

© 2017 Joseph M. Rutherford

COMPUTATION OF MULTISCALE TIME-HARMONIC
ELECTROMAGNETIC RADIATION

BY

JOSEPH M. RUTHERFORD

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Electrical and Computer Engineering
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2017

Urbana, Illinois

Doctoral Committee:

Professor Weng Cho Chew, Chair
Professor P. Scott Carney
Professor Erhan Kudeki
Professor Jonathan Makela

ABSTRACT

The effect of installation for an antenna is key to understanding its performance in the intended operating environments. While existing formulations support such analysis, their mapping of calculations to distributed computing hardware does not support simulating installation environments of arbitrary size. This work builds upon existing techniques to simulate installed antenna behavior using scattering analyses tailored to system components. The scattering operations reveal opportunities to introduce approximate techniques which form generalized hybrid solvers. The source antenna (with both subwavelength-scale and wavelength-scale features) is modeled with the electric field integral equation (EFIE), and it interfaces with the installation site using the equivalence principle algorithm (EPA) as a domain decomposition method (DDM). The use of EPA to enclose the EFIE-modeled antenna generalizes the method to arbitrary antenna models. The electrically large exterior structures are modeled with physical optics (PO) without loss of generality to other approximate or high-frequency asymptotic methods through a Schur complement analysis of the continuous and discretized equations. The proposed Schur complement EPA-PO hybrid introduces clear physics with applicability to other formulations for the individual domains. The proposed hybrid also maps the PO calculations efficiently to distributed parallel computing resources; the parallel computations are demonstrated by executing the simulations on a hybrid parallel distributed- and shared-memory computing cluster. The calculation of antenna interactions with electrically large structures implies transition from wave-physics to ray-physics behaviors, which raises questions of reduced rank in the discretized operators. These questions are addressed by identifying the wave-to ray-physics transition and observing reduced rank in the space of plane wave functions.

To Kate.

ACKNOWLEDGMENTS

I am grateful to the many people in my life who have so graciously supported me in this endeavor. Throughout this endeavor my family, friends, and colleagues have excused my absences, offered much-needed encouragement, and shouldered burdens so that I may pursue this goal. I owe much gratitude for their acts of selfless kindness; I would not have accomplished this without your sacrifices. Thank you, all.

My instructors have naturally had a profound influence on my methods. My doctoral advisor, Professor Weng Cho Chew, has had the greatest influence on my learning. It has been my honor to study in his research group, and the insights I have gained under his tutelage have been most helpful in both academia and industry. I am also grateful to Professor Paul Scott Carney for shaping my mind as my master's degree advisor; his investments in me have been most appreciated.

My colleagues in both industry and academia have been a consistent source of engaging ideas and valuable technical discussions. Thank you to all who have so patiently listened to me work out these ideas over time. Special thanks to Nathan Smith for his insightful questions and for always being willing to sit through one more practice run of a presentation.

My friends and family have been a bottomless well of support. Thank you for listening, empathizing, and encouraging as I have persevered. Special thanks to John and Sandy Brix for Nessie, Yoshi, and Llewey so that I have had a (canine) friend available at all times.

Most of all I thank my wife, Kate Rutherford, for walking with me through every challenge, frustration, and success. Our years of mutual self-sacrificial love have been a remarkable blessing.

TABLE OF CONTENTS

LIST OF ABBREVIATIONS	vii
CHAPTER 1 INTRODUCTION	1
1.1 Motivation for Addressing Computational Electromagnetics	1
1.2 Thesis Organization	2
CHAPTER 2 SURFACE INTEGRAL EQUATIONS	3
2.1 Differential Equations	3
2.2 Integral Equations and Equivalent Problems	4
2.3 Equivalence Principle for a Single Scatterer	7
2.4 Equivalence Principle for Multiple Scatterers	7
2.5 Figures	9
CHAPTER 3 CONSTRUCTING NUMERICAL SOLUTIONS	11
3.1 Discretized SIE	11
3.2 Obtaining Numerical Solutions	19
3.3 Fast Translator Matrix-Vector-Product Operations	20
3.4 Computing Solutions Using the Schur Complement	25
CHAPTER 4 EPA HYBRIDIZED WITH PO	26
4.1 Introduction	26
4.2 Previous Work	27
4.3 Physical Optics as a Scattering Operator	27
4.4 Eliminating Unknowns on an EPA Domain	29
4.5 Examples	32
4.6 Conclusions	35
4.7 Figures	36
CHAPTER 5 PARALLEL CALCULATIONS OF EPA-PO HYBRID SOLUTIONS	42
5.1 Introduction	42
5.2 Sequential and Parallel Calculations	43
5.3 MLFMA Operations	48
5.4 Schur Complement EPA-PO Hybrid	51
5.5 Results	53

5.6	Conclusions	55
5.7	Figures	56
CHAPTER 6 BEAM FORMATION USING PLANE WAVES		62
6.1	Introduction	62
6.2	Numerical Analysis of the Green's Function's Rank	63
6.3	FMM as Fourier-Legendre Series	66
6.4	Emergence of Ray Physics within FMM	68
6.5	Fast Multipole Methods and Rank Deficiency	70
6.6	Conclusions	75
6.7	Table and Figures	76
CHAPTER 7 CONCLUSION		91
7.1	Summary	91
7.2	Future Work	92
REFERENCES		94

LIST OF ABBREVIATIONS

DDM	Domain Decomposition Method
EPA	Equivalence Principle Algorithm
FaFFA	Fast Far-Field Approximation
FMM	Fast Multipole Method
GMRES	Generalized Minimal Residual
HPC	High-Performance Computing
MLFMA	Multilevel Fast Multipole Algorithm
MoM	Method of Moments
MPI	Message Passing Interface
MVM	Matrix-Vector-Multiply
PEC	Perfect Electric Conductor
PO	Physical Optics
RHS	Right-Hand Side
RPFMA	Ray Propagation Fast Multipole Algorithm
SIE	Surface Integral Equation
SVD	Singular Value Decomposition

CHAPTER 1

INTRODUCTION

1.1 Motivation for Addressing Computational Electromagnetics

Systems utilizing electrical and magnetic (EM) phenomena are engineered for specific operating environments. The simplest operating environment is an otherwise homogeneous (infinite) space supporting only a transmitter and a receiver. Such an environment is not suitable as an approximation for many cases because the surrounding medium is made inhomogeneous by the presence of a transmitter- or receiver-mounting structure, variable atmosphere, or other details not specifically planned as part of the system design. Therefore accurately modeling the inhomogeneous operating environment must be part of the system design process for accurate performance prediction. It is particularly important to develop simulation techniques which map efficiently to the distributed parallel computing hardware available in the foreseeable future.

EM phenomena are governed by Maxwell's equations [1, 2]. A computational electromagnetic (CEM) simulation computes a solution to these differential equations corresponding to the materials and sources characterizing the case of interest. Many cases of interest are accurately modeled by linear materials, which allows for a time-harmonic analysis and therefore supports composing distributional solutions via Green's functions. Such solutions are formed by solving problems of the form $\overline{\mathbf{A}}\mathbf{x} = \mathbf{b}$ where discrete linear operator $\overline{\mathbf{A}}$ is determined by the domain (environment and devices) under study, unknown vector \mathbf{x} is the fields in the domain, and known vector \mathbf{b} characterizes the scenario excitations.

Computing a solution is an intensive process, and may easily overwhelm available computational resources. Deploying parallel computing resources

to solve problems of interest demands parallel algorithms which map the computational workload efficiently to the available hardware. Past efforts to develop parallel algorithms for solving CEM problems have met with mixed results owing to the complexities of the wave equation. Approximate solutions are often suitably accurate for problem classes of interest with the advantage of high parallel scalability. This thesis develops foundational techniques to map approximated discretized linear EM problems to distributed parallel high-performance computing (HPC) hardware.

1.2 Thesis Organization

In this work, time-constant materials and time-harmonic linear fields are used to develop techniques for obtaining solutions to the governing equations using distributed parallel computing hardware. Chapter 2 explains the theory used to organize surface integral equation (SIE) problems using the surface equivalence theorem. Chapter 3 explains the discretization of the SIEs in support of computing results for arbitrary cases. Chapter 4 develops a novel approach to obtaining hybrid solutions combining SIE solutions and approximate scattering operators. Chapter 5 explores the distributed parallel solution to the hybrid solutions, especially using HPC resources. Finally, Chapter 6 explores beam formation in radiation calculations to identify opportunities to reduce the computational workload of accurate high-frequency radiation calculations.

CHAPTER 2

SURFACE INTEGRAL EQUATIONS

2.1 Differential Equations

In the following equations,

- \mathbf{E} (\mathbf{H}) is the electric (magnetic) field intensity,
- \mathbf{D} (\mathbf{B}) is the electric (magnetic) flux density,
- \mathbf{J} (\mathbf{M}) is the electric (magnetic) current density,
- ϵ is the electric permittivity, and
- μ is the magnetic permeability.

\mathbf{M} is generally regarded as a nonphysical quantity that must be zero in practice, but it balances the structure of the equations to simplify subsequent analysis through duality. The Maxwell-Heaviside equations [2] (often simply called the Maxwell equations [1]) are given in time-harmonic form, such that $f(\mathbf{r}, t) = \text{Re}[f(\mathbf{r})e^{-i\omega t}]$, as [3, Eqs. (1.1.12)-(1.1.15)]

$$\nabla \times \mathbf{E}(\mathbf{r}) = i\omega\mathbf{B}(\mathbf{r}) - \mathbf{M}(\mathbf{r}), \quad (2.1)$$

$$\nabla \times \mathbf{H}(\mathbf{r}) = -i\omega\mathbf{D}(\mathbf{r}) + \mathbf{J}(\mathbf{r}). \quad (2.2)$$

Simple algebra converts these into vector wave equations

$$\nabla \times \nabla \times \mathbf{E}(\mathbf{r}) - \omega^2\epsilon\mu\mathbf{E}(\mathbf{r}) = i\omega\mu\mathbf{J}(\mathbf{r}) - \nabla \times \mathbf{M}(\mathbf{r}), \quad (2.3)$$

$$\nabla \times \nabla \times \mathbf{H}(\mathbf{r}) - \omega^2\epsilon\mu\mathbf{H}(\mathbf{r}) = i\omega\epsilon\mathbf{M}(\mathbf{r}) + \nabla \times \mathbf{J}(\mathbf{r}), \quad (2.4)$$

under the assumptions that

1. the temporal frequency ω is not close to zero so that solutions may be obtained [4]

2. the constitutive parameters are independent of the fields; this defines linear equations so that
 - (a) weak solutions may be obtained using generalized functions,
 - (b) \mathbf{E} , \mathbf{H} , \mathbf{D} , \mathbf{B} , \mathbf{J} , and \mathbf{M} are asserted to have a common time-harmonic factor $e^{-i\omega t}$,
3. the constitutive relations are those of isotropic materials such that $\mathbf{D} = \epsilon\mathbf{E}$, $\mathbf{B} = \mu\mathbf{H}$,
4. the constitutive parameters $\epsilon \in \mathbb{C}$ and $\mu \in \mathbb{C}$ are both time-invariant and piecewise-constant, and
5. the second partial derivatives with respect to space and time $\frac{\partial^2}{\partial r \partial t}$, $\frac{\partial^2}{\partial t \partial r}$ of both \mathbf{E} and \mathbf{H} are continuous and equal.

If the tangential fields satisfy boundary conditions on the region interfaces and the media are lossy, then unique solutions are defined within each region. In this context, lossy media satisfy $\text{Im}(\epsilon) > 0$, $\text{Im}(\mu) > 0$, and the lossless case is supported in the limit approaching zero loss.

2.2 Integral Equations and Equivalent Problems

A dyadic Green's function $\overline{\mathbf{G}}(\mathbf{r}, \mathbf{r}')$ relates the fields to the sources by satisfying a PDE [3, Eqs. (8.1.16)-(8.1.17)] within a piecewise constant region V

$$\nabla \times \nabla \times \overline{\mathbf{G}}(\mathbf{r}, \mathbf{r}') - k^2 \overline{\mathbf{G}}(\mathbf{r}, \mathbf{r}') = \overline{\mathbf{I}} \delta(\mathbf{r} - \mathbf{r}'), \mathbf{r}, \mathbf{r}' \in V, \quad (2.5)$$

where $\overline{\mathbf{I}}$ is the identity dyad and $\delta(\mathbf{r})$ is the Dirac delta function. The vector calculus analysis of Eqs. (2.3), (2.4), and (2.5) which yields the electric field integral equation (EFIE) and the magnetic field integral equation (MFIE) [3, Sect. 8.1.2] also yields the concept of a SIE translator. Abbreviate the evaluated internal source fields using volume integrals as

$$\begin{aligned} \mathbf{E}^{\text{inc}}(\mathbf{r}) &= i\omega\mu \int_V \overline{\mathbf{G}}(\mathbf{r}, \mathbf{r}') \cdot \mathbf{J}(\mathbf{r}') dV' - \\ &\int_V \overline{\mathbf{G}}(\mathbf{r}, \mathbf{r}') \cdot \nabla \times \mathbf{M}(\mathbf{r}') dV', \end{aligned} \quad (2.6)$$

$$\begin{aligned} \mathbf{H}^{\text{inc}}(\mathbf{r}) &= i\omega\epsilon \int_V \overline{\mathbf{G}}(\mathbf{r}, \mathbf{r}') \cdot \mathbf{M}(\mathbf{r}') dV' + \\ &\int_V \overline{\mathbf{G}}(\mathbf{r}, \mathbf{r}') \cdot \nabla \times \mathbf{J}(\mathbf{r}') dV'. \end{aligned} \quad (2.7)$$

The fields anywhere in the volume under study are then calculated using the internal and boundary source fields as

$$\begin{bmatrix} \mathbf{E}(\mathbf{r}) \\ \mathbf{H}(\mathbf{r}) \end{bmatrix}_{\text{tan}} = \begin{bmatrix} \mathbf{E}^{\text{inc}}(\mathbf{r}) \\ \mathbf{H}^{\text{inc}}(\mathbf{r}) \end{bmatrix}_{\text{tan}} + \mathcal{T} \begin{bmatrix} \mathbf{M}^{\text{sca}} \\ \mathbf{J}^{\text{sca}} \end{bmatrix}, \quad \mathbf{r} \in V \quad (2.8)$$

with the integral operator abbreviations

$$\mathcal{L}\mathbf{X} = \int_S [\overline{\mathbf{G}}(\mathbf{r}, \mathbf{r}') \cdot \mathbf{X}(\mathbf{r}')] dS', \quad (2.9)$$

$$\mathcal{K}\mathbf{X} = \int_S [(\nabla' \times \overline{\mathbf{G}}(\mathbf{r}, \mathbf{r}')) \cdot \mathbf{X}(\mathbf{r}')] dS', \quad (2.10)$$

$$\mathcal{T} = \begin{bmatrix} \mathcal{K} & i\omega\mu\mathcal{L} \\ -i\omega\epsilon\mathcal{L} & -\mathcal{K} \end{bmatrix}. \quad (2.11)$$

The EFIE and MFIE in Eq. (2.8) express the fields at any point as the sum of the internal source fields and propagated boundary fields. Note that EFIE and MFIE are commonly evaluated at points on the surface of the scatterer itself.

SIEs depend upon knowing the Green's function for the problem under study, but inhomogeneous medium Green's functions are generally unavailable. Without loss of generality, a single antenna radiating in the presence of another antenna on a platform is analyzed to establish an equivalent problem using the homogeneous medium Green's function. The scenario is illustrated in Figure 2.1 as four connected non-overlapping domains: source antenna A, receiving antenna B, platform C, and exterior D. Domain D is bounded by domains A, B, and C on its interior and a closed exterior surface. In order to solve for the fields in any of domains A, B, or C, the tangential fields

bounding domain D must be obtained.

Domains B and C are replaced with material matching domain D. The sources within domain A are replaced by equivalent boundary surface currents radiating

- zero fields inside domain A, and
- isolated-antenna exterior fields outside domain A

as shown in Figure 2.2. These equivalent currents are the incident field source. The equivalent fields within domain A are zero, and its contents are also replaced with the exterior medium. This defines an equivalent scenario which may be solved using both Eq. (2.8) and the homogeneous medium Green's function [3, Eqs. (7.1.1) and (7.1.19)]

$$\overline{\mathbf{G}}(\mathbf{r}, \mathbf{r}') = \left(\overline{\mathbf{I}} + \frac{1}{k^2} \nabla \nabla \right) \frac{e^{ik|\mathbf{r}-\mathbf{r}'|}}{4\pi |\mathbf{r} - \mathbf{r}'|}. \quad (2.12)$$

The radiation boundary condition for the homogeneous equivalent problem is applied to the exterior surface. As the exterior surface expands in the limit of an infinite radius, the radiation condition forces the radiation integral to vanish. Using Eq. (2.12) with Eq. (2.8) satisfies this boundary condition and thus defines the unique solution for the equivalent problem.

The equivalent currents corresponding to the scattered fields are defined using the desired total field equivalent problem in Figure 2.3: the original problem's total fields in domain D (equal to the incident fields radiating from domain A plus the scattered fields radiated from all domains) and extinguished fields in domains A, B, and C. This equivalent problem satisfies Eqs. (2.3) and (2.4) with the correct solution in domain D and zero fields in the other domains. The equivalent sources yielding the exterior scattered fields are shown in Figure 2.4. Said surface currents radiate with the free space dyadic Green's function to calculate the scattered field everywhere exterior to the problem. Returning to EFIE and MFIE in Eq. (2.8), the definition of the Green's function completes the translator relating boundary sources to observed fields within a closed region.

2.3 Equivalence Principle for a Single Scatterer

Regardless of the SIE formulation for a particular scatterer with surface σ , the translator of Eq. (2.11) may be used with the equivalence principle to transform the scattering operator away from the skin σ to an enclosing surface Σ . Using labeling from [5, Eq. (12)]

- $\mathcal{T}_{\text{outside-in}}$ propagates the field from the exterior interface to the interior scatterer surface,
- $\mathcal{T}_{\text{inside-out}}$ propagates the field from the scatterer surface interior to the exterior interface, and
- $\mathcal{S}_{\text{current-solver}}$ is the interior structure scattering operator,

the new equivalent scattering operator is

$$\mathcal{S}_{\Sigma} = \mathcal{T}_{\text{inside-out}} \mathcal{S}_{\text{current-solver}} \mathcal{T}_{\text{outside-in}}. \quad (2.13)$$

The outside-in and inside-out operators translate between $\mathcal{S}_{\text{current-solver}}$ and the interface surface, often allowing coarser, more uniform sampling at the interface [5] because the evanescent field components near the antenna structures decay over distance. This abstraction facilitates hybridizing solutions by isolating physics-sensitive scattering operators to designated domains and relating fields using abstract translators.

2.4 Equivalence Principle for Multiple Scatterers

The equivalence principle algorithm (EPA) [5, 6, 7] provides a mechanism to define scattering operators per domain with a common interface. Equation (2.8) shows that the fields on any surface in domain D are determined by the radiation of both sources embedded in D and the boundary surface fields. One may compose a set of equations relating the fields on the surfaces of domains A, B, and C. The first step is to note that each domain is a boundary source for the others. The source fields \mathbf{M}^{inc} , \mathbf{J}^{inc} radiate from A to both B and C; the scattered fields of domains B and C are the only sources illuminating domain A. The sum of all sources in the scenario observed on

the interface surface of domain $l \in \{A, B, C\}$ is restated from Eq. (2.8) as

$$\begin{bmatrix} \mathbf{E}_l^{\text{src}} \\ \mathbf{H}_l^{\text{src}} \end{bmatrix}_{\text{tan}} = \mathcal{T}_{lA} \begin{bmatrix} \mathbf{M}_A^{\text{inc}} \\ \mathbf{J}_A^{\text{inc}} \end{bmatrix} + \sum_{k, k \neq l} \mathcal{T}_{lk} \begin{bmatrix} \mathbf{M}_k^{\text{sca}} \\ \mathbf{J}_k^{\text{sca}} \end{bmatrix}. \quad (2.14)$$

The scattered field equivalent currents bounding each scatterer shown in Figure 2.4 are defined by the scattering operators \mathcal{S}_A , \mathcal{S}_B , and \mathcal{S}_C applied to Eq. (2.14). The details of such scattering operators are developed later; it suffices for the moment to stipulate their existence. As in [7, Eq. (8)], the system of equations is generalized for any number of equivalence principle surfaces as

$$\begin{bmatrix} \mathbf{E}_l^{\text{sca}} \\ \mathbf{H}_l^{\text{sca}} \end{bmatrix} = \mathcal{S}_l \mathcal{T}_{lA} \begin{bmatrix} \mathbf{M}_A^{\text{inc}} \\ \mathbf{J}_A^{\text{inc}} \end{bmatrix} + \mathcal{S}_l \sum_{k, k \neq l} \mathcal{T}_{lk} \begin{bmatrix} \mathbf{M}_k^{\text{sca}} \\ \mathbf{J}_k^{\text{sca}} \end{bmatrix}. \quad (2.15)$$

The simultaneous equations of multiple scattering are expressed compactly in this form with a succinct abstraction of scattering operators and translation operators. Evaluating the fields on each domain l yields a system of equations

$$\left(\mathcal{I} - \begin{bmatrix} \mathcal{S}_A & 0 & 0 \\ 0 & \mathcal{S}_B & 0 \\ 0 & 0 & \mathcal{S}_C \end{bmatrix} \begin{bmatrix} 0 & \mathcal{T}_{AB} & \mathcal{T}_{AC} \\ \mathcal{T}_{BA} & 0 & \mathcal{T}_{BC} \\ \mathcal{T}_{CA} & \mathcal{T}_{CB} & 0 \end{bmatrix} \right) \begin{bmatrix} \mathbf{M}_A^{\text{sca}} \\ \mathbf{J}_A^{\text{sca}} \\ \mathbf{M}_B^{\text{sca}} \\ \mathbf{J}_B^{\text{sca}} \\ \mathbf{M}_C^{\text{sca}} \\ \mathbf{J}_C^{\text{sca}} \end{bmatrix} = \begin{bmatrix} \mathcal{S}_A & 0 & 0 \\ 0 & \mathcal{S}_B & 0 \\ 0 & 0 & \mathcal{S}_C \end{bmatrix} \begin{bmatrix} 0 \\ 0 \\ \mathcal{T}_{BA} \begin{bmatrix} \mathbf{M}_A^{\text{inc}} \\ \mathbf{J}_A^{\text{inc}} \end{bmatrix} \\ \mathcal{T}_{CA} \begin{bmatrix} \mathbf{M}_A^{\text{inc}} \\ \mathbf{J}_A^{\text{inc}} \end{bmatrix} \end{bmatrix}. \quad (2.16)$$

This expresses the EPA domain decomposition method (DDM) in continuous form. It is readily discretized and solved using the method of moments (MoM) just as any other surface integral equation problems. Effects are now isolated by domain so that methods suitable to electrically small regions and electrically large regions may be handled with tailored scattering algorithms.

The process of discretizing this SIE is developed next.

2.5 Figures

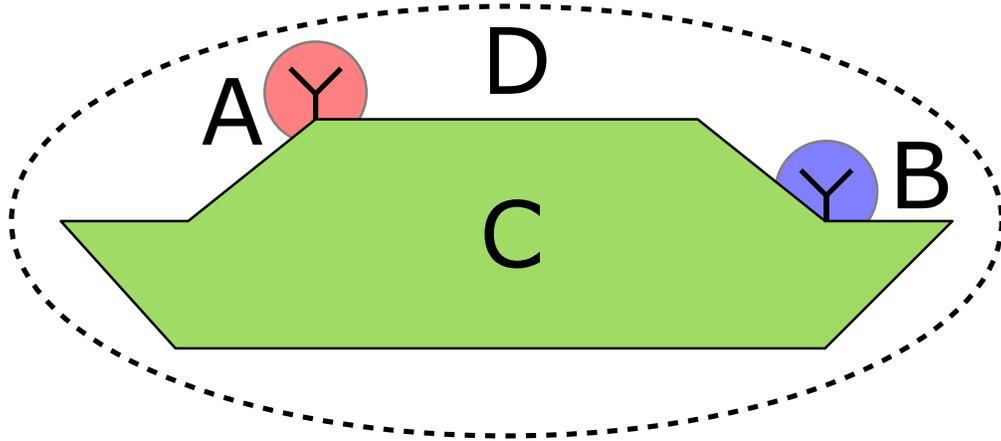


Figure 2.1: Development of SIE for exterior domain: original problem with four domains A, B, C, and D.

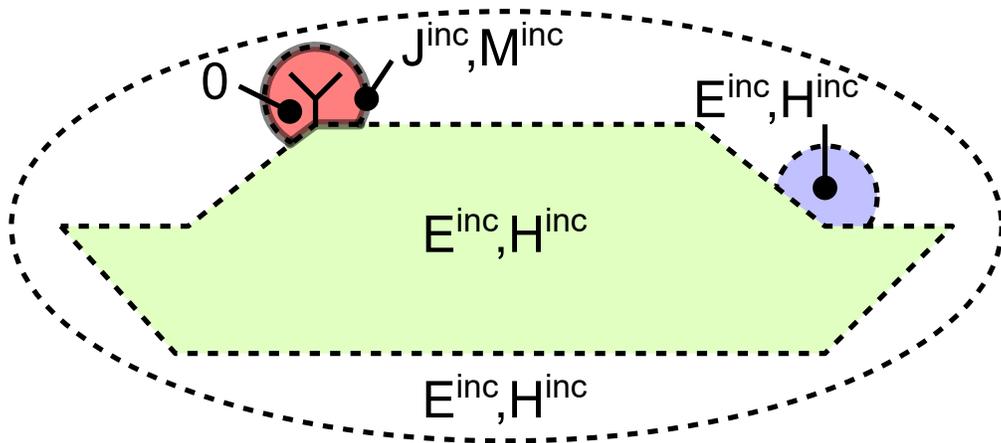


Figure 2.2: Development of SIE for exterior domain: homogeneous medium definition of incident field with equivalent sources bounding A.

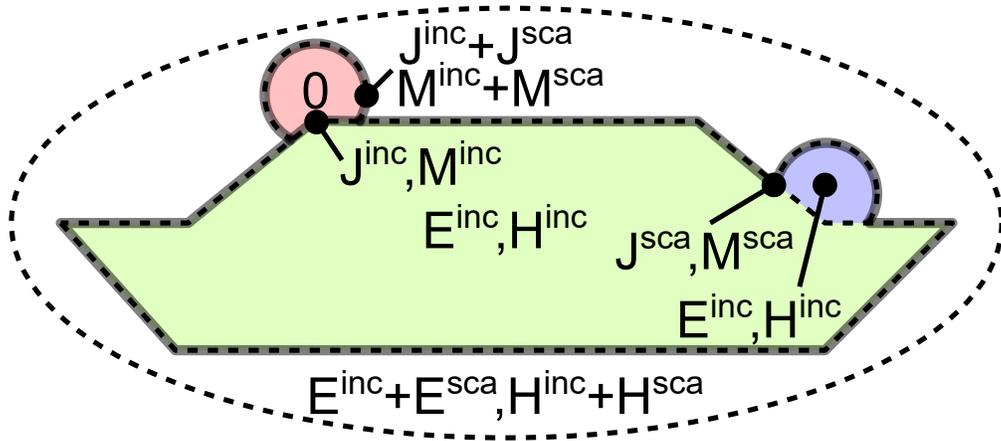


Figure 2.3: Development of SIE for exterior domain: homogeneous medium equivalent problem with same result as original in D but extinct fields in A, B, and C.

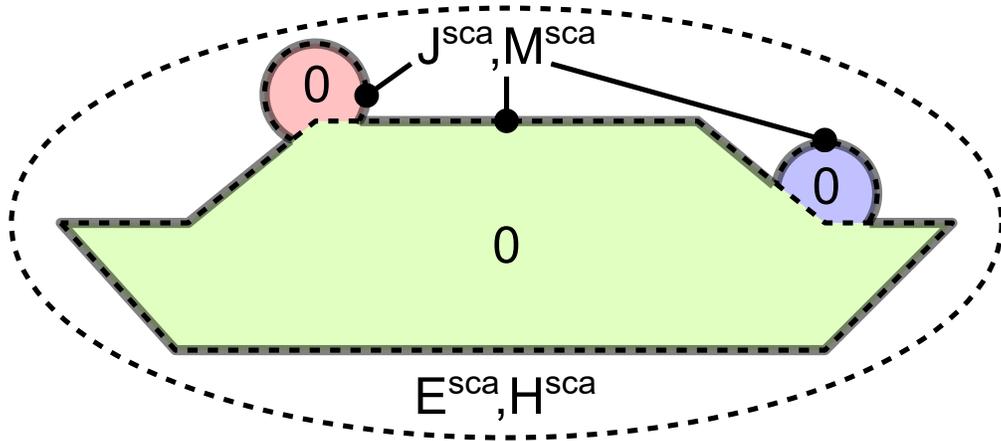


Figure 2.4: Development of SIE for exterior domain: definition of surface current densities in homogeneous medium to yield desired scattered fields.

