FDTD Simulation Techniques for Simulation of Very Large 2D and 3D Domains

Applied to Radar Propagation over the Ocean

by

Brandon Dowd

A Dissertation Presented in Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy

Approved July 2018 by the
Graduate Supervisory Committee:

Rodolfo E. Diaz, Chair
George Pan
Kevin Schmidt
James Aberle

ARIZONA STATE UNIVERSITY

August 2018

ABSTRACT

A domain decomposition method for analyzing very large FDTD domains, hundreds of thousands of wavelengths long, is demonstrated by application to the problem of radar scattering in the maritime environment. Success depends on the elimination of artificial scattering from the "sky" boundary and is ensured by an ultra-high-performance absorbing termination which eliminates this reflection at angles of incidence as shallow as 0.03 degrees off grazing. The two-dimensional (2D) problem is used to detail the features of the method. The results are cross-validated by comparison to a parabolic equation (PE) method and surface integral equation method on a 1.7km sea surface problem, and to a PE method on propagation through an inhomogeneous atmosphere in a 4km-long space, both at X-band. Additional comparisons are made against boundary integral equation and PE methods from the literature in a 3.6km space containing an inhomogeneous atmosphere above a flat sea at S-band. The applicability of the method to the three-dimensional (3D) problem is shown via comparison of a 2D solution to the 3D solution of a corridor of sea. As a technical proof of the scalability of the problem with computational power, a 5m-wide, 2m-tall, 1050m-long 3D corridor containing 321.8 billion FDTD cells has been simulated at X-band. A plane wave spectrum analysis of the (X-band) scattered fields produced by a 5m-wide, 225m-long realistic 3D sea surface, and the 2D analog surface obtained by extruding a 2D sea along the width of the corridor, reveals the existence of out-of-plane 3D phenomena missed by the traditional 2D analysis. The realistic sea introduces random strong flashes and nulls in addition to a significant amount of cross-polarized field. Spatial integration using a dispersion-corrected Green function is used to reconstruct the scattered fields outside of the computational FDTD space which would

impinge on a 3D target at the end of the corridor. The proposed final approach is a hybrid method where 2D FDTD carries the signal for the first tens of kilometers and the last kilometer is analyzed in 3D.

ACKNOWLEDGMENTS

The prior five years have comprised a path in my life unlike anything I could have ever before predicted. I cannot express enough gratitude to Professor Rodolfo Diaz for his insight and support which allowed the previously unimaginable to come to fruition. When I initially met Dr. Diaz, I was on a less than efficient academic path, filled with more directional uncertainty that I was comfortable with. He offered up a tremendous amount of advice and guidance, and when I proposed that I would like to volunteer in his lab over the summer to gather experience, he instead offered admittance into graduate school, a research assistant position in his lab group, and subsequently introduced me to the opportunities which, through diligent work and perseverance, will forever continue to shape the rest of my life. For this, and his unwavering dedication, I am forever grateful.

I would like to express my gratitude to my committee as well, Professors George Pan, Kevin Schmidt, and James Aberle. They have provided exceptional support both inside and outside of the classroom and have contributed tremendously to my learning.

I am very thankful for the patience and encouragement of my family and friends, and for their understanding of how difficult it is to attempt to maintain a social life while obtaining a PhD. They have never ceased to provide anything but assistance and encouragement, especially my parents, whose support has been limitless in every sense of the word.

Additionally, I would like to specially thank Glen Erler and the late Ryan Correy. Glen helped me assemble my first computer when I was 13 years old, and over the following several years dedicated unprecedented volumes of his time to share his knowledge and to aid me in troubleshooting a countless number of computer issues. He set

me up with the tools to pursue what would become one of my primary passions in life. Ryan's lifetime accomplishments extend well beyond someone twice his elder. From a young age, shortly after our family met Ryan, I began to model my personal definitions of the ideas of hard work, dedication, and focus based on what he exemplified. He's been a persistent inspiration about what it means to chase a goal.

I would also like to give thanks to Dr. Robert Burkholder of The Ohio State University for his method of moments contributions and collaboration on our initial cross-validation efforts with VTRPE. Additionally, to Dr. Frank Ryan and Dale Zolnick for providing very detailed VTRPE data used for our cross-validations. Finally, to our program manager, Dr. Steven Russell of the Office of Naval Research, for funding our research, and for the invaluable personal recommendations he has provided which have led to my first post-graduation career opportunities.

TABLE OF CONTENTS

Page

CHAPTER Page

CHAPTER Page

APPENDIX Page

LIST OF TABLES

Table Page

LIST OF FIGURES

Figure Page

Figure Page

Figure Page

Figure Page

Figure Page

Figure Page

Figure Page

Figure Page

Figure Page

Figure Page

Figure Page

CHAPTER 1

INTRODUCTION

1.1 Background of the Sea-Atmosphere Problem

There exist problems in applied electromagnetics in which the size of the domain of concern is not an incidental "technological" issue but rather a key parameter in the development of the physical phenomena of interest. These problems are usually characterized by a complex, inhomogeneous environment that, over a long distance, can transform the deterministic source fields produced by a localized radiator into a propagating phenomenon with stochastic properties. Examples are: the propagation of light in fog, radio waves in a forest, and radar signals over the sea. Of these, marine radar is an essential safety tool in the shipping industry, not only for the location of land and other ships but also for monitoring sea state conditions and forecasting the approach of adverse weather. [1, 2, 3]

Although physics-based models to estimate the propagation phenomenology in such cases have existed for a long time [4], eventually the need for ever more accurate characterization requires the application of fullwave computational electromagnetics modeling to the representative environments. Then Monte Carlo techniques can be used to rigorously derive the stochastic properties, validate the original models, and discover any new phenomenology that must be added to those models to maximize their fidelity. This is the situation with the problem of radar propagation over the sea, over distances measured in tens or hundreds of kilometers.

