE-based in vivo implementation of GMT, for which reducing memory requirements and computation time are both critical factors.

For cases in which the coil is placed at some distance from the imaging object, the coupling matrix is low-rank, thus can be compressed using ACA as UV^* . For coarse voxel resolutions, such compression strategy alone is effective and allows for a rapid CPU implementation of matrix-vector

products. However, as the voxel grid of the VIE domain is refined, the memory demand of U increases until the standard ACA becomes impractical for applications that need high accuracy, such as GMT. In fact, in order to keep using ACA, one would have to relax the tolerance, sacrificing accuracy. To avoid this, in this work we introduced an extension of ACA (Algorithm 3), for which the memory demand remains as low as the memory required by one column of the coupling matrix for any tolerance. Furthermore, Algorithm 3 can be executed in GPU for rapid computations also in the case of fine voxel resolutions.

An important aspect of the approach presented in this manuscript is that it can effectively compress the coupling matrix both when the coil is close to or far from the scatterer. Because of this, our method allows using GPUs to accelerate the matrix-vector products for most applications. For example, if the close-fitting coil geometry in Fig. 7 were integrated with a head gradient insert with a surrounding conductive shield at a certain distance, with our approach the coupling matrix would still be compressible and fit in the memory of a GPU. In fact, the interactions between the shield and the head would have lower Tucker ranks than the ones between the coil's conductors and the head (see Section IV.A.1). On the other hand, the pFFT+Tucker method would no longer be efficient because it would require extending the VIE domain to fully include the shield. Even though the Green's function tensors of the extended domain would still be compressible with Tucker decomposition, the unknowns that multiply such tensors element-wise would not be. In fact, their dimensions would have to significantly increase in order to perform the relevant FFT-based matrix-vector products.

For the previous example, in order to fully exploit the highly parallel architecture of the GPU with a minimal number of matrix-vector products operations, one could use a hybrid method that combines pFFT and Algorithm 3. To do that, the near and far interactions between the VIE domain and the conducting surfaces would need to be separated. Then, the pFFT could be used to model the near interactions between the VIE domain and the coil as in [33], while the arising Green's function operators could be compressed with the Tucker decomposition as in [10]. Finally, the remaining far interactions between the VIE domain and the coil could be modeled with a coupling matrix, which would be vastly compressed with Algorithm 3. Such hybrid method could enable us to rapidly execute the matrix-vector products in GPU for most cases, including complex coil-shield geometries, fine voxel resolutions, and PWL basis functions. The described hybrid method will be investigated in future work.

The main limitation of our proposed method is that it requires a considerable amount of time for the assembly of the coupling matrix, when the matrix is not compressible with ACA. It is especially slow because each element of the coupling matrix is a 5D integral. We showed that this could be addressed by implementing matrix assembly and compression in parallel, but such approach is not always possible due to memory limitations (IV.B.1). For such cases, one could alternatively employ 3D cross-Tucker approximation algorithms [18], [45], which are less efficient than HOSVD

for the tensor dimensions in this work, but do not suffer from memory constraints in case of large tensors. In fact, 3D cross-Tucker methods require only a small number of rows, columns, and fibers of the tensor they approximate, and they can be implemented with linear complexity (with respect to tensor's linear size). In future work, we will explore the execution of multiple 3D cross-Tucker steps in parallel to avoid memory overflows when assembling a compressed coupling matrix in the case of extremely fine resolutions. Furthermore, the application of the tensor train decomposition [22] on 4D reshapes of the coupling matrix will be investigated for remote geometries like the one appearing in section IV.B.2, as an alternative to Algorithm 3.

VI. CONCLUSION

We presented a memory compression technique for the coupling matrix in VSIE systems. Our method enables one to form and store the coupling matrix even when its full form size is prohibitively large (~ 80 TB). Specifically, in this work we were able to achieve a compression between ~ 0.5 (PWC) and ~ 2 (PWL) million times when simulating interactions between MRI coils and realistic body models with some distance between them, in the case of fine voxel resolutions of 1 mm^3 . The error was around one order of magnitude higher than the tolerance employed for the algorithm. The stored, compressed matrices could be used multiple times without the need to repeat the assembly. For example, this would allow one to rapidly perform EM simulations for the same coil with different body geometries, as far as they are contained in the original computational domain. For most cases, our compression method enables fitting large coupling matrices in GPUs, resulting in rapid execution of the VSIE matrix-vector product (from 1 to 56 seconds for the studied scenarios). Finally, the proposed method could facilitate the implementation of VSIE-based GMT for in vivo mapping of tissue electrical properties at clinically-relevant voxel resolutions.