CHAPTER 3

CONSTRUCTING NUMERICAL SOLUTIONS

3.1 Discretized SIE

Discretizing a function as a vector of basis function coefficients is combined with discretizing operators as matrices to obtain solutions via the method of moments (MoM) [8]; the technique is summarized here. A function $a(\mathbf{r})$ expanded in a basis set is expressed using Hilbert space vectors as

$$|a\rangle = \sum_n b_n |c_n\rangle, \quad (3.1)$$

$$|a\rangle = \mathbf{c}^\top \cdot \mathbf{b}, \quad (3.2)$$

with coefficients $[\mathbf{b}]_n = b_n \in \mathbb{C}$ and $c_n(\mathbf{r}) \in \mathbb{R}$. One may compose a system of equations by projecting onto testing functions $d_n(\mathbf{r}) \in \mathbb{R}$:

$$\langle \mathbf{d}, a \rangle = \langle \mathbf{d}, \mathbf{c}^\top \rangle \cdot \mathbf{b}, \quad (3.3)$$

$$\langle \mathbf{d}, \mathbf{c}^\top \rangle \cdot \mathbf{b} = \mathbf{e}. \quad (3.4)$$

This linear system is of the form $\overline{\mathbf{F}} \cdot \mathbf{b} = \mathbf{e}$. It may be solved for unknown coefficients $[\mathbf{b}]_n = b_n$ given (a) the Gramian matrix elements $[\overline{\mathbf{F}}]_{mn} = \langle d_m, c_n \rangle$ and (b) the projection of the known function $|a\rangle$ onto the testing functions $[\mathbf{e}]_m = \langle d_m, a \rangle$. Thus the operator projecting a function defined in one basis onto another basis is discretized in the Gramian matrix; this concept will be used extensively in discretizing other operators.

The discretization of a continuous linear operator \mathcal{A} is accomplished with this notation by analyzing a linear operator problem

$$\mathcal{A}x = b. \quad (3.5)$$

By expanding x and b in terms of respective basis functions $f_x\rangle$ and $f_b\rangle$,

$$x = \mathbf{f}_x^\top \cdot \mathbf{x}, \quad (3.6)$$

$$b = \mathbf{f}_b^\top \cdot \mathbf{b}, \quad (3.7)$$

the linear operator problem becomes

$$\mathcal{A} \mathbf{f}_x^\top \cdot \mathbf{x} = \mathbf{f}_b^\top \cdot \mathbf{b}. \quad (3.8)$$

One may pick a set of testing functions $f_t\rangle$ and test the linear operator problem with any testing function indexed i to obtain

$$\langle f_{t,i}, \mathcal{A} \mathbf{f}_x^\top \rangle \cdot \mathbf{x} = \langle f_{t,i}, \mathbf{f}_b^\top \rangle \cdot \mathbf{b}. \quad (3.9)$$

This testing process is repeated for each member of the set of testing functions. Therefore the discrete form of \mathcal{A} is seen to be

$$[\mathbf{A}]_{ij} = \langle f_{t,i}, \mathcal{A} f_{x,j} \rangle, \quad (3.10)$$

$$[\overline{\mathbf{F}}_{tb}]_{ij} = \langle f_{t,i}, f_{b,j} \rangle, \quad (3.11)$$

and the linear operator equation is discretized using the MoM as

$$\overline{\mathbf{A}} \cdot \mathbf{x} = \overline{\mathbf{F}}_{tb} \cdot \mathbf{b}. \quad (3.12)$$

The result is that all manner of continuous linear operators are readily expressed in terms of basis function coefficients on either side so that if $\overline{\mathbf{A}}$, \mathbf{x} , and $\overline{\mathbf{F}}_{tb}$ are known and \mathbf{b} is unknown, one obtains

$$\mathbf{b} = \overline{\mathbf{F}}_{tb}^{-1} \overline{\mathbf{A}} \cdot \mathbf{x}, \quad (3.13)$$

assuming that $\overline{\mathbf{F}}_{tb}$ is invertible. Likewise, if $\overline{\mathbf{A}}$, \mathbf{b} , and $\overline{\mathbf{F}}_{tb}$ are known and \mathbf{x} is unknown, then

$$\mathbf{x} = \overline{\mathbf{A}}^{-1} \overline{\mathbf{F}}_{tb} \cdot \mathbf{b}, \quad (3.14)$$

under the assumption that $\overline{\mathbf{A}}$ is invertible.

This method of discretizing an SIE also conveniently expresses a Poynting vector surface integral using surface finite element functions. If real-valued finite element functions $f_E\rangle$ and $f_H\rangle$ are used to express the electric and

magnetic fields on a closed surface S such that

$$\mathbf{E}(\mathbf{r}) = \mathbf{f}_E^\top \cdot \mathbf{e}, \quad (3.15)$$

$$\mathbf{H}(\mathbf{r}) = \mathbf{f}_H^\top \cdot \mathbf{h}, \quad (3.16)$$

then the complex Poynting vector

$$\mathbf{S}(\mathbf{r}) = \frac{1}{2} \mathbf{E}(\mathbf{r}) \times \mathbf{H}^*(\mathbf{r}) \quad (3.17)$$

is expressed as

$$\mathbf{S}(\mathbf{r}) = \frac{1}{2} (\mathbf{f}_E^\top \cdot \mathbf{e}) \times (\mathbf{f}_H^\top \cdot \mathbf{h}^*). \quad (3.18)$$

The complex Poynting vector is projected onto the surface normal $\hat{n}(\mathbf{r})$, which allows a trivial restatement using the scalar triple product

$$\hat{n}(\mathbf{r}) \cdot \mathbf{S}(\mathbf{r}) = \frac{1}{2} (\mathbf{f}_H^\top \cdot \mathbf{h}^*) \cdot \hat{n} \times \mathbf{f}_E^\top \cdot \mathbf{e}. \quad (3.19)$$

Performing the surface integral of this expression is no different from testing $\hat{n} \times \mathbf{E}(\mathbf{r})$ with $\mathbf{H}(\mathbf{r})$:

$$\oint_S d^2r \hat{n}(\mathbf{r}) \cdot \mathbf{S}(\mathbf{r}) = \frac{1}{2} \mathbf{h}^* \cdot \langle \mathbf{f}_H, \hat{n} \times \mathbf{f}_E \rangle \cdot \mathbf{e}. \quad (3.20)$$

Therefore the Poynting vector is integrated over a closed surface using only the Gramian $\langle \mathbf{f}_H, \hat{n} \times \mathbf{f}_E \rangle$ and the expansion coefficients of the electric and magnetic fields.

3.1.1 Identity in Non-Orthogonal Basis

If the basis functions are orthonormal, or $\langle d_m, c_n \rangle = \delta_{mn}$, then $b_m = \langle d_m, a \rangle$. If the basis functions are finite in support but nonorthogonal, this is a sparse system of equations and may be solved with little computational expense, allowing expression of $a \rangle$ using the basis functions:

$$a \rangle = \mathbf{c}^\top \cdot \overline{\mathbf{F}}^{-1} \cdot \mathbf{e}. \quad (3.21)$$

Inserting Eq. (3.21) into Eq. (3.3) reveals a discretized identity operator [9, Sect. 2.8]

$$\mathcal{I}_{cd} = \mathbf{c}^\top \cdot \overline{\mathbf{F}}^{-1} \cdot \langle \mathbf{d}. \quad (3.22)$$

In the process of discretizing compound continuous operators, this identity operation is inserted between continuous operators to obtain discretized matrix-matrix operations. Without loss of generality to greater counts of continuous operators, this technique is applied to a two-operator linear problem

$$\mathcal{A}\mathcal{B}x = b. \quad (3.23)$$

If the above process of discretization via MoM is applied to this equation, one obtains

$$\langle f_{t,i}, \mathcal{A}\mathcal{B}\mathbf{f}_x^\top \rangle \cdot \mathbf{x} = \langle f_{t,i}, \mathbf{f}_b^\top \rangle \cdot \mathbf{b}. \quad (3.24)$$

The difficulty posed by the adjacent continuous operators is overcome by evaluating the range of \mathcal{A} and the domain of \mathcal{B} on a shared interface S_{AB} on which an identity can be evaluated.

$$\langle f_{t,i}, \mathcal{A}\mathcal{I}_{S_{AB}}\mathcal{B}\mathbf{f}_x^\top \rangle \cdot \mathbf{x} = \langle f_{t,i}, \mathbf{f}_b^\top \rangle \cdot \mathbf{b}, \quad (3.25)$$

$$\langle f_{t,i}, \mathcal{A}\mathbf{f}_a^\top \rangle \cdot \overline{\mathbf{F}}_{ab}^{-1} \cdot \langle \mathbf{f}_b, \mathcal{B}\mathbf{f}_x^\top \rangle \cdot \mathbf{x} = \langle f_{t,i}, \mathbf{f}_b^\top \rangle \cdot \mathbf{b}, \quad (3.26)$$

$$\overline{\mathbf{A}}_{ta} \cdot \overline{\mathbf{F}}_{ab}^{-1} \cdot \overline{\mathbf{B}}_{bx} \cdot \mathbf{x} = \overline{\mathbf{F}}_{tb} \cdot \mathbf{b}, \quad (3.27)$$

$$\langle f_{t,i}, \mathcal{A}\mathcal{B}\mathbf{f}_x^\top \rangle = \overline{\mathbf{A}}_{ta} \cdot \overline{\mathbf{F}}_{ab}^{-1} \cdot \overline{\mathbf{B}}_{bx}. \quad (3.28)$$

Therefore the identity defined between \mathcal{A} and \mathcal{B} facilitates discretizing compound continuous operators in a simple manner.

The practical application of this method will require either factorization of the Gramian matrix or repeated iterative solution. The function spaces related by the identity may require an asymmetric Gramian matrix, so the Gramian inverse will be treated as a general matrix. In this work, finite elements with real functions are used with complex coefficients on mesh elements. The Gramians are therefore real and sparse, which will greatly reduce the computational burden.

3.1.2 Discretized EFIE for PEC Objects

A PEC object defined by surface Σ satisfies the boundary condition

$$\hat{n} \times (\mathbf{E}^{\text{inc}} + \mathbf{E}^{\text{sca}}) = \mathbf{0}, \quad (3.29)$$

which is to say that the tangential incident electric field is exactly cancelled by the scattered electric field on Σ with no restriction placed on the magnetic field. The equivalent radiating electric surface currents $\mathbf{J}_{\Sigma}^{\text{sca}} = \hat{n} \times \mathbf{H}_{\Sigma}^{\text{sca}}$ and incident fields satisfying EFIE are expanded in tangential functions on the surface of the scatterer:

$$\mathcal{L}_{EJ} \mathbf{f}^{\top} \cdot \mathbf{j}^{\text{sca}} = -\mathbf{f}^{\top} \cdot \mathbf{e}^{\text{inc}}. \quad (3.30)$$

Testing with the basis functions indexed m as $\langle f_m$ (Galerkin's method) yields a square matrix system

$$\sum_{n=1}^N \langle f_m, \mathcal{L}^{\text{EJ}} f_n \rangle j_n^{\text{sca}} = -\sum_{n=1}^N \langle f_m, f_n \rangle e_n^{\text{inc}} \quad (3.31)$$

$$\overline{\mathbf{L}}^{\text{EJ}} \cdot \mathbf{j}^{\text{sca}} = -\overline{\mathbf{F}} \cdot \mathbf{e}^{\text{inc}}. \quad (3.32)$$

Now the system of equations may be solved numerically; this is a familiar form of EFIE for PEC objects in terms of RWG function [10] currents and incident fields. Noting that the \mathbf{f} magnetic current basis corresponds to a $-\hat{n} \times \mathbf{f}$ electric field basis, the scattered magnetic fields are expressed in the manner of Eq. (3.14)

$$M_{\Sigma}^{\text{inc}} \rangle = -\sum_{n=1}^N \hat{n} \times f_n \rangle e_n^{\text{inc}}, \quad (3.33)$$

$$[\mathbf{m}^{\text{inc}}]_n = e_n^{\text{inc}}, \quad (3.34)$$

$$\mathbf{j}^{\text{sca}} = -\left(\overline{\mathbf{L}}^{\text{EJ}}\right)^{-1} \overline{\mathbf{F}} \cdot \mathbf{m}^{\text{inc}}. \quad (3.35)$$

Thus one defines a numerical relationship between coefficients for scattered electric currents in the f space and incident magnetic currents both expressed in the f space assuming that the discretized $\overline{\mathbf{L}}^{\text{EJ}}$ operator is invertible.

3.1.3 Translation Operator

Moving the observation surface away from the surface of the scatterer surface yields a finite volume. As a building block step, one can assert zero sources on one surface Σ and the only sources on a surface σ ; the topological relationship between the surfaces is not specified. Were no scatterer present in the region containing the sources, the field on the surface Σ would be expressed as

$$\begin{bmatrix} \mathbf{E}_\Sigma \\ \mathbf{H}_\Sigma \end{bmatrix}_{tan} = \mathcal{T}_{\Sigma\sigma} \begin{bmatrix} \mathbf{M}_\sigma \\ \mathbf{J}_\sigma \end{bmatrix}. \quad (3.36)$$

This expression of the tangential fields on Σ is expressed as equivalent currents radiating in the direction \hat{n} and expanded in basis functions

$$\begin{bmatrix} \langle (\mathbf{f}_\Sigma^M)^\top \cdot \mathbf{m}_\Sigma \rangle \\ \langle (\mathbf{f}_\Sigma^J)^\top \cdot \mathbf{j}_\Sigma \rangle \end{bmatrix} = \begin{bmatrix} -\hat{n} \times & 0 \\ 0 & \hat{n} \times \end{bmatrix} \mathcal{T}_{\Sigma\sigma} \begin{bmatrix} \langle (\mathbf{f}_\sigma^M)^\top \cdot \mathbf{m}_\sigma \rangle \\ \langle (\mathbf{f}_\sigma^J)^\top \cdot \mathbf{j}_\sigma \rangle \end{bmatrix}. \quad (3.37)$$

The equations for the translated fields may be tested on the observation surface using a set of functions on that surface to discretize them. This discretized form is converted into a set of linear equations suitable to computerized solution by testing with a number of functions on the inner surface to form a square system of equations

$$\left[\overline{\mathbf{K}}_{\Sigma\sigma}^{EM} \right]_{mn} = \langle \hat{n} \times \mathbf{f}_\Sigma^M, \mathcal{K}^{EM} (\mathbf{f}_\sigma^M)^\top \rangle, \quad (3.38)$$

$$\left[\overline{\mathbf{L}}_{\Sigma\sigma}^{EJ} \right]_{mn} = \langle \hat{n} \times \mathbf{f}_\Sigma^M, \mathcal{L}^{EJ} (\mathbf{f}_\sigma^J)^\top \rangle, \quad (3.39)$$

$$\left[\overline{\mathbf{K}}_{\Sigma\sigma}^{HJ} \right]_{mn} = -\langle \hat{n} \times \mathbf{f}_\Sigma^J, \mathcal{K}^{HJ} (\mathbf{f}_\sigma^J)^\top \rangle, \quad (3.40)$$

$$\left[\overline{\mathbf{L}}_{\Sigma\sigma}^{HM} \right]_{mn} = -\langle \hat{n} \times \mathbf{f}_\Sigma^J, \mathcal{L}^{HM} (\mathbf{f}_\sigma^M)^\top \rangle, \quad (3.41)$$

$$\begin{bmatrix} \overline{\mathbf{F}}_\Sigma^M & 0 \\ 0 & \overline{\mathbf{F}}_\Sigma^J \end{bmatrix} \begin{bmatrix} \mathbf{m}_\Sigma \\ \mathbf{j}_\Sigma \end{bmatrix} = \begin{bmatrix} \overline{\mathbf{K}}_{\Sigma\sigma}^{EM} & \overline{\mathbf{L}}_{\Sigma\sigma}^{EJ} \\ \overline{\mathbf{L}}_{\Sigma\sigma}^{HM} & \overline{\mathbf{K}}_{\Sigma\sigma}^{HJ} \end{bmatrix} \begin{bmatrix} \mathbf{m}_\sigma \\ \mathbf{j}_\sigma \end{bmatrix}. \quad (3.42)$$

If the Gramians for the observation surface functions are invertible, then this expression of the equivalent currents on the observation surface Σ with respect to the source currents is readily converted into a discrete translation

operator

$$\bar{\mathbf{T}}_{\Sigma\sigma} = \begin{bmatrix} (\bar{\mathbf{F}}_{\Sigma}^{\text{M}})^{-1} & 0 \\ 0 & (\bar{\mathbf{F}}_{\Sigma}^{\text{J}})^{-1} \end{bmatrix} \begin{bmatrix} \bar{\mathbf{K}}_{\Sigma\sigma}^{\text{EM}} & \bar{\mathbf{L}}_{\Sigma\sigma}^{\text{EJ}} \\ \bar{\mathbf{L}}_{\Sigma\sigma}^{\text{HM}} & \bar{\mathbf{K}}_{\Sigma\sigma}^{\text{HJ}} \end{bmatrix}, \quad (3.43)$$

$$\begin{bmatrix} \mathbf{m}_{\Sigma} \\ \mathbf{j}_{\Sigma} \end{bmatrix} = \bar{\mathbf{T}}_{\Sigma\sigma} \begin{bmatrix} \mathbf{m}_{\sigma} \\ \mathbf{j}_{\sigma} \end{bmatrix}. \quad (3.44)$$

Thus the process of translating fields between surface functions is discretized using basis and test functions on the respective source and observation surfaces. This is noteworthy because the results of a translation operation may be used as the sources to another translation operation or scattering operation.

3.1.4 Equivalent Scattering Operator

Moving the surface Σ for evaluating the scattering operator away from the surface of the scatterer σ leads to the construction of an equivalent scattering operator [5, 6, 11] in a sequence of operations:

- translation from the outer surface to the inner scatterer $\mathcal{T}_{\sigma\Sigma}$,
- localized scattering effects on the scatterer surface \mathcal{S}_{σ} , and
- translation from the inner scatterer to the outer surface $\mathcal{T}_{\Sigma\sigma}$.

For some subcomponents \mathcal{S}_{σ} may be comparatively expensive to compute as it requires solution for equivalent currents. The outside-in and inside-out operators translate its effects to the interface surface, possibly allowing coarser sampling at the interface [5] because the evanescent field components near the antenna structures decay over distance. The net effect of these three operations is a new scattering operator at the surface of the exterior surface

$$\mathcal{S}_{\Sigma} = \mathcal{T}_{\Sigma\sigma} \mathcal{S}_{\sigma} \mathcal{T}_{\sigma\Sigma}, \quad (3.45)$$

$$\begin{bmatrix} (\mathbf{f}_{\Sigma}^{\text{M}})^{\text{T}} \cdot \mathbf{m}_{\Sigma}^{\text{sca}} \\ (\mathbf{f}_{\Sigma}^{\text{J}})^{\text{T}} \cdot \mathbf{j}_{\Sigma}^{\text{sca}} \end{bmatrix} = \mathcal{T}_{\Sigma\sigma} \mathcal{S}_{\sigma} \mathcal{T}_{\sigma\Sigma} \begin{bmatrix} (\mathbf{f}_{\Sigma}^{\text{M}})^{\text{T}} \cdot \mathbf{m}_{\Sigma}^{\text{inc}} \\ (\mathbf{f}_{\Sigma}^{\text{J}})^{\text{T}} \cdot \mathbf{j}_{\Sigma}^{\text{inc}} \end{bmatrix}. \quad (3.46)$$

These equations are discretized by testing on the outer surface Σ with testing functions:

$$\begin{bmatrix} \overline{\mathbf{F}}_{\Sigma}^{\text{M}} \mathbf{m}_{\Sigma}^{\text{sca}}, \\ \overline{\mathbf{F}}_{\Sigma}^{\text{J}} \mathbf{j}_{\Sigma}^{\text{sca}} \end{bmatrix} = \begin{bmatrix} \langle \mathbf{f}_{m,\Sigma}^{\text{M}} & 0 \\ 0 & \langle \mathbf{f}_{m,\Sigma}^{\text{J}} \end{bmatrix} \mathcal{T}_{\Sigma\sigma} \mathcal{S}_{\sigma} \mathcal{T}_{\sigma\Sigma} \begin{bmatrix} (\mathbf{f}_{\Sigma}^{\text{M}})^{\top} \cdot \mathbf{m}_{\Sigma}^{\text{inc}} \\ (\mathbf{f}_{\Sigma}^{\text{J}})^{\top} \cdot \mathbf{j}_{\Sigma}^{\text{inc}} \end{bmatrix}. \quad (3.47)$$

A discretized identity expressed in a basis set on σ , per Eq. (3.22), is applied between the translation and scattering operators for both \mathbf{E} and \mathbf{H} in the manner of Eq. (3.28) to obtain

$$\overline{\mathbf{S}}_{\sigma} = \begin{bmatrix} \left(\overline{\mathbf{F}}_{\sigma}^{\text{M}}\right)^{-1} & 0 \\ 0 & \left(\overline{\mathbf{F}}_{\sigma}^{\text{J}}\right)^{-1} \end{bmatrix} \begin{bmatrix} \langle \mathbf{f}_{\sigma}^{\text{M}}, \mathcal{S}_{\sigma}^{\text{MM}} (\mathbf{f}_{\sigma}^{\text{M}})^{\top} \rangle & \langle \mathbf{f}_{\sigma}^{\text{M}}, \mathcal{S}_{\sigma}^{\text{MJ}} (\mathbf{f}_{\sigma}^{\text{J}})^{\top} \rangle \\ \langle \mathbf{f}_{\sigma}^{\text{J}}, \mathcal{S}_{\sigma}^{\text{JM}} (\mathbf{f}_{\sigma}^{\text{M}})^{\top} \rangle & \langle \mathbf{f}_{\sigma}^{\text{J}}, \mathcal{S}_{\sigma}^{\text{JJ}} (\mathbf{f}_{\sigma}^{\text{J}})^{\top} \rangle \end{bmatrix}, \quad (3.48)$$

$$\begin{bmatrix} \mathbf{m}_{\Sigma}^{\text{sca}} \\ \mathbf{j}_{\Sigma}^{\text{sca}} \end{bmatrix} = \overline{\mathbf{T}}_{\Sigma\sigma} \overline{\mathbf{S}}_{\sigma} \overline{\mathbf{T}}_{\sigma\Sigma} \begin{bmatrix} \mathbf{m}_{\Sigma}^{\text{inc}} \\ \mathbf{j}_{\Sigma}^{\text{inc}} \end{bmatrix}. \quad (3.49)$$

In this way, the coefficients for the scattered field are expressed as a function of the coefficients for the incident field.

3.1.5 Equivalence Principle Algorithm

The above development of discrete scattering and translation operators is readily applied to Eq. (2.16) to form a discretized EPA system of equations

$$\begin{bmatrix} \overline{\mathbf{I}}_A & -\overline{\mathbf{S}}_A \overline{\mathbf{T}}_{AB} & -\overline{\mathbf{S}}_A \overline{\mathbf{T}}_{AC} \\ -\overline{\mathbf{S}}_B \overline{\mathbf{T}}_{BA} & \overline{\mathbf{I}}_B & -\overline{\mathbf{S}}_B \overline{\mathbf{T}}_{BC} \\ -\overline{\mathbf{S}}_C \overline{\mathbf{T}}_{CA} & -\overline{\mathbf{S}}_C \overline{\mathbf{T}}_{CB} & \overline{\mathbf{I}}_C \end{bmatrix} \begin{bmatrix} \mathbf{m}_A^{\text{sca}} \\ \mathbf{j}_A^{\text{sca}} \\ \mathbf{m}_B^{\text{sca}} \\ \mathbf{j}_B^{\text{sca}} \\ \mathbf{m}_C^{\text{sca}} \\ \mathbf{j}_C^{\text{sca}} \end{bmatrix} = \begin{bmatrix} 0 \\ \overline{\mathbf{S}}_B \overline{\mathbf{T}}_{BA} \\ \overline{\mathbf{S}}_C \overline{\mathbf{T}}_{CA} \end{bmatrix} \begin{bmatrix} \mathbf{m}_A^{\text{inc}} \\ \mathbf{j}_A^{\text{inc}} \\ \mathbf{m}_A^{\text{inc}} \\ \mathbf{j}_A^{\text{inc}} \end{bmatrix}. \quad (3.50)$$

This discretization of Eq. (2.16) illustrates the relationship between the incident and scattered fields in terms of abstract operators. The required operators are the identity $\overline{\mathbf{I}}$, dense translators $\overline{\mathbf{T}}_{\{ABC\}\{ABC\}}$, and (generally) dense

scattering operators $\overline{\mathbf{S}}_{\{ABC\}}$. This system is readily constructed in this form for numerical solution, although the operators need not be filled as dense operators, as developed subsequently.

3.2 Obtaining Numerical Solutions

Whatever the choice of basis and test functions, Eq. (3.50) is a dense system of linear equations of the form $\overline{\mathbf{A}}\mathbf{x} = \mathbf{b}$, where

- $\overline{\mathbf{A}}$ is the dense EPA system matrix defined by geometry and boundary conditions,
- \mathbf{x} is a vector of unknown coefficients for basis functions modeling tangential scattered fields enclosing problem domains, and
- \mathbf{b} is a vector of known coefficients for basis functions modeling tangential incident fields enclosing problem domains.

Filling a dense matrix $\overline{\mathbf{A}}$ of N equations and N unknowns demands $\mathcal{O}(N^2)$ complexity in computational time and storage. Direct solution methods (e.g. LU [12, Sect. 3.2] or QR [12, Sect. 5.2] factorizations) for a densely filled matrix require $\mathcal{O}(N^3)$ computation times.

Iterative solver methods construct an approximate solution at each iteration, with the goal of obtaining an approximate solution \mathbf{x}_n at iteration n such that $\overline{\mathbf{A}}\mathbf{x}_n - \mathbf{b} \approx \mathbf{0}$. Because the computational time complexity of the simple matrix-vector-product (MVP) is $\mathcal{O}(N^2)$, an iterative solver performing a single MVP at each of n iterations would bear computational complexity $\mathcal{O}(nN^2)$; assuming $n < N$ this is an improvement over direct solution. The discretized system in Eq. (3.50) is generally asymmetric, and the generalized minimal residual (GMRES) algorithm [13] provides a method to compute solutions within a specified tolerance. Next we will develop an abstract translator operator with controllable error and both lower computational time and storage complexity than the $\mathcal{O}(N^2)$ of the original dense matrix.

3.3 Fast Translator Matrix-Vector-Product Operations

Equation (3.50) reveals that the discretized translator defined in Eq. (3.43) is critical to efficient simulations. The discretized $\bar{\mathbf{L}}$, $\bar{\mathbf{K}}$, and $\bar{\mathbf{F}}^{-1}$ comprise the discretized $\bar{\mathbf{T}}_{\Sigma\sigma}$ in Eq. (3.43). $\bar{\mathbf{L}}$ and $\bar{\mathbf{K}}$ are both dense operators, but they are readily implemented as abstract operators using methods developed for accelerating EFIE and MFIE solutions, respectively. $\bar{\mathbf{F}}^{-1}$ need not be computed explicitly: because the Gramian is very sparse for finite elements, this inversion is efficiently effected by solving the sparse system after applying the other operators in the equation. Therefore $\bar{\mathbf{T}}_{\Sigma\sigma}$ may be implemented as a composite operator of accelerated MVP operations for $\bar{\mathbf{L}}$ and $\bar{\mathbf{K}}$ operators followed by a direct sparse solution of Eq. (3.42) to normalize for the basis-function weighting coefficients.

If surfaces Σ and σ have N_Σ and N_σ field samples, respectively, then the computational and memory expense to fill $\bar{\mathbf{L}}$ and $\bar{\mathbf{K}}$ operators, and therefore $\bar{\mathbf{T}}_{\Sigma\sigma}$, is $\mathcal{O}(N_\Sigma N_\sigma)$. The computational complexity to perform a MVP using one of the densely filled operators is $\mathcal{O}(N_\Sigma N_\sigma)$. Various techniques are available to solve SIE formulations based upon these operators. Owing to the operators' low-rank off-diagonal blocks, adaptive cross approximation (ACA) offers an algebraic method to reduce the memory and computational complexity to $\mathcal{O}(N^{\frac{4}{3}} \log N)$ for wave physics EM problems [14]. If the Green's function may be diagonalized with a similarity transformation, the fast multipole method (FMM) offers polynomial $\mathcal{O}(N^{\frac{3}{2}})$ computational complexity [15, Sect. 3.2.2], and its hierarchical counterpart the multilevel fast multipole algorithm (MLFMA) has linearithmic $\mathcal{O}(N \log N)$ storage and computational complexity [15, Sect. 3.3] with controllable error. Because of the superior asymptotic costs and physical insight, MLFMA is to be applied to $\bar{\mathbf{L}}$ and $\bar{\mathbf{K}}$ (and thus to $\bar{\mathbf{T}}_{\Sigma\sigma}$) to reduce the computational and memory costs of both the abstract operator construction and the MVM from polynomial $\mathcal{O}(N_\Sigma \log N_\Sigma + N_\sigma \log N_\sigma)$.

3.3.1 Fast Multipole Method

For brevity, only the key points of the FMM and MLFMA salient to computing a matrix-vector product in Eq. (3.43) with algorithmic complexity less than $\mathcal{O}(N_\Sigma N_\sigma)$ (brute force quadrature) are stated here. The details are

available in the literature [15]. This analysis assumes partitioning of all test and basis functions into uniform cubic subdivision of voxels. The essential operation is to compute the effect on a single voxel's contained operator-range test functions due to a distant voxel's contained operator-domain basis functions. Aggregating effects between groups by factorizing the scalar Green's function in k -space both elucidates the salient physics and reduces the work done from the brute force implementation.