The standard tool available today for analyzing this problem is the Variable Terrain Radio Parabolic Equation method (VTRPE) applied to a two-dimensional (2D) slice of space, up to hundreds of kilometers downrange and hundreds of meters in height [5]. Fullwave alternatives in 2D include the surface integral equation (e.g. the forwards-backwards Method of Moments (MoM) [6, 7]) and the Boundary Integral Equation (BIE) methods [8]. However, the computational load of these other methods eventually limits them to domains of the order of a few kilometers long. Furthermore, they can only handle inhomogeneous atmospheres if they take the form of one of the special profiles admitting an exact Green function.

More importantly, all 2D approaches ignore out-of-plane phenomena. Yet these phenomena are known to cause flashes of increased power to be delivered to a target, well above the 2D expectation [9]. Similarly, transverse inhomogeneities in the dielectric constant of the atmosphere must be expected to provide focusing and defocusing paths to the propagating signal that cannot be seen in a 2D slice simulation. Although three-dimensional (3D) extensions of the parabolic equation method have been proposed and may contain some of the relevant physics [10, 11], the use of a scalar field to emulate a 3D vector field, that must exhibit time varying polarization properties, constitutes an approximation of unknown fidelity.

To correctly evaluate the radar-sea interaction, the strength of these out-of-plane effects must be assessed using rigorous fullwave methods, readily scalable from 2D to 3D. The time domain nature of the computational electromagnetics Finite-Difference Time-Domain (FDTD) method, and its straight-forward parallelization, make it a natural choice for the extreme domain decomposition that is required to bring the fullwave computational

burden to manageable levels. The key problem to be overcome is the accumulation of numerical artefacts over the extent of the simulation. In this work, an FDTD domain decomposition method is reported for this purpose. A 2D scenario exhibiting the stochastic nature of the scatter from the sea, simulated in FDTD at X-band as one of the first cross-validation cases, is illustrated in Figure 1.1.

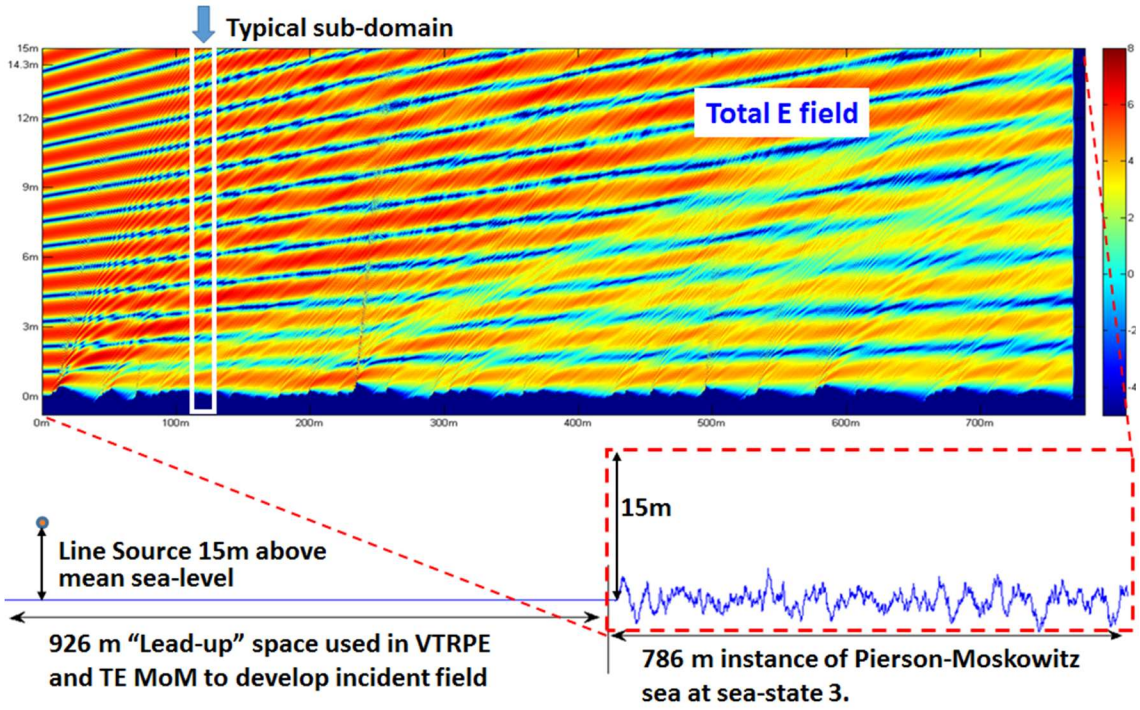

Figure 1.1: The baseline 2D computational space > 1km long used to cross-validate FDTD to VTRPE and MoM in X-band. (TE case shown.)

1.2 Placing This Work in Context with the Literature

The work reported here is similar in many respects to the excellent work by Lai et al on scattering from large (3D) rough surfaces (including the sea surface at 1.25GHz) [12] but it also differs in significant ways. The most important difference is that Lai et al consider the scattering of a wave incident from a distant source onto the large surface. This assumption allows them to limit domain-to-domain coupling to only adjacent domains

because the principal incident field exciting every domain of the problem is the direct illumination by the distant source.

From the viewpoint of the plane wave spectrum representation of electromagnetic fields (with the coordinate x-y plane assumed to be parallel to the rough surface) [13], we can say that their source field and their observation direction (located at a distant z-plane) both lie in the visible region, while the invisible, evanescent, part of the spectrum is only involved in adjacent domain interaction. In our case, our source is a low directivity radiator (9GHz dipole) suspended a few meters above the surface of the sea. And the field of interest is the propagating field that exists also near the surface of the sea (within about 20 meters), several kilometers away. Therefore, in our problem we are mostly concerned with long range propagation along the horizon, that is, the part of the plane wave spectrum existing right at the boundary between the visible and evanescent regions of the spectrum.