REFERENCES

- [1] S. W. Anderson *et al.*, "Effect of disease progression on liver apparent diffusion coefficient values in a murine model of NASH at 11.7 Tesla MRI," *Journal of Magnetic Resonance Imaging*, vol. 33, no. 4, pp. 882–888, 2011.
- [2] J. Jin and J. Chen, "On the SAR and field inhomogeneity of birdcage coils loaded with the human head," *Magnetic resonance in medicine*, vol. 38, no. 6, pp. 953–963, 1997.
- [3] R. Lattanzi *et al.*, "Electrodynamic constraints on homogeneity and radiofrequency power deposition in multiple coil excitations," *Magnetic resonance in medicine*, vol. 61, no. 2, pp. 315–334, 2009.
- [4] X. Zhang *et al.*, "From complex B1 mapping to local SAR estimation for human brain MR imaging using multi-channel transceiver coil at 7T," *IEEE transactions on medical imaging*, vol. 32, no. 6, pp. 1058–1067, 2013.
- [5] M. Cosottini *et al.*, "Short-term side-effects of brain MR examination at 7 T: a single-centre experience," *European radiology*, vol. 24, no. 8, pp. 1923–1928, 2014.
- [6] A. Taflov and K. R. Umashankar, "Review of FD-TD numerical modeling of electromagnetic wave scattering and radar cross section," *Proceedings of the IEEE*, vol. 77, no. 5, pp. 682–699, 1989.
- [7] R. Lee and A. C. Cangellaris, "A study of discretization error in the finite element approximation of wave solutions," *IEEE Transactions on Antennas and Propagation*, vol. 40, no. 5, pp. 542–549, 1992.
- [8] J. R. Phillips and J. K. White, "A precorrected-FFT method for electrostatic analysis of complicated 3-D structures," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 16, no. 10, pp. 1059–1072, 1997.
- [9] A. A. Tambova *et al.*, "On the generalization of directfn for singular integrals over quadrilateral patches," *IEEE Transactions on Antennas and Propagation*, vol. 66, no. 1, pp. 304–314, 2017.
- [10] I. I. Giannakopoulos, M. S. Litsarev, and A. G. Polimeridis, "Memory footprint reduction for the fft-based volume integral equation method via tensor decompositions," *IEEE Transactions on Antennas and Propagation*, vol. 67, no. 12, pp. 7476–7486, 2019.
- [11] J. F. Villena *et al.*, "Fast electromagnetic analysis of MRI transmit RF coils based on accelerated integral equation methods," *IEEE Transactions on Biomedical Engineering*, vol. 63, no. 11, pp. 2250–2261, 2016.
- [12] A. G. Polimeridis *et al.*, "Stable FFT-JVIE solvers for fast analysis of highly inhomogeneous dielectric objects," *Journal of Computational Physics*, vol. 269, pp. 280–296, 2014.
- [13] S. Rao, D. Wilton, and A. Glisson, "Electromagnetic scattering by surfaces of arbitrary shape," *IEEE Transactions on antennas and propagation*, vol. 30, no. 3, pp. 409–418, 1982.
- [14] I. P. Georgakis *et al.*, "A fast volume integral equation solver with linear basis functions for the accurate computation of electromagnetic fields in MRI," *IEEE Transactions on Antennas and Propagation, Early Access*, 2020.