The scalar Green's function for homogeneous media may be expressed for non-overlapping source and observation points as an integral over the Ewald sphere and a truncated convergent summation. Express the vector between any two points \mathbf{r}_i and \mathbf{r}_j as a vector between their respective representative points (e.g. voxel centroids) \mathbf{r}_m and $\mathbf{r}_{m'}$ plus two (shorter) vectors from the reference points to the nearby original points [15, Eq. (3.16)]

$$\mathbf{r}_{ji} = \mathbf{r}_{jm} + \mathbf{r}_{mm'} - \mathbf{r}_{im'}. \quad (3.51)$$

This allows one to split the analysis into factors characterizing propagation along the main ray from voxel m' to m and factors defining local source and observation patterns about their respective reference points; this obtains [15, Eqs. (3.17)-(3.18)]

$$\frac{\exp(ik|\mathbf{r}_{ji}|)}{|\mathbf{r}_{ji}|} = \int d^2\hat{k} \exp(i\mathbf{k} \cdot (\mathbf{r}_{jm} - \mathbf{r}_{im'})) \alpha(\mathbf{k}, \mathbf{r}_{mm'}), \quad (3.52)$$

$$\alpha(\mathbf{k}, \mathbf{r}_{mm'}) = \frac{ik}{4\pi} \sum_{l=0}^L i^l (2l+1) h_l^{(1)}(k|\mathbf{r}_{mm'}|) P_l(\hat{k} \cdot \hat{\mathbf{r}}_{mm'}). \quad (3.53)$$

The scalar Green's function is thus an integral over all plane waves from the source to the observation. This representation of the scalar Green's function is inserted into the dyadic Green's function to obtain [15, Eq. (3.19)]

$$\overline{\mathbf{G}}(\mathbf{r}_j, \mathbf{r}_i) \approx \frac{1}{4\pi} \int d^2\hat{k} (\overline{\mathbf{I}} - \hat{k}\hat{k}) \exp(i\mathbf{k} \cdot (\mathbf{r}_{jm} - \mathbf{r}_{im'})) \alpha(\mathbf{k}, \mathbf{r}_{mm'}). \quad (3.54)$$

This expression reflects a plane wave relationship between a source point and observation point. We appear to be worse off than when we started because we have replaced a direct evaluation of the dyadic Green's function with an integral over the energy shell. The relationships between collections of source and observation points are connected in the integral kernel,

and only the plane wave components relating the source and observation points to their respective representative points \mathbf{r}_m and $\mathbf{r}_{m'}$ change with each source/observation point. If the k -space integral is evaluated numerically, then it may be performed at a specified quadrature points; the kernel values are evaluated between each voxel-representative point pair at all k -space positions.

A single term of $\overline{\mathbf{L}}$ is obtained by discretizing Eq. (3.54) using a basis function $\mathbf{b}_i(\mathbf{r}')$ and a test function $\mathbf{t}_j(\mathbf{r})$.

$$[\overline{\mathbf{L}}]_{ji} = [\langle \mathbf{t}_j, \overline{\mathbf{G}} \mathbf{b}_i \rangle]_{ji}, \quad (3.55)$$

$$= \iint_{\Sigma_j} d^2r \mathbf{t}_j(\mathbf{r}) \cdot \iint_{\Sigma_i} d^2r' \overline{\mathbf{G}}(\mathbf{r}, \mathbf{r}') \cdot \mathbf{b}_i(\mathbf{r}'), \quad (3.56)$$

$$\approx \frac{1}{4\pi} \int d^2\hat{k} \alpha(k\hat{k}, \mathbf{r}_{mm'}) \iint_{\Sigma_j} d^2r \left[\mathbf{t}_j(\mathbf{r}) - \hat{k} (\hat{k} \cdot \mathbf{t}_j(\mathbf{r})) \right] \exp(i k \hat{k} \cdot \mathbf{r}_{jm}) \cdot \iint_{\Sigma_i} d^2r' \left[\mathbf{b}_i(\mathbf{r}') - \hat{k} (\hat{k} \cdot \mathbf{b}_i(\mathbf{r}')) \right] \exp(-i k \hat{k} \cdot \mathbf{r}'_{im'}), \quad (3.57)$$

where the orders of integration are interchangeable because the fields are non-singular. This is simply the plane wave representation of the basis field as if radiating from its reference point $\mathbf{r}'_{m'}$ projected onto the test function observing from its reference point \mathbf{r}_m . Therefore every basis function is evaluated as outbound plane waves; the test functions are evaluated as inbound far-fields in plane waves (or “outbound” with a direction opposite the source plane wave direction). Note that $\alpha(k\hat{k}, \mathbf{r}_{mm'})$ relates the outbound and inbound plane waves as a diagonal k -space operator. The discretized \mathcal{K} operator is expressed in the same manner, and differs from $\overline{\mathbf{L}}$ only by a polarization

rotation applied to test function plane waves.

$$[\mathbf{K}]_{ji} = [\langle \mathbf{t}, \nabla' \times \overline{\mathbf{G}} \mathbf{b} \rangle]_{ji}, \quad (3.58)$$

$$= \iint_{\Sigma_j} d^2 r \mathbf{t}_j(\mathbf{r}) \cdot \iint_{\Sigma_i} d^2 r' \nabla' \times \overline{\mathbf{G}}(\mathbf{r}, \mathbf{r}') \cdot \mathbf{b}_i(\mathbf{r}'), \quad (3.59)$$

$$\begin{aligned} &\approx \frac{1}{4\pi} \int d^2 \hat{k} \alpha(k \hat{k}, \mathbf{r}_{mm'}) \\ &\quad \iint_{\Sigma_j} d^2 r \left(-\hat{k} \right) \times \left[\mathbf{t}_j(\mathbf{r}) - \hat{k} \left(\hat{k} \cdot \mathbf{t}_j(\mathbf{r}) \right) \right] \exp \left(ik \hat{k} \cdot \mathbf{r}_{jm} \right) \cdot \\ &\quad \iint_{\Sigma_i} d^2 r' \left[\mathbf{b}_i(\mathbf{r}') - \hat{k} \left(\hat{k} \cdot \mathbf{b}_i(\mathbf{r}') \right) \right] \exp \left(-ik \hat{k} \cdot \mathbf{r}'_{im'} \right). \quad (3.60) \end{aligned}$$

Therefore the block translator of Eq. (3.43) is expressed for distant basis and test functions as a diagonal k -space outbound-to-inbound (O2I) operator with domain of basis function far-fields and range of test function far-fields. This allows the expression of distant translations as diagonalized blocks in k -space and is fundamental to fast MVP operations.

These FMM operators are then a cascade of sparse operators:

- Basis functions are converted to sets of plane waves emanating from the containing domain voxel,
- Each domain-to-range voxel translation is a diagonal operator in the space of plane waves, and
- Incoming plane waves at each range voxel are converted to the test functions in the voxel.

The plane wave transformation at the range and domain voxels are specific to the plane waves defined emanating to/from the voxel reference point. This technique achieves polynomial computational complexity $\mathcal{O}\left(N^{\frac{3}{2}}\right)$ for the MVP [15, Sect. 3.2], which then brings the complexity of an n -iteration solution to $\mathcal{O}\left(nN^{\frac{3}{2}}\right)$.

3.3.2 Multilevel Fast Multipole Algorithm

The analysis is readily extended to hierarchical form by noting that the plane wave transformations for a voxel can be nested in the same way that basis functions were grouped. MLFMA is defined by a cascade of sparse operations,

and achieves linearithmic $\mathcal{O}(N \log N)$ computational complexity and storage size [15, Sect. 3.3]. Therefore, this brings the computational complexity of an n -iteration solution to $\mathcal{O}(nN \log N)$.

The development from FMM to MLFMA [15, Sect. 3.3] is summarized as follows. Because the fields are spatially band-limited, the radiation patterns computed within smaller voxels can be locally interpolated to the denser sampling of a larger voxel's pattern. Likewise, the radiation patterns computed for larger voxels can be locally antepolated to the coarser sampling of a smaller voxel's pattern. The operation radiating outbound-to-outbound (O2O) from the smallest basis voxels up to higher-level basis voxels is a series of phase shifts and local interpolations, which can be composed with sparse matrices. The same diagonal O2I operator concept from single-level FMM is used here, but at the largest possible voxel level not violating the addition theorem. The operation radiating inbound-to-inbound (I2I) from the higher-level test voxels down to the smallest test voxels is a series of phase shifts and local antepolations like O2O, and may be composed with sparse matrices. In the case of translating between different meshes σ and Σ with function counts N_σ and N_Σ , respectively, the computational and storage complexity of the MVP is $\mathcal{O}(N_\Sigma \log N_\Sigma + N_\sigma \log N_\sigma)$. This acceleration is fundamental to solving large problems.

3.3.3 Iterative Solver Computational Complexity

The existence of fast solvers warrants reconsideration of the discrete operators in Eq. (3.50) to determine the computational complexity of solving the system. The identity is discretized as a Gramian matrix; for finite elements in which $\mathcal{O}(1)$ such functions have a nontrivial Gramian entry with any other function, the cost of discretizing and performing a MVM with $\mathcal{O}(N)$ computational and memory usage. The expense of solving for the on-surface coefficients may be achieved by iterative sparse solver with complexity $\mathcal{O}(N)$ (e.g. GMRES [13]) or by a direct sparse solver with complexity $\mathcal{O}(N^2 \log N)$ (e.g. parallel sparse LU factorization [16]). (The direct sparse solver's complexity in many cases is lower than the given worst case complexity [16], but such a risk must be acknowledged.) The translators are comprised by $\overline{\mathbf{L}}$ and $\overline{\mathbf{K}}$ operators, and thus have $\mathcal{O}(N \log N)$ storage and computational

complexity. The scattering operators are generally composed using iterative solvers and fast MVM routines so that they yield $\mathcal{O}(N \log N)$ storage and computations. The operations required to define the right-hand side (RHS) of Eq. (3.50) are the same as those used in the block matrix. Therefore the discrete solution of Eq. (3.50) may be achieved with $\mathcal{O}(N \log N)$ storage and computational complexity.

3.4 Computing Solutions Using the Schur Complement

The solution of a linear system of equations is sometimes intractable for various reasons. If the solution on only part of the problem is required, then one may reformulate the original system of equations to solve for only the required domain. If the system is partitioned as

$$\begin{bmatrix} \mathcal{A} & \mathcal{B} \\ \mathcal{C} & \mathcal{D} \end{bmatrix} \begin{bmatrix} \mathbf{q} \\ \mathbf{r} \end{bmatrix} = \begin{bmatrix} \mathbf{s} \\ \mathbf{t} \end{bmatrix}, \quad (3.61)$$

then the Schur complement allows the restatement of the system without the unknowns \mathbf{r} as

$$(\mathcal{A} - \mathcal{B}\mathcal{D}^{-1}\mathcal{C}) \mathbf{q} = \mathbf{s} - \mathcal{B}\mathcal{D}^{-1}\mathbf{t}. \quad (3.62)$$

This incomplete solution is available if \mathcal{D}^{-1} is viable for calculation. The new matrix system is smaller than the original; its dimensions match those of the \mathcal{A} matrix.

CHAPTER 4

EPA HYBRIDIZED WITH PO

4.1 Introduction

Solving the discretized equations of Chapter 3 for arbitrary cases is limited by the total number of basis and test functions N used in all the domains. As the required computational storage and time scales with asymptotic complexity $\mathcal{O}(N \log N)$ when using MLFMA, the scale of N is the parameter that determines whether a problem may be solved with the available computational hardware. As the problem size kd increases with respect to the wavelength,

$$N = \mathcal{O}((kd)^2), \quad (4.1)$$

and the computational storage and time scale with respect to the linear dimension

$$\mathcal{O}(N \log N) = \mathcal{O}((kd)^2 \log(kd)). \quad (4.2)$$

This points to a challenge solving general problems where the objects have arbitrarily large size scales. Hybridizing EPA with high-frequency asymptotic methods provides a means to characterize otherwise intractable problems for SIE analysis.

In this chapter, a brief overview of the physical optics (PO) high-frequency asymptotic model is developed to form an approximate scattering operator for use within an EPA framework. Then a novel EPA-PO hybrid solution is derived, and computed results are presented, verifying the accuracy of the technique. The proposed hybrid solution establishes an equation structure that generalizes from PO to other approximate methods.

4.2 Previous Work

Hybridizing solutions between MoM and approximate/asymptotic methods is a topic with a long history. A noteworthy early hybrid expands PO currents in finite element functions as if it were part of the MoM region with a sparse self-interaction matrix [17]. A similar approach was applied to hybridizing EPA with PO for electrically connected objects [18]. These approaches are most comparable to evaluating the method developed here before performing the Schur complement stage. In application with an iterative solver, this is also comparable to the MoM-PO hybrid described in [19] and EI-MoM-PO hybrid described in [20]. However, these methods described techniques specifically computing the interactions between a PEC MoM region and a PEC PO region. The method described here is amenable to solution with “black-box” iterative solver such as GMRES or use with a compressed matrix representation framework [14].

The developed method also may be compared to the MLFMA-UTD hybrid described in [21]. This approach modifies the diagonal $\bar{\alpha}$ translation operator with UTD-calculated fields. Because the reflected rays from the approximately modeled domain will not generally align with pre-computed source/observation pattern samples, the effects of the reflected rays impose additional interpolation/interpolation calculation, beyond FMM/MLFMA conventions. The method proposed here has the advantage of supporting implementation as either a whole operator or as an augmentation to an existing homogeneous medium implementation.

4.3 Physical Optics as a Scattering Operator

To obtain a suitable approximate scattering operator, one may look to high-frequency asymptotic scattering techniques in general and to the PO approximation in particular. High-frequency asymptotic methods are developed in the literature [22, 23, 24, 25]. Among the common methods, PO stands out for its simplicity. PO approximates scattered fields using physical equivalence principle sources corresponding locally to a tangent plane approximation to a large smooth scatterer. Other methods may be used to compute local surface fields, and then the equivalent PO surface currents radiate to obser-

vation points [26]. The argument developed for PO as a hybrid retains the generality of the greater EPA solver framework so that further research may apply other asymptotic methods beyond PO.

In the local tangent plane approximation to a smooth surface, one may analyze a finite surface as a bounded portion of an infinite plane. A surface with a PEC boundary condition requires that the total tangential electric field be zero. If the surface is a part of an infinite flat plate, then phase matching the plane wave solutions on the surface requires that, given the vector orientations in Figure 4.1a the complex coefficients characterizing the fields must be

- $\hat{n} \times \mathbf{E}^{\text{sca}} \times \hat{n} = -\hat{n} \times \mathbf{E}^{\text{inc}} \times \hat{n}$ and $\hat{n} \times \mathbf{H}^{\text{sca}} \times \hat{n} = \hat{n} \times \mathbf{H}^{\text{inc}} \times \hat{n}$ on the illuminated side of the scatterer, and
- $\hat{n} \times \mathbf{E}^{\text{sca}} \times \hat{n} = -\hat{n} \times \mathbf{E}^{\text{inc}} \times \hat{n}$ and $\hat{n} \times \mathbf{H}^{\text{sca}} \times \hat{n} = -\hat{n} \times \mathbf{H}^{\text{inc}} \times \hat{n}$ on the dark side of the scatterer.

The physical equivalent model of the scenario replaces the PEC scatterer with the exterior medium to form a single homogeneous space with radiating current sheets enclosing the scatterer. As illustrated in Figure 4.1b, the equivalent currents on the illuminated side are $\mathbf{J} = 2\hat{n} \times \mathbf{H}^{\text{inc}}$, $\mathbf{M} = \mathbf{0}$ and on the dark side they are $\mathbf{J} = \mathbf{0}$, $\mathbf{M} = \mathbf{0}$. In this manner, the local tangent scattered fields are defined in terms of the incident fields. Point-wise on the surface of a planar PEC scatterer one sees that

$$\begin{bmatrix} \mathbf{E}^{\text{sca}} \\ \mathbf{H}^{\text{sca}} \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ 0 & -2\hat{n} \times \mathcal{I}_i \end{bmatrix} \begin{bmatrix} \mathbf{M}^{\text{inc}} \\ \mathbf{J}^{\text{inc}} \end{bmatrix}, \quad (4.3)$$

$$\mathbf{J}^{\text{sca}} = 2\mathbf{J}^{\text{inc}}. \quad (4.4)$$

Therefore a trivial continuous approximate scattering operator is available for an electrically large planar surface. Applying the operator discretization of Eq. (3.13) yields a trivial numerical operator if the incident and scattered fields are expressed in the same basis. The PEC surface case is stated here without loss of generality to other material types. The illuminated/shadowed distinction is applied as a scalar weight on fields radiating to the PO scattering surface. In this way, one defines a new domain type to be incorporated in an EPA-based solver framework. Moreover, the availability of a sparse scattering operator and trivially modified translators suggests

that high-frequency asymptotic scattering may facilitate solving otherwise intractable problems using Eq. (3.50).

4.4 Eliminating Unknowns on an EPA Domain

An efficiently evaluated scattering operator for one domain of Eq. (3.50) has significant impact on how one may obtain a solution. A domain approximated by a PO model may be discretized directly with the approximate scattering operator, as reported in [18]. Within the context of this analysis, such a method will be described as a “Direct PO” hybrid. Further analysis suggests a more general approach to hybridizing approximate scattering algorithms with EPA; this is developed next.

The Schur complement technique described in Section 3.4 can be applied to eliminate one of the domains. This is a straightforward manipulation, especially because of the identity on the system diagonal. If the domain to be eliminated is indexed C , then eliminating those fields from Eq. (3.50) yields

$$\begin{aligned}
& \left(\begin{bmatrix} \bar{\mathbf{I}}_A & -\bar{\mathbf{S}}_A \bar{\mathbf{T}}_{AB} \\ -\bar{\mathbf{S}}_B \bar{\mathbf{T}}_{BA} & \bar{\mathbf{I}}_B \end{bmatrix} \right. \\
& \left. - \begin{bmatrix} \bar{\mathbf{S}}_A \bar{\mathbf{T}}_{AC} \bar{\mathbf{S}}_C \bar{\mathbf{T}}_{CA} & \bar{\mathbf{S}}_A \bar{\mathbf{T}}_{AC} \bar{\mathbf{S}}_C \bar{\mathbf{T}}_{CB} \\ \bar{\mathbf{S}}_B \bar{\mathbf{T}}_{BC} \bar{\mathbf{S}}_C \bar{\mathbf{T}}_{CA} & \bar{\mathbf{S}}_B \bar{\mathbf{T}}_{BC} \bar{\mathbf{S}}_C \bar{\mathbf{T}}_{CB} \end{bmatrix} \right) \begin{bmatrix} \mathbf{m}_A^{\text{sca}} \\ \mathbf{j}_A^{\text{sca}} \\ \mathbf{m}_B^{\text{sca}} \\ \mathbf{j}_B^{\text{sca}} \end{bmatrix} \\
& = \begin{bmatrix} 0 \\ \bar{\mathbf{T}}_{BA} \begin{bmatrix} \mathbf{m}_A^{\text{inc}} \\ \mathbf{j}_A^{\text{inc}} \end{bmatrix} \end{bmatrix} + \begin{bmatrix} \bar{\mathbf{T}}_{AC} \bar{\mathbf{S}}_C \bar{\mathbf{T}}_{CA} \\ \bar{\mathbf{T}}_{BC} \bar{\mathbf{S}}_C \bar{\mathbf{T}}_{CA} \end{bmatrix} \begin{bmatrix} \mathbf{m}_A^{\text{inc}} \\ \mathbf{j}_A^{\text{inc}} \\ \mathbf{m}_A^{\text{inc}} \\ \mathbf{j}_A^{\text{inc}} \end{bmatrix}. \tag{4.5}
\end{aligned}$$

The tangential scattered fields on the surface of all objects are readily solved in this form with the caveat that they are all taken to be radiating in the presence of the platform instead of the homogeneous background medium. In this expression of the EPA multi-scatterer system, the physics of scattering within an inhomogeneous medium are made plain.

Most importantly, $\bar{\mathbf{S}}_C$ is surrounded by translation operators from the

other domains, which corresponds to radiation-scattering-radiation interaction chains that may be characterized by myriad asymptotic scattering calculations. In this view, using PO as the local scattering formulation becomes an implementation detail suitable for replacement by other formulations. The early MoM-PO hybrid method [17] was extended to include additional UTD scattering enhancements to the PO scattering regions [27]. Such an enhancement is immediately applicable to the Schur complement method proposed here. But the radiation-scattering-radiation calculations in the Schur-complement PO method are amenable to using any other scattering formulation, not strictly equivalent-current-based methods. The eliminated domain may be so electrically large that meshing it and tracking unknowns may be prohibitive; in such cases field-modeling methods such as UTD are viable candidates for integration into the hybridization framework. Subject to the restrictions of the addition theorem, MLFMA decomposition may be utilized to perform such analyses in the space of plane waves.

4.4.1 Computational Considerations

Going forward, the problem is split into N_{EPA} functions in the EPA-modeled domains and N_{PO} functions in the single PO domain. Operations on the N_{PO} functions spanning the PO surface are expected to dominate the calculations as $N_{\text{PO}} \gg N_{\text{EPA}}$. The computational steps required are discussed next in order to understand the anticipated computational complexity. The equivalent current basis functions used to model all structures are RWG [10] functions using a seven-point Gaussian quadrature rule per triangle and double precision floating point values.

In a “Direct PO” hybrid EPA-PO system (without the Schur complement applied), the cost per MLFMA MVM is $\mathcal{O}((N_{\text{EPA}} + N_{\text{PO}}) \log(N_{\text{EPA}} + N_{\text{PO}}))$. It is noteworthy that because the scattering operator is a simple scalar, one need not apply the inverse Gramian after translation to the PO surface if block rows are scaled by the Gramian. This saves an expensive sparse solver step with a cost between $\mathcal{O}(N_{\text{PO}})$ (GMRES) and $\mathcal{O}(N_{\text{PO}}^2 \log N_{\text{PO}})$ (sparse LU). In exchange for avoiding this sparse solver step, the Direct PO method explicitly solves for current unknown coefficients across the PO surface. If GMRES is used to solve the system, then $\mathcal{O}(N_{\text{PO}})$ time and storage is re-

quired in the iterative solver step. Moreover, the iterative solver will treat approximate PO unknown weights as equally important to the final solution as the fields on the more accurate antennas, leading to challenges satisfying the target solution residual.

The Schur complement EPA-PO hybrid uses the approximate PO fields only to obtain the scattered fields at other surfaces. This offers the opportunity for reduced result dataset sizes from the Schur complement method because the PO surface unknowns are not retained. In the ‘‘Schur PO’’ hybrid EPA-PO system, the same MLFMA implementation is used as in the Direct PO system, which leads to the same net complexity. Unlike the Direct PO algorithm, the Schur PO algorithm requires normalization with the inverse Gramian after translation to the N_{PO} functions because translation to the PO surfaces is immediately followed by translation back out to the other domains. Consequently, the expense of normalizing by the Gramian across N_{PO} functions is added to the Schur PO algorithm in comparison to Direct PO. Therefore, an efficient sparse solver is critical to practical application of this expression of the method. This normalization is required only because the scattering operator is discretized as other MoM-modeled domains. In future work, other asymptotic methods may avoid this expense. The Schur complement EPA-PO hybrid restates the equations in a manner suggesting that another high-frequency asymptotic (approximate) method may be used in place of PO without loss of generality to the formulation.

The solution of Eq. (4.5) differs from that of Eq. (3.50) by having a dramatically smaller space of unknown and RHS values. Evaluating a solution’s residual in the latter equation treats the fields across the PO-modeled domain as first-class quantities to be considered in evaluating a solution. As a matter of practicality, the fields on the asymptotically modeled domain are not important for evaluating the accuracy of a solution because the local boundary conditions are only approximately satisfied. Rather, the effect of radiation in the presence of the approximately modeled domain need only accurately satisfy boundary conditions on the surfaces modeled with ‘‘exact’’ methods such as EFIE or EPA. The reduced set of basis and test functions analyzed simplifies the work to be done by the iterative solver from $\mathcal{O}(N_{\text{EPA}} + N_{\text{PO}})$ to $\mathcal{O}(N_{\text{EPA}})$, which helps offset the expense of normalizing coefficients on the PO surface in this implementation.

The new hybrid also requires that the iterative solver’s residual test for

convergence be constructed with the restated problem in mind. The RHS of Eq. (4.5) contains radiative effects of both the incident field and the incident field’s scattering from the PO surface, so initial residual error is expected to be much lower than it would be for Eq. (3.50). Likewise, the effect of each iteration includes both direct radiation from non-Schur-complement domains and their scattering from the PO-modeled domains. The reduced initial residual error and the more productive iterations lead to an expected reduction in number of iterations.

4.5 Examples

The viability of the proposed method is evaluated through application to various cases. EFIE is used as the reference solution in each case to validate the computed solution. EPA hybridized with PO without the Schur-complement (labeled “Direct PO”) is also computed to verify the correct function of the Schur-complement PO hybrid (labeled “Schur PO”). In every case, the Schur PO hybrid is verified to function correctly. The comparison to the EFIE validation data indicates reliable behavior, but the PO physics results in expected deviations from EFIE. The cases are chosen so that the PO-modeled surfaces are physically separated so that radiative physics is the only coupling mechanism. Moreover, the PO-modeled surfaces are large and locally smooth so that the local tangent plane approximation is a suitable approximation to the field interactions.

4.5.1 Dipole-Plate Case

A single square plate illuminated by a single half-wavelength dipole serves as the core validation case. The case is derived from [28, Fig. 8] and is illustrated in Figure 4.2. The antenna is driven at its center with a 1 V (peak-to-peak) source in the center 3 mm of the cylinder at 8 GHz. The plate is a square in the $x = 0$ plane with side length 177.8 mm in the y - and z -directions, or about 4.745λ . The dipole is a wire aligned with the z -axis offset from the plate center in the x -direction by 130 mm (about 3.469λ). The dipole wire structure itself is not described in the reference, but it is taken to be a right circular cylinder with radius 0.75 mm. The dipole and plate are both

modeled with a PEC boundary condition.

The antenna’s installed far-field patterns are computed using multiple algorithms. The validation case “EFIE:EFIE” applies EFIE to both the dipole and the square plate. The verification case “EPA:Direct PO” applies EPA (enclosing the dipole modeled with EFIE) to the source antenna and applies PO to the plate without using the Schur complement approach to the PO hybrid. The test case “EPA:Schur PO” applies the same EPA model to the antenna and applies the Schur complement modified PO hybrid to the plate. The far-field relative electric field intensity for the validation, verification, and test cases are plotted in the horizontal ($z = 0$) plane in Figure 4.3 and in the vertical ($y = 0$) plane in Figure 4.4. The computed results for all three algorithms show strong agreement with both published UTD and measured data [28, Figs. 8(a)-(b)]. The “EFIE:EFIE” algorithm validated here is used as the reference data for the remaining scenarios, for which external data is not available. The agreement between the EFIE and EPA-PO hybrid calculations indicates that the PO calculations are an accurate approximation to the fields on the plate. The overlaid plots of the Direct PO and Schur PO results indicate that the two methods are computing similar solutions.

4.5.2 Dipole-Dihedral Case

The Dipole-Plate validation case is extended by adding a much larger rectangular plate to the scenario and observing the effects on the antenna pattern. The perpendicular plates form a dihedral, leading to a shorthand name “Dipole-Dihedral”. The added plate is a rectangle parallel to the $z = 0$ plane and spanning 40λ in the x -direction and 30λ in the y -direction. Relative to the center of the small plate, the large plate is offset in the x - and y -directions by 10λ and 7.5λ , respectively. The z -direction offset is 5λ below the bottom edge of the small plate. The large plate is also modeled using a PEC boundary condition. This scenario is visualized in Figure 4.5.

The antenna’s installed far-field patterns are computed using multiple algorithms. The validation case “EFIE:EFIE” applies EFIE to both the dipole and dihedral plates. The verification case “EPA:Direct PO” applies EPA (enclosing the dipole modeled with EFIE) to the source antenna and applies PO to the dihedral plates without using the Schur complement approach

to the PO hybrid. The test case “EPA:Schur PO” applies the same EPA model to the antenna, applies Direct PO to the small plate, and applies the Schur-complement modified PO hybrid to the large plate. The far-field relative electric field intensity for the validation, verification, and test cases are plotted in the horizontal ($z = 0$) plane in Figure 4.6, and in the vertical ($y = 0$ and $x = 0$) planes in Figures 4.7 and 4.8. The presence of two PO-modeled plates necessitates MLFMA-accelerated interactions between the antenna and the plates. The computed results for both PO algorithms agree with reference EFIE data. The overlaid plots of Direct PO and Schur PO results indicate that the two methods are computing similar solutions.

4.5.3 Dipole-Curved-Plate Case

The Dipole-Plate validation case is altered by replacing the vertical plate with a much larger horizontal curved plate and observing the effects on the antenna pattern. The perpendicular plates form a dihedral, leading to a shorthand name “Dipole-Curved-Plate”.

The curved plate is a quarter section of a right circular cylinder perpendicular to the $x = 0$ plane and spanning 40λ in the x -direction. The cylinder radius in the YZ plane is 30λ , and the quarter section is oriented so that the mid-point has a surface normal vector pointing in the $z > 0$ direction. The quarter cylinder is offset from the dipole as illustrated in Figure 4.9. This geometry exercises the PO blockage check as part of the curved plate is blocked from the antenna. The lack of surface wave support and PO shadow boundary current truncation error is expected to cause increased error relative to the flat plate cases considered.

The antenna’s installed far-field patterns are computed using multiple algorithms. The validation case “EFIE:EFIE” applies EFIE to both the dipole and curved plate. The verification case “EPA:Direct PO” applies EPA (enclosing the dipole modeled with EFIE) to the source antenna and applies PO to the curved plate without using the Schur complement approach to the PO hybrid. The test case “EPA:Schur PO” applies the same EPA model to the antenna and applies the Schur-complement modified PO hybrid to the curved plate. The far-field relative electric field intensities for the validation, verification, and test cases are plotted in the horizontal ($z = 0$) plane in

Figure 4.10, and in the vertical ($y = 0$ and $z = 0$) planes in Figures 4.11 and 4.12. The computed results for both PO algorithms generally agree with reference EFIE data, with less-reliable agreement in the far-field region shadowed by the curved plate. The overlaid plots of Direct PO and Schur PO results indicates that the two methods are computing similar solutions.