Furthermore, the principal incident field in our case is not the direct illumination from the source. Back and forth interaction of the X-band wave between successive crests of the sea surface is critical to the phenomena we seek to investigate. And it is clear that after a few hundred meters of propagation the principal field entering a computational domain down range is not just the direct ray from the source but the continuously diffracted signal accumulating along the entire surface.

We are thus working in the anomalous range of extreme grazing incidence [14, 15]. Therefore, the problem we address can be considered complementary to that reported by Lai et al and, as such, many parallels will be found between the techniques employed here and in [12] to solve it; but each is uniquely suited to the particulars of its target problem. For instance: In [12] the domain-to-domain coupling is performed via near field radiation

by equivalent currents collected on the reference plane above the surface. In this work the domain-to-domain coupling is effected by direct transference of the FDTD fields from one domain to the next using the exact discrete Schelkunoff currents implemented via the field teleportation algorithm on the boundary connecting the domains [16].

Domain-to-domain coupling in [12] is limited to adjacent domains because this excitation is secondary to the principal, incident field from the remote source and because the observed field of interest is dominated by the power radiated to the far field by every domain. In this work, domain-to-domain coupling is limited to adjacent sequential domains because each domain is made large enough to contain the essential forward and backward interaction between successive wave crests [7] and because the observed field of interest is the forward scattered total accumulated field reaching a distant plane close to the surface of the sea in the final sub-domain. Whereas in [12] the typical sub-domain is up to 13 wavelengths on the side, in this work the typical sub-domain is 500 wavelengths long in 2D and 100 wavelengths long in 3D.

In [12] CPML boundaries are used to terminate the domains and prevent any unphysical scattering from reaching the far field. In this work, because of the extreme shallow angles of propagation involved, we cannot use PMLs, nor can we use absorbing boundaries optimized to a set of angles [17, 18] since the effective angle of incidence on the sky boundary becomes shallower and shallower with distance (of the order of 0.2 degrees off-grazing at 4 kilometers and as low as 0.03 degrees in special cases as will be shown below). Instead we use a new, extremely high-performance re-Radiating Absorbing Termination (rRAT) derived from the self-teleportation algorithm of the re-Radiating Boundary Condition (rRBC) [19].

The success of the method in [12] is validated by comparing their results on a 32λ by 32λ surface (an estimated 100 to 500 million 3D FDTD cells) to integral equation and more canonical classic approximations [20]. (The method is then demonstrated on a 66λ by 66λ surface.) In this work we validate the 2D result against VTRPE and MoM X-band solutions on a 1.7km-long problem (8.8 billion FDTD cells), then against a boundary integral equation S-band solution [8] of a 3.6km-long flat sea with an inhomogeneous atmosphere (65.1 billion cells), and another X-band 2D cross-validation against a VTRPE solution of a 4km realistic sea-atmosphere space (43.3 billion cells). As a proof of scalability, we apply our techniques to a 15km (450,000λ) long, 41m (1,240λ) tall 2D inhomogeneous sea-atmosphere space comprised of 178.3 billion cells and 11.4 million time steps. We then extend the techniques into three dimensions in order to ascertain the existence of what relevant 3D out-of-plane physics is currently missed by 2D methodologies.

In summary, we are reporting the fullwave solution and cross-validation of the electromagnetic scattering problem between a source and a nearby nominal 122,500 wavelengths long lossy surface in 2D. Additionally, to demonstrate that the techniques which apply to the 2D grazing incidence and grazing observation problem translate directly to 3D, we show a cross-validation with the 2D problem by solving it in 3D for a "corridor" of sea 110m (3,300λ) long, 1.3m (40λ) wide, and 2.25m (68λ) tall. Upon verification of functionality in 3D, we analyze the plane wave spectrum of the scatter from a large corridor of realistic 3D sea surface in comparison to a 2D-analogous "extruded" sea corridor. Through application of special techniques, we are able to generate the FDTD scattered fields over a relevantly large spatial region in front of a sea-based radar target, so as to

subsequently produce angle of arrival information for the fields impinging upon the target. The plane wave spectrum and angle of arrival will serve to demonstrate the 3D physics missed by current methods. They are, in fact, the quantities which determine whether current 2D methods are "enough." This information is gathered from 3D FDTD simulations at X-band frequencies of corridors of realistic 3D sea in upwards of 1,050m (31,500λ) long, 5m (150λ) wide, and 2m (60λ) tall containing 315.3 billion FDTD cells; this appears to be the largest FDTD problem reported in the literature to date.

Novel features of the work reported here include:

- A stackable, stable, and ultra-high-performance FDTD absorbing boundary that eliminates artificial reflections from the termination of the domain out to 0.03 degrees from grazing.
- Domain decomposition by seamless transference of fields by use of the exact discrete Schelkunoff currents (teleportation algorithm.)
- Dispersion-corrected tracking of the propagating (and spreading) time domain pulse to maintain fidelity over the run.
- Dispersion-corrected analytic field injection FDTD implementation over the same domain to cross-validate the results against a purely scattered field formulation.
- Cross-validation of 2D sea scatter prediction fidelity against parabolic equation and fullwave MoM solvers; and cross-validation of the application to inhomogeneous atmosphere problems against a BIE result in the literature.
- Cross-validation of the 2D code with a 3D implementation on a 110m long sea corridor.

