

Adaptive and Parallel Surface Integral Equation Solvers for Very Large-Scale Electromagnetic Modeling and Simulation

Brian MacKie-Mason¹, Andrew Greenwood², and Zhen Peng^{1, *}

(Invited Paper)

Abstract—This work investigates an adaptive, parallel and scalable integral equation solver for very large-scale electromagnetic modeling and simulation. A complicated surface model is decomposed into a collection of components, all of which are discretized independently and concurrently using a discontinuous Galerkin boundary element method. An additive Schwarz domain decomposition method is proposed next for the efficient and robust solution of linear systems resulting from discontinuous Galerkin discretizations. The work leads to a rapidly-convergent integral equation solver that is scalable for large multi-scale objects. Furthermore, it serves as a basis for parallel and scalable computational algorithms to reduce the time complexity via advanced distributed computing systems. Numerical experiments are performed on large computer clusters to characterize the performance of the proposed method. Finally, the capability and benefits of the resulting algorithms are exploited and illustrated through different types of real-world applications on high performance computing systems.

1. INTRODUCTION

Surface integral equation (SIE) methods [1–7] have shown to be effective in solving electromagnetic (EM) radiation and scattering problems. Applications range from advanced antenna design [8, 9], stealth technologies [10, 11], integrated circuits [12] to optics and photonics [13, 14]. One advantage of SIE is that both the analysis and unknowns reside only on the boundary surfaces of the targets. It often requires fewer unknowns to solve when compared to differential equation methods, where the unknowns scale volumetrically. Since surface-based modeling is used, it is easier to prepare analysis-suitable models from Computer Aided Design (CAD) geometries. However, application of SIE methods to Maxwell's Equations leads to a dense, complex and indefinite matrix equation to solve, therefore, robust preconditioning techniques and scalable computational algorithms are needed.

Over the past decades, with the synergy between the advancement of fast algorithms and increasingly capable computing hardwares, many very large-scale EM problems can be simulated in a reasonable timeframe. Fast and parallel algorithms such as the fast multipole method (FMM) [15–18], hybrid FMM and fast Fourier transform (FFT) [19], and parallel adaptive integral method [20, 21] have been developed to accelerate the dense matrix-vector product (MVP). Other significant developments include the parallel higher-order method of moments [22] and fast direct solver of the SIE linear system [23–25]. With these advancements, the solutions with over several hundred millions and a billion unknowns have been possible [17, 18, 26].

Reciprocally, the growing sophistication in real-world EM applications and the increasing demand on enhancements of the model fidelity have driven the need for more advanced mathematical models

Received 30 November 2015, Accepted 22 December 2015, Scheduled 29 December 2015

Invited paper for the Commemorative Collection on the 150-Year Anniversary of Maxwell's Equations.

* Corresponding author: Zhen Peng (pengz@unm.edu).

¹ Department of Electrical and Computer Engineering, University of New Mexico, USA. ² US Air Force Research Laboratory, Kirtland AFB, USA.

and numerical algorithms. One representative EM application considered in this work is the full wave analysis of a modern full-scale aircraft carrier. As illustrated in Fig. 1, this high-definition platform consists of hundreds of EM sub-systems and over 100,000 parts. In the Ku band, it requires more than 20 billion mesh points to discretize the entire platform, and leads to a dense matrix equation with approximately 30 billion degrees of freedom (DoFs) to solve. Another major difficulty comes from the multi-scale nature of the high-definition geometry, in which the spatial scales differ by several orders of magnitude. The coexistence of an electrically large aircraft carrier platform and electrically small fine features reflects into disparate mesh sizes. It may cause the so-called “mixed-frequency” problem [12, 27] and result in an extremely ill-conditioned matrix equation.

Such very large multi-scale EM problems present significant mathematical and computational challenges. They can hardly ever be completed within a reasonable amount of time on a single computer. This puts a high premium on investigation into parallel SIE simulators which are scalable on massively parallel supercomputers. Furthermore, to fully exploit the potential of petaflops and exaflops high performance computing (HPC) systems, all essential elements in the IE modeling and simulation process, including model preparation, mesh generation, preconditioning, solution and post-processing, must be designed with scalable parallelism in mind.

The objective of this research is to investigate an adaptive and parallel SIE solver for the solution of time-harmonic Maxwell’s Equations. Both the simulation capability and modeling fidelity of the proposed solver are expected to scale with the exponential growth in computing power. To realize this objective, advances have been made on three fronts: (i) a geometry-adaptive discontinuous Galerkin boundary element method (DG-BEM), which permits the use of non-conformal surface discretizations and facilitates the mesh generation task for high-definition objects; (ii) a non-overlapping additive Schwarz domain decomposition (DD) method for the iterative solution of the DG-BEM matrix equations, which leads to scalable convergence in the DD iterations; (iii) parallel and adaptive computational algorithms to reduce the time complexity of very large-scale simulations via distributed memory HPC architectures.

The work is shown to be a suitable paradigm not only for the parallel, scalable solution of the SIE matrix equation on advanced HPC computers, but also for new simulation-aided design tools capable of component oriented optimization, discretization and simulation. As depicted in Fig. 1, many EM applications today involve the interaction of multiple bodies and the integration of multiple components. At conceptual design stage, individual components are often modified on a daily basis. The proposed work allows generating analysis-suitable models per-component, analyzing individual components independently, and automating assembly of multiple components to obtain the virtual prototyping of entire product.




Figure 1. A high-definition aircraft carrier.

The remainder of the paper is organized as follows. In Section 2 we introduce the model problem of EM scattering from non-penetrable PEC objects. We then introduce the DG-BEM variational formulation and Schwarz DD preconditioning. In Section 3, we present the major ingredients of the parallel computing methodology. Based on the proposed algorithms, we have developed a hybrid Message Passing Interface (MPI)/OpenMP parallel implementation. Section 4 includes representative numerical experiments to validate the analysis and to illustrate the performance of the proposed method. Concluding remarks are summarized in Section 5.

2. THEORY AND FORMULATION

2.1. Model Problem

We focus on the solution of time-harmonic EM scattering from a finite three-dimensional PEC object Ω with its exterior boundary \mathcal{S} , as illustrated in Fig. 2. The exterior region Ω_{ext} is assumed to be free space. Two surface trace operators on \mathcal{S} are employed in this work. They are the tangential trace operator $\pi_\tau(\mathbf{f}) := \hat{\mathbf{n}} \times (\mathbf{f} \times \hat{\mathbf{n}})|_{\mathcal{S}}$, and the twisted tangential trace operator $\pi_\times(\mathbf{f}) := \hat{\mathbf{n}} \times (\mathbf{f})|_{\mathcal{S}} = \hat{\mathbf{n}} \times \pi_\tau(\mathbf{f})$. The plane wave incident electric and magnetic fields are denoted by \mathbf{E}^{inc} and \mathbf{H}^{inc} , respectively.




Figure 2. EM scattering from a non-penetrable PEC object.

The surface integral equation method is a natural choice for such an EM scattering problem. To do so, we first introduce the electric field boundary potential, \mathcal{L} , and magnetic field boundary potential \mathcal{K} , defined as,

$$\mathcal{L}(\mathbf{f}; \mathcal{S})(\mathbf{r}) := -ik_0 \Psi_A(\mathbf{f}; \mathcal{S})(\mathbf{r}) + \frac{1}{ik_0} \nabla \Psi_F(\nabla_\tau \cdot \mathbf{f}; \mathcal{S})(\mathbf{r}), \quad (1)$$

$$\mathcal{K}(\mathbf{f}; \mathcal{S})(\mathbf{r}) := \nabla \times \Psi_A(\mathbf{f}; \mathcal{S})(\mathbf{r}), \quad (2)$$

where Ψ_A and Ψ_F are the single-layer vector and scalar potential, defined by:

$$\Psi_A(\mathbf{f}; \mathcal{S})(\mathbf{r}) := \int_{\mathcal{S}} \mathbf{f}(\mathbf{r}') G(\mathbf{r}, \mathbf{r}') d\mathbf{r}', \quad (3)$$

$$\Psi_F(\rho; \mathcal{S})(\mathbf{r}) := \int_{\mathcal{S}} \rho(\mathbf{r}') G(\mathbf{r}, \mathbf{r}') d\mathbf{r}', \quad (4)$$

and $G(\mathbf{r}, \mathbf{r}') := \frac{\exp^{-ik_0|\mathbf{r}-\mathbf{r}'|}}{4\pi|\mathbf{r}-\mathbf{r}'|}$ is the free-space Green's function.