4.6 Conclusions

The Schur PO algorithm is demonstrated to correctly compute the hybridization of EPA and PO in radiative coupling cases. The algorithm computes solutions over other domains radiating in the presence of the Schur complement PO domain. This shrinks the solution size for the primary solver at the expense of an additional (sparse) Gramian solution over the PO domain. This added expense is mitigated by the physically intuitive equation structure obtained, which will support alternative scattering operators with the same operator domain and range interfaces as the Schur complement without the cost of the sparse Gramian solution. Moreover, the obtained equations reveal the opportunity for parallel execution discussed next.

4.7 Figures

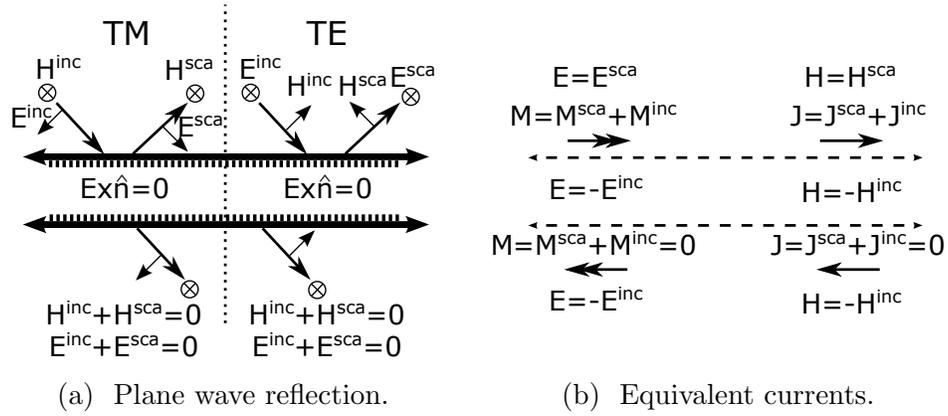


Figure 4.1: Plane wave reflection and physically equivalent surface current model for two parallel PEC planes separated by a finite distance.

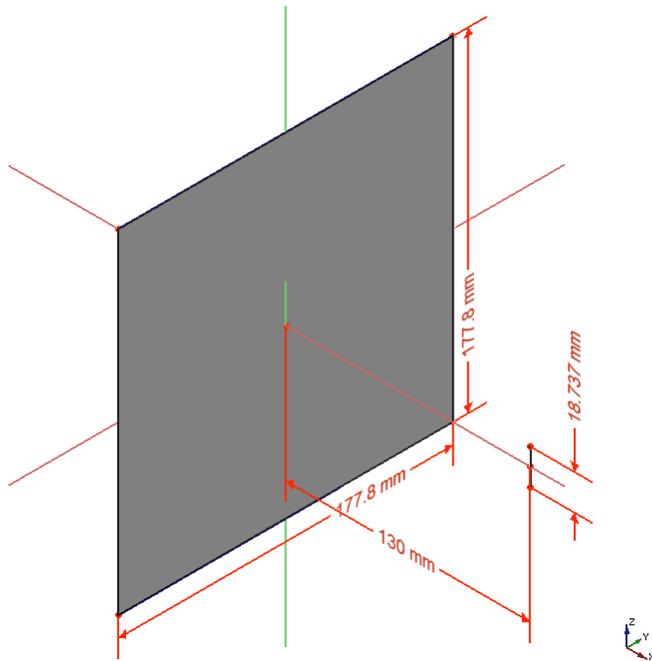


Figure 4.2: Half-wavelength dipole parallel to square plate following [28, Fig. 8].

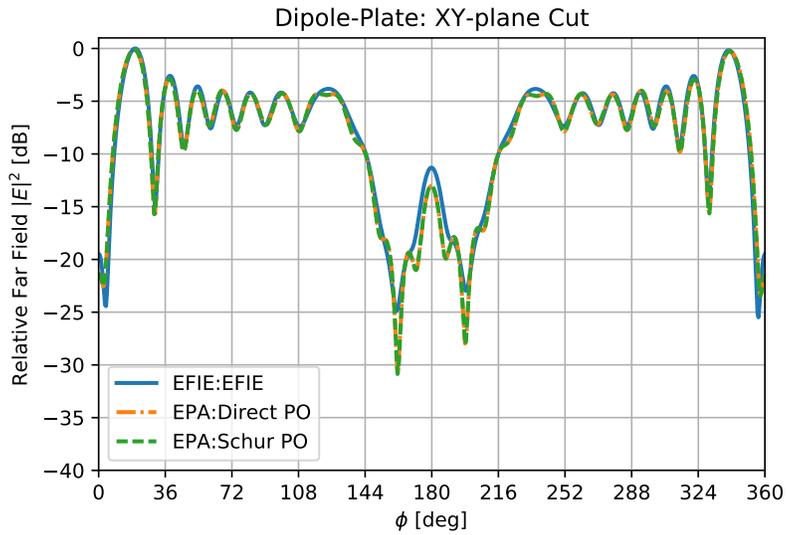


Figure 4.3: Horizontal plane ($z = 0$) cut of far-field intensity for a half-wavelength dipole parallel to a square plate.

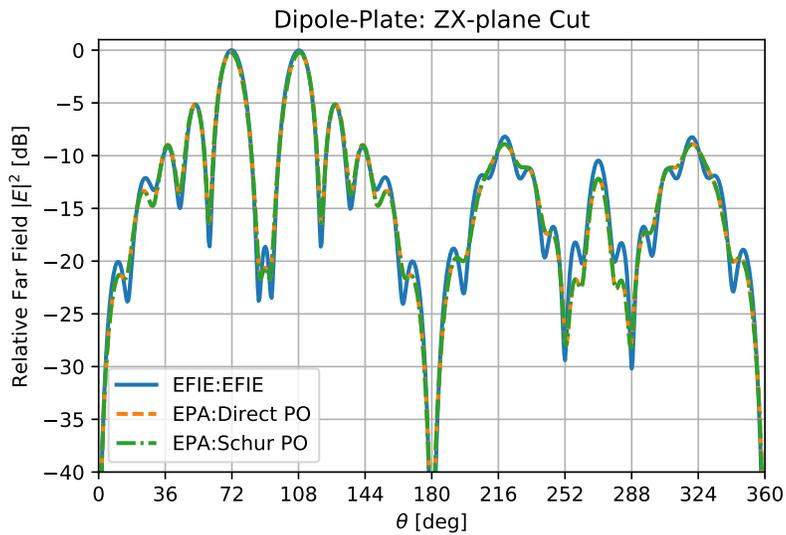


Figure 4.4: Vertical plane ($y = 0$) cut of far-field intensity for a half-wavelength dipole parallel to a square plate.

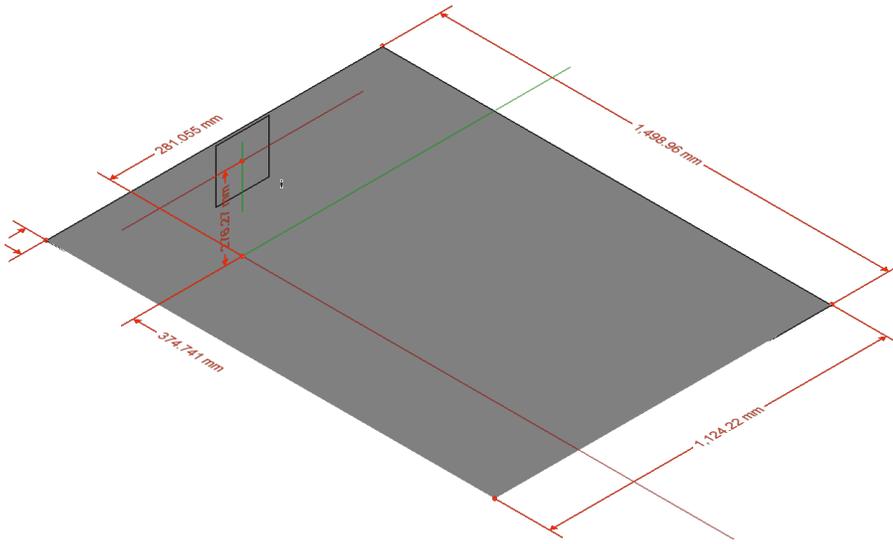


Figure 4.5: Half-wavelength dipole parallel to square plate and perpendicular to a large rectangular plate.

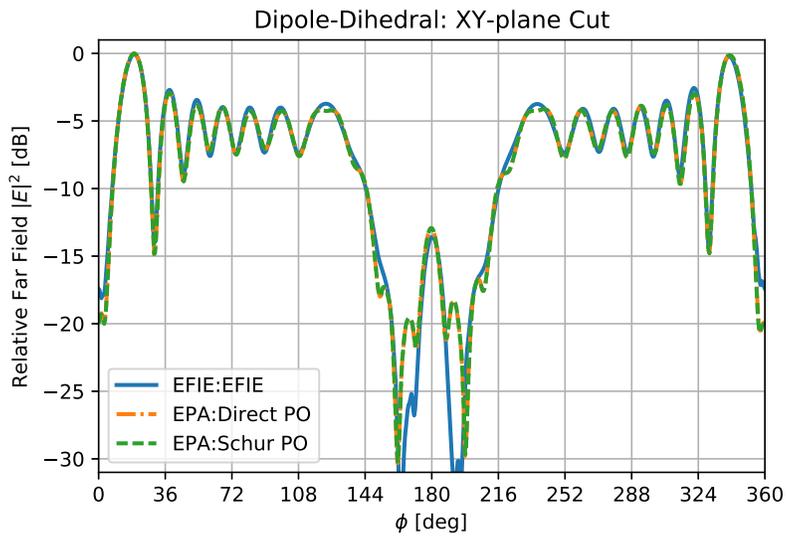


Figure 4.6: Horizontal plane ($z = 0$) cut of far-field intensity for a half-wavelength dipole parallel to a square plate and perpendicular to a large rectangular plate.

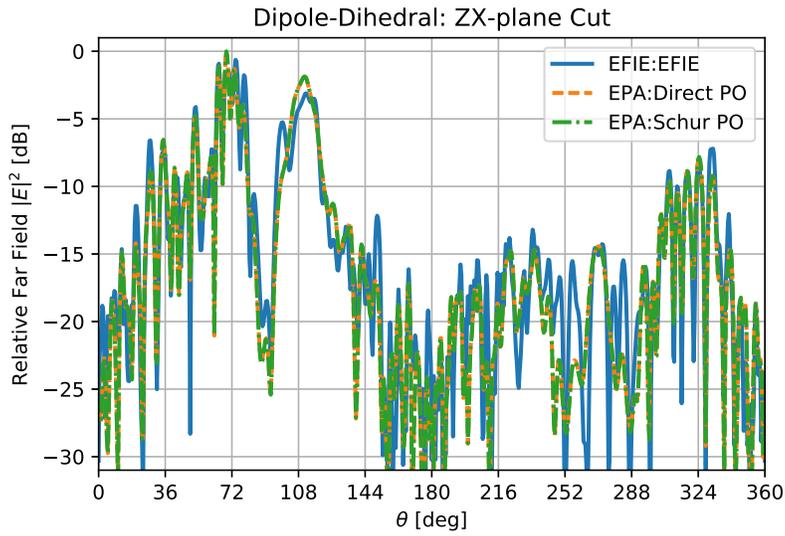


Figure 4.7: Vertical plane ($y = 0$) cut of far-field intensity for a half-wavelength dipole parallel to a square plate and perpendicular to a large rectangular plate.

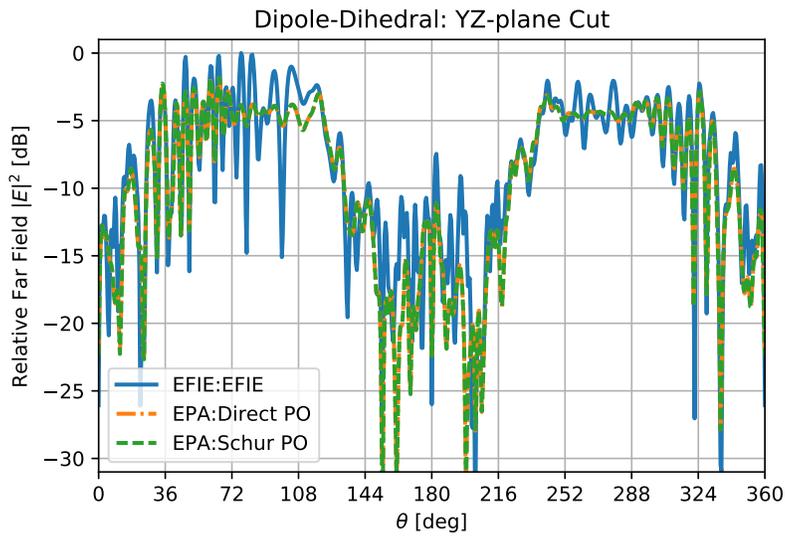


Figure 4.8: Vertical plane ($x = 0$) cut of far-field intensity for a half-wavelength dipole parallel to a square plate and perpendicular to a large rectangular plate.

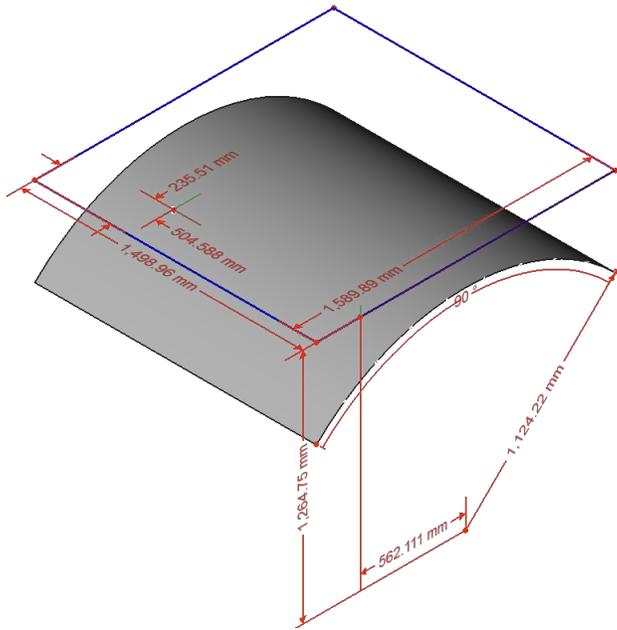


Figure 4.9: Half-wavelength dipole above a large curved plate; the quadrilateral outline is the projection of the curved plate outline into the $z = 0$ plane.

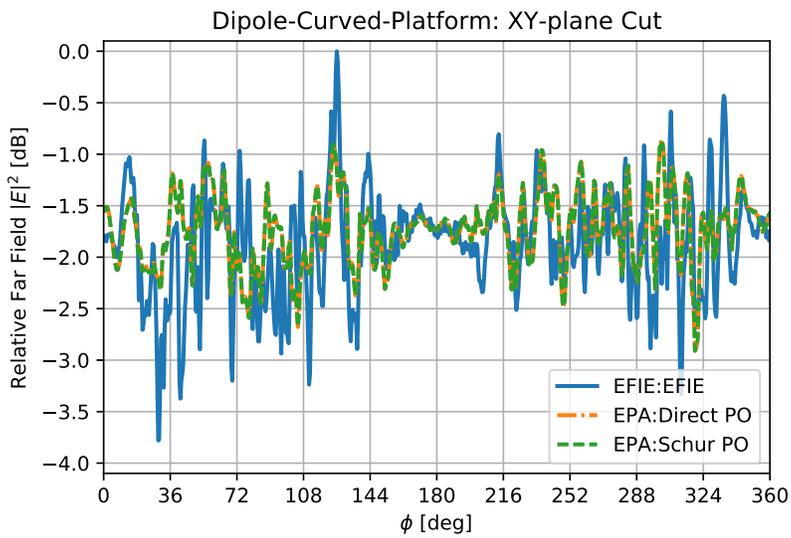


Figure 4.10: Horizontal plane ($z = 0$) cut of far-field intensity for a half-wavelength dipole above a large curved plate.

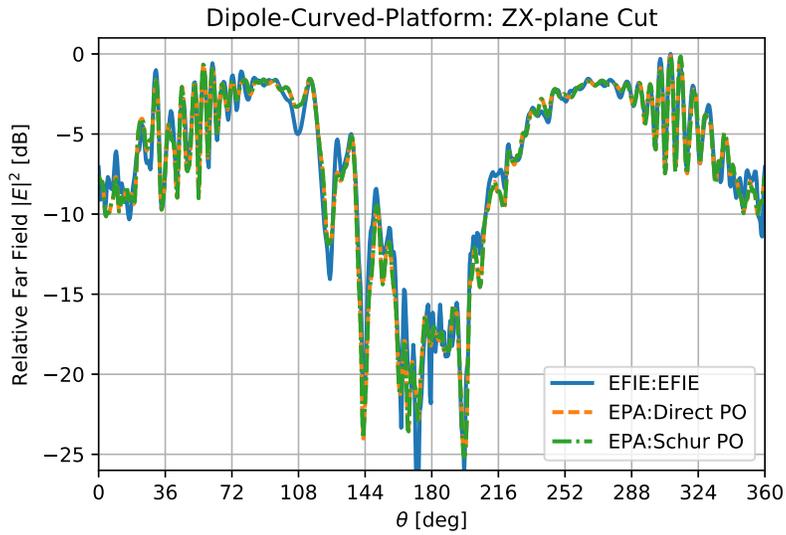


Figure 4.11: Vertical plane ($y = 0$) cut of far-field intensity for a half-wavelength dipole above a large curved plate.

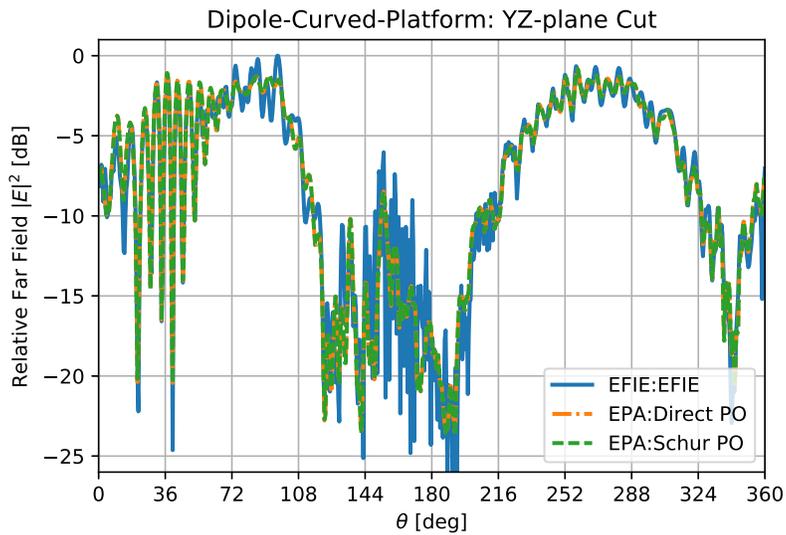


Figure 4.12: Vertical plane ($x = 0$) cut of far-field intensity for a half-wavelength dipole above a large curved plate.

CHAPTER 5

PARALLEL CALCULATIONS OF EPA-PO HYBRID SOLUTIONS

5.1 Introduction

The core idea of the Schur complement EPA-PO hybridization developed in Chapter 4 is the exchange of solving for the PO current sample coefficients for a somewhat more complicated linear system. This method, like all other computational methods, can be evaluated in light of mapping computational needs to available computational resources for ever-larger problems. This characteristic is broadly referred to as a method’s “parallel scaling” as applied to its memory-utilization or time-consumption with respect to the problem size and computational resources deployed. The parallel scaling is driven by both the underlying algorithms and the deployed computational resources.

Integral equation based approaches, including the EPA-PO hybrid developed in Chapter 4, are particularly challenging because N basis functions interact with N test functions to yield a square dense system. The acceleration of MVM operations reduces the computational and memory costs of each solver iteration to $\mathcal{O}(N \log N)$ for discretized SIEs, but as objects increase in electrical size scaling kd , the number of unknowns scales per Eq. (4.1) as $N \propto (kd)^2$. The computational storage and time complexity scales per Eq. (4.2) as $N \log N \propto (kd)^2 \log(kd)$. The expense of such algorithms can quickly overwhelm a single computer’s fixed resources. Therefore algorithms must be chosen judiciously to adaptively scale the computational resources used with the problem under study. In the context of parallel computing this is called “weak scaling,” and the degree to which a problem of fixed size reduces time requirements with increased resources is the “strong scaling” [29, 30]. The community has already developed suitable parallel distribution of the homogeneous medium MLFMA workload [31, 32]. The purpose of the approach implemented here is to apply distributed-memory parallel compu-

tations to the (inhomogeneous medium) Schur complement EPA-PO hybrid system. Because the PO scattering calculation is highly parallel-scalable, this offers an opportunity to transfer the computational burden away from an MLFMA-specific parallelization strategy to leverage readily parallelizable calculations.

Linearity and local scattering enable a simple use of parallel calculations for the EPA-PO hybrid. The system of equations is linear per the premises underlying Eqs. (2.3) and (2.4); these premises follow through to the continuous and discretized SIE operators in Eqs. (3.50) and (4.5), respectively. The SIE operators' contributions to each observed function weight are the linear sum of contributions due to each source function, which facilitates term-wise analysis of the MVP. The PO tangent plane approximation for curved surfaces is a point-wise local approximation of the object's scattering operator per Section 4.3. The equivalent currents characterizing the scattered fields are therefore validly partitioned into separate geometric regions which radiate independently. This combination of linear operators and point-wise analysis provides flexibility in composing parallel evaluation strategies.

To support arbitrarily large platform problems (scalable to the available HPC resources distributed across P compute nodes), we seek to split the platform calculations across distributed computational resources. The analysis begins by exploring theoretical parallel execution of the constituent operations. Then the implemented solver is documented in its execution on an HPC cluster using various node counts. Lastly, opportunities for future improvement are described.

5.2 Sequential and Parallel Calculations

The iterative solution of a linear system necessitates some sequencing to the calculations, e.g. the solution update for iteration n follows iterations 1 through $n - 1$ where n is in the set of natural numbers (positive integers) \mathbb{N} . Operations within each iteration may be partitioned into tasks with unique input variables and output results so that they may be partitioned into multiple tasks which may be evaluated with zero required intercommunication between units performing tasks; compute resources must be synchronized (thus requiring communication) between rounds of these tasks, i.e. at the

end of one step of the sequence. The consequence of such a task partitioning is that

- multiple independent tasks may be evaluated asynchronously on multiple cores to reduce the time a user waits for a simulation to complete, and
- more computing resources may be added to support analysis of problems of increasing size.

“Embarrassingly parallel” calculations are those which are uniquely split (including intermediate results) across computing resources and therefore execute on assigned resources with high parallel efficiency. Such embarrassingly parallel problems are targets for parallel execution partitioning because they minimize the overhead required for communication between processors.

The required operations for the EPA-PO hybrid will be discussed next in the context of supporting parallel evaluations. Parallel execution results will be presented for example problems to verify the suitability of the proposed parallelization strategy as applied to high-performance computing (HPC) hardware. For the purposes of this work, HPC hardware is defined to be a group of $P \in \mathbb{N}$ distributed nodes with an internode network which is relatively expensive (in time) to use in comparison to the intranode connection among $p \in \mathbb{N}$ local cores per node. Increasing P by one adds both distributed memory and p cores. Large values of P support problems with large distributed parallel components, provided the chosen algorithm is implemented in a manner to exploit hybrid distributed/shared-memory parallel execution.

Calculations may be grouped by how they are mapped to nodes and cores within HPC systems with the aforementioned P nodes each with p cores. It is assumed here that a calculation’s evaluation on one node takes time T and memory M ; the time for completion on all nodes and the total memory required across all nodes are distinguished in this taxonomy.

- Sequential calculations are evaluated in the same way on each node. This requires $\mathcal{O}(T)$ time and $\mathcal{O}(MP)$ total memory usage across all nodes.
- Distributed-memory parallel calculations are evaluated differently on

each node. Distributed parallel efficiency e_P implies that the required time is $\mathcal{O}(T/Pe_P)$. The memory used is $\mathcal{O}(M)$.

- Shared-memory parallel calculations are evaluated in the same way on each node but leverage multiple cores per node. Assuming local parallel efficiency e_p , the required time for the shared-memory parallel execution is $\mathcal{O}(T/pe_p)$, and the memory used is $\mathcal{O}(MP)$.
- Hybrid parallel calculations are evaluated differently on each node and leverage multiple cores per node. Shared-memory parallel efficiency e_p and distributed-memory parallel efficiency e_P are both utilized in hybrid parallelism. The required time is $\mathcal{O}(T/pe_pPe_P)$, and the memory used is $\mathcal{O}(M)$.

The calculations to be evaluated next are categorized according to their designated parallel execution mode using these definitions.

In this work, a data-parallel paradigm is all parallel tasks. This implies that with the exception of file writing, all P distributed processing nodes execute the operations coordinating via the message passing interface (MPI). Each distributed process computes an independent contribution to the final result utilizing up to p local task execution units, or cores. The local tasks are orchestrated using the OpenMP model. Both of these are implemented in conjunction with the open-source Trilinos [33] software family’s Tpetra [34] and Kokkos [35] libraries. Tpetra supports developing abstract operators with distinct MPI-distributed domain and range spaces, and provides supporting data structures to support distributed linear algebra. Kokkos maps computations to abstracted shared-memory parallel computing hardware, including multiple multicore processors using OpenMP among other possible methods. These packages provide the foundation of developing a hybrid parallel linear problem solver.

5.2.1 Hybrid Parallel Calculations

A node’s assigned computations are comprised of cascaded compound operators. Abstract operations such as block operators and composite “A X Plus Y” (or “AXPY”) operators comprise the system structure, and each of these is built upon concrete instances of “CrsMatrix” objects. These CrsMatrix

instances provide the core MVM operations within the MLFMA implementation, and they are implemented using hybrid-memory parallel execution provided by the underlying Tpetra/Kokkos combination. The population of the CrsMatrix instances is also performed by applying the Kokkos library to a data-parallel evaluation of the sparse matrix rows filled because each row's operations differ principally in the function indices used. The subdivided operations are of sufficient scale that the portions allocated to each node efficiently execute in shared-memory parallel operation. Ubiquitous multicore processors and increasingly common many-core accelerators dictate that a hybrid distributed- and shared-memory parallel evaluation mode is fundamental to an effective parallel implementation. The distributed-parallel operations discussed above are implemented to function in a hybrid parallel manner.

5.2.2 Shared-Memory Parallel Calculations

The number of matrix rows to be filled in each CrsMatrix is generally much greater than the number of local processors p . Consequently, data-parallel analysis is an effective strategy for evaluating the elements of the comprising matrices: Gramian matrices, the FMM-voxel near-field terms, the FMM-voxel plane wave patterns, the local interpolators, the O2O propagators, the O2I plane wave translators, the local antinterpolators, and the I2I propagators are all populated in this way to leverage shared-memory parallelism at initialization. Once these matrices are populated, they automatically gain shared memory parallel execution during MVM operations as provided by the Tpetra CrsMatrix data structure.

5.2.3 Distributed-Memory Parallel Calculations

The largest scale operations in the solution of the linear systems are orchestrated to leverage independent calculations. These operations are

- the governing GMRES iterative solver,
- field coefficient normalization with Gramians,
- the per-MVM data flow management, and

- MLFMA-accelerated source-to-field MVMs.

They are explored here in light of their time and memory requirements in a distributed-memory parallel execution.

Executing the GMRES [13] algorithm on the composed linear system is the coarsest level operation executed. It is implemented within the Trilinos Belos library [36], and relies on an abstract MVM operator interface. The composed linear system is implemented to satisfy this interface. Beyond the expense of the MVM it invokes, GMRES itself has $\mathcal{O}(N)$ storage requirement split across all distributed nodes.

Projecting fields onto testing functions is a key element of forming discretized solutions via the MoM. This projection is readily split across many distributed-memory parallel nodes. To obtain the basis function coefficients corresponding to the tested values, the effect of a Gramian MVM must be inverted. This may be performed by either iteratively solving every time such coefficients are needed during the simulation, or by direct solution. In this work, the Gramian is inverted by applying a sparse LU through the SuperLU_DIST library [16] via the Trilinos Amesos2 system of sparse direct solvers [36].

Within each MVM required by the iterative solver, communication operations are undertaken. To facilitate the iterative solution process, the operator domain and range are each a non-replicated distribution of basis and test functions, respectively. Because the MLFMA accelerates a dense matrix-vector multiply, this necessitates some degree of communication. The MLFMA operator domain and range functions are distributed by partitioning MLFMA voxels across the P distributed nodes, but each node’s implementation of the operator requires a complete copy of the domain coefficients. Therefore the $\mathcal{O}(N)$ coefficients are gathered from all nodes to all nodes, sometimes called an “allgather” operation, in $\mathcal{O}(\log P)$ communication steps among the nodes for a total time of $\mathcal{O}(N \log P)$.