- Plane wave spectrum and angle of arrival analysis on simulations of large corridors of realistic 3D sea up to 31,500λ long, 150λ wide, and 60λ tall, containing over 315.3 billion cells.
- Dispersion-corrected 2D and 3D scattered field reconstruction, which drastically reduces the computational requirements of sea spaces containing a homogeneous atmosphere.
- Multi-GPU implementation of the computer code.
- Proof-of-concept techniques which may allow for hybrid 2D-3D FDTD modeling to produce solutions of very large inhomogeneous sea-atmosphere spaces which contain 3D physics.

Insofar as the solution of extremely large problems in FDTD is concerned, the literature stretches back to the rise of modern FDTD computation:

In 1985, Taflove et al [21] reported using a Cray-1S supercomputer with 8MB of RAM to simulate a 112 x 48 x 48 (258,048 cells) FDTD space containing a 3-D scatterer with a maximum dimension of 9λ with 16-bit precision; thirty times larger than any previous simulation. In 1987, Kriegsmann et al, used the On-Surface Radiation Condition to enable the computation of the scattering from objects greater than 10λ in size [22]. By 1992, a domain decomposition method was demonstrated by Umashankar et al which increased the size of the object addressed to more than 100λ [23].

Motivated by the distribution of the computational load over multiple compute nodes, in 1994 Nguyen et al [24] formulated an algorithm to divide an FDTD domain into a collection of sub-domains which would be sent to individual nodes for processing at each

time step. A multi-threaded solitary machine implementation was demonstrated by Yang et al (2003) [25] for arbitrarily shaped divisions of the domain. In 2014, a 2D single-domain rough surface, 6554λ long, was solved using a GPU implementation of FDTD [26].

To solve much larger problems requires both parallelization and decomposition at the physical problem level. In 2006, Abd-El-Raouf et al [27] published a novel way to decompose large electromagnetics problems into computationally manageable chunks. The domains were linked together by overlapping grids, using the time history of the tangential E-fields to create a radiation boundary to propagate the fields into the next domain, and processed serially. Issues with that implementation included: Overlap regions of the grid could take up as much as 50% of the sub-domain size, increased simulation time over non-domain-decomposed approaches, and imperfections from the overlap region truncation introduces errors. The work continued in [28] culminating in [12].

This present work represents the next steps in this long history of FDTD domain decomposition development, showing the solution of a 15km stretch of sea at X-band, a scattering surface 450,000 wavelengths long. Additionally, the methods are extended into three dimensions, encompassing corridors of realistic 3D sea more than 280 million cubic wavelengths in size.

Organization of the dissertation: The mechanics of the method are outlined in 0, including the key FDTD update equations used for the implementation of the field transfer, shallow angle attenuation, and scattered field formulation. Cross-validation against other techniques follows in Chapter 3 against VTRPE and Forward-Backward MoM, followed by BIE, and against the scattered field VEP formulation. Chapter 4 explores the initial 3D implementation cross-validated against the 2D results. Chapter 5 demonstrates the validity

of the techniques when applied to solve a relevantly-sized 3D sea-corridor problem, with 0 exploring the development of the realistic sea surface which fills such a sea corridor. Chapter 7 qualifies and quantifies the observable out-of-plane phenomena existing in 3D over 2D using a plane wave spectrum analysis. Chapter 8 develops a spatial integration reconstruction method to augment the size of the solution space and builds upon the quantification of relevant out-of-plane phenomena of Chapter 7 by investigation of the angle of arrival characteristics of the scattered fields at several points down range. Lastly, in Chapter 9, methods are proposed to hybridize 2D and 3D solutions to efficiently solve the 3D inhomogeneous atmosphere problem.

CHAPTER 2

SPECIAL TECHNIQUES FOR PROBLEMS OF EXTREME SIZE

2.1 The Mechanics of the Method

Figure 2.1 illustrates the typical problem of concern: An instance of the sea surface (vertical scale exaggerated), constructed using a modified Pierson-Moskowitz [29] spectrum, and an inhomogeneous atmosphere with a realistic ducting profile [30], are excited by an X-band source suspended several meters from the surface. The total field arriving at the last sub-domain is to be found. As the figure suggests the total field propagating in this environment consists of the direct ray (black arrow – which may be refracted by the atmosphere or even trapped in a duct), rays repeatedly reflected and diffracted by the sea surface (red arrows), rays redirected by the inhomogeneous atmosphere (green arrow), and rays exiting the domain at the sky boundary many of which can impinge on this boundary at extremely shallow angles (blue arrows) either because of refraction in the atmosphere or simply because of the very long distance traveled.

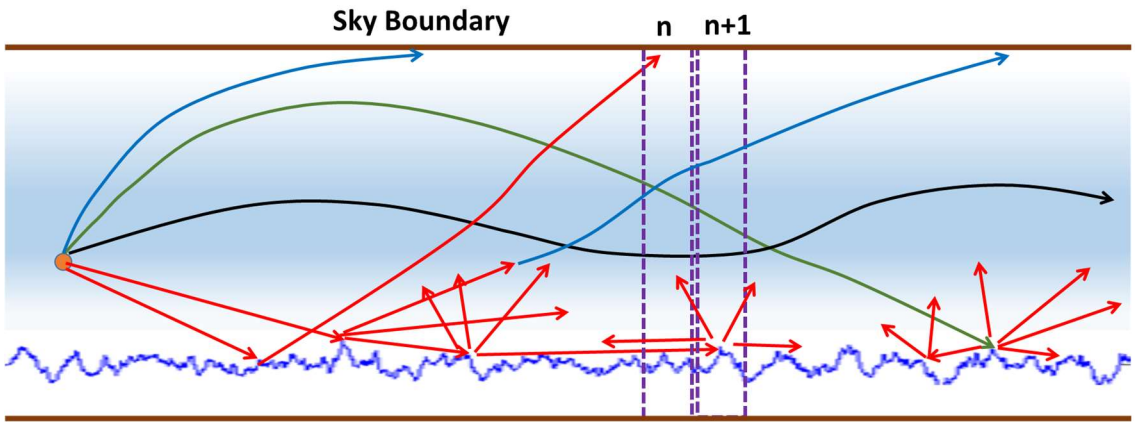

Figure 2.1: Possible ray paths for a source above the sea (exaggerated vertical scale) in the presence of an inhomogeneous atmosphere.