We proceed to introduce an auxiliary variable on \mathcal{S} , $\mathbf{j} = \frac{1}{-ik_0} \pi_\times(\nabla \times \mathbf{E})$, which represents the scaled surface electric current. To avoid internal resonance solutions, we herein employ a combined field integral equation (CFIE) formulation, which can be written as:

$$\frac{\alpha}{2} \mathbf{j} - \alpha \pi_\times(\bar{\mathcal{K}}(\mathbf{j}; \mathcal{S})) - (1 - \alpha) \pi_\tau(\mathcal{L}(\mathbf{j}; \mathcal{S})) = \alpha \mathbf{j}^{inc} + (1 - \alpha) \mathbf{e}^{inc} \text{ on } \mathcal{S}, \quad (5)$$

where $\bar{\mathcal{K}}$ stands for the *principle value* of the magnetic field boundary potential, \mathcal{K} , on \mathcal{S} , and $\mathbf{j}^{inc} = \eta_0 \pi_\times(\mathbf{H}^{inc})$ and $\mathbf{e}^{inc} = \pi_\tau(\mathbf{E}^{inc})$. The coupling parameter α is chose to be $\frac{1}{2}$ in the following study.

2.2. Discontinuous Galerkin Formulation

Equation (5) is often solved via the Galerkin method, which is based on a variational formulation in suitable trial and testing function spaces. Therefore, conforming boundary element spaces defined on a conformal discretization of the target's surface are commonly required. However, generating conformal discretizations for engineering system-level simulations is far from trivial, as the complexity of modern engineering applications increases at a fast pace. Among the previous works addressing the above mentioned deficiencies, we mention recent works [28–32].

In this work, we investigate an interior penalty discontinuous Galerkin boundary element method, which allows the solution of CFIE (5) using square-integrable, \mathbf{L}^2 , trial and testing functions spaces. Due to the local characteristics of \mathbf{L}^2 vector functions, it admits discontinuities of the elements in the discrete space and allows the use of non-conformal surface discretizations of the targets. In this way, the troublesome mesh generation task of complex, multi-scale targets can be facilitated dramatically. Initial work has been done to demonstrate the potential of DG-BEM in solving the EM wave scattering from PEC objects in [33]. We have shown that the DG-BEM provides the same order of accuracy and convergence behavior compared to that of the traditional Galerkin CFIE method, which uses standard $\mathbf{H}(\text{div})$ -conforming boundary element spaces. Moreover, the stability and convergence have been validated for targets with sharp edges and corners. Along the line of this research, we herein employ a non-symmetric interior penalty formulation [34]. Although it results in non-symmetric coupling matrices after the finite dimensional discretization, it enjoys the remarkable property of discrete ellipticity of the weak formulation and well-posedness of the discrete formulation. The consistency of the discrete weak formulation is illustrated in this section.

Let us consider a simplicial mesh $\mathcal{S}^h = \bigcup_{m=1}^M \{\mathcal{T}_m\}$ of \mathcal{S} , i.e., a tessellation of \mathcal{S} made of M triangles. In the DG-BEM formulation, the triangles are neither required to be conformal nor matching. Furthermore, for any $\mathcal{T}_m \in \mathcal{S}^h$ we define the boundary of the triangle to be \mathcal{C}_m and the in-plane unit normal $\hat{\mathbf{t}}_m$ outer to the boundary \mathcal{C}_m . The approximation of the current \mathbf{j}^h can be written as:

$$\mathbf{j}^h(\mathbf{r}) = \bigoplus_{m=1}^M \mathbf{j}_m^h(\mathbf{r}) \quad (6)$$

where $\mathbf{j}_m^h \in \mathbf{X}_m^h$ is the local approximation of the surface current within each element, and \mathbf{X}_m^h is taken as the space spanned by the vector basis functions introduced in [35] with the DOFs defined independently for each triangle. Moreover, let \mathcal{T}_m and \mathcal{T}_n to be two adjacent triangles sharing the contour boundary \mathcal{C}_{mn} . The jump of the surface electric current across adjacent elements is denoted by

$$\llbracket \mathbf{j}^h \rrbracket_{mn} := \hat{\mathbf{t}}_m \cdot \mathbf{j}_m^h + \hat{\mathbf{t}}_n \cdot \mathbf{j}_n^h \quad \text{on } \mathcal{C}_{mn}. \quad (7)$$

We define the following bilinear form on $\mathbf{X}^h \times \mathbf{X}^h$:

$$\begin{aligned} a_h(\mathbf{v}^h, \mathbf{j}^h) &:= \frac{ik_0}{2} \sum_{m=1}^M \left\langle \mathbf{v}_m^h, \sum_{n=1}^M \Psi_A(\mathbf{j}_n^h; \mathcal{T}_n) \right\rangle_{\mathcal{T}_m} + \frac{1}{2ik_0} \sum_{m=1}^M \left\langle \nabla_\tau \cdot \mathbf{v}_m^h, \sum_{n=1}^M \Psi_F(\nabla_\tau \cdot \mathbf{j}_n^h; \mathcal{T}_n) \right\rangle_{\mathcal{T}_m} \\ &+ \frac{1}{4ik_0} \sum_{\mathcal{C}_{mn} \in \mathcal{C}} \left\langle \llbracket \mathbf{j}^h \rrbracket_{mn}, \sum_{n=1}^M \Psi_F(\nabla_\tau \cdot \mathbf{v}_n^h; \mathcal{T}_n) \right\rangle_{\mathcal{C}_{mn}} - \frac{1}{4ik_0} \sum_{\mathcal{C}_{mn} \in \mathcal{C}} \left\langle \llbracket \mathbf{v}^h \rrbracket_{mn}, \sum_{n=1}^M \Psi_F(\nabla_\tau \cdot \mathbf{j}_n^h; \mathcal{T}_n) \right\rangle_{\mathcal{C}_{mn}} \\ &+ \frac{\beta}{2k_0} \sum_{\mathcal{C}_{mn} \in \mathcal{C}} \left\langle \llbracket \mathbf{v}^h \rrbracket_{mn}, \llbracket \mathbf{j}^h \rrbracket_{mn} \right\rangle_{\mathcal{C}_{mn}} + \frac{1}{4} \sum_{m=1}^M \left\langle \mathbf{v}_m^h, \mathbf{j}_m^h \right\rangle_{\mathcal{T}_m} + \frac{1}{2} \sum_{m=1}^M \left\langle \pi_\times(\mathbf{v}_m^h), \sum_{n=1}^M \bar{\mathcal{K}}(\mathbf{j}_n^h; \mathcal{T}_n) \right\rangle_{\mathcal{T}_m} \end{aligned}$$

where the stabilization parameter [33] is chosen to be $\beta = |\log \bar{h}|/10$, where \bar{h} is the average element size over the entire discretization.

The variational weak formulation of CFIE (5) using discontinuous Galerkin discretizations can be formally stated as: Seek $\mathbf{j}^h = \bigoplus_{m=1}^M \mathbf{j}_m^h$, $\mathbf{j}_m^h \in \mathbf{X}_m^h$ such that

$$a_h(\mathbf{v}^h, \mathbf{j}^h) = \frac{1}{2} \langle \mathbf{v}^h, \mathbf{e}^{inc} \rangle_{\mathcal{S}^h} + \frac{1}{2} \langle \mathbf{v}^h, \mathbf{j}^{inc} \rangle_{\mathcal{S}^h} \quad (8)$$

$\forall \mathbf{v}^h = \bigoplus_{m=1}^M \mathbf{v}_m^h$, $\mathbf{v}_m^h \in \mathbf{X}_m^h$. After expanding the current in terms of the basis functions, the discrete system can then be cast as a matrix equation as: $\mathbf{A}\mathbf{x} = \mathbf{b}$, where \mathbf{x} is the unknown coefficient vector and \mathbf{b} is the right hand side vector.