The MLFMA-accelerated integral operators defining radiation between domains incur a large expense in the overall simulation. Both constructing an MLFMA abstract operator, discussed in Section 5.3, and the evaluation of the MVM using the same operator are amenable to distributed parallel evaluation. The calculations to be undertaken may be partitioned using the operator domain and range function voxels, the k -space quantities, or a com-

bination of the two; strategies are summarized in [31]. Motivated by the Schur complement EPA-PO hybrid, in this work the MLFMA operator is partitioned across nodes by distributing the coordinate space voxels. The domain vector is copied to all nodes so that each node has all the inputs required to completely calculate a unique subset of the range vectors. Because the domain vector is replicated across all nodes, the usual $\mathcal{O}(N \log N)$ cost is incurred on each node. The range vector is partitioned across the cluster with time and memory usage $\mathcal{O}((N/P)(\log N - \log P))$. Clearly, the replicated work on the domain side dominates the MLFMA MVM expenses, so the total memory used across the entire cluster is $\mathcal{O}(PN \log N)$.

Finally, some distributed memory calculations are reproduced on each node in an inconsistent manner and are implemented as needed on each distributed node as if they are shared-memory parallel calculations. Such calculations represent additional computational burden for the completed system while mitigating the need for additional communication between nodes, thereby trading the time expense of computing these values to avoid the communication expense of transmitting them between nodes. An example of this is found in the MLFMA translator operator construction. The definition of μ , α , and ν matrices consists largely of relative calculations which benefit greatly from cached calculations easily available in shared-memory parallel calculations but dependent upon expensive internode communications in distributed-memory parallel implementations. This particular problem has been addressed in the literature [32], but for the purposes of this work, the simpler spatial-partitioning-only is accepted on each node because it maps well to the proposed partitioning of the Schur complement EPA-PO hybrid.

5.3 MLFMA Operations

The core calculation of evaluating the MVM of Eqs. (3.50) and (4.5) is the radiation calculations accelerated by MLFMA. Consequently, MLFMA's structure will inform the choices to be made regarding parallel calculations. The MLFMA abstract MVM operator is defined by composing sequences of sparse matrix operations. Relative to the number of unknowns in a problem N , the composition of the sparse matrices is an $\mathcal{O}(N \log N)$ computational time and storage cost to be incurred before invoking the iterative solver. Within

each iteration the MVM is another $\mathcal{O}(N \log N)$ computational time expense. These two expenses are naturally connected, so the initial composition will use the same distributed parallelism strategy as the per-iteration MVM evaluation. Before examining parallelization strategies, the core structure of an MLFMA operator will be illustrated. This structure motivates the choices to be made in parallelizing the Schur complement EPA-PO system of equations.

5.3.1 MLFMA Data Structure Illustration

These operations are illustrated without loss of generality by considering an arbitrary integral operator for a case with geometry partitioned at MLFMA level three. Such a two-dimensional case's second and third level filled voxels are shaded in Figures 5.1a and 5.1b, respectively. In these illustrations each tile is labeled $A : BC$ with $A \in \{2, 3\}$ being the tree subdivision count and $B, C \in \{1, \dots, 8\}$ being the row and column indices of the subdivided tiles, respectively. The steps of defining an MLFMA operator for the voxel-partitioned functions are broadly

- computing block-wise plane wave pattern matrices for basis and test functions,
- computing block-to-block propagators between levels,
- applying interpolation and anterpolation matrices with propagator matrices,
- computing near-term radiation between coordinate space functions, and
- constructing block abstract operators.

Using an abstract operator $\overline{\mathbf{A}}$ with source function polarization $\mathbf{f}(\hat{k}_l)$, observation function polarization $\mathbf{g}(\hat{k}_l)$, interpolator $\gamma_i(\hat{k}_l, \hat{k}_{l'})$, and anterpolator

$\gamma_a(\hat{k}_l, \hat{k}_{l'})$, one defines the key block-wise operators as

$$a, b \in \{3 : 23, 3 : 24, 3 : 56, 3 : 57\}, \quad (5.1)$$

$$[\bar{\beta}_3]_{lb} = \iint_{S_b} d^2r \mathbf{f}(\hat{k}_l) \exp(-ik\hat{k}_l \cdot \mathbf{r}_{bm'}), \quad (5.2)$$

$$[\bar{\mu}_{a,b}]_{ll'} = \gamma_i(\hat{k}_l, \hat{k}_{l'}) \exp(ik\hat{k}_l \cdot (\mathbf{r}_a - \mathbf{r}_b)), \quad (5.3)$$

$$[\bar{\alpha}_{a,b}]_{ll'} = \delta_{ll'} \alpha(k\hat{k}_l, \mathbf{r}_a - \mathbf{r}_b), \quad (5.4)$$

$$[\bar{\nu}_{a,b}]_{ll'} = \gamma_a(\hat{k}_l, \hat{k}_{l'}) \exp(ik\hat{k}_l \cdot (\mathbf{r}_a - \mathbf{r}_b)), \quad (5.5)$$

$$[\bar{\beta}_3^\top]_{al} = \iint_{S_a} d^2r \mathbf{g}(\hat{k}_l) \exp(ik\hat{k}_l \cdot \mathbf{r}_{am'}). \quad (5.6)$$

Owing to the Kronecker delta function, the $\bar{\alpha}_{a,b}$ operator is diagonal in the space of plane waves. Because local interpolation/antepolation is used, the $\bar{\mu}_{a,b}$ and $\bar{\nu}_{a,b}$ operators are sparse. The $\bar{\beta}$ operators are dense in their own blocks of coordinate space function and plane waves. These block operators are readily assembled into a greater system of block operators with many zero blocks

$$\bar{\mu}_{2,3} = \begin{bmatrix} \bar{\mu}_{2:11,3:23} & \bar{0} & \bar{0} & \bar{0} \\ \bar{0} & \bar{\mu}_{2:12,3:24} & \bar{0} & \bar{0} \\ \bar{0} & \bar{0} & \bar{\mu}_{2:23,3:56} & \bar{\mu}_{2:23,3:57} \end{bmatrix} \quad (5.7)$$

$$\bar{\alpha}_3 = \begin{bmatrix} \bar{0} & \bar{0} & \bar{\alpha}_{3:23,3:56} & \bar{\alpha}_{3:23,3:57} \\ \bar{0} & \bar{0} & \bar{\alpha}_{3:24,3:56} & \bar{\alpha}_{3:24,3:57} \\ \bar{\alpha}_{3:56,3:23} & \bar{\alpha}_{3:56,3:24} & \bar{0} & \bar{0} \\ \bar{\alpha}_{3:57,3:23} & \bar{\alpha}_{3:57,3:24} & \bar{0} & \bar{0} \end{bmatrix} \quad (5.8)$$

$$\bar{\nu}_{32} = \begin{bmatrix} \bar{\nu}_{3:23,2:11} & \bar{0} & \bar{0} \\ \bar{0} & \bar{\nu}_{3:24,2:12} & \bar{0} \\ \bar{0} & \bar{0} & \bar{\nu}_{3:56,2:23} \\ \bar{0} & \bar{0} & \bar{\nu}_{3:57,2:23} \end{bmatrix} \quad (5.9)$$

$$\bar{\alpha}_2 = \begin{bmatrix} \bar{0} & \bar{0} & \bar{\alpha}_{2:11,2:23} \\ \bar{0} & \bar{0} & \bar{0} \\ \bar{\alpha}_{2:23,2:11} & \bar{0} & \bar{0} \end{bmatrix} \quad (5.10)$$

$$\overline{\mathbf{A}}_3^{\text{far}} = \overline{\boldsymbol{\beta}}_3^{\text{T}} \cdot (\overline{\boldsymbol{\alpha}}_3 + \overline{\boldsymbol{\nu}}_{32} \cdot \overline{\boldsymbol{\alpha}}_2 \cdot \overline{\boldsymbol{\mu}}_{23}) \cdot \overline{\boldsymbol{\beta}}_3 \quad (5.11)$$

$$\overline{\mathbf{A}}_3^{\text{near}} = \begin{bmatrix} \overline{\mathbf{A}}_{3:23,3:23} & \overline{\mathbf{A}}_{3:23,3:24} & \overline{\mathbf{0}} & \overline{\mathbf{0}} \\ \overline{\mathbf{A}}_{3:24,3:23} & \overline{\mathbf{A}}_{3:24,3:24} & \overline{\mathbf{0}} & \overline{\mathbf{0}} \\ \overline{\mathbf{0}} & \overline{\mathbf{0}} & \overline{\mathbf{A}}_{3:56,3:56} & \overline{\mathbf{A}}_{3:56,3:57} \\ \overline{\mathbf{0}} & \overline{\mathbf{0}} & \overline{\mathbf{A}}_{3:57,3:56} & \overline{\mathbf{A}}_{3:57,3:57} \end{bmatrix} \quad (5.12)$$

$$\overline{\mathbf{A}}_3 = \overline{\mathbf{A}}_3^{\text{near}} + \overline{\mathbf{A}}_3^{\text{far}} \quad (5.13)$$

This case is a simple example illustrating the nature of the problem under study; it is discussed without loss of generality to larger objects calling for larger numbers of MLFMA octree levels. The completed operator is demonstrated to be a cascaded multiply-add operation of abstract block operators with many zero blocks. Parallel evaluation of an MVM will depend upon the distribution of the responsibility computing coordinate-space data (split by voxel) and k -space data for the required plane waves (radiating between voxels).

5.3.2 MLFMA Parallel Evaluation

The abstract MLFMA operator is constructed independently on parallel computing hardware. The parallel evaluation is performed by using the constructed operators in-place on the resources utilized to construct them. If the responsibility for particular voxels' effects is distributed across parallel computing resources, then the fill and MVM operations are readily distributed in the same way. Every row output requires all row inputs, so the distributed responsibility requires serialization of the set of input coefficients.

5.4 Schur Complement EPA-PO Hybrid

The $\overline{\mathbf{T}} \overline{\mathbf{S}}_{\text{PO}} \overline{\mathbf{T}}$ operator of Eq. (4.5) is an independent bistatic PO simulation relating domain-to-range voxel plane wave translators. The number of local PO scattering calculations is both the dominant physical effect and a candidate for low-communication parallel evaluation if the EPA-modeled antenna region geometric feature scale d_{EPA} is designated such that the Schur

complement PO-modeled size scale d_{PO} satisfies

$$d_{\text{EPA}} \ll d_{\text{PO}}. \quad (5.14)$$

Consequently the number of samples on the EPA-modeled region $N_{\text{EPA}} = \mathcal{O}((kd_{\text{EPA}})^2)$ is overwhelmed by the sample count on the Schur complement PO-modeled $N_{\text{PO}} = \mathcal{O}((kd_{\text{PO}})^2)$. Therefore by inspecting Eq. (4.5), one readily sees that the largest part of any MVM is computing fields radiated to and from the N_{PO} functions. This drives the strategy for distributing the workload across distributed parallel resources.

To support arbitrarily large platform problems (scalable to the available HPC resources distributed across P compute nodes), the chosen approach is to split the platform calculations across distributed computational resources for efficient MLFMA and PO calculations. The Schur complement hybrid partitions the PO basis functions uniquely across each of the compute nodes for $\mathcal{O}(N_{\text{PO}}/P)$ PO functions per node. To minimize repeated MLFMA calculations, the MLFMA tree is split across the distributed-memory nodes at the coarsest level possible; for load-balancing purposes, the tree split was performed at the level at which each node holds at least two unique voxels assigned to each node. The mean number of basis functions per node is then

$$N_{\text{mean}} = N_{\text{EPA}} + \frac{N_{\text{PO}}}{P}. \quad (5.15)$$

The principal costs of parallel execution of the Schur complement hybrid are driven by the MLFMA radiation to/from the PO surface and normalizing the PO coefficients using the factorized Gramian matrix. These are considered next.

The per-node computational and memory complexity of MLFMA radiation for the PO surface on each node is reduced to $\mathcal{O}(N_{\text{mean}} \log N_{\text{mean}})$. In the asymptotic limit of increasing P , this reaches its minimum when N_{PO}/P approaches N_{EPA} . This implies a shared-memory parallel $\mathcal{O}(N_{\text{EPA}} \log N_{\text{EPA}})$ computational and memory expense on each of the P nodes.

Because the PO scattering in Eq. (4.4) is a scaled identity, the scattering operator itself is local. However, the discretized form interfaces with the MLFMA translators, which require normalization of tested functions by inverting a Gramian for functions on the PO surface. The distributed Gramian

for the entire Schur complement PO surface is LU factorized using the SuperLU_DIST library [16] within the Trilinos Amesos2 library [36]. Assuming the number of nonzeros in the matrix is proportional to the number of rows or columns N , SuperLU_DIST provides $\mathcal{O}(N^2 \log N)$ serial computational time complexity; the scaling of time complexity for a variety of problems has been reported to be approximately $\mathcal{O}(N^{4/3})$ in practice [16, Fig. 2]. The library also demonstrated high computational efficiency for large numbers of distributed processors when applied to a discretized Laplacian (efficiency fell to approximately 50 percent when split across 128 distributed processors [16, Table IX]), suggesting that factorizing the Gramian matrix will also exhibit suitable weak parallel scaling of computational time. The total distributed memory required for the sparse solver is demonstrated for several cases to scale with the parallel processor count P as $\mathcal{O}(\sqrt{P})$ for $P = 2^n, n \in \{1, 2, 3\}$ [16, Table VIII], so the expected memory expense for the sparse solver portion of the solver is $\mathcal{O}(N_{\text{PO}}\sqrt{P})$.

5.5 Results

The parallel scaling performance of the chosen implementation is evaluated by executing example cases using a cluster of nodes running Scientific Linux 6.1 as the operating system. Each node provided a pair of Intel X5650 processors (six cores per processor) operating between 2.66 and 3.06 GHz and with varying amounts of memory (12 GB minimum) per node. The compute nodes were connected using “quad data rate” (QDR) InfiniBand and the distributed-memory parallel calculations are orchestrated using the MPI standard implemented by OpenMPI 2.0.1. Distributed-memory and shared-memory parallel calculations are both implemented using Tpetra and Kokkos libraries, respectively, provided by Trilinos 12.10.1. The shared-memory parallel calculations are performed using the OpenMP module within Kokkos library. The research code and performance-critical libraries were compiled using the GNU Compiler Collection (GCC) 6.2.0.

The cases analyzed are the Dipole-Dihedral and Dipole-Curved-Plate cases analyzed in Section 4.5. Each was run three times using $1 \leq P \leq 12$ distributed-memory nodes and $p = 12$ shared-memory parallel (OpenMP) threads per node. The times reported are the “wall clock” time in seconds.

The node-parallel efficiency is computed from the times and node counts as the ratio of the single node run time to the observed run time scaled by the node count. The memory usage is that reported by the TORQUE job queuing system upon job completion. All jobs were run on nodes of sufficient memory to avoid forcing the node to swap virtual memory between memory and disk.

5.5.1 Dipole-Dihedral

The Dipole-Dihedral case uses 2,136 functions on the EPA surface and 4,332 functions on the smaller PEC plate modeled with PO. The Schur complement PO model of the large plate uses 229,272 functions. The run time and node-parallel efficiency of the solver process are illustrated in Figures 5.2 and 5.3, respectively. The efficiency is observed to be approximately 20 to 30 percent over the full range of nodes. The iterative solver's time spent calling the MVM operation is illustrated in Figure 5.4, and the node-parallel efficiency is illustrated in Figure 5.5. The MVM portion of the simulation is less efficient than the simulation in total, which is expected because the MVM includes most of the required communication among the nodes and duplicated MLFMA calculations. The increase in memory usage with respect to number of distributed processors in Figure 5.6 illustrates that the memory consumption is approximately linear with respect to node count P ; the per-node memory burden is approximately 1.5 GB. The node-parallel efficiency is relatively low, but the memory scaling supports distributing the problem across multiple nodes with less memory than that required for the single node execution so that a distributed-memory parallel HPC system can solve large problems by utilizing more nodes.

5.5.2 Dipole-Curved-Plate

The Dipole-Curved-Plate case uses 2,136 functions on the EPA surface. The Schur complement PO model of the curved plate uses 361,915 functions for the PEC PO surface. The run time and node-parallel efficiency of the solver process are illustrated in Figures 5.7 and 5.8, respectively. The efficiency is observed to be approximately 40 to 70 percent over the full range of nodes.

The iterative solver’s time spent calling the MVM operation is illustrated in Figure 5.9, and the node-parallel efficiency is illustrated in Figure 5.10. The MVM portion of the simulation is less efficient than the simulation in whole, which is expected because the MVM includes most of the required communication among the nodes. The increase in memory usage with respect to number of distributed processors in Figure 5.11 illustrates that the memory consumption is approximately linear with respect to node count P ; the per-node memory burden is approximately 1.2 GB. The memory scaling supports distributing the problem across multiple nodes with less memory than that required for the single node execution so that a distributed-memory parallel HPC system can solve large problems by utilizing more nodes.

5.5.3 Analysis

The intent of the parallel execution strategy is to distribute the PO calculations for the single largest part of the model. Consequently, the ratio of high-parallel-efficiency Schur complement unknowns relative to lower-parallel-efficiency unknowns drives the parallel scaling. Therefore, the parallel efficiency is expected to be lower for the Dipole-Dihedral than the Dipole-Curved-Plate because the Schur complement region is a lower fraction of unknowns in the former case than in the latter. The increased memory cost per node in the Dipole-Dihedral case is due to the additional plate included in the scene, which is included in the replicated calculations on every node.

5.6 Conclusions

The EPA-PO hybrid is observed to be effectively distributed across nodes of an HPC cluster system. The workload is distributed by partitioning the Schur complement EPA-PO hybrid geometry across nodes; the MLFMA octree voxels are the fundamental mechanism of partitioning the calculations for both the MLFMA radiation and LU factorizing the Gramian on the surface. In the presented cases, it was observed that the parallel efficiency and memory scaling were such that distributed-parallel HPC nodes were added to accumulate sufficient memory to execute the scenarios with acceptable performance.

The parallel efficiency of the implemented method is observed in a consistent range as the node count P varies, and the total memory required increases at a problem-dependent rate per node added to a simulation. Consequently, it is concluded that the calculation dominating the simulation time appears to be the MLFMA implementation instead of the sparse LU factorization. This suggests that further improving the MLFMA implementation for better load balancing is a viable path forward to speeding the parallel execution of the EPA-PO hybrid. Alternatively, the high-frequency asymptotic model applied to the EPA-PO hybrid may be exploited to reduce the number of plane waves utilized.

5.7 Figures

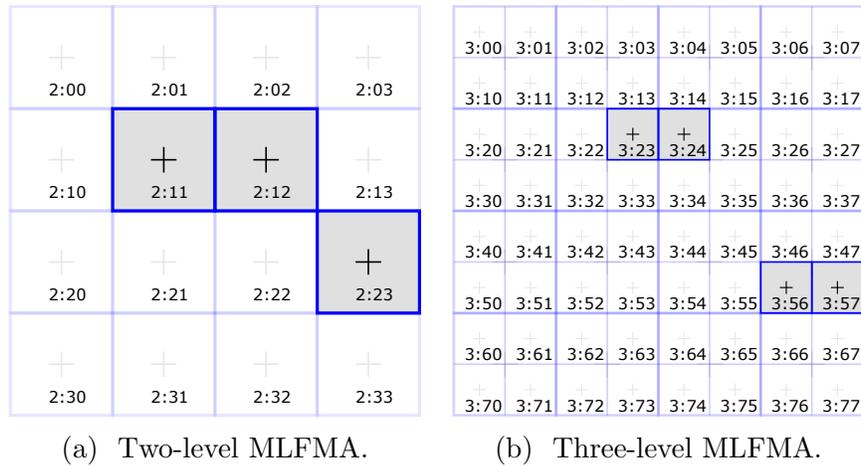


Figure 5.1: Two-dimensional illustration of MLFMA filled voxels tracking with unique voxel identifiers.

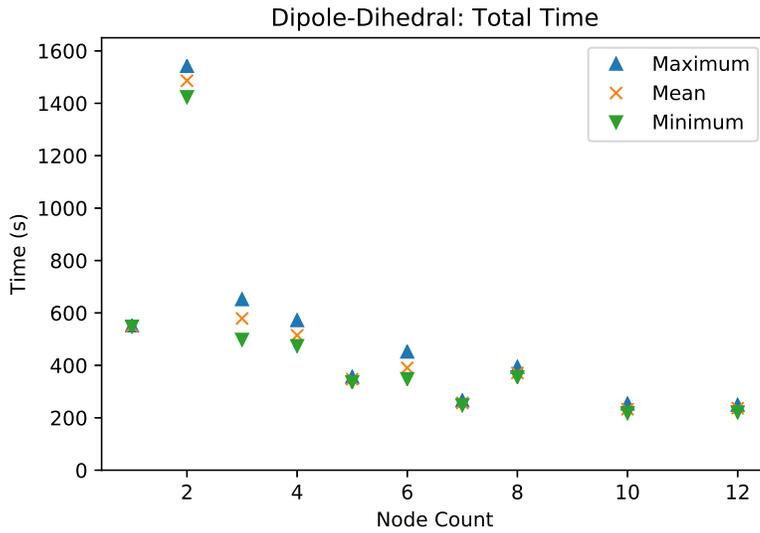


Figure 5.2: Wall clock time spent solving EPA-PO hybrid for a half-wavelength dipole parallel to a square plate and perpendicular to a large rectangular plate.

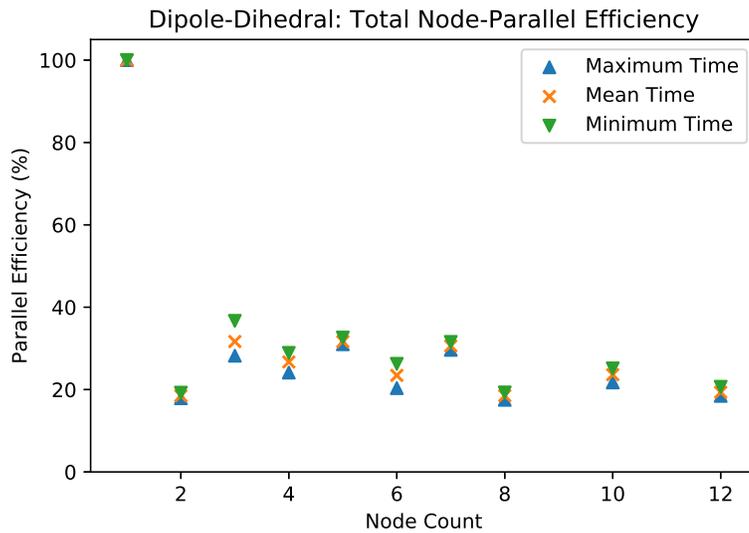


Figure 5.3: Node-parallel efficiency solving EPA-PO hybrid for a half-wavelength dipole parallel to a square plate and perpendicular to a large rectangular plate.

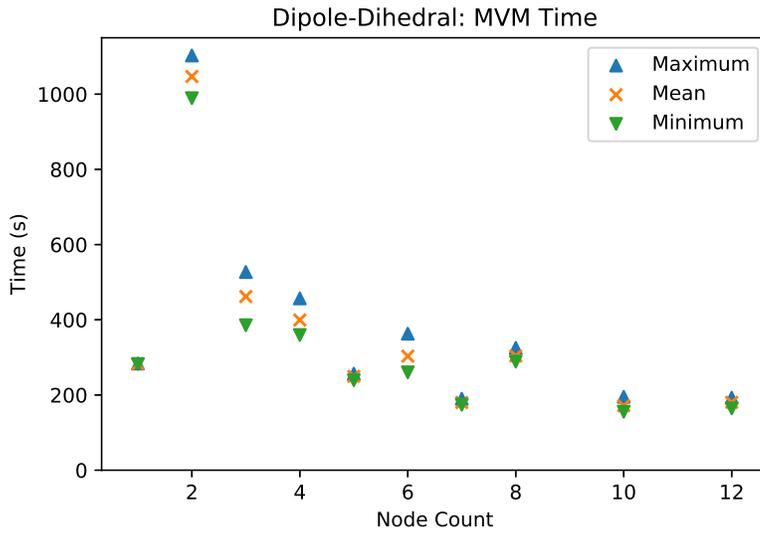


Figure 5.4: Wall clock time spent performing MVM while solving EPA-PO hybrid for a half-wavelength dipole parallel to a square plate and perpendicular to a large rectangular plate.

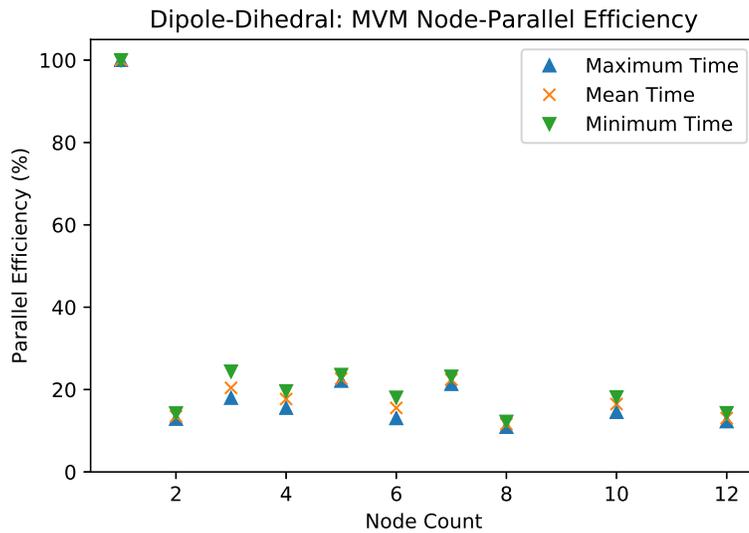


Figure 5.5: Node-parallel efficiency of MVM while solving EPA-PO hybrid for a half-wavelength dipole parallel to a square plate and perpendicular to a large rectangular plate.

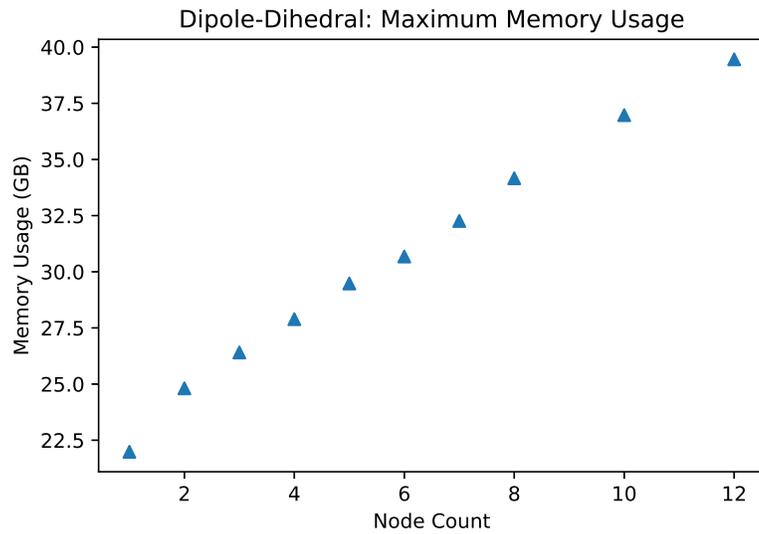


Figure 5.6: Memory usage reported after solving EPA-PO hybrid for a half-wavelength dipole parallel to a square plate and perpendicular to a large rectangular plate.

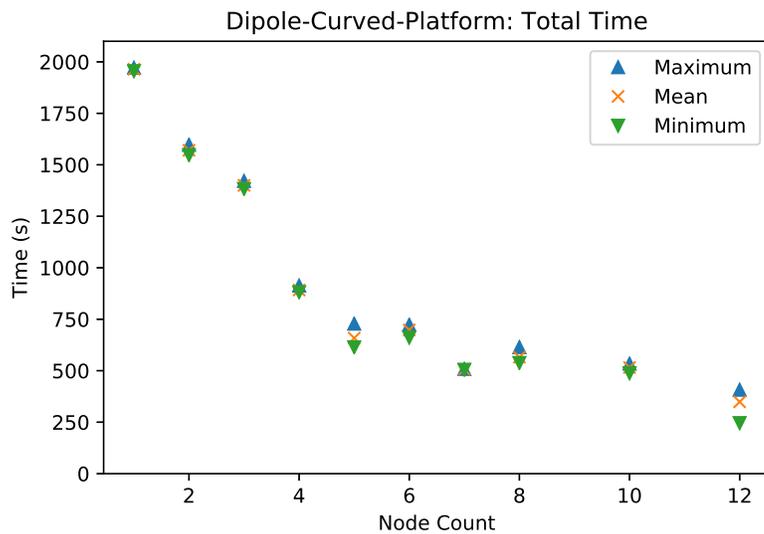


Figure 5.7: Wall clock time spent solving EPA-PO hybrid for a half-wavelength dipole above a large curved plate.

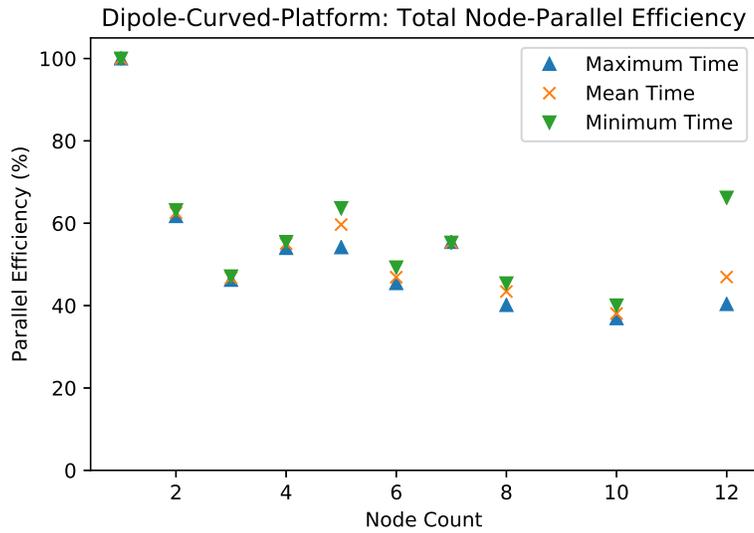


Figure 5.8: Node-parallel efficiency solving EPA-PO hybrid for a half-wavelength dipole above a large curved plate.