- [15] J. E. Serrallés *et al.*, "Noninvasive Estimation of Electrical Properties from Magnetic Resonance Measurements via Global Maxwell Tomography and Match Regularization," *IEEE Transactions on Biomedical Engineering*, vol. 67, no. 1, pp. 3–15, 2019.
- [16] I. Giannakopoulos *et al.*, "Magnetic-resonance-based electrical property mapping using Global Maxwell Tomography with an 8-channel head coil at 7 Tesla: a simulation study," *IEEE Transactions on Biomedical Engineering*, vol. 68, no. 1, pp. 236–246, 2019.
- [17] L. R. Tucker, "Some mathematical notes on three-mode factor analysis," *Psychometrika*, vol. 31, no. 3, pp. 279–311, 1966.
- [18] I. I. Giannakopoulos, M. S. Litsarev, and A. G. Polimeridis, "3D cross-Tucker approximation in FFT-based volume integral equation methods," in *2018 IEEE International Symposium on Antennas and Propagation & USNC/URSI National Radio Science Meeting*. IEEE, 2018, pp. 2507–2508.
- [19] J. Zhang, Y. Han, and J. Jiang, "Tucker decomposition-based tensor learning for human action recognition," *Multimedia Systems*, vol. 22, no. 3, pp. 343–353, 2016.
- [20] M. Wang *et al.*, "VoxCap: FFT-Accelerated and Tucker-Enhanced Capacitance Extraction Simulator for Voxelized Structures," *arXiv preprint arXiv:2004.02609*, 2020.
- [21] C. Qian and A. C. Yucel, "On the Compression of Translation Operator Tensors in FMM-FFT-Accelerated SIE Simulators via Tensor Decompositions," *arXiv preprint arXiv:2010.00520*, 2020.
- [22] I. V. Oseledets, "Tensor-train decomposition," *SIAM Journal on Scientific Computing*, vol. 33, no. 5, pp. 2295–2317, 2011.
- [23] B. N. Khoromskij and I. Oseledets, "Quantics-TT collocation approximation of parameter-dependent and stochastic elliptic PDEs," *Computational Methods in Applied Mathematics Comput. Methods Appl. Math.*, vol. 10, no. 4, pp. 376–394, 2010.
- [24] L. Grasedyck, "Hierarchical singular value decomposition of tensors," *SIAM Journal on Matrix Analysis and Applications*, vol. 31, no. 4, pp. 2029–2054, 2010.
- [25] E. Tyrtshnikov, "Mosaic-skeleton approximations," *Calcolo*, vol. 33, no. 1-2, pp. 47–57, 1996.
- [26] E. Tyrtshnikov, "Mosaic ranks and skeletons," in *International Workshop on Numerical Analysis and Its Applications*. Springer, 1996, pp. 505–516.
- [27] S. A. Goreinov, E. E. Tyrtshnikov, and N. L. Zamarashkin, "A theory of pseudoskeleton approximations," *Linear algebra and its applications*, vol. 261, no. 1-3, pp. 1–21, 1997.
- [28] S. Kurz, O. Rain, and S. Rjasanow, "The adaptive cross-approximation technique for the 3D boundary-element method," *IEEE transactions on Magnetics*, vol. 38, no. 2, pp. 421–424, 2002.
- [29] M. Bebendorf and S. Rjasanow, "Adaptive low-rank approximation of collocation matrices," *Computing*, vol. 70, no. 1, pp. 1–24, 2003.
- [30] L. Greengard and V. Rokhlin, "A fast algorithm for particle simulations," *Journal of computational physics*, vol. 73, no. 2, pp. 325–348, 1987.
- [31] R. Coifman, V. Rokhlin, and S. Wandzura, "The fast multipole method for the wave equation: A pedestrian prescription," *IEEE Antennas and Propagation Magazine*, vol. 35, no. 3, pp. 7–12, 1993.