We proceed to show that the DG-BEM weak formulation is consistent, i.e., the exact solution \mathbf{j} of CFIE (5) solves the finite dimensional weak form:

$$a_h(\mathbf{v}^h, \mathbf{j}) = \frac{1}{2} \langle \mathbf{v}^h, \mathbf{e}^{inc} \rangle_{S^h} + \frac{1}{2} \langle \mathbf{v}^h, \mathbf{j}^{inc} \rangle_{S^h} \quad \forall \mathbf{v}^h \in \mathbf{X}^h \quad (9)$$

To prove this statement, we first recognize that, for exact solution \mathbf{j} , the jump of the surface electric current is equal to zero, $[\mathbf{j}]_{mn} = 0$. Thus, we obtain that for $\mathbf{v}^h \in \mathbf{X}^h$

$$\begin{aligned} a_h(\mathbf{v}^h, \mathbf{j}) &:= \frac{ik_0}{2} \sum_{m=1}^M \left\langle \mathbf{v}_m^h, \sum_{n=1}^M \Psi_A(\mathbf{j}_n; \mathcal{T}_n) \right\rangle_{\mathcal{T}_m} + \frac{1}{2ik_0} \sum_{m=1}^M \left\langle \nabla_\tau \cdot \mathbf{v}_m^h, \sum_{n=1}^M \Psi_F(\nabla_\tau \cdot \mathbf{j}_n; \mathcal{T}_n) \right\rangle_{\mathcal{T}_m} \\ &- \frac{1}{4ik_0} \sum_{C_{mn} \in \mathcal{C}} \left\langle \llbracket \mathbf{v}^h \rrbracket_{mn}, \sum_{n=1}^M \Psi_F(\nabla_\tau \cdot \mathbf{j}_n; \mathcal{T}_n) \right\rangle_{C_{mn}} + \frac{1}{4} \sum_{m=1}^M \langle \mathbf{v}_m^h, \mathbf{j}_m \rangle_{\mathcal{T}_m} + \frac{1}{2} \sum_{m=1}^M \left\langle \pi_\times \mathbf{v}_m^h, \sum_{n=1}^M \bar{\mathcal{K}}(\mathbf{j}_n; \mathcal{T}_n) \right\rangle_{\mathcal{T}_m} \end{aligned}$$

Next, we recall the definition of jump operator, and the boundary integral term involving discontinuous testing functions, $\frac{-1}{4ik_0} \sum_{C_{mn} \in \mathcal{C}} \langle \llbracket \mathbf{v}^h \rrbracket_{mn}, \sum_{n=1}^M \Psi_F(\nabla_\tau \cdot \mathbf{j}_n; \mathcal{T}_n) \rangle_{C_{mn}}$, can be rewritten as:

$$-\frac{1}{4ik_0} \sum_{C_{mn} \in \mathcal{C}} \left\langle \llbracket \mathbf{v}^h \rrbracket_{mn}, \sum_{n=1}^M \Psi_F(\nabla_\tau \cdot \mathbf{j}_n; \mathcal{T}_n) \right\rangle_{C_{mn}} = -\frac{1}{2ik_0} \sum_{n=1}^M \left\langle \mathbf{v}_m \cdot \hat{\mathbf{t}}_m, \sum_{n=1}^M \Psi_F(\nabla_\tau \cdot \mathbf{j}_n; \mathcal{S}_n) \right\rangle_{C_m} \quad (10)$$

After applying a surface Green's formula, we arrive at the following expression:

$$\begin{aligned} a_h(\mathbf{v}^h, \mathbf{j}) &= \frac{ik_0}{2} \sum_{m=1}^M \left\langle \mathbf{v}_m^h, \sum_{n=1}^M \Psi_A(\mathbf{j}_n; \mathcal{T}_n) \right\rangle_{\mathcal{T}_m} - \frac{1}{2ik_0} \sum_{m=1}^M \left\langle \mathbf{v}_m^h, \sum_{n=1}^M \nabla_\tau \Psi_F(\nabla_\tau \cdot \mathbf{j}_n; \mathcal{T}_n) \right\rangle_{\mathcal{T}_m} \\ &+ \frac{1}{4} \sum_{m=1}^M \langle \mathbf{v}_m^h, \mathbf{j}_m \rangle_{\mathcal{T}_m} + \frac{1}{2} \sum_{m=1}^M \left\langle \pi_\times(\mathbf{v}_m^h), \sum_{n=1}^M \bar{\mathcal{K}}(\mathbf{j}_n; \mathcal{T}_n) \right\rangle_{\mathcal{T}_m} \\ &= \sum_{m=1}^M \left\langle \mathbf{v}_m^h, \frac{1}{4} \mathbf{j} \right\rangle_{\mathcal{T}_m} - \sum_{m=1}^M \left\langle \mathbf{v}_m^h, \frac{1}{2} \pi_\times(\bar{\mathcal{K}}(\mathbf{j}; \mathcal{S})) \right\rangle_{\mathcal{T}_m} - \sum_{m=1}^M \left\langle \mathbf{v}_m^h, \frac{1}{2} \pi_\tau(\mathcal{L}(\mathbf{j}; \mathcal{S})) \right\rangle_{\mathcal{T}_m} \\ &= \sum_{m=1}^M \left\langle \mathbf{v}_m^h, \frac{1}{2} \mathbf{j}^{inc} \right\rangle_{\mathcal{T}_m} + \sum_{m=1}^M \left\langle \mathbf{v}_m^h, \frac{1}{2} \mathbf{e}^{inc} \right\rangle_{\mathcal{T}_m} \end{aligned}$$

Note that we have used Eq. (5) in the final step of derivation. The above proves the consistency of the DG-BEM weak formulation. Furthermore, since the DG-BEM solution $\mathbf{j}^h \in \mathbf{X}^h$ also solve the discrete weak formulation, we have:

$$a_h(\mathbf{v}^h, \mathbf{j}^h) = \langle \mathbf{v}^h, \mathbf{e}^{inc} \rangle_{\mathcal{K}^h} \quad \forall \mathbf{v}^h \in \mathbf{X}^h \quad (11)$$

Subtracting (9) and (11) and making use of the linearity in the bilinear form, we arrive at:

$$a_h(\mathbf{v}^h, \mathbf{j} - \mathbf{j}^h) = 0 \quad \forall \mathbf{v}^h \in \mathbf{X}^h \quad (12)$$

Namely, the projection of the solution error, $\mathbf{r} := \mathbf{j} - \mathbf{j}^h$, onto the finite dimensional subspace \mathbf{X}^h is zero with respect to the energy inner product defined by the bilinear form $a_h(\cdot, \cdot)$. Eq. (12), also known as Galerkin orthogonality property, states that \mathbf{j}^h is the best approximation of \mathbf{j} in \mathbf{X}^h with respect to the energy inner product, as the error is orthogonal to the space we are seeking for the solution. Therefore, the DG-BEM discrete formulation will be converging, and the solution error decreases with increasingly bigger finite dimensional subspaces.

2.3. Domain Decomposition Solver

This section investigates efficient and robust solution techniques for the linear system resulting from DG-BEM discretizations. For many problems of practical interest, the number of DOFs, and therefore the dimension of the matrix equation to be solved, is too large to permit the use of direct solution techniques. Instead, one usually turns to preconditioned Krylov subspace iterative methods. When such iterative solutions are employed, the efficient and parallel computation of effective preconditioners poses an immense challenge [6, 12, 36–42].

This work proposes a one-level non-overlapping additive Schwarz DD preconditioner [43] for the solution of the DG-BEM linear system equation, $\mathbf{Ax} = \mathbf{b}$. The main advantages of the method are twofold: (i) it results in a highly efficient and naturally parallelizable computational algorithm on distributed memory many-core parallel computing architectures; (ii) it enables the design of SIE based “divide-and-conquer” scheme in which appropriate SIE solvers can be seamlessly integrated to solve individual subdomain problems.

To apply the proposed method, we first partition the triangular mesh \mathcal{S}^h into N subdomains, \mathcal{S}_i^h , $i = 1, \dots, N$. The sub-meshes associated to i th subdomain are denoted by $\mathcal{S}_i^h = \bigcup_{m=1}^{M_i} \{\mathcal{T}_m^i\}$. The matrix equation for the decomposed problem is given by:

$$\begin{bmatrix} \mathbf{A}_1 & \mathbf{C}_{12} & \dots & \mathbf{C}_{1N} \\ \mathbf{C}_{21} & \mathbf{A}_2 & \dots & \mathbf{C}_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{C}_{N1} & \mathbf{C}_{N2} & \dots & \mathbf{A}_N \end{bmatrix} \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \\ \vdots \\ \mathbf{x}_N \end{bmatrix} = \begin{bmatrix} \mathbf{b}_1 \\ \mathbf{b}_2 \\ \vdots \\ \mathbf{b}_N \end{bmatrix} \quad (13)$$

In (13), \mathbf{A}_i denotes the subdomain CFIE impedance matrix on \mathcal{S}_i^h , and \mathbf{C}_{ij} is the coupling matrix between subdomains \mathcal{S}_i^h and \mathcal{S}_j^h . We note that these subdomains are obtained by a direct partitioning of the object’s surface without the introduction of artificial interfaces [11, 44] or auxiliary unknowns [45, 46]. We proceed to introduce \mathbf{R}_i , a rectangular restriction matrix that returns the vector of coefficients in the subdomain \mathcal{S}_i^h , i.e., $\mathbf{x}_i = \mathbf{R}_i \mathbf{x}$. The subdomain matrix \mathbf{A}_i can be written as:

$$\mathbf{A}_i = \mathbf{R}_i \mathbf{A} \mathbf{R}_i^T \quad (14)$$

The one-level additive Schwarz preconditioner \mathbf{P}^{-1} can then be expressed as the following:

$$\mathbf{P}^{-1} = \sum_{i=1}^N \mathbf{R}_i \mathbf{A}_i^{-1} \mathbf{R}_i^T \quad (15)$$