To propagate this electromagnetic field multiple kilometers at high frequency (X-band) with high fidelity requires (i) seamless transfer of the fields from sub-domain to sub-domain (e.g. from n to n+1 in Figure 2.1) and (ii) prevention of artificial reflections from the terminating boundaries of the sub-domains. The key to accomplishing both requirements is the application of the teleportation algorithm [16]. Because this algorithm creates exact discrete Schelkunoff equivalent currents in the FDTD space, the fields so generated only propagate into the intended space and exhibit zero leakage backwards (to within the machine precision truncation error.)

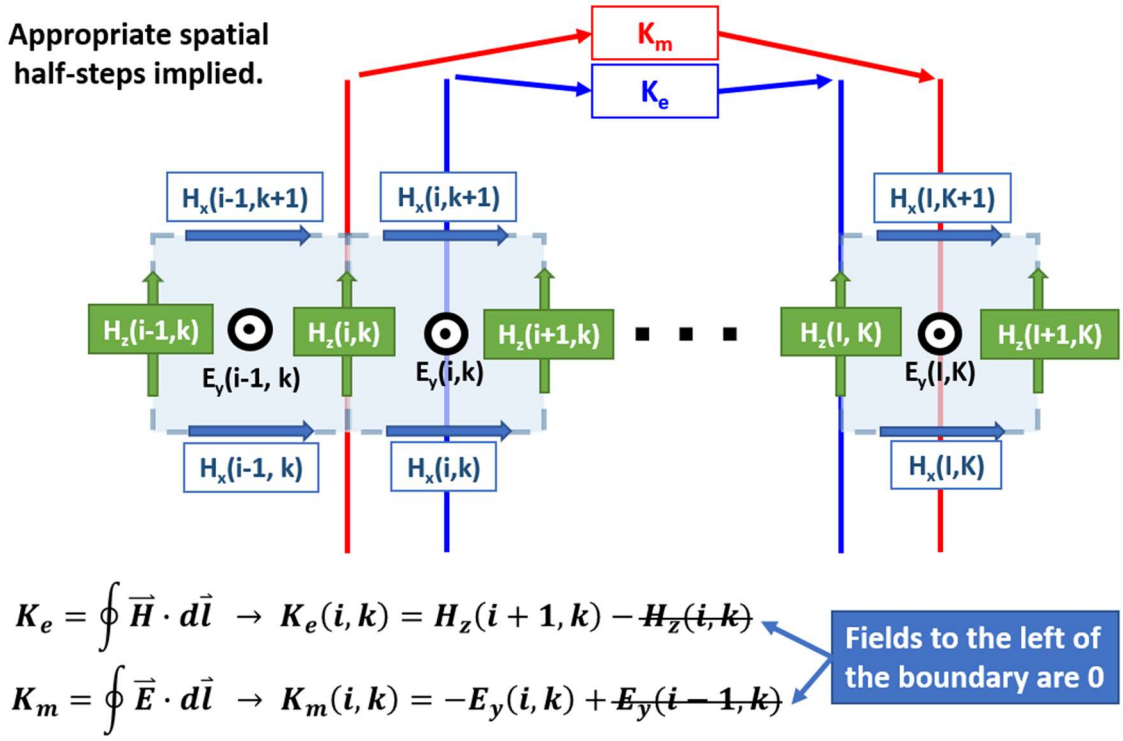

Figure 2.2: The teleportation algorithm creates equivalent currents from the collection plane E and H fields (left) and these are used to inject E and H fields at the teleportation plane (right).

2.2 The Transfer of the Signal from Sub-Domain to Sub-Domain

Because we are working in the time domain, each sub-domain needs to be long enough to allow sufficient back-and-forth interaction with sea features to accumulate into

the forward traveling pulse from the source before passing that signal on to the next sub-domain (detailed in Chapter 5). The time history of the pulse is converted into a time history of Schelkunoff surface currents (at the collection window) as it exits the sub-domain, is stored, and then used to teleport the fields to the next sub-domain (Figure 2.2.)

This is the simplest instance of the teleportation algorithm of Watts [16] because it need only be done over a plane. The typical update loop scheme for the 2D TEy case, teleporting from i_0 to i_{tele}, is shown in equations 2.1 and 2.2. Let $K = k + \frac{1}{2}$, $I_{0E} = i_0 + \frac{1}{2}$, and $I_{teleE} = i_{tele} + \frac{1}{2}$ for simplicity.

1. Standard E update
2. Teleport E:

$$E_y(I_{teleE}, K) = E_y(I_{teleE}, K) - A \cdot H_z(i_0, K) \quad 2.1$$

3. Standard H update
4. Teleport H:

$$H_z(i_{tele}, K) = H_z(i_{tele}, K) - B \cdot E_y(I_{0E}, K) \quad 2.2$$

Where $A = \left(\frac{2\Delta t}{1+\sigma_e \Delta t}\right) \cdot \frac{1}{\Delta s}$ and $B = \left(\frac{2\Delta t}{1+\sigma_m \Delta t}\right) \cdot \frac{1}{\Delta s}$ are standard FDTD update coefficients.