- [32] B. Shanker and H. Huang, "Accelerated Cartesian expansions—a fast method for computing of potentials of the form $R^{-\nu}$ for all real ν ," *Journal of Computational Physics*, vol. 226, no. 1, pp. 732–753, 2007.
- [33] G. G. Guryev *et al.*, "Fast field analysis for complex coils and metal implants in MARIE 2.0," in *Proc. ISMRM*, 2019, p. 1035.
- [34] Y. Brick and A. Boag, "Multilevel nonuniform grid algorithm for acceleration of integral equation-based solvers for acoustic scattering," *IEEE transactions on ultrasonics, ferroelectrics, and frequency control*, vol. 57, no. 1, pp. 262–273, 2009.
- [35] R. F. Harrington, *Field computation by moment methods*. Wiley-IEEE Press, 1993.
- [36] M. F. Catedra, E. Gago, and L. Nuno, "A numerical scheme to obtain the RCS of three-dimensional bodies of resonant size using the conjugate gradient method and the fast Fourier transform," *IEEE transactions on antennas and propagation*, vol. 37, no. 5, pp. 528–537, 1989.
- [37] P. Zwamborn and P. M. Van Den Berg, "The three dimensional weak form of the conjugate gradient FFT method for solving scattering problems," *IEEE Transactions on Microwave Theory and Techniques*, vol. 40, no. 9, pp. 1757–1766, 1992.
- [38] H. Gan and W. C. Chew, "A discrete BCG-FFT algorithm for solving 3D inhomogeneous scatterer problems," *Journal of Electromagnetic Waves and Applications*, vol. 9, no. 10, pp. 1339–1357, 1995.
- [39] J. Jin *et al.*, "Computation of electromagnetic fields for high-frequency magnetic resonance imaging applications," *Physics in Medicine & Biology*, vol. 41, no. 12, p. 2719, 1996.
- [40] M. Van Beurden and S. Van Eijndhoven, "Well-posedness of domain integral equations for a dielectric object in homogeneous background," *Journal of Engineering Mathematics*, vol. 62, no. 3, pp. 289–302, 2008.
- [41] J. Markkanen *et al.*, "Analysis of volume integral equation formulations for scattering by high-contrast penetrable objects," *IEEE Transactions on Antennas and Propagation*, vol. 60, no. 5, pp. 2367–2374, 2012.
- [42] P. Yla-Oijala *et al.*, "Surface and volume integral equation methods for time-harmonic solutions of Maxwell's equations," *Progress In Electromagnetics Research*, vol. 149, pp. 15–44, 2014.
- [43] C.-T. Tai, *Dyadic Green functions in electromagnetic theory*. Institute of Electrical & Electronics Engineers (IEEE), 1994.
- [44] L. De Lathauwer, B. De Moor, and J. Vandewalle, "A multilinear singular value decomposition," *SIAM journal on Matrix Analysis and Applications*, vol. 21, no. 4, pp. 1253–1278, 2000.
- [45] I. V. Oseledets, D. Savostianov, and E. E. Tyrtshnikov, "Tucker dimensionality reduction of three-dimensional arrays in linear time," *SIAM Journal on Matrix Analysis and Applications*, vol. 30, no. 3, pp. 939–956, 2008.
- [46] S. A. Goreinov and E. E. Tyrtshnikov, "The maximal-volume concept in approximation by low-rank matrices," *Contemporary Mathematics*, vol. 280, pp. 47–52, 2001.
- [47] W. Chai and D. Jiao, "Theoretical study on the rank of integral operators for broadband electromagnetic modeling from static to electrodynamic frequencies," *IEEE Transactions on Components, Packaging and Manufacturing Technology*, vol. 3, no. 12, pp. 2113–2126, 2013.
- [48] M. A. Francavilla *et al.*, "Maxwell parallel imaging," *arXiv preprint arXiv:2008.09042*, 2020.
- [49] K. Zhao, M. N. Vouvakis, and J.-F. Lee, "The adaptive cross approximation algorithm for accelerated method of moments computations of EMC problems," *IEEE transactions on electromagnetic compatibility*, vol. 47, no. 4, pp. 763–773, 2005.
- [50] A. Christ *et al.*, "The Virtual Family-development of surface-based anatomical models of two adults and two children for dosimetric simulations," *Physics in Medicine & Biology*, vol. 55, no. 2, p. N23, 2009.
- [51] SIEMENS Healthineers, "Magnetom skyra." [Online]. Available: <https://www.siemens-healthineers.com/magnetic-resonance-imaging/3t-mri-scanner/magnetom-skyra>
- [52] E. Milshteyn *et al.*, "Individualized SAR calculations using computer vision-based MR segmentation and a fast electromagnetic solver," *Magnetic Resonance in Medicine*, vol. 85, no. 1, pp. 429–443, 2021.