The matrix Eq. (13) can then be solved with a right preconditioned form, $(\mathbf{A}\mathbf{P}^{-1})\mathbf{P}\mathbf{x} = \mathbf{b}$. Explicitly, the preconditioned matrix equation can be formulated as:

$$\begin{bmatrix} \mathbf{I} & \mathbf{C}_{12}\mathbf{A}_2^{-1} & \dots & \mathbf{C}_{1N}\mathbf{A}_N^{-1} \\ \mathbf{C}_{21}\mathbf{A}_1^{-1} & \mathbf{I} & \dots & \mathbf{C}_{2N}\mathbf{A}_N^{-1} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{C}_{N1}\mathbf{A}_1^{-1} & \mathbf{C}_{N2}\mathbf{A}_2^{-1} & \dots & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{u}_1 \\ \mathbf{u}_2 \\ \vdots \\ \mathbf{u}_N \end{bmatrix} = \begin{bmatrix} \mathbf{b}_1 \\ \mathbf{b}_2 \\ \vdots \\ \mathbf{b}_N \end{bmatrix} \quad (16)$$

where $\mathbf{u}_i = \mathbf{A}_i \mathbf{x}_i$. Thus, the system matrix equation is solved in two steps. The first step is to solve the preconditioned matrix equation Eq. (16) by a Krylov subspace iterative method, in which the residual norm, $\|\mathbf{A}\mathbf{u} - \mathbf{b}\|_2$, is the same as the original matrix equation. Once the unknown vector \mathbf{u}_i is computed, the solution for each subdomain can be recovered via $\mathbf{x}_i = \mathbf{A}_i^{-1} \mathbf{u}_i$.

The proposed work can be viewed as an effective preconditioning scheme that reduces the condition number of very large systems of equations. Numerical experiments illustrate that the iteration count of preconditioned systems depends logarithmically on the number of sub-domains and the electrical size of the object. Therefore, the method is very effective in solving very large-scale EM engineering applications. Furthermore, it is shown to be a suitable computing paradigm not only for the scalable parallel solution of the IE matrix equation on advanced computing architectures, but also for a flexible hybrid solution strategy where fast direct solvers and fast multipole methods based iterative solvers can be applied to individual sub-domain problems.

3. PARALLELIZATION STRATEGY

The mathematical advancements of the proposed work enable an adaptive, parallel and scalable SIE-based analysis framework well-suited for advanced distributed computing systems. The basic units of the parallel analysis framework consist of pre-processing, parallel computing and post-processing, all of which are formulated with proposed mathematical ingredients towards scalable parallelism. This section provides an overview of the parallel implementation of the proposed algorithms.

3.1. Pre-processing

3.1.1. Geometry-Based Domain Decomposition

The simulation flowchart begins with a decomposition of the CAD-generated surface model into a collection of components based on simulation-aided design aspects. Each component is a surface module and considered as one sub-domain, and all components in the modeling and simulation processes are interchangeable. Subsequently, the virtual prototyping of the entire product is obtained by automating assembly of multiple components. As a representative example, the high-definition aircraft carrier shown in Fig. 1 is first decomposed into three simulation modules, (i) aircrafts on deck, (ii) an integrated island, and (iii) the flight deck and ship body. Those modules are further decomposed into a collection of components directly based on the design description embodied in the CAD models. Each component has its own geometric representation and represents one subdomain in the simulation framework. The resulting geometry-based decomposition is depicted in Fig. 3.




Figure 3. Geometry-based domain decomposition of the aircraft carrier.

3.1.2. Parallel Geometry-Adaptive Mesh Generation

Next, individual components are discretized independently and concurrently into a tessellation of triangular elements. The largest discretization size employed in the computation is determined by the wavelength of the operating frequency. Fine surface discretizations are generated locally to accurately represent geometrically complex regions. Each triangulated subdomain contains its own collection of triangles, edges, and vertices. As a result, it enables a trivially parallel mesh generation and allows engineers to rapidly generate high-fidelity models of complex geometries. Moreover, since every component and associated discretization are interchangeable, it provides unprecedented flexibility and convenience for simulation-based design and parameter studies, i.e., antenna in-situ performance characterization, antenna co-site interference analysis, etc.

3.1.3. Algorithmic Graph Partitioning

The goal of this part is to develop an automatic spatial decomposition that will balance the computation among the processors in the HPC system. To do so, we first employ an algorithmic graph partitioning algorithm, METIS [47], to partition the DG discretizations associated with each component into modular work units based on an a-priori complexity analysis and HPC architectural parameters. Subsequently, work units are then aggregated into balanced computational subdomains based on the number of processors available and the local memory each processor can access. Taking the aircraft carrier shown in Fig. 3 for example, the subdomains obtained from CAD geometry-based partitioning are highly unbalanced for parallel computation. After applying the load balancing strategy, we show in Fig. 4 the new balanced computational partitioning. The proposed method provides a much-needed adaptive load balancing strategy for massively parallel HPC computers.




Figure 4. Load-balanced computational partitioning.

3.2. Parallel Computing

To fully exploit the recent success of multi-core processors and massively parallel distributed memory supercomputers, we have considered a hybrid MPI/OpenMP implementation of proposed algorithms. The computation in solving the preconditioned DD matrix equation Eq. (16) consists of three parts: (i) a parallel preconditioned Krylov subspace iterative method; (ii) the application of the additive Schwarz preconditioner involving local subdomain solutions; and (iii) the EM *radiation coupling* among subdomains corresponding to off-diagonal coupling matrices, \mathbf{C}_{ij} , in Eq. (16).

3.2.1. Parallel Krylov Iterative Method

Since the system matrix equation is non-symmetric due to the use of CFIE formulation, we have employed the truncated GCR method [48] for the solution of Eq. (16). The objective of the GCR method is to obtain an approximation \mathbf{u}^n to the solution \mathbf{u} in the Krylov subspace $\mathcal{K}^n((\mathbf{A}\mathbf{P}^{-1}), \mathbf{r}^0)$, where the Krylov subspace is defined as $\mathcal{K}^n(\mathbf{B}, \mathbf{y}) := \text{Span}\{\mathbf{y}, \mathbf{B}\mathbf{y}, \dots, \mathbf{B}^{n-1}\mathbf{y}\}$ and the initial preconditioned residual is denoted by $\mathbf{r}^0 = (\mathbf{b} - \mathbf{A}\mathbf{P}^{-1}\mathbf{u}^0)$. The GCR method can be easily extended to the proposed work with M subdomains. In a distributed memory MPI programming model, each subdomain Krylov subspace vector will be stored only in the local memory associated with individual MPI processes. The main algorithms are summarized in the following.

3.2.2. Task-based Parallelism of Subdomain Solutions

As illustrated in the previous section, the application of the additive Schwarz preconditioner, \mathbf{P}^{-1} , requires the solution of individual subdomain problems. One advantage of the proposed preconditioner

Algorithm 1 Computation of matrix-vector multiplication $\mathbf{r} = (\mathbf{A}\mathbf{P}^{-1})\mathbf{v}$

```

1: Subdomain solution  $\mathbf{t}_i = \mathbf{A}_i^{-1}\mathbf{v}_i, \forall i = 1, M$ 
2: for  $i = 1, M$  do
3:    $\mathbf{r}_i = \sum \mathbf{C}_{ij}\mathbf{t}_j, \quad \forall j = 1, M$  and  $j \neq i$ 
4:    $\mathbf{r}_i = \mathbf{v}_i + \mathbf{r}_i$ 
5: end for

```

Algorithm 2 Computation of solution vector $\mathbf{x} = \mathbf{P}^{-1}\mathbf{u}$

```

1: Initialize  $\mathbf{x}_i = 1, \forall i = 1, M$ 
2: for  $i = 1, M$  do
3:    $\mathbf{x}_i = \mathbf{A}_i^{-1}\mathbf{u}_i$ 
4: end for

```

is the ability to solve all the subdomain problems simultaneously in each DD iteration. Each subdomain is allowed to choose its own appropriate subdomain solver based on local EM wave characteristics and geometrical features. In the parallel implementation, we have employed a task-based parallelism for subdomain solutions. Namely, those subdomain solutions are considered as independent tasks in an MPI programming model. Individual MPI processes will execute different tasks simultaneously. Furthermore, we have developed a queue-based task balancing strategy. The time spent for individual tasks in the previous DD iteration is measured, and the tasks are sorted and recorded into a task queue. Each MPI process will be assigned with a group of tasks based on the timing data, leading to a dynamic load balancing environment.