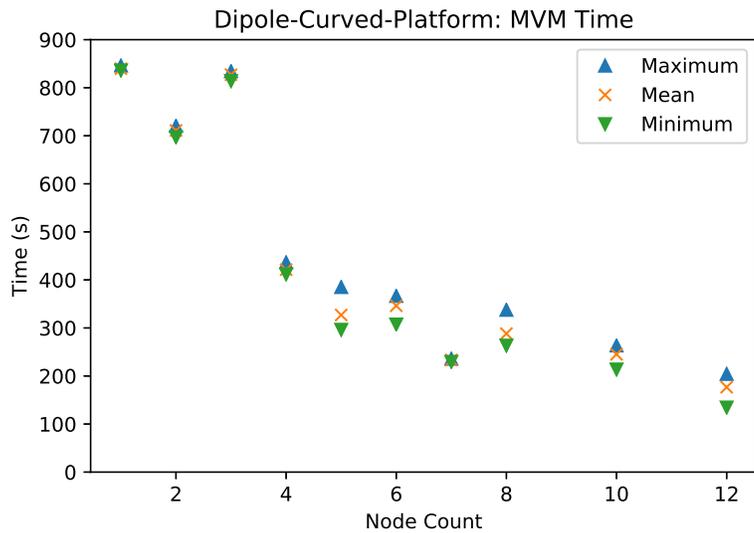


Figure 5.9: Wall clock time spent performing MVM while solving EPA-PO hybrid for a half-wavelength dipole above a large curved plate.

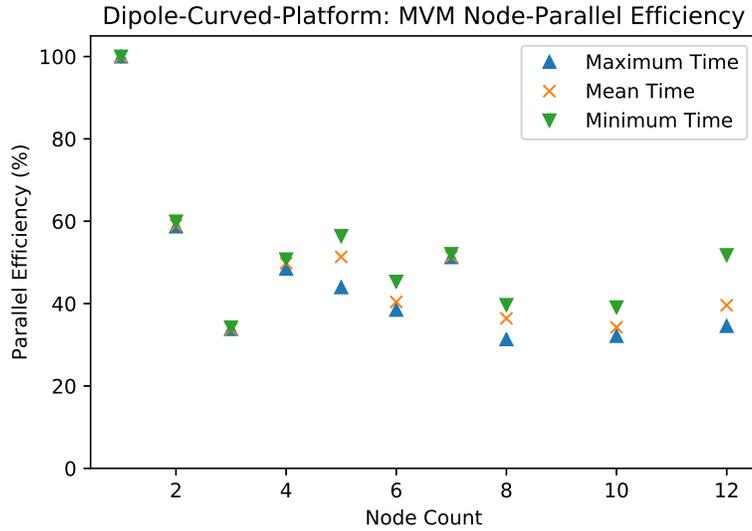


Figure 5.10: Node-parallel efficiency of MVM while solving EPA-PO hybrid for a half-wavelength dipole above a large curved plate.

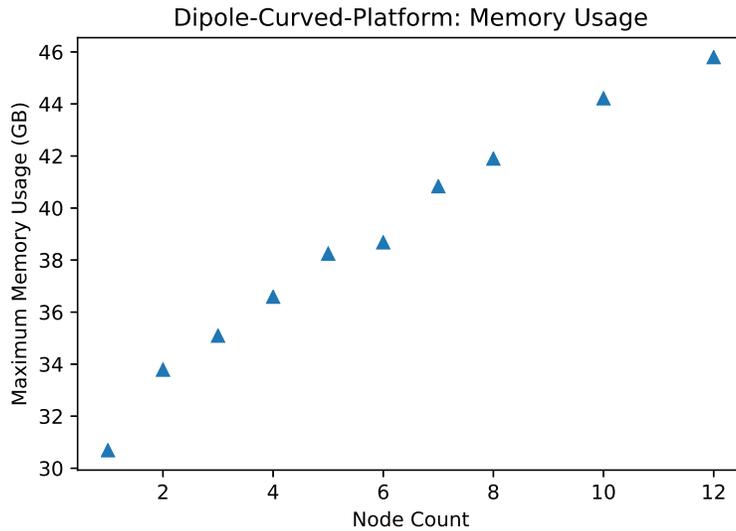


Figure 5.11: Memory usage reported after solving EPA-PO hybrid for a half-wavelength dipole above a large curved plate.

CHAPTER 6

BEAM FORMATION USING PLANE WAVES

6.1 Introduction

The core of the integral equations studied to this point is the discretization by MoM of linear operators convolving a Green's function with solution sources. In all such cases, the scalar Green's function is the core of the calculation. In the previously developed Schur-complement-based EPA-PO hybrid, MLFMA was applied to all radiation calculations to obtain fast MVM operations, but this same radiation to/from PO surfaces remains the dominant calculation. Because MLFMA requires a k -space surface integral over the Ewald sphere, and the integral kernel function order scales with the electrical size of the objects analyzed, electrically large objects are expected to involve a large number of plane waves. Historically effective high-frequency asymptotic models lead one to expect the radiation between electrically large components to be dominated by ray-physics behavior. Such propagation corresponds to select narrow sectors of integration on the Ewald sphere in an FMM analysis, and thus a reduced computational workload for the dominant calculation of radiation. A more efficient EPA-PO hybrid may leverage the naturally expected beam formation physics to reduce the workload of the k -space integral.

In this chapter, a method of analyzing the FMM translator is developed which reveals the relationship between objects' electrical sizes, their separation distances, and asymptotic scaling of operator numerical rank for increasing sizes. The FMM (and by extension MLFMA) provides a robust accelerated radiation MVM by diagonalizing the Green's function in the space of plane waves. Recent observation of low matrix rank in discretized radiation operators [37] raises questions regarding physical mechanisms driving the rank deficiency; the FMM-diagonalized translator provides a mechanism

for such an investigation. Because the time-harmonic electromagnetic vector case is a straightforward scaling of polarization effects onto the scalar case, the scalar Green’s function is analyzed without loss of generality. The results obtained indicate mechanisms for source and observation pattern-independent reduced operator rank relative to conventional FMM operators for electrically large translation operators for electrically large objects.

6.2 Numerical Analysis of the Green’s Function’s Rank

Using wavenumber k , surfaces of characteristic scale d require $\mathcal{O}((kd)^2)$ samples in coordinate space when discretized by basis and testing functions, assuming a uniform sampling rate in a linear direction. Thus a discretized translation operator $\overline{\mathbf{T}}$ between any two such objects satisfies

$$\text{rank}(\overline{\mathbf{T}}) \leq \mathcal{O}((kd)^2). \quad (6.1)$$

As an example of such analysis, one may consider the case where both geometries are spheres of identical diameter d and center-to-center separation $2d$. This case is diagrammed in Figure 6.1 showing the voxels around each sphere and the buffer voxel between them. Because of the surface equivalence principle discussed previously, this case is explored without loss of generality to other shapes within closed surfaces. As in the analysis of the same scalar case in [37, Sect. IV.D], the translator is discretized using Dirac delta basis and test functions. Consequently, the translator to observation surface A ’s function index i from source surface B ’s function index j is [37, Eq.(72)]

$$[\overline{\mathbf{T}}_{AB}]_{ij} = \frac{e^{ik|\mathbf{r}_i - \mathbf{r}_j|}}{|\mathbf{r}_i - \mathbf{r}_j|}. \quad (6.2)$$

This problem supports a simple implementation for numerical study using both low-rank approximation and FMM-based methods.

6.2.1 SVD Analysis of Translator

The SVD for a complex matrix relating translation between functions in voxels indexed a and b , $\overline{\mathbf{T}}_{ab} \in \mathbb{C}^{m \times n}$, requires that there exist unitary matrices

$\bar{\mathbf{U}}$ and $\bar{\mathbf{V}}$ such that [12, Sect. 2.5.6]

$$\bar{\mathbf{U}}^T \bar{\mathbf{T}} \bar{\mathbf{V}} = \text{diag}(\sigma_1, \dots, \sigma_p) \in \mathbb{R}^{m \times n}. \quad (6.3)$$

Alternatively, one may restate this as

$$\bar{\mathbf{U}} \text{diag}(\sigma_1, \dots, \sigma_p) \bar{\mathbf{V}}^T = \bar{\mathbf{T}} \in \mathbb{C}^{m \times n}. \quad (6.4)$$

The computation of the SVD of Eq. (6.2) is a numerical exercise performed with a “black box” procedure, but the low rank approximants are readily used to reconstitute a facsimile of the original dense matrix. There is little physical insight gained from this procedure, except to note that there are orthonormal domain function modes which will be observed as range function modes with positive definite eigenvalues.

The operator rank scaling with size is evaluated by considering a variety of object scales and their computed numerical rank. Paired spheres of diameter 1, 2, 4, 8, 16, and 32 λ are sampled at a linear sample distance of approximately $\lambda/2$ with respective per-sphere sample counts of 18, 67, 283, 900, 3,617, and 14,565. The operator rank is determined by using the singular value decomposition (SVD) to determine the number of eigenvalues greater than a specified tolerance [12, Sect. 2.5.5]. The ranks with respect to particular tolerances are shown in Figure 6.2, and agree with [37, Fig. 11(b)]. The linear increase in operator rank with respect to geometry feature scale means that as the matrix dimension increases quadratically with feature size, the matrix has a quadratically increasing number of effectively zero eigenvalues. This effect was documented for other three-dimensional object simulations using the ACA+ algorithm for low-rank factorization and SVD for rank analysis [37, Sect. IV.E].

6.2.2 FMM Diagonalized Translator

The FMM diagonalization of Eq. (6.2) is expressed as

$$\bar{\mathbf{T}}_{ab} = \bar{\boldsymbol{\beta}}_a^\top \bar{\boldsymbol{\alpha}}_{ab} \bar{\boldsymbol{\beta}}_b, \quad (6.5)$$

$$\left[\bar{\boldsymbol{\beta}}_a^\top\right]_{jl} = \exp\left(ik\hat{k}_l \cdot \mathbf{r}_{ja}\right), \quad (6.6)$$

$$\left[\bar{\boldsymbol{\beta}}_b\right]_{li} = \exp\left(ik\hat{k}_l \cdot \mathbf{r}_{bi}\right), \quad (6.7)$$

$$\left[\bar{\boldsymbol{\alpha}}_{ab}\right]_{ll'} = \delta_{ll'} q_l \alpha\left(k\hat{k}_l, \mathbf{r}_{ab}\right), \quad (6.8)$$

where $\delta_{ll'}$ is the Kronecker delta function, and q_l is the Ewald sphere quadrature weight corresponding to plane wave direction \hat{k}_l . As both the number of source and observation points is quadratic with respect to feature scale, and the number of plane waves used in the Ewald sphere surface integral is linear with respect to the number of source and observation points, one sees that a simple analysis yields

$$\text{rank}(\bar{\mathbf{T}}_{ab}) = \min\left(\text{rank}(\bar{\boldsymbol{\beta}}_a), \text{rank}(\bar{\boldsymbol{\alpha}}_{ab}), \text{rank}(\bar{\boldsymbol{\beta}}_b)\right), \quad (6.9)$$

$$\text{rank}(\bar{\mathbf{T}}) \leq \mathcal{O}\left((kd)^2\right). \quad (6.10)$$

This perspective on the translation operator does not immediately reveal the source of the reduced rank, but it provides a mechanism to isolate the physical effects in question.

A short comparison of Eqs. (6.4) and (6.5) shows that FMM points to the contributors to reduced operator rank. Firstly, the observation and source patterns in Eqs. (6.6) and (6.7) are not unitary and therefore have nonuniform eigenvalue spectra. These are far-field antenna patterns which are highly dependent upon the observation and source function positions relative to their corresponding reference ray end points. Due to coherent field effects these will have peaks and nulls in the k -space weighted pattern dot-products, which will give rise to source- and observation-specific ranks. The diagonal ‘‘O2I’’ translator $\bar{\boldsymbol{\alpha}}_{ab}$ weights the pattern dot-products with terms given in Eq. (6.8). This diagonal operator is trivially subjected to SVD analysis to reveal its rank. As the translator’s effect is independent of the source and observation patterns, it is key to understanding the core physics under consideration.

6.3 FMM as Fourier-Legendre Series

The three-dimensional scalar Green's function is expressed in terms of a vector between two reference points and then two (shorter) vectors from the reference points to their respective nearby original points as [15, Eq. (3.12)]

$$\mathbf{r}_{ji} = \mathbf{r}_j - \mathbf{r}_i, \quad (6.11)$$

$$= \mathbf{r}_{mm'} + \mathbf{r}_\Delta, \quad (6.12)$$

$$\mathbf{r}_\Delta = \mathbf{r}_{jm} - \mathbf{r}_{im'}, \quad (6.13)$$

$$|\mathbf{r}_\Delta| < |\mathbf{r}_{mm'}|, \quad (6.14)$$

$$\frac{\exp(ik|\mathbf{r}_{ji}|)}{|\mathbf{r}_{ji}|} = ik \sum_{l=0}^{\infty} (-1)^l (2l+1) j_l(k|\mathbf{r}_\Delta|) h_l^{(1)}(k|\mathbf{r}_{mm'}|) P_l(\hat{\mathbf{r}}_\Delta \cdot \hat{\mathbf{r}}_{mm'}). \quad (6.15)$$

This summation can be thought of as a truncated projection of one infinite-length vector onto another using the function order l as the vector index. The summation is clearly convergent, and the conditions of its truncation are explained in [15, Sect. 3.4.1]. The vectors comprising the dot product can be chosen in any number of ways. The physical intuition of the k -space integral form in [15, Eqs. (3.17) and (3.18)] is utilized next to identify physically interesting vector quantities.

For brevity, the k -space integral's azimuthal integral may be expressed in terms of the Bessel function $J_0(\cdot)$. One may restate the zenithal integral in terms of Legendre polynomials by using the Dirac delta function

$$f(b) = \int_a^c dx f(x) \delta(x-b) : a < b < c \quad (6.16)$$

$$f(\cos \psi) = \int_d^{d+\pi} d\psi \sin \psi (\cos \psi) \delta(\cos \psi - \cos \psi) : d, \psi \in \mathbb{R} \quad (6.17)$$

Insertion of the completeness relation for Legendre polynomials,

$$\delta(\cos \theta_k - \cos \theta'_k) = \sum_{l=0}^{\infty} \frac{2l+1}{2} P_l(\cos \theta_k) P_l(\cos \theta'_k), \quad (6.18)$$

splits the integral over the zenithal angle into two separate zenithal angle integrals. The integrals over zenithal integrals then reveal the series repre-

sentation of the Green's function in terms of Legendre polynomial orders

$$\frac{\exp(ik|\mathbf{r}_{ji}|)}{|\mathbf{r}_{ji}|} = \mathbf{a} \cdot \mathbf{b}, \quad (6.19)$$

$$[\mathbf{a}]_l = \frac{2l+1}{2} \int_{-1}^1 d\chi P_l(\chi) \alpha(\chi), \quad (6.20)$$

$$= \frac{ik}{4\pi} i^l (2l+1) h_l^{(1)}(k|\mathbf{r}_{mm'}|), \quad (6.21)$$

$$[\mathbf{b}]_l = \int_{-1}^1 d\chi' P_l(\chi') \exp(ik\hat{r}_{mm'} \cdot \mathbf{r}_\Delta \chi') \quad (6.22)$$

$$2\pi J_0(k|(\mathcal{I} - \hat{r}_{mm'}\hat{r}_{mm'}) \cdot \mathbf{r}_\Delta| \sin\theta'_k),$$

$$= \int d^2\hat{k} P_l(\hat{k} \cdot \hat{r}_{mm'}) \exp(ik\hat{k} \cdot \mathbf{r}_\Delta), \quad (6.23)$$

$$= 4\pi i^l j_l(k|\mathbf{r}_\Delta|) P_l(\hat{r}_\Delta \cdot \hat{r}_{mm'}). \quad (6.24)$$

Thus the scalar Green's function is expressed as the dot product of two vectors of infinite length. The first vector, \mathbf{a} , is the divergent term-wise scaling and phase shift of radiation between a source and observer voxel with respect to the Legendre polynomial order $l \geq 0$. The second vector, \mathbf{b} , is the convergent term-wise projection of local source and observer far-field patterns onto the Legendre polynomial. The convergence of \mathbf{b} overwhelms the divergence of \mathbf{a} , bounding the Legendre polynomial order and Hankel function order to $L \in \mathbb{N}$. The Green's function is therefore expressed as a Fourier-Legendre series with the form of the addition theorem in Eq. (6.16). This illustrates a transform pair relationship between function orders and spatial representation of source and observation fields. In one view, the projection of source and observer patterns represented as plane waves is projected onto Legendre polynomials, and the translation between reference points is in closed-form. The effect of propagation from one point to another with respect to the source and observer patterns in the space of spherical modes is also revealed. By simple inspection, Eq. (6.19) is clearly a restatement of Eq. (6.16), but the physical role of each vector's terms is made clearer. Owing to the clearly convergent behavior of the vector dot product, windowing may be included in the analysis.

An FMM (or a MLFMA) implementation requires a number of modes L calculated by the excess bandwidth formula [15, Eq. (3.47)]; for objects of electrical size $kd > 1$, $L = \mathcal{O}(kd)$. Moreover, the plane wave sampling across

the Ewald sphere requires $\mathcal{O}(L^2)$ samples for accurate integration (and interpolation/interpolation). Consequently the FMM-diagonalized propagation operator satisfies

$$\overline{\mathbf{T}}_{ab} = \overline{\boldsymbol{\beta}}_a^\top \overline{\boldsymbol{\alpha}}_{ab} \overline{\boldsymbol{\beta}}_b, \quad (6.25)$$

$$\text{rank}(\overline{\mathbf{T}}_{ab}) = \min(\text{rank}(\overline{\boldsymbol{\beta}}_a), \text{rank}(\overline{\boldsymbol{\alpha}}_{ab}), \text{rank}(\overline{\boldsymbol{\beta}}_b)), \quad (6.26)$$

$$\text{rank}(\overline{\mathbf{T}}) \leq \mathcal{O}((kd)^2). \quad (6.27)$$

6.4 Emergence of Ray Physics within FMM

As a translator's source and observation groups increase in characteristic size, the quadratic sample count on the Ewald sphere becomes overwhelmingly large if all plane wave samples are used. Far-field analysis suggests that plane waves paraxial to rays from the source to the observation surface's points should be dominant when the groups are electrically large. Direct analysis of the diagonal $\overline{\boldsymbol{\alpha}}_{ab}$ operator does not make the ray physics sufficiently clear; the alternate expressions for propagation explained here clarify the ray behavior.

The key observation is that the $\overline{\boldsymbol{\alpha}}_{ab}$ operator expresses radiation as a Legendre-Fourier series. The summation index l is the Fourier mode index into the countably infinite number of spatial modes. The plane waves' conic coordinate $\cos\theta = \hat{k} \cdot \hat{r}_{mn'}$ serves as a continuous coordinate in the space of ray cones defining the beam. The mode index and the conic coordinate are transform pair domains. When the summation is truncated at index L , it has the effect of a discrete Fourier transform evaluated with a rectangular window function [38], including the side lobes due to Gibbs' phenomenon in conic coordinate space.

6.4.1 Ray Propagation Fast Multipole Algorithm

The ray propagation fast multipole algorithm (RPFMA) applies the common discrete Fourier analysis practice of windowing to the series at hand. The effect of truncating the convergent sum with a vector dot product between Eqs. (6.21) and (6.24) is to apply a square window to the transformed quantities. Owing to the terms of Eq. (6.24) converging to zero at a higher rate than the terms of Eq. (6.21) diverge, there is latitude to introduce a

simple windowing into the terms of the summation without impacting accuracy. Consequently the diverging terms in Eq. (6.21) may be augmented with a window to suppress the higher order modes' influence over the spatial behavior of the translator. Per the findings of [38, 39, 40], the conventional rectangular window may be replaced with a flat window with cosine-tapered end of the form [38, Eq. (24)]

$$w_n = \begin{cases} 1, & n \leq L - K, \\ \frac{1}{2} (1 + \cos(\frac{n-L+K}{K}\pi)), & L - K < n < L, \\ 0, & n \geq L, \end{cases} \quad (6.28)$$

where the circumstance of L even and $K = L/2$ is named here a ‘‘Flat-Hanning’’ window.

6.4.2 Fast Far-Field Approximation

As the source and observation blocks become farther separated, one intuitively expects the far-field approximation to become more applicable. The fast far-field approximation (FaFFA) introduces the far-field approximation into the α operator summation [41, 42, 40]. This is manifested in two expressions of FaFFA. In the proper FaFFA method the far-field approximation is utilized directly in place of the addition theorem. In so doing, the α operator reduces to a spatial filter passing a single plane wave from source to observer [42, 40]. Alternatively, one may insert the far-field approximation to the spherical Hankel function $h_l^{(1)}(k_0 r_{mm'})$ into the addition theorem definition of α . Applying the windowing method of RPFMA to the approximated translator suppresses side lobes and provides a practical limit to the behavior of RPFMA approximation. This is defined here as ‘‘SemiFaFFA’’. The applicability of the far-field approximation is explored next in a short study.

The accuracy of SemiFaFFA is made clear by comparison to RPFMA in a chosen set of example translator cases. Figures 6.3 to 6.17 plot the amplitude of the translator with respect to the ray-zenithal angle for a object sizes $2^p, p \in [4, 5, 6, 7, 8]$ separated by distances $2^{2p+q}, q \in [-1, 0, 1]$. With constant object size and increasing separation, SemiFaFFA has a fixed beam width, but RPFMA's beam contracts. RPFMA is seen to approach Semi-

FaFFA when the separation is the square of the object scale. This finding is consistent with the criteria of applicability for the Fraunhofer approximation to the scalar diffraction integral [43, Sect. 8.3] [44, Sect. 4.3] and the closely related antenna far-field radiation [45, Sect. 6.8.2]. The consistent transition from RPFMA into SemiFaFFA illustrates a transition from Fresnel radiation into Fraunhofer radiation. It is important to note that the unwindowed FMM addition theorem is applicable in the near-field, so that one sees FMM-based methods being applicable at a very broad set of size scales.

The relationship between SemiFaFFA and FaFFA also points out a key relationship between wave and ray physics without resorting to classic Fresnel zone analysis. As object size increases, the SemiFaFFA translator shrinks in angle more rapidly than the sample density in angle increases; the consequence is an $\mathcal{O}(1)$ plane wave sample count asymptotic scaling with respect to object size. FaFFA's use of the far-field approximation in the untruncated addition theorem is a translator defined by a Dirac delta function in k -space; this has meaning only when used in an integral over the Ewald sphere. SemiFaFFA expresses a windowed summation which may be evaluated numerically at arbitrary directions in k -space. This may be rightly understood as paraxial rays about the primary ray direction. Because of the similarity between SemiFaFFA and FaFFA, one may conclude that where SemiFaFFA is accurate the number of plane waves required has converged to $\mathcal{O}(1)$ in keeping with the equivalent Dirac delta function representation.

6.5 Fast Multipole Methods and Rank Deficiency

The question of rank deficiency is addressed in light of the FMM family of methods. In particular, we seek a physical argument for the number of plane waves to be included in the calculation. As discussed above, the shape of the translator as a function of conic angle off the reference ray axis provides insight into the problem.

6.5.1 Monolithic Translator Beams

The translator beams for object diameters 8λ to $2,048\lambda$, with respective reference ray lengths of double the object diameters, are plotted in Figures 6.18

to 6.26. Along with translators computed using FMM, RPFMA, and Semi-FaFFA, two vertical lines mark the angles defining the cones enclosing the farthest-faces of the FMM voxels (the narrowest cone in which rays are expected) and the nearest-faces of the FMM voxels (the widest cone in which rays are expected). These cones are comparable to the stationary-point cone of [46, Fig. 1]. These translators provide insight into the rank of the off-diagonal blocks in SIE-based simulations. It is clear that the beam contracts with increasing object size and separations, with the largest objects showing little contraction. The narrowest cone angle is near what appears to be a limit to the contraction of the RPFMA beam, which is predictable from a high-frequency asymptotic ray-based picture of radiation between components of the objects under study.

Within an FMM approach, the solid angle spanned by such beams is sampled in a O2I operator at quadrature points in k -space. If the beam is truncated at amplitudes below a threshold, then a particular beam-axis-angle ψ_{\max} is the limit of the zenithal angle of integration in the beam. The zenithal angle sample rate is the usual $(L + 2) / \pi$ so that the sample count is

$$N_{\psi} = \frac{\psi_{\max} (L + 2)}{\pi} \quad (6.29)$$

The azimuthal angle χ is sampled at the same rate as ψ at the widest part of the Ewald sphere sampled. The radius of the Ewald sphere at the widest part of the beam is at the maximum values of $\sin \psi$. Therefore the number of azimuthal samples is

$$N_{\chi} = \begin{cases} 2L, & \psi_{\max} > \frac{\pi}{2}, \\ 2L \sin \psi_{\max}, & \psi_{\max} \leq \frac{\pi}{2}. \end{cases} \quad (6.30)$$

The number of samples for the beam is the product of the azimuthal and zenithal sample counts.

In a far-field analysis the separation D and the size d satisfy the usual far-field criteria

$$D > 2d^2, \quad (6.31)$$

then the conic angle spanned is

$$\psi_{\max} \approx \sin \psi_{\max} \approx \frac{d}{D}. \quad (6.32)$$

In this narrow-beam case, one obtains a sample count

$$N_\chi N_\psi \approx \frac{2}{\pi} \psi_{\max}^2 L(L+2). \quad (6.33)$$

The case of a fixed ψ_{\max} but increasing object size d such that $L = \mathcal{O}(d)$ yields $N_\chi N_\psi = \mathcal{O}(d^2)$. If the object size is proportional to the separation using a subdivision count s , then

$$d = 2^{-s} D, s \in \mathbb{N}, \quad (6.34)$$

and the number of required quadrature points is

$$N_\chi N_\psi = \mathcal{O}(2^{-4s} D^2). \quad (6.35)$$

Far-field analysis puts a lower limit on the number of far-field samples

$$N_\chi N_\psi \geq \mathcal{O}(1) \quad (6.36)$$

The shrinking angle analysis reveals the approach of RPFMA to FaFFA using the estimated sample count. As the angle becomes ever smaller and the far-field approximation holds, one observes a transition from the Fresnel zone into the Fraunhofer zone of radiation analysis.

6.5.2 Multiple Beams

The analysis of beams between a single source-observation pair may be subdivided into multiple pairs; such an approach leads to a spectrum of O2I operator structures. If the characteristic size d surface is uniformly split $s \in \mathbb{N}$ times, then one obtains a 2^{2s} source and observation voxels of characteristic size $d2^{-s}$. This gives rise to a dense block operator (all blocks populated) with each block populated by diagonal O2I operators between the newly obtained voxels. Each observation voxel in the operator range accumulates contributions from all source voxels in the operator domain. If one defines an expression for the subdivision count

$$s = u + \log_2 d^t, 0 \leq t \leq 1, \quad (6.37)$$

where u and t are chosen to satisfy $s \in \mathbb{N}$, then the subdivision s gives rise to $\mathcal{O}(2^{2u}d^{2t})$ voxels. The number of plane waves required for a single observation voxel per source voxel is $N_\chi N_\psi = \mathcal{O}(2^{-4u}d^{2-4t})$. Plane waves are needed per domain voxel per range voxel. If chosen subdivision level s supports a far-field approximation, that is $t > \frac{1}{2}$, then the beams are paraxial groups of $\mathcal{O}(1)$ plane wave samples about rays from each source voxel to the observation voxel as discussed in Section 6.4.2. The total number of plane waves required in this circumstance is $\mathcal{O}(d^2)$.

The geometry of subdivided observation and source voxels reveals physical mechanisms for reduced rank in the radiation calculation. The rank of the translator is limited by the dimension of the dense block-diagonal operator; the upper limit of this rank is the number of plane waves per operator range voxel multiplied by the number of range voxels. Both large ($t = 1$) and small ($t \approx 0$) subdivision counts give rise to voxel-plane wave count products which are quadratic with regard to the original object size. Between these extremes one finds overlapping beams from the observation voxel to neighboring source voxels. Because each observation voxel's results are the sum of all k -space contributions, replicated k -space contributions from aligned and adjacent source voxels do not contribute to the rank. The other key effect driving the dimension of per-voxel-pair radiation calculations is the conic-angle width of the translator beam which contracts as the distance between voxels increases but the voxel sizes are held constant. The required sample count density is constant but the subtended angle contracts as the distance to the nearest source voxels increases for each observation voxel. This narrowing conic beam in k -space pushes the required rank downward in combination with the shared spatial bandwidth. These effects push the net rank downward, but do not change the asymptotic scaling of the rank as $\mathcal{O}(d^2)$.

6.5.3 Subdivided RPFMA Rank Calculations

The viability of reduced rank due to far-field effects has been explored by counting the number of plane waves used in the k -space quadrature. Given a range voxel and a domain voxel of coordinate space functions with respective k -space representations, Algorithm 1 is a pseudocode computing the plane waves contained in a beam about the ray connecting the range voxel to the

domain voxel. Each voxel has an internal reference point (usually the voxel center) used for this calculation, and it has an external reference point (usually a greater geometry centroid). An example conic beam between a domain and range voxel using a fixed angle is illustrated in Figure 6.27a. Using the same range voxel, another domain voxel will require a cone positioned at a different ray direction with partial overlap with the original beam cone; the overlapping beams are illustrated in Figure 6.27b along with a red-highlighted greater beam covering both smaller beams. The greater containing beam is computed with a relative zenithal direction from the range voxel's internal reference point to the domain voxels' common external reference point. The minimal cone enclosing all the plane waves observed at the range voxel is the k -space region which must support accurate quadrature.