While in [12], the fields are tapered in the overlapping region of the sub-domains, in our case the surface of each sub-domain is tapered to flatness outside the incoming and outgoing teleportation windows. The effect is the same: to minimize backscatter through the connecting boundary. The tapering is done by reducing the deviation of the sea surface from the edge value down to the local mean sea level over 3 wavelengths and then

terminating the domain with flat sea into FDTD boundary termination (Figure 2.3 top). Because the sea surface topography is continuous across the outgoing window of domain "n" and the receiving window of domain "n+1" the transfer is seamless (Figure 2.3 bottom), that is, backscatter noise is at the -140dB word truncation level for single precision arithmetic.

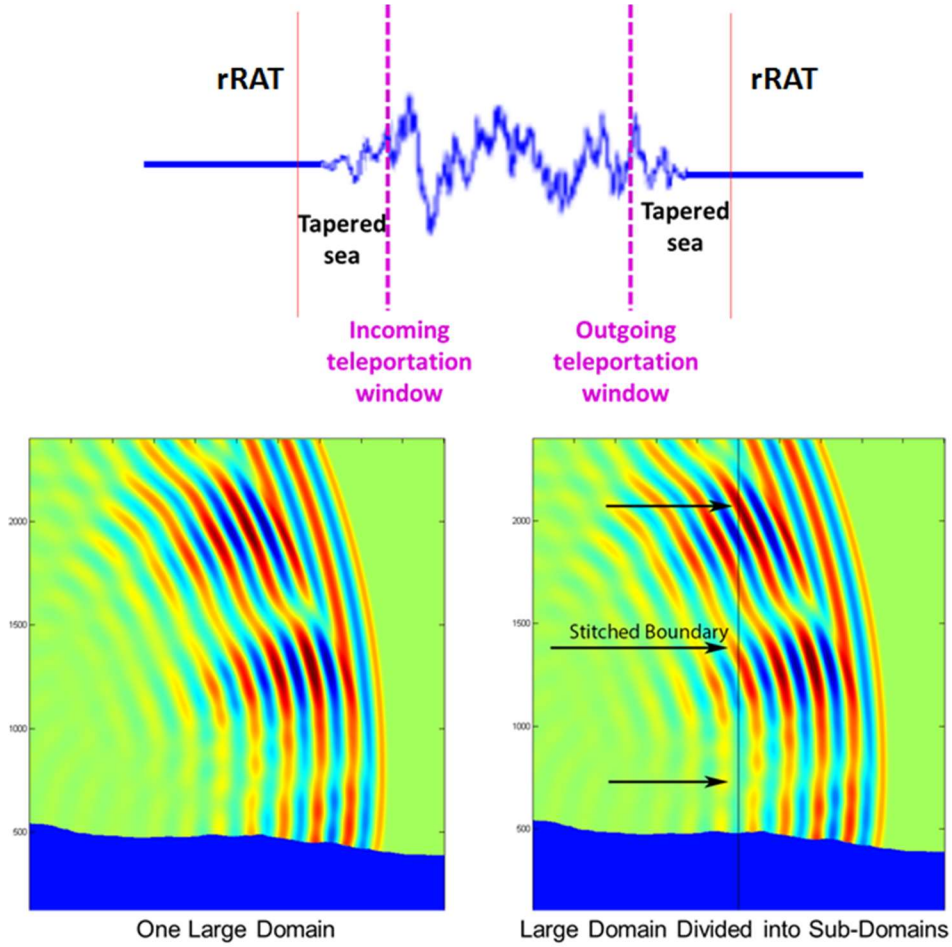

Figure 2.3: The stitching together of sub-domains is seamless, and the same forward-travelling information will be observed in a single large domain (left) as well as in a subdivided large domain (right).

Even though the source pulse is short (nominally a 128-cycle sinusoidal wave packet), the propagating field is continually accumulating more and more diffracted energy. As a result, the propagating pulse to be tracked becomes increasingly longer with

distance. (The rough maritime environment, like fog, behaves as a diffusing medium.) Therefore, it is necessary to account for this growth over the length of the computer run.

As a minimum, the time history recorded at the window, must last long enough to contain the leading-edge ray and a ray that has traversed one round trip in that domain: from the bottom corner horizontally to the opposite bottom corner and back again to the opposite top corner (approximately 38,000 time steps in a 500λ long, 1,100λ tall domain). We find that after 4km, we must actually record 10.2 million time steps in the window time history.

In addition, to accurately track the leading edge we must take into account the dispersive effect of the grid on the signal's group velocity (eq. 2.3 [31]) along the principal axis (x-axis) which is not quite the speed of light. We find that with a Courant number of 0.707 in 2D, and a resolution of 17.8 cells per wavelength, the effective speed of propagation of this leading edge along the x-axis is $0.9921c_0$.

$$v_g = \sqrt{1 - \left[\frac{\Delta s}{c_0 \Delta t}\sin\left(\frac{1}{2}\omega\Delta t\right)\right]^2} \cdot \left[\cos\left(\frac{1}{2}\omega\Delta t\right)\right]^{-1} \qquad 2.3$$

2.3 The Truncation of the FDTD Domains

The use of the teleportation algorithm to implement an FDTD boundary termination (in 2D and 3D) is reported in [19, 32, 33]. This termination, called the re-Radiating Boundary Condition (rRBC), seeks to eliminate outgoing waves by teleporting a negative copy of the wave onto itself as it travels outwards. The field information (in the form of Schelkunoff equivalent currents) is recorded at a collection plane, stored for two time steps and then teleported with a negative sign one step forward in space (ds = dx, dz). This

scheme avoids feedback between the copied and the injected fields; and, for the case when the time-step $\Delta t = 0.5\Delta s/c_0$, it results in near synchrony between the injected fields and the propagating fields, thus resulting in a strongly subtracted electromagnetic field beyond the teleportation plane. Because it uses the same update equations as FDTD, there is no requirement for the subtracted field to be a plane wave; in fact, reactive fields in close proximity to a source are just as easily teleported and (approximately) cancelled.