3.2.3. Radiation Coupling among Subdomains

The final part in the parallel computing is the radiation coupling among multiple subdomains. It requires that the surface currents in each independent subdomain be radiated to all other subdomains. The computation is accomplished with two mathematical ingredients: (i) a hierarchical multi-level fast multipole method (H-MLFMM) [49, 50], which leads to the seamless integration of multi-level skeletonization technique [51] into the FMM framework and results in an effective matrix compression for non-uniform DG discretizations; (ii) a primal-dual octree partitioning algorithm for separable subdomain coupling [52]. Namely, instead of partitioning the entire computational domain into a single octree as in the traditional FMM, we have created independent octrees for all subdomains. Those octrees are allowed to be overlapping or intersecting. We have again employed a task-based parallelism for the radiation coupling. To illustrate, we consider the case of M subdomains and M distributed computing nodes. At the begin of the radiation coupling, we use MPI to broadcast the input data, $\mathbf{t}_i, \forall i = 1, M$, among all computing nodes. Individual subdomains proceed to calculate the MLFMM radiation pattern simultaneously and independently using their own octrees. Next, the interactions between subdomains are also distributed among MPI processes. For example, the i th MPI process will execute the task: $\mathbf{r}_i = \sum \mathbf{C}_{ij}\mathbf{t}_j, \forall j = 1, M$; and $j \neq i$. The skeletonized near-field and FMM based far-field compressions are applied to \mathbf{C}_{ij} with the help of a primal tree partitioning. We note that individual tasks are performed without any communication or synchronization among the computing nodes. The parallelization within each task is attained using OpenMP, which exploits the fast memory access in the shared-memory multi-core processors. Finally, we remark that, for a very large number of subdomains, a highly scalable parallel performance in radiation coupling can be achieved by a hybrid MLFMA-FFT algorithm, which has been demonstrated in [53, 54].

3.3. Post-Processing

Once the simulation is completed, the surface electric current for each subdomain will be calculated and translated into separate Silo files. Silo is a popular mesh and field I/O library and scientific database developed by Lawrence Livermore National Laboratory. In order to process large-scale Silo data files efficiently, we have employed the Visit visualization tool [55], which is a distributed, parallel visualization

and graphical analysis tool. In addition, an embarrassingly parallel far-field calculator has been written to calculate the scattered far-field from the surface electric currents.

4. NUMERICAL EXPERIMENTS

In this section, we study the performance of the proposed work via numerical experiments. We first present the convergence analysis and study of the scalability of the proposed method. It is followed by parallel strong scaling and weak scaling experiments. Finally, several examples of practical interest are included to demonstrate the capability of the proposed method.

4.1. Numerical Study

4.1.1. Eigenspectrum

We first examine the eigenvalue distribution and condition number of the preconditioned system matrix (16). To do so, we use a PEC spherical target with radius 0.5 m, and the eigenspectrum is examined at two difference frequencies: 180 MHz and 240 MHz. The discretization size is kept fixed at $h = \lambda_0/10$. The resulting triangular meshes at two frequencies are partitioned into 2 subdomains and 3 subdomains, respectively. As depicted in Fig. 5, the subdomains are formed by a direct partitioning of the triangular mesh, which leads to jagged boundaries between adjacent subdomains. The eigenvalue distributions for the preconditioned DD matrix are shown in Fig. 5. We notice that all the eigenvalues are within the shifted unit circle. Moreover, they are well separated from the origin. The results indicate the effectiveness of the additive Schwarz preconditioner. Thus, we expect a rapid convergence of Krylov solvers for the preconditioned matrix equation, see [56, 57].

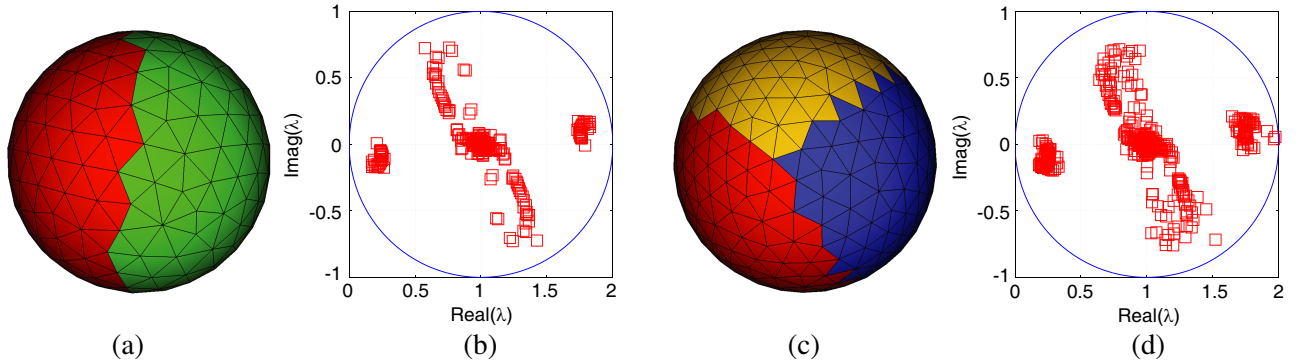


Figure 5. Eigenspectrum for a PEC sphere with irregular subdomain boundaries. (a) Domain partitioning, $N = 2$. (b) Eigenspectrum, $f = 180$ MHz. (c) Domain partitioning, $N = 3$. (d) Eigenspectrum, $f = 240$ MHz.

4.1.2. Convergence Study

Next, we investigate the convergence in iterative matrix solutions of Eq. (16) with respect to the problem size. Namely, we keep the size of each subdomain constant and increase the number of subdomains for higher operating frequencies. The studies performed here demonstrate the potential of the proposed method to solve very large-scale problems. A PEC cube of dimensions $1 \text{ m} \times 1 \text{ m} \times 1 \text{ m}$ is used and the performance of the method is evaluated at five different frequencies, 0.6 GHz, 1.2 GHz, 2.4 GHz, 4.8 GHz and 9.6 GHz. The electrical size of the cube increase from $1\lambda_0$ to $16\lambda_0$, and the number of subdomain increases from 2 to 512, correspondingly. The mesh size is fixed at $h = \lambda_0/10$. The iterative solver convergences for various simulations are plotted in Fig. 6. We observe that the convergence behavior is quite insensitive to the increase in the number of subdomains. Regarding to the cases of 2 and 512 subdomains, the iteration counts increase from 4 to 7 to reach a relative residual 10^{-2} , and from 24 to 35 to reach a relative residual 10^{-6} . Fig. 7 gives electric current distributions at 9.6 GHz using the




Figure 6. Iterative solver convergence with respect to operating frequency.

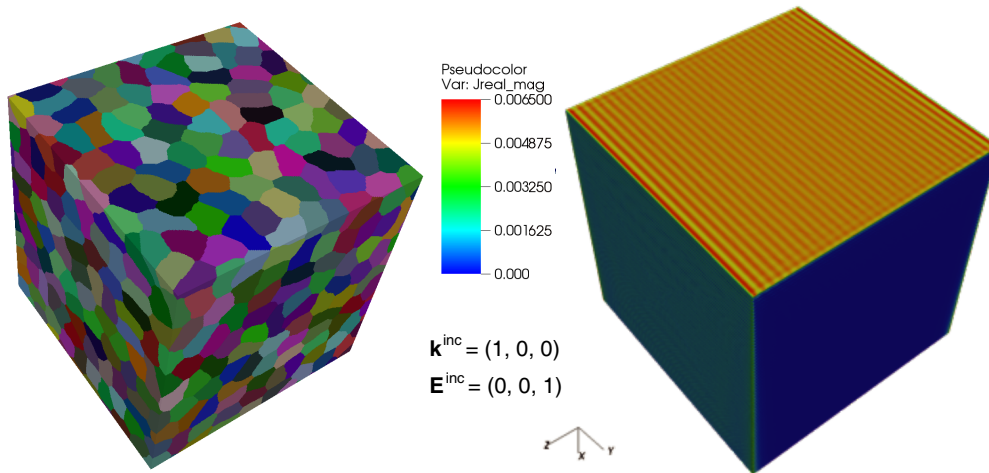


Figure 7. Electric current on a PEC cube at 9.6 GHz using 512 subdomains.

proposed work method with 512 subdomains. We notice that the normal continuity of electric current is enforced correctly at the surface of PEC cube.