Algorithm 1 Beam-forming O2I operator plane wave directions.

Require: $t \ll 1$, $d > 0$, $|\mathbf{D}| > \sqrt{3}d$

Ensure: $p \neq \emptyset$

 Compute FMM zenithal samples $Z(\mathbf{D}, d)$

 Compute RPFMA O2I $\alpha(Z)$

 Find beam zenithal samples $z \in Z : |\alpha(z)| \leq t$

 Compute zenithal sample directions $\hat{k}(z)$

if $\min(\hat{k}(z) \cdot \mathbf{D} > 0)$ **then**

 Compute RPFMA zenithal sample density d_z

 Compute uniformly-distributed azimuthal samples a with density d_z

else

 Compute FMM zenithal sample density d_z

 Compute uniformly-distributed azimuthal samples a with density d_z

end if

 Compute plane wave directions $p = \hat{k}(a, z)$

return p

The same zenithal and azimuthal sampling strategy used in Algorithm 1 is applied to each range voxels' greater beam cone; the number of plane waves used is the sum of the range voxels' conic beam plane waves. If this analysis is applied with subdivision count $s = 0$, then the analysis is completed. If $s > 0$, then the O2I translator goes from being diagonal to being a dense block matrix. Each block is a diagonal O2I translator between small objects separated by a large distance. The rank of the block matrix is limited by the cumulative number of domain or range plane waves used by all ranges or domains. The vast majority of the plane waves between subvoxels lie within

the narrower cone angles marked in Figures 6.18 to 6.26, and virtually all the plane waves are within the marked wider cone angles. Therefore, if one were to consolidate all voxels and patterns into one, one expects the case of $s > 0$ to become that of $s = 0$ by physical intuition.

This picture of the translator rank has been tested for spheres with diameters 16, 32, 64, 128 and 256 λ and a translator tolerance of $1e - 4$. Each has been evaluated with zero subdivisions, and the plane wave count from both FMM and RPFMA serves as reference values for the operator rank. Each object was then subdivided to obtain block-wise RPFMA analyses with one and two subdivisions. The source- and observation-functions are uniformly split on each of the three Cartesian axes so that at most a factor of 8 subblocks are filled with increasing subdivision level, implying exponential growth in the number of voxels. However, the number of voxel interactions inspected is quadratic with respect to the exponentially growing number of voxels, so the number of subdivisions utilized is small as a matter of computational necessity. The results are reported in Table 6.1.

The results show the effects of object scale on propagator rank. The FMM rank is clearly quadratic in nature with respect to characteristic size d . For sufficiently large objects, the beam rank is demonstrated to be reduced by subdividing the domain and range voxels and applying RPFMA. For smaller objects ($d \leq 64 \lambda$) the rank is increased by applying subdivided RPFMA. At $d = 128 \lambda$ a single subdivision has a small rank-reducing effect for the beam. At $d = 256 \lambda$ a single subdivision provides a noteworthy rank reduction, and a second subdivision reduces the rank even further. This trend illustrates reduced rank by examining the emerging ray physics, but the rank's asymptotic scaling $\mathcal{O}(d^2)$ is not observed to change. The emerging ray physics is due to applying the analyses to smaller source and observation zone objects. This transitions the analyses from near-field radiation to far-field radiation, which is analogous to a transition from Fresnel zone to Fraunhofer zone radiation.

6.6 Conclusions

The FMM O2I operator supports multiple expressions based upon partitioning the source and observation functions of interest into voxels. The usual diagonal operator was replaced with a densely filled block operator relating

multiple voxels; each block is a diagonal operator relating outbound plane waves at source voxels to inbound plane waves at each block row’s observation voxel. Analyzing the interaction between subdivided components revealed ray physics when applied to electrically large voxels. The FMM radiation calculation’s rank is reduced using physically clear overlapping beams and narrowing beams for various observation voxel choices. This was tested in a numerical experiment applied to meshed spheres of various size. The rank was reduced, but it was not pushed below quadratic with respect to object diameter. The observed accurate transition from FMM analysis over the entire Ewald sphere to beams paraxial to rays between voxels suggests that high-frequency asymptotic methods may be incorporated into FMM-based analyses using the k -space sampling approaches developed here.

6.7 Table and Figures

Table 6.1: Table of computed ranks using RPFMA relative tolerance 1e-4.

Size	Separation	Samples	FMM Rank	Splits	Beam Rank
16	32	3,617	89,042	0	72,584
16	32	3,617	89,042	1	145,080
16	32	3,617	89,042	2	315,798
32	64	14,565	310,472	0	165,480
32	64	14,565	310,472	1	357,132
32	64	14,565	310,472	2	680,464
64	128	57,807	1,134,018	0	680,600
64	128	57,807	1,134,018	1	776,244
64	128	57,807	1,134,018	2	1,101,322
128	256	232,120	4,280,738	0	1,970,496
128	256	232,120	4,280,738	1	1,943,932
128	256	232,120	4,280,738	2	2,190,577
256	512	912,418	16,508,258	0	6,393,300
256	512	912,418	16,508,258	1	5,694,767
256	512	912,418	16,508,258	2	5,320,139

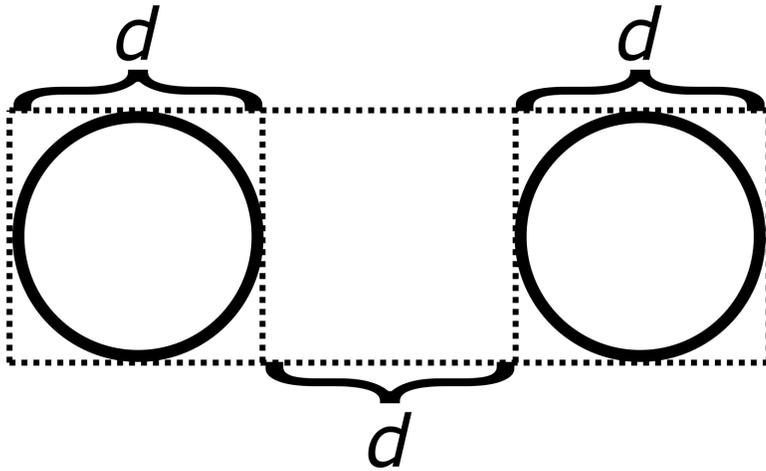


Figure 6.1: Sketch of two spheres case with sphere diameter d used to mark the characteristic dimensions of the problem; FMM enclosure voxels are marked with dotted lines.

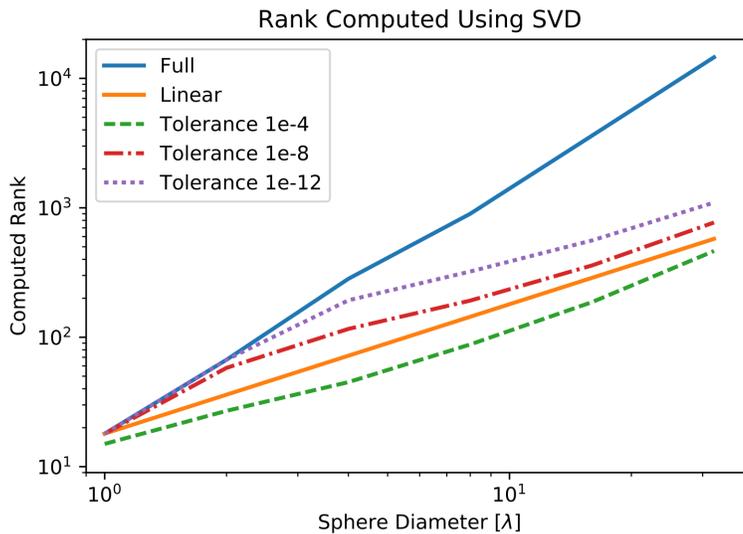


Figure 6.2: Rank computed for translation operator between two spheres separated by twice the characteristic size for various tolerances.

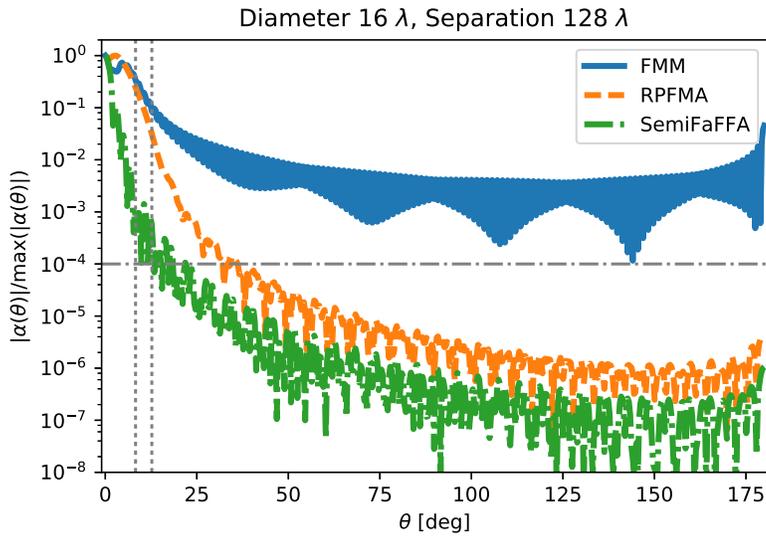


Figure 6.3: Translator amplitude for two spheres of diameter 16λ and separation 128λ .

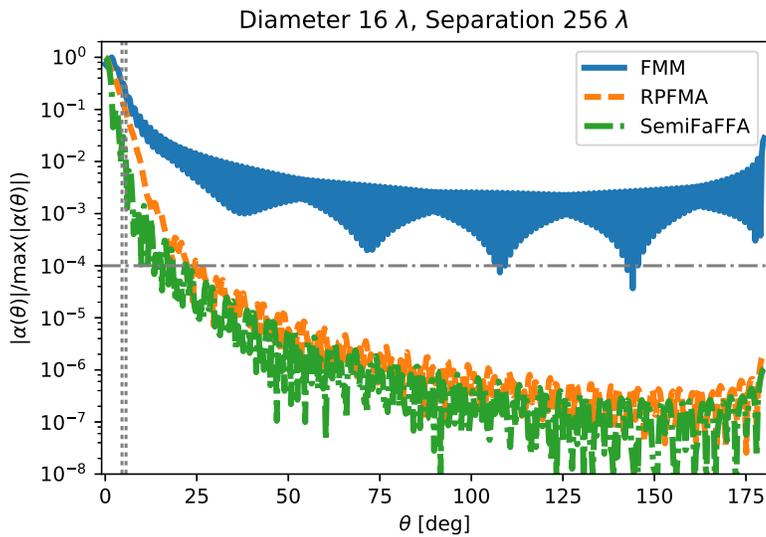


Figure 6.4: Translator amplitude for two spheres of diameter 16λ and separation 256λ .

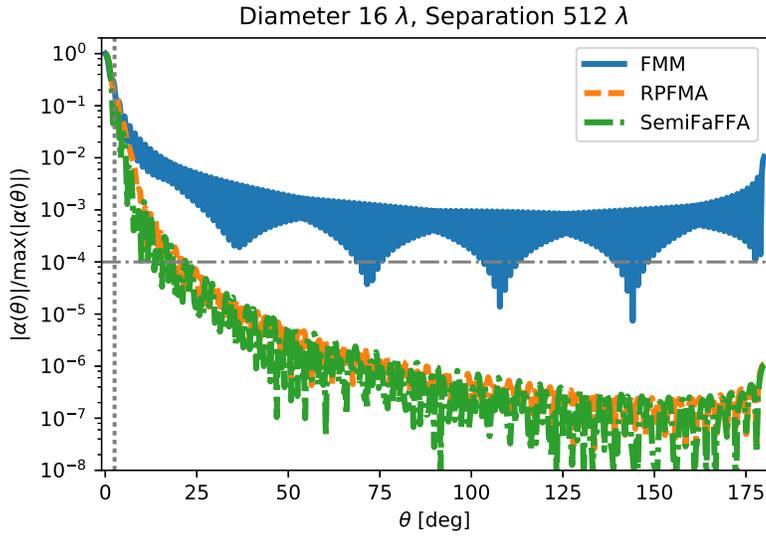


Figure 6.5: Translator amplitude for two spheres of diameter 16λ and separation 512λ .

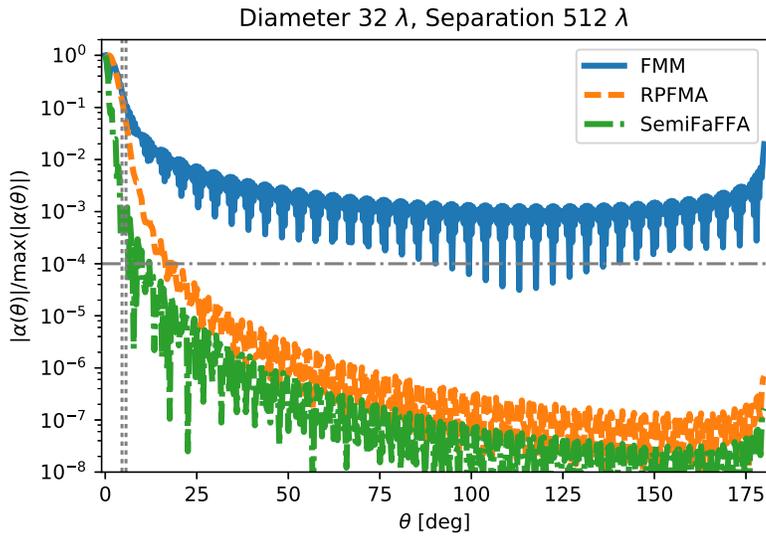


Figure 6.6: Translator amplitude for two spheres of diameter 32λ and separation 512λ .

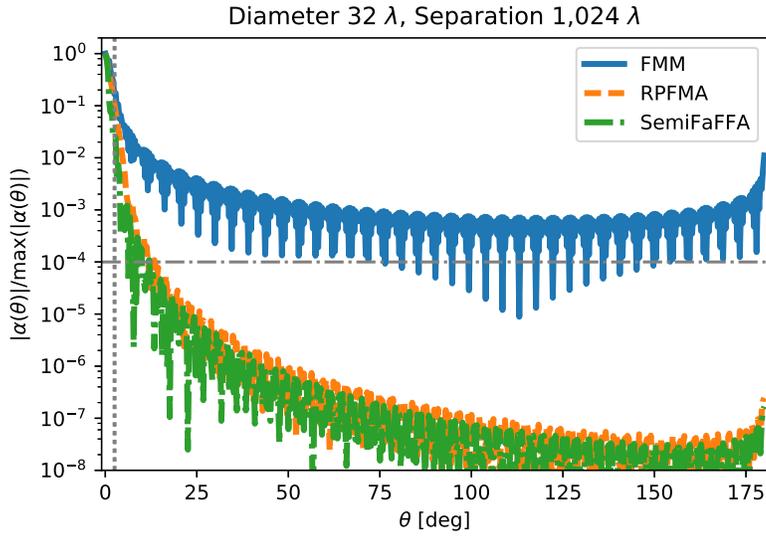


Figure 6.7: Translator amplitude for two spheres of diameter 32λ and separation $1,024 \lambda$.

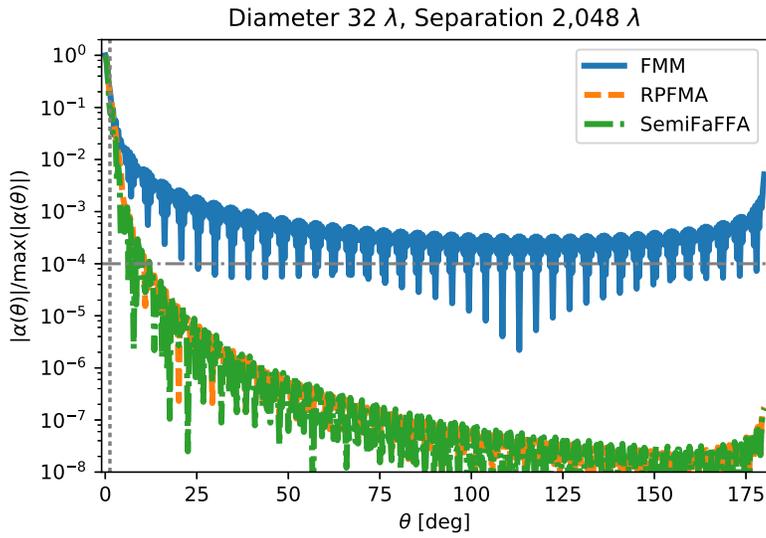


Figure 6.8: Translator amplitude for two spheres of diameter 32λ and separation $2,048 \lambda$.

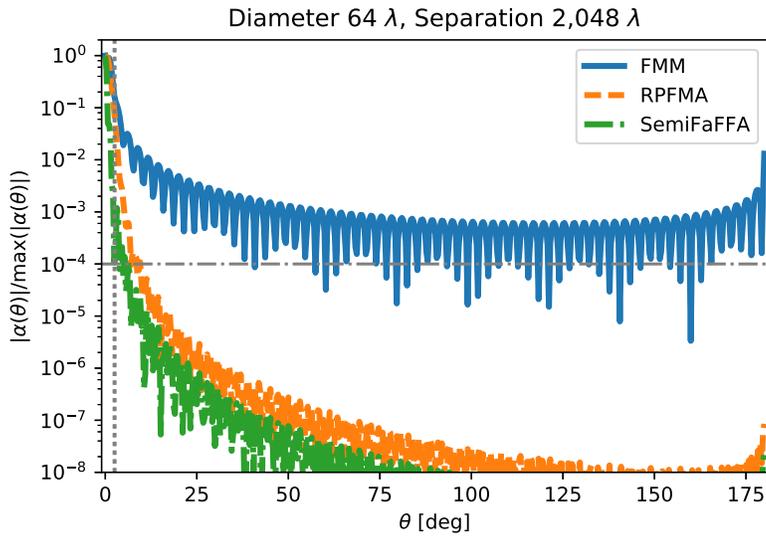


Figure 6.9: Translator amplitude for two spheres of diameter 64λ and separation $2,048 \lambda$.

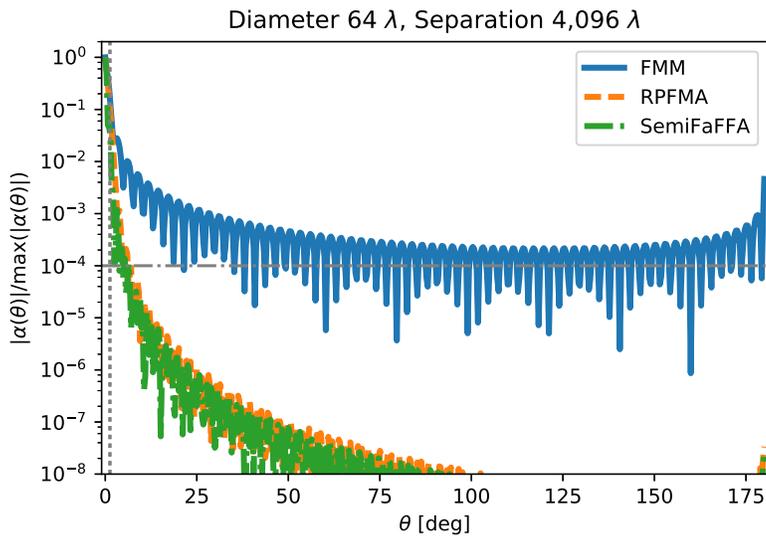


Figure 6.10: Translator amplitude for two spheres of diameter 64λ and separation $4,096 \lambda$.

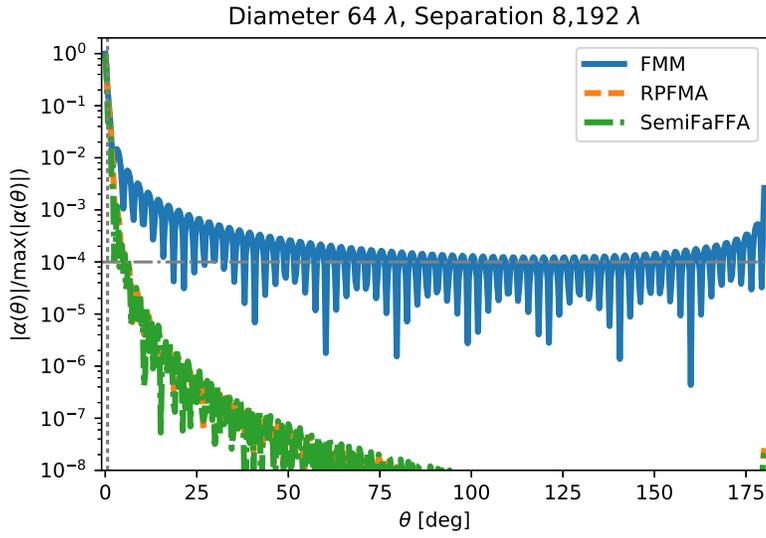


Figure 6.11: Translator amplitude for two spheres of diameter 64λ and separation $8,192 \lambda$.

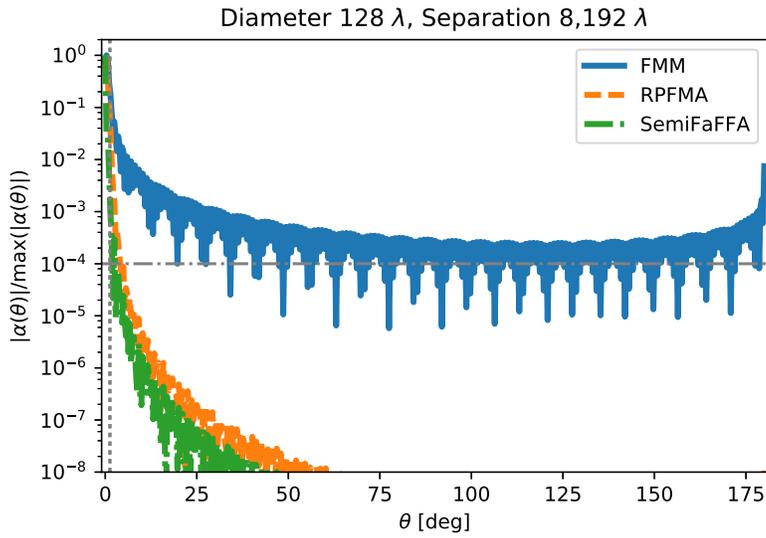


Figure 6.12: Translator amplitude for two spheres of diameter 128λ and separation $8,192 \lambda$.

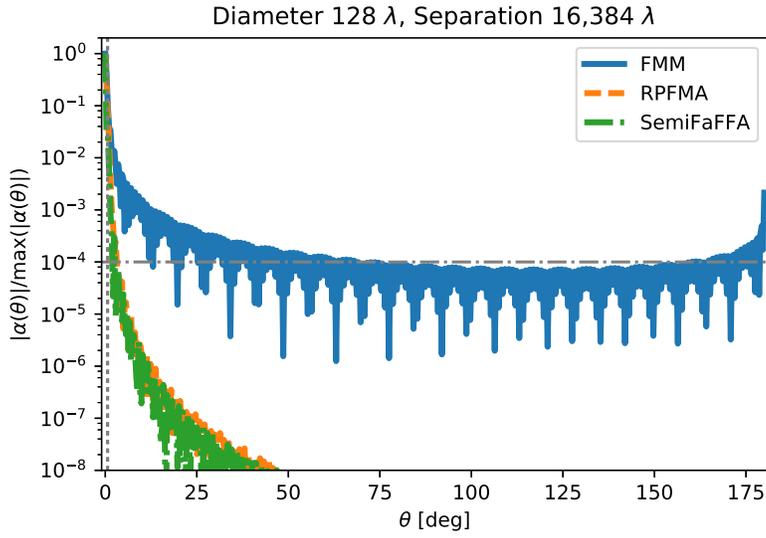


Figure 6.13: Translator amplitude for two spheres of diameter 128λ and separation $16,384 \lambda$.

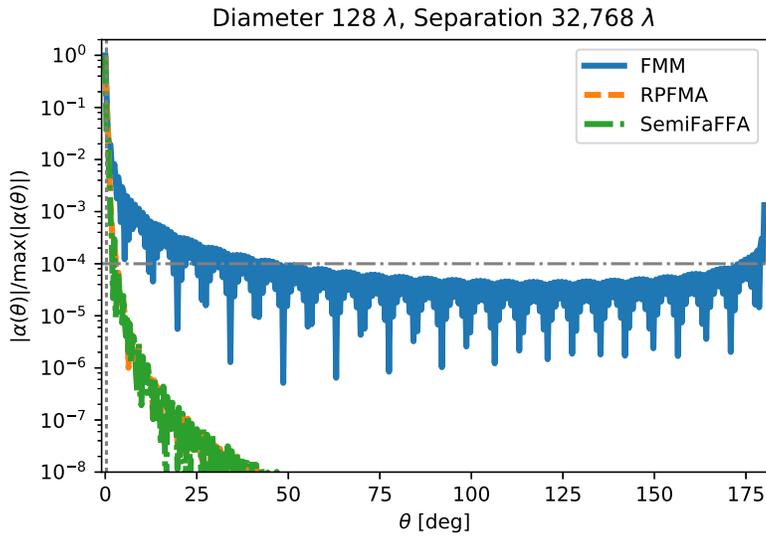


Figure 6.14: Translator amplitude for two spheres of diameter 128λ and separation $32,768 \lambda$.

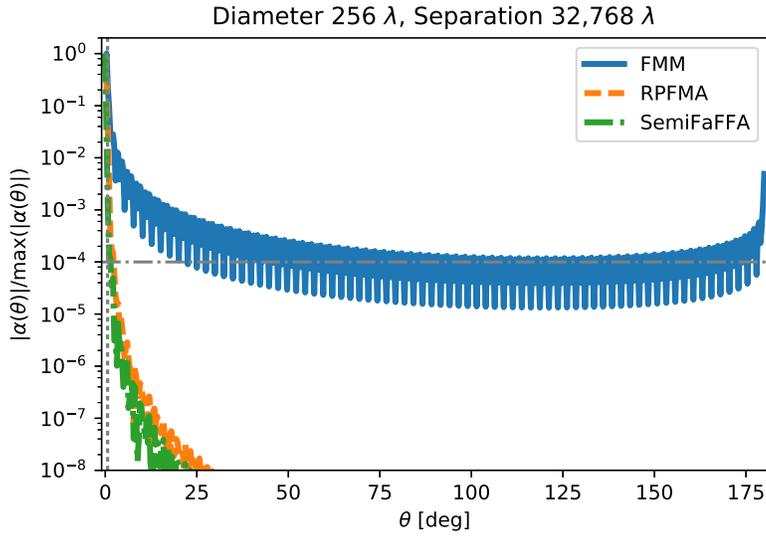


Figure 6.15: Translator amplitude for two spheres of diameter 256λ and separation $32,768 \lambda$.

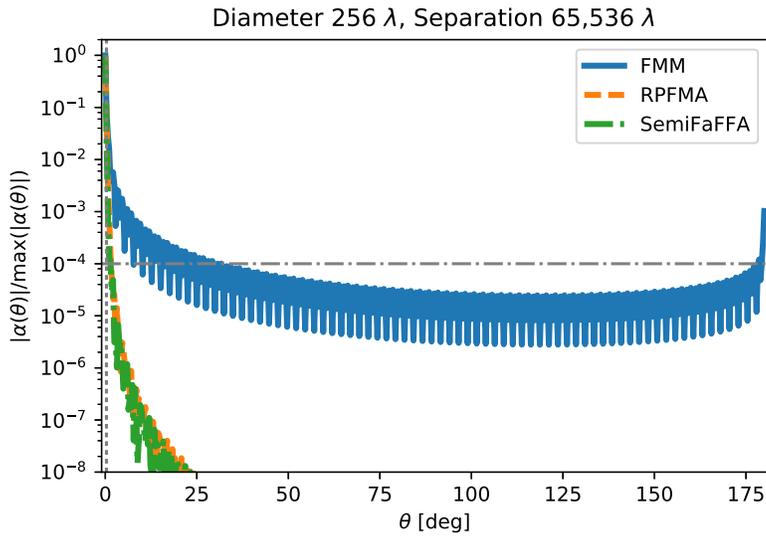


Figure 6.16: Translator amplitude for two spheres of diameter 256λ and separation $65,536 \lambda$.

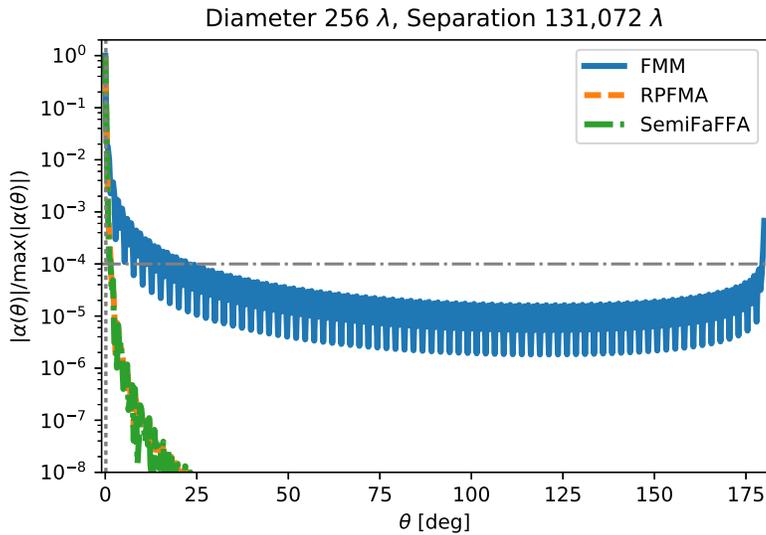


Figure 6.17: Translator amplitude for two spheres of diameter 256λ and separation $131,072 \lambda$.