To understand the effect of the rRBC as an operator, consider the case of a nearly plane wave exiting the domain: The greatest difference between the negative teleported version of the wave and the propagating wave one step forwards is the slight error in synchrony, call it δt, a quantity that is smaller than Δt, the time-step, and certainly much smaller than the temporal variation of the wave. Then the subtracted field is of the form:

$$f(t) - f(t - \delta t) \approx \delta t \cdot \frac{\Delta f}{\Delta t} \qquad 2.4$$

And we see that the rRBC is an *active surface* that differentiates the field crossing it and scales it down (attenuates it) by the synchrony error.

Because the rRBCs update is identical to the conventional FDTD update, and is implemented on a plane (not a volume), it has very low computational burden and can cut seamlessly through inhomogeneous media. However, because it is an active surface it is not impervious to instability. For instance, to avoid instability due to mutual feedback between successive rRBCs, they are placed no closer than 3 cells apart.

Of more concern in a computational method is the sensitivity to numerical (truncation) noise. As reported in [32] the *attenuating differentiator* property of the rRBC for a wave traveling forward through it means that a wave traveling backwards through it

is integrated and therefore amplified. Nevertheless, because the backwards amplification is less than the forward attenuation, a stack of three rRBCs terminated in a trivial one-cell Salisbury screen (377Ohm termination) operates as an excellent absorber for most conventional problems. The typical FDTD rRBC boundary termination echo has been reported to be less than -80dB up to 10 degrees from grazing incidence.

The sensitivity to noise is manifested in the problem of diminishing returns: Although rRBCs in principle can seamlessly cross each other (x-plane boundaries cross y-plane boundaries and z-plane boundaries at the corners of the domain), they are non-linear, active boundaries. This means that numerical truncation noise can eventually interfere with their stability at those crossings. Thus, even though rRBCs can be stacked in layers to improve total absorption, stacking too many layers will eventually cause the rRBCs to feedback on each other, and introduce late-time spurious noise into the FDTD domain.

It is for this reason that in the published work the stack has been traditionally limited to three layers. To reduce the sensitivity to noise, the proportion of the stored field to be used for injection must be reduced. Typically, rRBCs store fields at the collection plane with a 0.99 multiplier. To attain noise-free operation with a stack of 7 rRBCs requires a multiplier of 0.96 or lower.

The above detailed discussion is necessary because the rRBC is the essential building block for the new rRAT termination, and because the rRBC has been classified as a "Huygens boundary" [34] even though its properties as an active surface are more complicated than that name implies.

For instance, the rRBC is angle-tunable. If the field information at the collection boundary is stored both one time step in the past and two time steps in the past, and these

are used in a linear combination for the teleported field, the attenuation of waves crossing the boundary can be made stronger for a selectable angle of incidence. This follows intuitively from realizing that the synchrony between the fields stored two time steps in the past and the field reaching one space step forward depends on the angle the wave makes with the plane that defines the term "forward".

Thus, it can be shown that the rRBC can be tuned to be an excellent absorber at a particular angle of incidence. But, as already stated, in our problem of concern the angle of incidence onto the sky boundary gets shallower and shallower as the propagation progresses, so selecting any one angle is never optimal, and changing the recipe from sub-domain to sub-domain would result in scattering from the discontinuity of the recipe at the teleportation window.

In summary, although the rRBC constitutes a simple, robust terminating condition for many conventional FDTD problems, it is not an adequate termination when dealing with extremely shallow angles of incidence. The inability to derive better performance from stacking more layers also hinders its use in problems requiring extremely low boundary reflections.

Although much progress has been made in optimizing the PML recipe for wide angle performance, there has been no absorbing termination reported that is (theoretically) perfectly reflectionless at arbitrarily shallow angles of incidence. The progression in the literature goes from a 75 degree incidence (15 degrees off grazing) in 1996 to 88 degrees incidence (2 degrees off grazing) by 2011 [17, 35, 36, 37, 38]. But to model realistic lengths of sea surface we need an absorbing termination which can handle tenths (and eventually hundredths) of degrees from grazing.

To remedy the two cited inadequacies of the rRBC and to attain reflectionless domain truncation down to true grazing incidence, we have implemented a new attenuating surface: the re-Radiating Absorbing Termination (rRAT).