4.1.3. Accuracy Study

Finally, we examine the solution accuracy of proposed method. To do so, we consider the EM scattering from a PEC sphere with radius 0.5 m, for which analytic solution is available in the form of Mie-series. The operating frequency varies from 1 GHz to 16 GHz, and the mesh size is chosen as $h = \lambda_0/12$ at each frequency. The generated surface discretizations are partitioned into a number of subdomains with a fixed number of DOFs per subdomain. The numerical solutions obtained from the proposed work are compared to the Mie-series analytical solutions. The error of the numerical solution is evaluated in L^2 norm, defined by:

$$\left\| \mathbf{j}^h - \mathbf{j}^{ref} \right\|_{L^2(S^h)} = \left[\int_{S^h} \left(\mathbf{j}^h(\mathbf{r}) - \mathbf{j}^{ref}(\mathbf{r}) \right) \cdot \left(\mathbf{j}^h(\mathbf{r}) - \mathbf{j}^{ref}(\mathbf{r}) \right)^* d\mathbf{r} \right]^{\frac{1}{2}}$$

The results are shown in Fig. 8. We notice that, for a fixed mesh size, the solution error in L^2 norm remains nearly constant with increasing higher operating frequencies. Fig. 9 presents the bistatic radar cross section (RCS) of the sphere at 16 GHz. The plane wave illuminates the sphere from $-\hat{\mathbf{z}}$ direction




Figure 8. Accuracy study of a PEC sphere with respect to operating frequency.

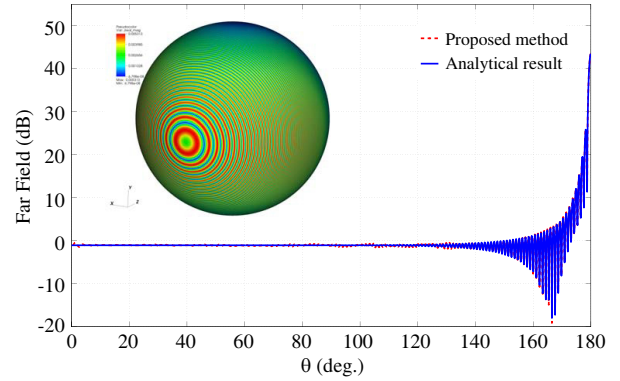


Figure 9. Bistatic RCS of a PEC sphere at 16 GHz.

and the electric field is polarized in the \hat{x} direction. We observe that the computational result agrees very well with the analytical result obtained by a Mie-series solution.

4.2. Parallel Performance Study

In this section, we present the parallel performance experiments of the parallel implementation of proposed algorithms. Results were obtained on Copper, a Cray XE6m DoD HPCMP Open Research System. The login and compute nodes are populated with AMD Opteron processors. Copper uses a dedicated Cray Gemini high-speed network for MPI messages and IO traffic. Copper uses Lustre to manage its parallel file system that targets arrays of SAS disk drives. Copper has 460 compute nodes that share memory only on the node; memory is not shared across the nodes. Each compute node has two sixteen-core processors (32 cores per node) that operate under a Cray Linux Environment (CLE) sharing 64 GBytes of DDR3 memory, with no user-accessible swap space. Copper is rated at 138 peak TFLOPS and has 239 TBytes (formatted) of parallel disk storage.

In the following experiments, the speedup and efficiency are displayed in terms of number of cores. All the elapsed time are obtained using the MPI routine `MPI_Wtime()`. Precisely, we concern with two cases of scalability [58]: (i) weak scalability, defined as how the solution time varies with the number of cores for a fixed problem size per core. Ideally the elapsed time is constant for a fixed ratio between the problem size and the number of cores. (ii) Strong scalability, defined as how the solution time varies with the number of cores for a fixed total problem size. Ideally, the elapsed time is inversely proportional to the number of cores.

4.2.1. Weak Scalability Test

We consider a plane wave scattering from a high-definition mockup jet aircraft. The dimensions of the jet aircraft are approximately 14.5 m long, 9.6 m wide, and 5.0 m high. We would like to analyze the EM wave scattering from this high definition multi-scale aircraft at four different frequencies, 3 GHz, 6 GHz, 9 GHz and 12 GHz. The plane wave incidents upon the aircraft from the nose (radome), $-\hat{x}$ direction. The electric field is polarized in the \hat{z} direction. The surface discretization is first generated at 3 GHz using DG boundary element methods. The discretization is chosen to accurately represent the arbitrarily-shaped geometrical surfaces, which usually results in distinct mesh sizes for different parts of the aircraft. At higher operating frequencies, individual local meshes is refined independently to satisfy the Nyquist sampling rate and to reach a desired level of accuracy.

In the scaling experiment, we use a nearly constant number of DOFs, 3.5 million, per subdomain. The global discretization at each operating frequency is then partitioned into a number of subdomains using METIS. The number of subdomains increases from 2 subdomains at 3 GHz to 32 subdomains at 32 GHz, correspondingly. The resulting partitions are displayed in Fig. 10. In the parallel computing,




Figure 10. Domain partitioning for jet aircraft simulations. (a) 3 GHz, 2 subdomains. (b) 6 GHz, 8 subdomains. (c) 9 GHz, 18 subdomains. (d) 12 GHz, 32 subdomains.





Figure 11. Parallelization efficiency in various simulations.

each subdomain is assigned with one MPI process in which 32 OpenMP threads are used. The parallelization efficiency relative to 32 cores is plotted in Fig. 11. It is evident that the efficiency remains 90% for all the computations. The details of computational statistics for individual operating frequencies are given in Table 1.

We first notice that the iteration counts to reach a relative residual of 10^{-2} only increase from 4 to 8 when comparing the cases of 3 GHz with 6.95 million DOFs and 12 GHz with 111.3 million DOFs. As a comparison, the GCR iterative solver with the conventional near-field preconditioner takes 223 iterations, 505 iterations and 1349 iterations for the cases of 3 GHz, 6 GHz and 9 GHz, respectively. The promising iterative solver convergence in the proposed work is one of the crucial factors to reach scalability for very large-scale EM simulations. Moreover, when using 1024 computing cores, the EM scattering from the jet aircraft at Ku-band can be solved within 7 and a half hours. Finally, we plot in Fig. 12 the electric current distributions on the exterior surfaces of the jet aircraft. From these figures, we have witnessed smooth current distributions without noticeable discontinuities across subdomain boundaries.

Table 1. Computational statistics for scaling experiments of EM scattering from a full-scale aircraft.

Num. of cores	Num. of subdomains	Subdomain solution (minutes)	Coupling (minutes)	Krylov method (seconds)	Total time per iteration (minutes)	Num. of iterations	Num. of DOFs (million)
64	2	50	7	3	57	4	6.95
256	8	47	8	15	55	5	27.4
576	18	40	14	28	65	7	62.3
1024	32	35	21	39	57	8	111.3

**Figure 12.** Surface electric currents distributions on jet aircraft at different operating frequencies. (a) 3 GHz, 2 subdomains. (b) 6 GHz, 8 subdomains. (c) 9 GHz, 18 subdomains. (d) 12 GHz, 32 subdomains.

4.2.2. Strong Scalability Test

In the strong scaling experiment, we examine the solution time required for a fixed problem size by increasing the number of cores. To do so, we consider EM scattering from the aircraft at 3 GHz and increase the number of subdomain from 2 to 32. In the parallel computing, each subdomain is assigned with one MPI process in which 32 OpenMP threads are used. Thus, the total computing cores increase from 64 to 1024. The timings of various simulations are given in Fig. 13. We see that the speedup is super-linear between 64 and 256 cores and then sub-linear from 256 cores to 1024 cores due to the slight growth of iteration numbers. At peak performance on 1024 cores, the speedup relative to 64 cores is approximately a twelve times decrease in simulation time.

4.3. Application: A Plane Wave Scattering from a Full-Scale Aircraft Carrier

We conclude this section with a very large-scale simulation, that of a plane wave scattering from a full-scale aircraft carrier as shown in Fig. 3. The carrier platform geometry, shown in Fig. 14, is approximately 250 m long, 65 m wide and 40 m high. This large high-definition platform model comprises of a number of complex smaller structures such radar antenna, mast, weapons and landing system located on the integrated island. Moreover, a variety of mockup aircrafts, including one V-22 osprey aircraft, three full-scale Apache helicopters, two full-scale stealth aircraft and three Unmanned Aerial




Figure 13. Parallel performance tests with respect to strong scalability, N : number of MPI processes (subdomains), K : number of iterations.




Figure 14. Aircraft carrier model.




Figure 15. Geometry-based decomposition of an aircraft carrier with a variety of jet aircraft.




Figure 16. Load-balanced computational partitioning of the aircraft carrier.




Figure 17. Surface electric current distributions on a full-scale aircraft carrier, $f = 400$ MHz.

Vehicles (UAVs) are located on the flight deck. The multi-scale nature of the underlying geometry and very large-scale problem sizes present significant computational difficulties for conventional IE methods.

We first simulate the plane wave scattering from the aircraft carrier at $f = 400$ MHz as depicted in Fig. 15. A total number of $M = 24$ subdomains are used, resulting in 35 million DOFs. The computational partitioning resulted from the Metis partitioner is shown in Fig. 16. Each subdomain is then assigned to one MPI process with 16 computing cores. The simulation requires 8 iterations to reach a relative residual 10^{-2} . The computation takes half an hour per DDM iteration. The surface electric current distribution for the large multi-scale simulation is given in Fig. 17.