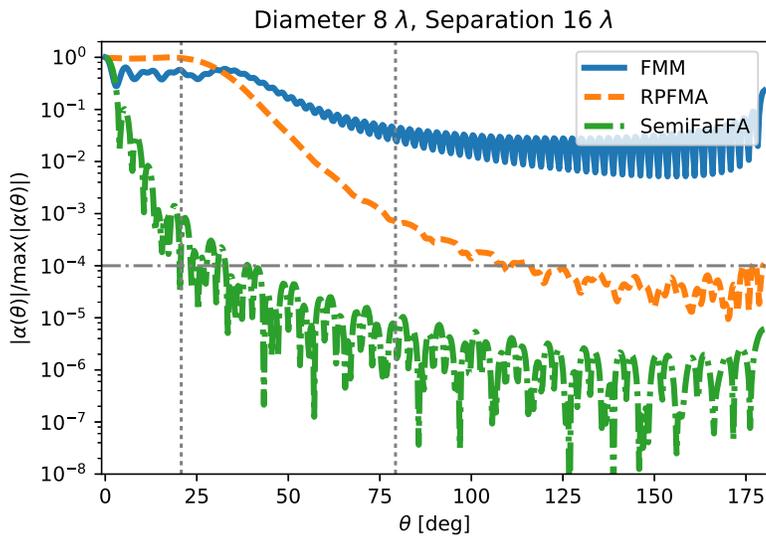


Figure 6.18: Translator amplitude for two spheres of diameter 8λ and separation 16λ .

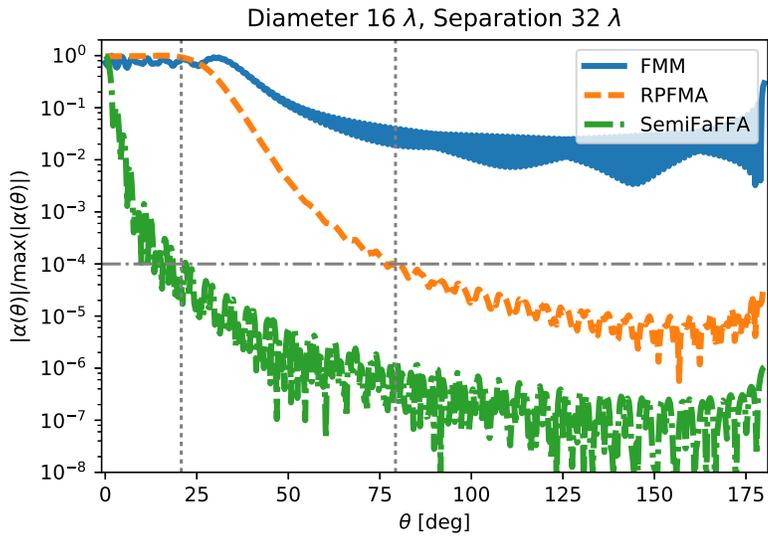


Figure 6.19: Translator amplitude for two spheres of diameter 16λ and separation 32λ .

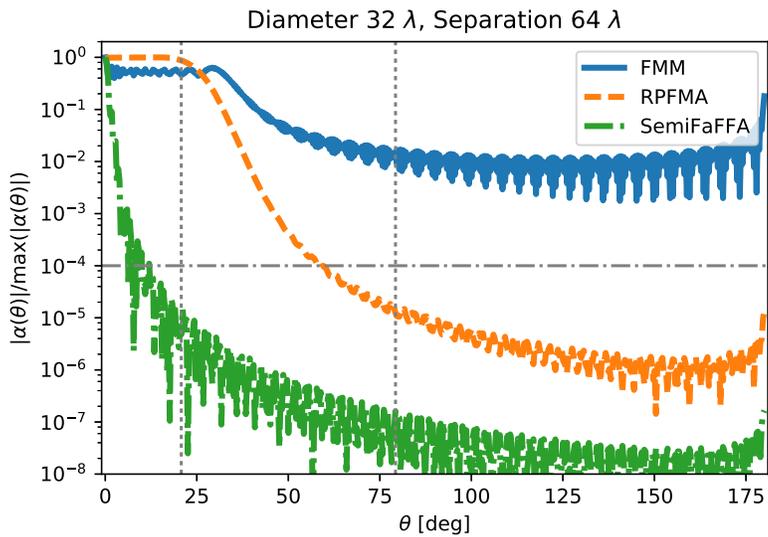


Figure 6.20: Translator amplitude for two spheres of diameter 32λ and separation 64λ .

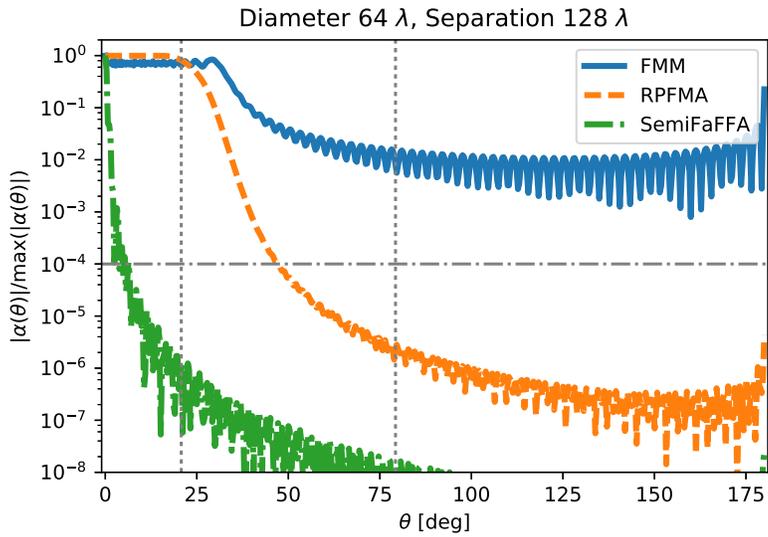


Figure 6.21: Translator amplitude for two spheres of diameter 64λ and separation 128λ .

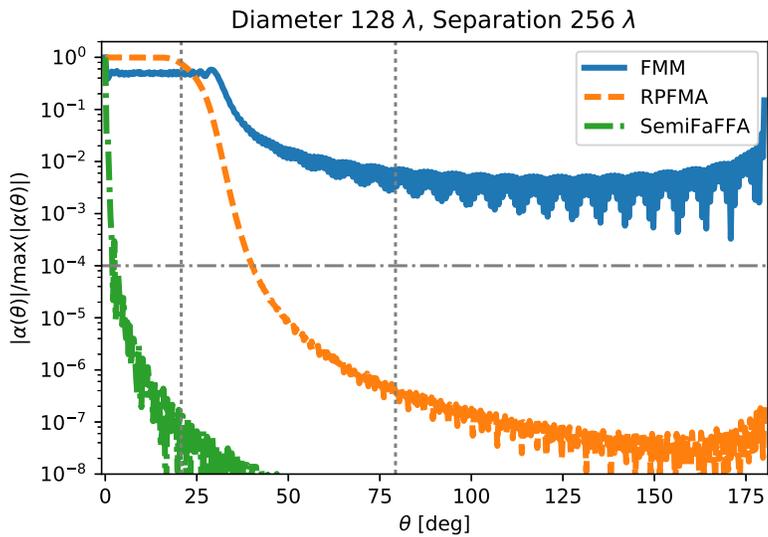


Figure 6.22: Translator amplitude for two spheres of diameter 128λ and separation 256λ .

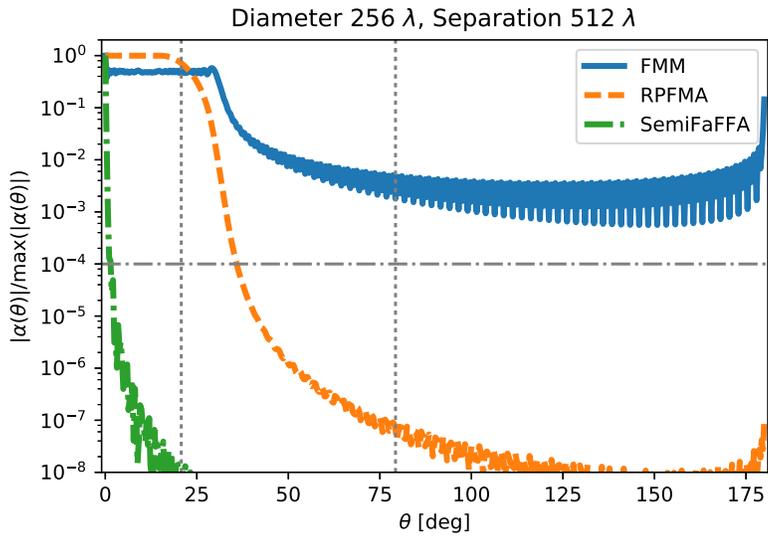


Figure 6.23: Translator amplitude for two spheres of diameter 256λ and separation 512λ .

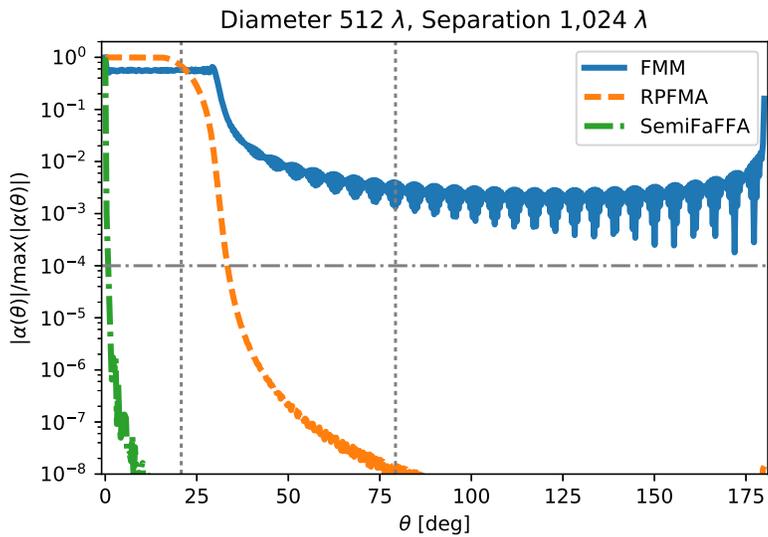


Figure 6.24: Translator amplitude for two spheres of diameter 512λ and separation $1,024 \lambda$.

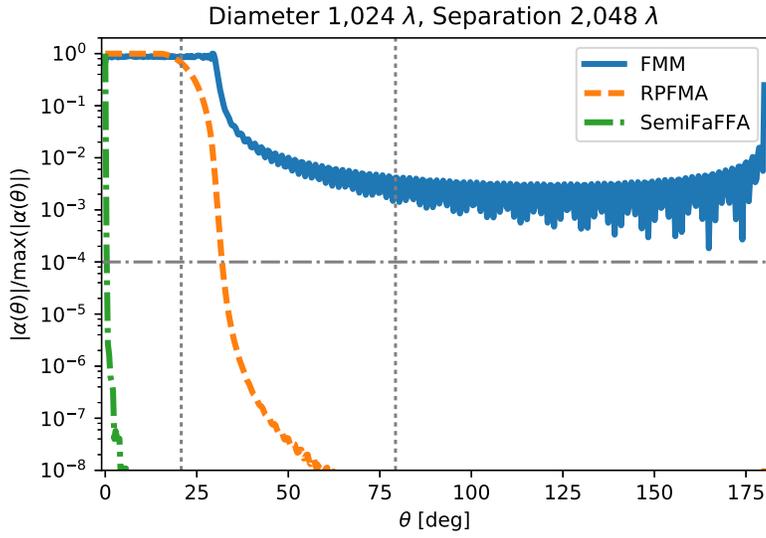


Figure 6.25: Translator amplitude for two spheres of diameter $1,024 \lambda$ and separation $2,048 \lambda$.

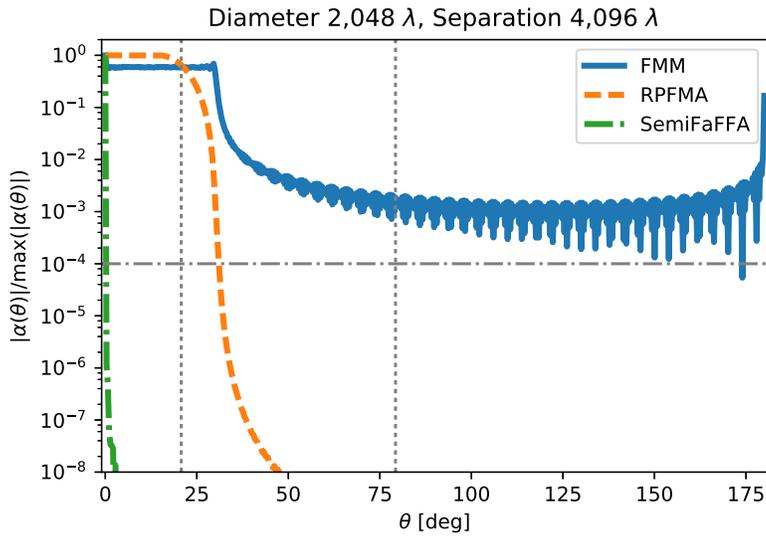


Figure 6.26: Translator amplitude for two spheres of diameter $2,048 \lambda$ and separation $4,096 \lambda$.

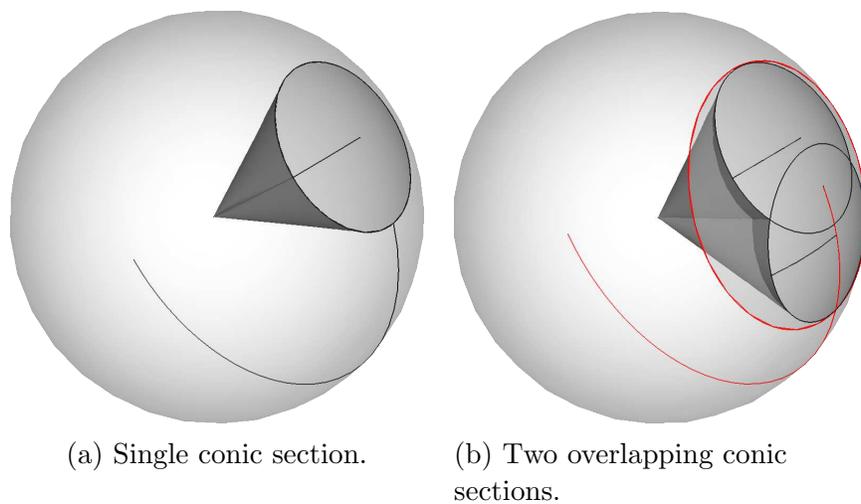


Figure 6.27: Conic sections of the Ewald sphere used to perform k -space analysis.

CHAPTER 7

CONCLUSION

7.1 Summary

The purpose of this work has been to identify solutions to a class of problems which overwhelm existing CEM approaches. EPA provides a framework well suited to a physically intuitive analysis of complicated multiple-object scattering cases using a generalized SIE method expressed in terms of identity operators, MLFMA-accelerated radiation, and object-tailored scattering analyses. A single large object overwhelms the other domains in the system of equations by requiring a large discretized scattering operator along with MLFMA radiation between it and all other domains. The Schur complement of this domain's identity block in the EPA system results in a much smaller system of equations with the large scattering operator isolated between MLFMA radiation calculations. The scattering operator is not feasible to compute in general, but broad classes of problems support the use of approximate scattering operators such that this method becomes more viable.

The PO approximation was applied as a high-frequency asymptotic model of this object's scattering operator, which yielded two SIE hybrids of EPA and PO. The original system of equations (with PO scattering operator) served as a reference to validate the performance of a novel Schur complement EPA-PO hybrid. The PO approximation was fully coupled into the greater SIE system without modification beyond the PO-specific radiation and scattering operators in both cases. The new Schur complement hybrid was found to be in agreement with the direct application of PO as a discretized EPA domain scattering operator. The Schur complement EPA-PO hybrid was analyzed without loss of generality to scattering operators derived from other high-frequency asymptotic models.

The PO calculations were further utilized as a mechanism for orchestrat-

ing massively parallel simulations. The locality of the PO approximation was leveraged to distribute the MLFMA and PO workload across an HPC cluster using a hybrid combination of distributed-memory parallel execution on separate cluster nodes and shared-memory parallel execution for portion of the work done on each cluster node. The observed consequence of partitioning the PO surface regions across cluster nodes was reduced execution times and a slow increase in required memory across the cluster so that multiple lower-memory nodes were added to the resource pool to complete a simulation more quickly than fewer nodes with higher per-node memory.

Discretized operators defining radiation between sets of basis and test functions comprise the SIE methods utilized in CEM. Reducing the workload in an FMM-based method requires reducing the number of plane wave samples required to radiate between the basis and test functions. To this end, the radiation between electrically large source and observation functions was analyzed in the space of plane waves. The radiation over large distances was observed to form beams independent of the source and observation block far-field patterns. Moreover, at sufficiently large size scale, the number of plane waves was observed to be reduced by subdividing the source and observation regions to form multiple overlapping narrower beams. This marks a transition from wide beams to ray physics, and it demonstrates size scales at which the far-field approximation and ray models may be accurately applied to accelerate radiation calculations. This observation will help accelerate hybridized SIE and high-frequency asymptotic methods such as the previously discussed Schur complement EPA-PO hybrid.

7.2 Future Work

The Schur complement EPA-PO hybrid was defined without loss of generality to other scattering operators. The approximated scattering operator is utilized only through MLFMA-accelerated radiation operators. These operators are diagonalized in the space of plane waves radiating from the source domain and received at the observation domain. The compound radiation-scattering-radiation sequence may be discretized in the space of plane waves as a dense operator in general. This dense operator is potentially expensive to evaluate, but it offers opportunities for hybridization of EPA and other

approximation scattering operator formulations. The EPA-PO hybrid implemented in this work can be seen to define such a dense operator using a composite product of discretized abstract operators. If FMM is utilized to perform a similarity transformation in the space of plane waves radiating to and from the approximate scattering operator domain and range, then the problem is expressed in the FMM plane wave space instead of coordinate space functions. The MLFMA-accelerated operators utilized in this work implicitly perform this transformation. In such a statement of the problem other methods of high-frequency asymptotic analysis defined for plane wave source and far-field observation become viable. The study on the formation of beams for electrically large source and observation regions will help inform which input and observation angles are expected to be the dominant contributions for each range- and domain-voxel pairing.

Exploring other hybrids of EPA with high-frequency asymptotic methods will also provide further opportunities to identify embarrassingly parallel calculations. Each iteration of the system solver requires solution of a bistatic scattering problem relating sources and fields from all domains, which may be evaluated explicitly by filling the dense Schur complement operator, or by evaluating the bistatic scattering problem at each iteration using an abstract operator. Either method provides opportunities to explore parallel calculations. If the method does not require normalization by a Gramian inverse on the scatterer surface, then the need for a distributed parallel sparse solver is also removed.

The study of subdivided RPFMA illustrated the formation of beams as objects and separations increase in linear size scale. This physical mechanism was identified in large objects, and does not explain the observed linear rank increase observed for smaller objects as obtained by SVD analysis. Because the mechanism for rank reduction is not found by analyzing the $\bar{\alpha}$ operator, the source and observation patterns are the remaining candidates for examination. Past efforts to predict the pattern characteristics of aperiodic antenna arrays may be applied to the source and observation voxel plane wave patterns and combined with the beam formation observed in this work to develop reduced rank models.

REFERENCES

- [1] J. C. Maxwell, “A dynamical theory of the electromagnetic field,” *Philos. Trans. R. Soc. London*, vol. 165, pp. 459–512, 1865.
- [2] O. Heaviside, *Electromagnetic Waves*. London, UK: Taylor and Francis, 1889.
- [3] W. C. Chew, *Waves and Fields in Inhomogeneous Media*. New York, NY: IEEE Press, 1995.
- [4] W. C. Chew, “Vector potential electromagnetics with generalized gauge for inhomogeneous media: Formulation,” *Prog. Electromagn. Res.*, vol. 149, pp. 69–84, 2014.
- [5] M.-K. Li, W. C. Chew, and L. J. Jiang, “A domain decomposition scheme based on equivalence theorem,” *Microwave Opt. Technol. Lett.*, vol. 48, no. 9, pp. 1853–1857, 2006.
- [6] M.-K. Li and W. C. Chew, “Wave-field interaction with complex structures using equivalence principle algorithm,” *IEEE Trans. Antennas Propag.*, vol. 55, no. 1, pp. 130–138, Jan 2007.
- [7] P. Ylä-Oijala and M. Taskinen, “Electromagnetic scattering by large and complex structures with surface equivalence principle algorithm,” *Waves Random Complex Media*, vol. 19, no. 1, pp. 105–125, 2009.
- [8] R. F. Harrington, “Matrix methods for field problems,” *Proc. IEEE*, vol. 55, no. 2, pp. 136–149, Feb 1967.
- [9] W. C. Chew, M. S. Tong, and B. Hu, *Integral Equation Methods for Electromagnetic and Elastic Waves*. Morgan and Claypool Publishers, 2009.
- [10] S. M. Rao, D. R. Wilton, and A. W. Glisson, “Electromagnetic scattering by surfaces of arbitrary shape,” *IEEE Trans. Antennas Propag.*, vol. 30, no. 3, pp. 409–418, May 1982.

- [11] M.-K. Li and W. C. Chew, "Using tap basis to implement the equivalence principle algorithm for domain decomposition in integral equations," *Microwave Opt. Technol. Lett.*, vol. 48, no. 11, pp. 2218–2222, 2006.
- [12] G. H. Golub and C. F. Van Loan, *Matrix Computations*, 3rd ed. Baltimore: The Johns Hopkins University Press, 1996.
- [13] Y. Saad and M. H. Schultz, "GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems," *SIAM J. Sci. Stat. Comput.*, vol. 7, no. 3, pp. 856–869, Jul 1986.
- [14] K. Zhao, M. N. Vouvakis, and J.-F. Lee, "The adaptive cross approximation algorithm for accelerated method of moments computations of emc problems," *IEEE Trans. Electromagn. Compat.*, vol. 47, no. 4, pp. 763–773, Nov 2005.
- [15] J. Song and W. C. Chew, "FMM and MLFMA in 3d and Fast Illinois Solver Code," in *Fast and Efficient Algorithms in Computational Electromagnetics*, W. C. Chew, J.-M. Jin, E. Michielssen, and J. Song, Eds. Boston, MA: Artech House, 2001.
- [16] X. S. Li and J. W. Demmel, "SuperLU_DIST: A scalable distributed-memory sparse direct solver for unsymmetric linear systems," *ACM Trans. Math. Softw.*, vol. 29, no. 2, pp. 110–140, June 2003.
- [17] U. Jakobus and F. M. Landstorfer, "Improved PO-MM hybrid formulation for scattering from three-dimensional perfectly conducting bodies of arbitrary shape," *IEEE Trans. Antennas Propag.*, vol. 43, no. 2, pp. 162–169, 1995.
- [18] K. Zhang, J. Ou Yang, F. Yang, J. Zhang, and Y. Li, "Analysis of multi-scale problem about antenna mounted on electrically large platform by using connected EPA-PO," *Prog. Electromagn. Res.*, vol. 126, pp. 169–183, 2012.
- [19] F. Obelleiro, J. M. Taboada, J. L. Rodriguez, J. O. Rubios, and A. M. Arias, "Hybrid moment-method physical-optics formulation for modeling the electromagnetic behavior of on-board antennas," *Microwave Opt. Technol. Lett.*, vol. 27, no. 2, pp. 88–93, 2000.
- [20] Z.-L. Liu and C.-F. Wang, "Efficient iterative method of moments-physical optics hybrid technique for electrically large objects," *IEEE Trans. Antennas Propag.*, vol. 60, no. 7, pp. 3520–3525, 2012.

- [21] A. Tzoulis and T. F. Eibert, “Hybridization of multilevel fast multipole method and uniform geometrical theory of diffraction for radiation and scattering computations,” in *2005 IEEE Antennas and Propagation Society International Symposium*, vol. 4A, July 2005, pp. 159–162 vol. 4A.
- [22] R. G. Kouyoumjian, “Asymptotic high-frequency methods,” *Proc. IEEE*, vol. 53, no. 8, pp. 864–876, Aug 1965.
- [23] G. L. James, *Geometrical Theory of Diffraction for Electromagnetic Waves*, revised ed. Stevenage, UK: Peter Peregrinus, 1980.
- [24] P. Pathak, “High frequency techniques for antenna analysis,” *Proceedings of the IEEE*, vol. 80, no. 1, pp. 44–65, Jan 1992.
- [25] P. Y. Ufimtsev, *Fundamentals of the Physical Theory of Diffraction*, 2nd ed. Hoboken, NJ: John Wiley and Sons, 2014.
- [26] S. W. Lee, H. Ling, and R. Chou, “Ray-tube integration in shooting and bouncing ray method,” *Microwave Opt. Technol. Lett.*, vol. 1, no. 8, pp. 286–289, 1988.
- [27] U. Jakobus and F. M. Landstorfer, “Improvement of the PO-MoM hybrid method by accounting for effects of perfectly conducting wedges,” *IEEE Trans. Antennas Propag.*, vol. 43, no. 10, pp. 1123–1129, 1995.
- [28] W. D. Burnside and R. J. Marhefka, “Antennas on aircraft, ships, or any large, complex environment,” in *Antenna Handbook: Theory, Applications, and Design*, Y. T. Lo and S. W. Lee, Eds. New York, NY: New York: Springer Science+Business Media, 1988.
- [29] G. M. Amdahl, “Validity of the single processor approach to achieving large scale computing capabilities,” in *Proceedings of the April 18-20, 1967, Spring Joint Computer Conference*, ser. AFIPS ’67 (Spring). New York, NY, USA: ACM, 1967, pp. 483–485.
- [30] J. L. Gustafson, “Reevaluating Amdahl’s law,” *Comm. ACM*, vol. 31, no. 5, pp. 532–533, 1988.
- [31] B. Michiels, J. Fostier, I. Bogaert, and D. D. Zutter, “Weak scalability analysis of the distributed-memory parallel mlfma,” *IEEE Trans. Antennas Propag.*, vol. 61, no. 11, pp. 5567–5574, Nov 2013.
- [32] B. Michiels, I. Bogaert, J. Fostier, and D. D. Zutter, “A well-scaling parallel algorithm for the computation of the translation operator in the mlfma,” *IEEE Trans. Antennas Propag.*, vol. 62, no. 5, pp. 2679–2687, May 2014.

- [33] M. A. Heroux and J. M. Willenbring, “A new overview of the Trilinos project.” *Scientif. Program.*, vol. 20, no. 2, pp. 83–88, 2012.
- [34] C. G. Baker and M. A. Heroux, “Tpetra, and the use of generic programming in scientific computing.” *Scientif. Program.*, vol. 20, no. 2, pp. 115–128, 2012.
- [35] H. C. Edwards, C. R. Trott, and D. Sunderland, “Kokkos: Enabling manycore performance portability through polymorphic memory access patterns,” *J. Parallel Distr. Comput.*, vol. 74, no. 12, pp. 3202–3216, 2014.
- [36] E. Bavier, M. Hoemmen, S. Rajamanickam, and H. Thornquist, “Amesos2 and Belos: Direct and iterative solvers for large sparse linear systems,” *Scientif. Program.*, vol. 20, no. 3, pp. 241–255, 2012.
- [37] W. Chai and D. Jiao, “Theoretical study on the rank of integral operators for broadband electromagnetic modeling from static to electrodynamic frequencies,” *IEEE Trans. Compon. Packag. Manuf. Technol.*, vol. 3, no. 12, pp. 2113–2126, 2013.
- [38] W. Chew, Y. Wang, and L. Gürel, “Recursive algorithm for wave-scattering solutions using windowed addition theorem,” *J. Electromagn. Waves*, vol. 6, no. 7, pp. 1537–1560, 1992.
- [39] R. L. Wagner and W. C. Chew, “A ray-propagation fast multipole algorithm,” *Microwave Opt. Technol. Lett.*, vol. 7, no. 10, pp. 435–438, 1994.
- [40] T.-J. Cui, W. C. Chew, G. Chen, and J. Song, “Efficient MLFMA, RPFMA, and FAFFA algorithms for EM scattering by very large structures,” *IEEE Trans. Antennas Propag.*, vol. 52, no. 3, pp. 759–770, March 2004.
- [41] C.-C. Lu and W. C. Chew, “Fast far-field approximation for calculating the rcs of large objects,” *Microwave Opt. Technol. Lett.*, vol. 8, pp. 238–241, 1995.
- [42] W. C. Chew, T. J. Cui, and J. M. Song, “A FAFFA-MLFMA algorithm for electromagnetic scattering,” *IEEE Trans. Antennas Propag.*, vol. 50, no. 11, pp. 1641–1649, Nov 2002.
- [43] M. Born and E. Wolf, *Principles of Optics*, 7th ed. Cambridge, UK: Cambridge University Press, 1999.
- [44] J. W. Goodman, *Introduction to Fourier Optics*, 2nd ed. Boston, MA: McGraw-Hill, 1996.

- [45] C. A. Balanis, *Advanced Engineering Electromagnetics*. Hoboken, NJ: Wiley, 1989.
- [46] T. B. Hansen, “Exact plane-wave expansion with directional spectrum: application to transmitting and receiving antennas,” *IEEE Trans. Antennas Propag.*, vol. 62, no. 8, pp. 4187–4198, 2014.