Lastly, to demonstrate the capability of the proposed method, we consider the large-scale high-fidelity aircraft carrier simulation at 1.2 GHz. The geometry-based decomposition and computational partitioning are displayed in Fig. 3 and Fig. 4, respectively. The simulation requires 254 million DOFs with $M = 62$ subdomains. Each subdomain is assigned to one MPI process with 32 computing cores. A total number of 1984 computing cores are used in the computation. The simulation takes 1 hour




Figure 18. Surface electric current distributions on a full-scale aircraft carrier, $f = 1.2$ GHz. (a) Aircrafts on deck and front view. (b) Zoom-in view of integrated island.

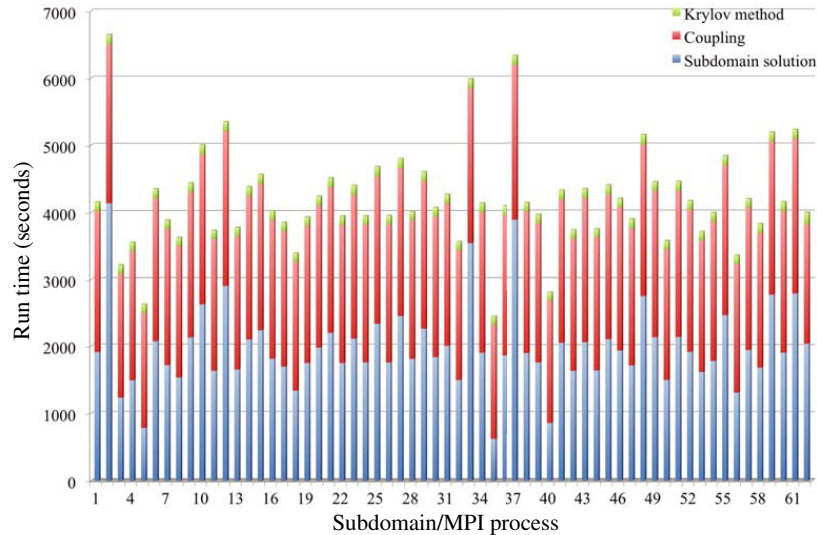


Figure 19. Comparison of the time spent in various steps for individual subdomains.

51 minutes for one DDM iteration and 12 iterations are required to reach a relative residual 10^{-2} . The surface electric currents on the aircraft carrier are depicted in Fig. 18. Finally, we show in Fig. 19 the runtime spent for individual subdomains in the first DDM iteration. Moreover, a breakdown of the time into, respectively, subdomain solution, coupling, and Krylov method is presented. We notice that the time spent in Krylov method, which includes the distribution of Krylov vector, computation of vector norms and message passing across MPI processes, is almost negligible comparing to subdomain solution and coupling times. Also, there is still room for improvement towards a perfect load balancing among computing nodes, which will be investigated in the future work.

5. CONCLUSIONS

The numerical results illustrate the capability of proposed algorithms for massively parallel HPC architectures. The technical advancements enable appealing parallel simulation capabilities including: (1) trivially parallel mesh generation, (2) scalable convergence in the DD iterations, (3) automatic load-balanced computational partitioning, and (4) embarrassingly parallel sub-domain solutions. We expect that continued work in this field will lead to a parallel EM simulation environment, where the fidelity and complexity of the EM problems that can be solved will scale with the exponential growth of supercomputing power. The longer-term goal is to create a reconfigurable, portable and maintainable IE solution suite based on the proposed framework. The already existing SIE solvers in academics, industry and government can be easily integrated with minor modifications.

ACKNOWLEDGMENT

The work was supported in part by the National Science Foundation (NSF) through award #CCF-1526605, and in part by the U.S. Department of Defense HPC Modernization Program, PP-CEA-KY06-003-P3.

REFERENCES

1. Harrington, R. F., "Boundary integral formulations for homogeneous material bodies," *Journal of Electromagnetic Waves and Applications*, Vol. 3, No. 1, 1–15, 1989.
2. Nédélec, J.-C., *Acoustic and Electromagnetic Equations: Integral Representations for Harmonic Problems*, Springer-Verlag, Berlin, Germany, 2001.

3. Buffa, A., R. Hiptmair, T. von Petersdorff, and C. Schwab, "Boundary element methods for Maxwell equations on Lipschitz domains," *Numer. Math.*, Vol. 95, 459–485, 2003.
4. Buffa, A. and R. Hiptmair, "Regularized combined field integral equations," *Numer. Math.*, Vol. 100, 1–19, 2005.
5. Tzoulis, A. and T. Eibert, "A hybrid FEBI-MLFMM-UTD method for numerical solutions of electromagnetic problems including arbitrarily shaped and electrically large objects," *IEEE Trans. Antennas and Propagation*, Vol. 53, 3358–3366, Oct. 2005.
6. Andriulli, F., K. Cools, F. Olyslager, A. Buffa, S. Christiansen, and E. Michielssen, "A multiplicative Calderon preconditioner for the electric field integral equation," *IEEE Trans. Antennas and Propagation*, Vol. 56, 2398–2412, Aug. 2008.
7. Chew, W. C., M. S. Tong, and B. Hu, *Integral Equation Methods for Electromagnetic and Elastic Waves*, Synthesis Lectures on Computational Electromagnetics, Morgan & Claypool Publishers, San Rafael, CA, USA, 2008.
8. Wang, X.-C., Z. Peng, and J.-F. Lee, "Multi-solver domain decomposition method for modeling EMC effects of multiple antennas on a large air platform," *IEEE Trans. Electromagnetic Compatibility*, Vol. 54, 375–388, Apr. 2012.
9. Hesford, A. J. and W. C. Chew, "On preconditioning and the eigensystems of electromagnetic radiation problems," *IEEE Trans. Antennas and Propagation*, Vol. 56, 2413–2420, Aug. 2008.
10. Ylä-Oijala, P., M. Taskinen, and J. Sarvas, "Surface integral equation method for general composite metallic and dielectric structures with junctions," *Progress In Electromagnetics Research*, Vol. 52, 81–108, 2005.
11. Peng, Z., K.-H. Lim, and J.-F. Lee, "Non-conformal domain decomposition methods for solving large multi-scale electromagnetic scattering problems," *Proceedings of IEEE*, Vol. 101, No. 2, 298–319, 2013.
12. Li, M.-K. and W. C. Chew, "Multiscale simulation of complex structures using equivalence principle algorithm with high-order field point sampling scheme," *IEEE Trans. Antennas and Propagation*, Vol. 56, 2389–2397, Aug. 2008.
13. Solis, D. M., J. M. Taboada, F. Obelleiro, and L. Landesa, "Optimization of an optical wireless nano link using directive nanoantennas," *Optics Express*, Vol. 21, No. 2, 2369–2377, 2013.
14. Solis, D. M., J. Taboada, F. Obelleiro, L. M. Liz-Marzán, and F. J. G. de Abajo, "Toward ultimate nanoplasmonics modeling," *ACS Nano*, Vol. 8, No. 8, 7559–7570, 2014.
15. Velamparambil, S., W. C. Chew, and J. Song, "10 million unknowns: Is it that big? [Computational electromagnetics]," *IEEE Antennas and Propagation Magazine*, Vol. 45, 43–58, Apr. 2003.
16. Ergül, Ö. and L. Gürel, "Efficient parallelization of the multilevel fast multipole algorithm for the solution of large-scale scattering problems," *IEEE Trans. Antennas and Propagation*, Vol. 56, No. 8, 2335–2345, 2008.
17. Pan, X.-M., W.-C. Pi, M.-L. Yang, Z. Peng, and X.-Q. Sheng, "Solving problems with over one billion unknowns by the MLFMA," *IEEE Trans. Antennas and Propagation*, Vol. 60, 2571–2574, May 2012.
18. Michiels, B., J. Fostier, I. Bogaert, and D. de Zutter, "Full-wave simulations of electromagnetic scattering problems with billions of unknowns," *IEEE Trans. Antennas and Propagation*, Vol. 63, No. 2, 796–799, 2015.
19. Taboada, J., L. Landesa, F. Obelleiro, J. Rodriguez, J. Bertolo, M. Araujo, J. Mourino, et al., "High scalability FMM-FFT electromagnetic solver for supercomputer systems," *IEEE Antennas and Propagation Magazine*, Vol. 51, No. 6, 20–28, 2009.
20. Wei, F. and A. E. Yilmaz, "A hybrid message passing/shared memory parallelization of the adaptive integral method for multi-core clusters," *Parallel Computing*, Vol. 37, No. 6, 279–301, 2011.
21. Wei, F. and A. Yilmaz, "A more scalable and efficient parallelization of the adaptive integral method — Part I: Algorithm," *IEEE Trans. Antennas and Propagation*, Vol. 62, 714–726, Feb. 2014.
22. Zhang, Y., Z. Lin, X. Zhao, and T. Sarkar, "Performance of a massively parallel higher-order method of moments code using thousands of CPUs and its applications," *IEEE Trans. Antennas*

- and Propagation*, Vol. 62, 6317–6324, Dec. 2014.
23. Adams, R., Y. Xu, X. Xu, S. Gedney, and F. Canning, “Modular fast direct electromagnetic analysis using local-global solution modes,” *IEEE Trans. Antennas and Propagation*, Vol. 56, 2427–2441, Aug. 2008.
 24. Wei, J.-G., Z. Peng, and J.-F. Lee, “A fast direct matrix solver for surface integral equation methods for electromagnetic wave scattering from non-penetrable targets,” *Radio Science*, Vol. 47, No. RS5003, 2012.
 25. Heldring, A., J. M. Tamayo, J. Rius, et al., “Accelerated direct solution of the method-of-moments linear system,” *Proceedings of the IEEE*, Vol. 101, No. 2, 364–371, 2013.
 26. Ergul, O. and L. Gurel, “Accurate solutions of extremely large integral-equation problems in computational electromagnetics,” *Proceedings of the IEEE*, Vol. 101, No. 2, 342–349, 2013.
 27. Vipiana, F., M. A. Francavilla, and G. Vecchi, “EFIE modeling of high-definition multiscale structures,” *IEEE Trans. Antennas and Propagation*, Vol. 58, 2362–2374, Jul. 2010.
 28. Nair, N. and B. Shanker, “Generalized method of moments: A novel discretization technique for integral equation,” *IEEE Trans. Antennas and Propagation*, Vol. 59, 2280–2293, Jun. 2011.
 29. Tong, M. S. and W. C. Chew, “A novel mesh less scheme for solving surface integral equations with flat integral domain,” *IEEE Trans. Antennas and Propagation*, Vol. 60, 3285–3293, Jul. 2012.
 30. Bendali, A., F. Collino, M. Fares, and B. Steif, “Extension to nonconforming meshes of the combined current and charge integral equation,” *IEEE Trans. Antennas and Propagation*, Vol. 60, 4732–4744, Oct. 2012.
 31. Chouly, F. and N. Heuer, “A Nitsche-based domain decomposition method for hypersingular integral equations,” *Numer. Math.*, Vol. 121, 705–729, Aug. 2012.
 32. Ubeda, E., J. Rius, and A. Heldring, “Nonconforming discretization of the electric-field integral equation for closed perfectly conducting objects,” *IEEE Trans. Antennas and Propagation*, Vol. 62, 4171–4186, Aug. 2014.
 33. Peng, Z., K.-H. Lim, and J.-F. Lee, “A discontinuous Galerkin surface integral equation method for electromagnetic wave scattering from nonpenetrable targets,” *IEEE Trans. Antennas and Propagation*, Vol. 61, No. 7, 3617–3628, 2013.
 34. Peng, Z., R. Hiptmair, Y. Shao, and B. MacKie-Mason, “Domain decomposition preconditioning for surface integral equations in solving challenging electromagnetic scattering problems,” *IEEE Trans. Antennas and Propagation*, doi: 10.1109/TAP.2015.2500908, 2015.
 35. Rao, S. M., D. R. Wilton, and A. W. Glisson, “Electromagnetic scattering by surfaces of arbitrary shape,” *IEEE Trans. Antennas and Propagation*, Vol. 30, 409–418, May 1982.
 36. Vecchi, G., “Loop-star decomposition of basis functions in the discretization of the EFIE,” *IEEE Trans. Antennas and Propagation*, Vol. 47, No. 2, 339–346, 1999.
 37. Lee, J. F., R. Burkholder, and R. Lee, “Loop star basis functions and a robust preconditioner for EFIE scattering problems,” *IEEE Trans. Antennas and Propagation*, Vol. 51, 1855–1863, Aug. 2003.
 38. Adams, R. J., “Physical and analytical properties of a stabilized electric field integral equation,” *IEEE Trans. Antennas and Propagation*, Vol. 52, 362–372, Feb. 2004.
 39. Stephanson, M. and J.-F. Lee, “Preconditioned electric field integral equation using Calderon identities and dual loop/star basis functions,” *IEEE Trans. Antennas and Propagation*, Vol. 57, 1274–1279, Apr. 2009.
 40. Bruno, O., T. Elling, R. Paffenroth, and C. Turc, “Electromagnetic integral equations requiring small numbers of Krylov-subspace iterations,” *J. Comput. Phys.*, Vol. 228, 6169–6183, Sep. 2009.
 41. Wiedenmann, O. and T. F. Eibert, “A domain decomposition method for boundary integral equations using transmission condition based on the near-zone coupling,” *IEEE Trans. Antennas and Propagation*, Vol. 62, 4105–4114, Aug. 2014.
 42. Echeverri Bautista, M., F. Vipiana, M. Francavilla, J. Tobon Vasquez, and G. Vecchi, “A nonconformal domain decomposition scheme for the analysis of multi-scale structures,” *IEEE Trans. Antennas and Propagation*, doi: 10.1109/TAP.2015.2430873, 2015.

43. Toselli, A. and O. Widlund, *Domain Decomposition Methods — Algorithms and Theory*, Springer, Berlin, 2005.
44. Peng, Z., X.-C. Wang, and J.-F. Lee, “Integral equation based domain decomposition method for solving electromagnetic wave scattering from non-penetrable objects,” *IEEE Trans. Antennas and Propagation*, Vol. 59, 3328–3338, Sep. 2011.
45. Langer, U., G. Of, O. Steinbach, and W. Zulehner, “Inexact data-sparse boundary element tearing and interconnecting methods,” *SIAM J. Sci. Comput.*, Vol. 29, No. 1, 290–314, 2007.
46. Steinbach, O. and M. Windisch, “Stable boundary element domain decomposition methods for the Helmholtz equation,” *Numer. Math.*, Vol. 118, No. 1, 171–195, 2011.
47. Karypis, G. and V. Kumar, “A fast and high quality multilevel scheme for partitioning irregular graphs,” *SIAM J. Sci. Comput.*, Vol. 20, No. 1, 359–392 (electronic), 1998.
48. Parks, M. L., E. D. Sturler, G. Mackey, D. D. Johnson, and S. Maiti, “Recycling Krylov subspaces for sequences of linear systems,” *SIAM J. Sci. Comput.*, Vol. 28, No. 5, 1651–1674, 2006.
49. Wei, J.-G., Z. Peng, and J.-F. Lee, “Multi-scale electromagnetic computations using a hierarchical multi-level fast multipole algorithm,” *Radio Science*, Vol. 49, No. 11, 1022–1040, 2014.
50. Pan, X.-M., J.-G. Wei, Z. Peng, and X.-Q. Sheng, “A fast algorithm for multiscale electromagnetic problems using interpolative decomposition and multilevel fast multipole algorithm,” *Radio Science*, Vol. 47, No. RS1011, 2012.
51. Ho, K. L. and L. Greengard, “A fast direct solver for structured linear systems by recursive skeletonization,” *SIAM J. Sci. Comput.*, Vol. 34, No. 5, A2507–A2532, 2012.
52. Mahaffey, J. V., “A direct approach at field computation using the FMM framework,” Master’s Thesis, The Ohio State University, Columbus, Ohio, 2012.
53. Taboada, J. M., M. G. Araújo, J. M. Bertolo, L. Landesa, F. Obelleiro, and J. L. Rodríguez, “MLFMA-FFT parallel algorithm for the solution of large-scale problems in electromagnetic (invited paper),” *Progress In Electromagnetics Research*, Vol. 105, 15–30, 2010.
54. Taboada, J., M. Araújo, F. Basteiro, J. Rodríguez, and L. Landesa, “MLFMA-FFT parallel algorithm for the solution of extremely large problems in electromagnetics,” *Proceedings of the IEEE*, Vol. 101, 350–363, Feb. 2013.
55. Lawrence Livermore National Laboratory, “Visit users manual,” <https://wci.llnl.gov/content/assets/docs/simulation/computer-codes/visit/VisItUsersManual1.5.pdf>, 2005.
56. Campbell, S. L., I. C. F. Ipsen, C. T. Kelley, and C. D. Meyer, “GMRES and the minimal polynomial,” *BIT Numer. Math.*, Vol. 36, 664–675, 1996.
57. Meurant, G. and J. D. Tebbens, “The role eigenvalues play in forming GMRES residual norms with non-normal matrices,” *Numer. Algorithms*, Vol. 68, No. 1, 143–165, 2015.
58. Dolean, V., P. Jolivet, and F. Nataf, “An introduction to domain decomposition methods: Algorithms, theory and parallel implementation,” Master, France, 2015.