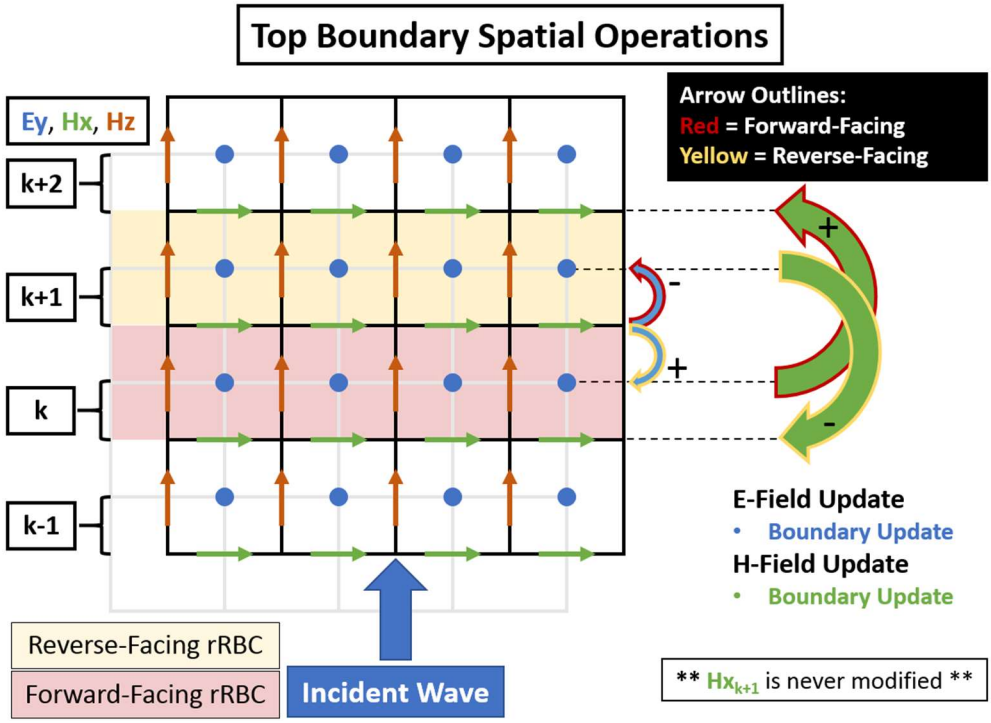

Figure 2.4: Spatial algorithm of the re-radiating absorbing termination (The time-delay operations are not explicitly shown here.)

The origin of the rRAT can be traced to the observation that backwards amplification through the rRBC is always less than the forwards attenuation. It follows that when a wave passes through a combination of two rRBCs – one in the forward and one in the reverse direction – the net result is a decrease in wave amplitude. Now, even though two rRBCs facing the same direction cannot be placed closer than 3 cells because their "copy" and "teleport" boundaries overlap and cause a feedback instability, a reverse rRBC can be placed one cell directly behind the forward rRBC and their back-to-back boundaries never modify the same cells simultaneously, thus this instability is avoided. For instance, in the illustration of a portion of the top boundary of a 2D TEy domain (Ey, Hx, Hz) in

Figure 2.4, even though the Hx field at (i, k+1) is used by the reverse rRBC, it is not addressed by the forward one.

This forward-backward rRBC boundary configuration is the rRAT. It is implemented the same way in 2D and 3D. We give the update equations in 2D for the sake of simplicity, but as will be shown in Chapter 4, they operate just as well in 3D, here shown in the TEy (x-z plane) case for the top-side wall positioned at z = K_{rRAT}, which is also diagrammed in Figure 2.4. Let $K_E \equiv K_{rRAT} + \frac{1}{2}$, $K_H \equiv K_{rRAT}$, and $I \equiv i + \frac{1}{2}$, also the parameter S_c controls the amount of field to teleport, usually S_c = 0.99. Additionally, $Ey_{present}$ is the E_y field at the current time-step, Ey_{past} is E_y at the previous time-step, and Ey_{stored} is the E_y field which will be stored and used as Ey_{past} in the n+1[st] time-step. The coefficients $A = \left(\frac{2\Delta t}{2\varepsilon+\sigma_e\Delta t}\right)\cdot\frac{1}{ds}$ and $B = \left(\frac{2\Delta t}{2\mu+\sigma_m\Delta t}\right)\cdot\frac{1}{ds}$ are standard FDTD update coefficients.

1. Standard E update
2. Update E in rRAT boundaries:
    a. Recall stored 'past' field and store new 'past' field:

Forward wall:

$$\begin{aligned} Hx_{past}(I, K_H+1) &= Hx_{stored}(I, K_H+1) \\ Hx_{present}(I, K_H+1) &= S_c \cdot H_x(I, K_H+1) \\ Hx_{stored}(I, K_H+1) &= Hx_{present}(I, K_H+1) \end{aligned} \quad (2.5)$$

Reverse wall:

$$\begin{aligned} Hx_{past}(I, K_H+1) &= Hx_{stored}(I, K_H+1) \\ Hx_{present}(I, K_H+1) &= S_c \cdot H_x(I, K_H+1) \\ Hx_{stored}(I, K_H+1) &= Hx_{present}(I, K_H+1) \end{aligned} \quad (2.6)$$

b. Update boundary fields

Forward wall:

$$E_y(I, K_E + 1) = E_y(I, K_E + 1) - A \cdot Hx_{past}(I, K_H + 1) \quad 2.7$$

Reverse wall:

$$E_y(I, K_E) = E_y(I, K_E) + A \cdot Hx_{past}(I, K_H + 1) \quad 2.8$$

3. Standard H update

4. Update H in rRAT boundaries:

a. Recall stored 'past' field and store new 'past' field:

Forward wall:

$$\begin{aligned} Ey_{past}(I, K_E) &= Ey_{stored}(I, K_E) \\ Ey_{present}(I, K_E) &= S_c \cdot E_y(I, K_E) \\ Ey_{stored}(I, K_E) &= Ey_{present}(I, K_E) \end{aligned} \quad 2.9$$

Reverse wall:

$$\begin{aligned} Ey_{past}(I, K_E + 1) &= Ey_{stored}(I, K_H + 1) \\ Ey_{present}(I, K_E + 1) &= S_c \cdot E_y(I, K_E + 1) \\ Ey_{stored}(I, K_E + 1) &= Ey_{present}(I, K_E + 1) \end{aligned} \quad 2.10$$

b. Update boundary fields

Forward wall:

$$H_x(I, K_H + 2) = H_x(I, K_H + 2) + B \cdot Ey_{past}(I, K_E) \quad 2.11$$

Reverse wall:

$$H_x(I, K_H) = H_x(I, K_H) - B \cdot Ey_{past}(I, K_E + 1) \quad 2.12$$

Even though the one-way attenuation through one rRAT plane is less than would be obtained going through a single forward rRBC, the round-trip attenuation is much greater, and the fact that it has no self-feedback instability means we can stack as many as

we need to attain the desired attenuation level. This is critical in dealing with problems in which waves travel at grazing incidence to the boundaries of the computational domain.

All numerical experiments to date have shown the rRAT-terminated FDTD domain to be stable (the largest stack tested so far has been 1024 layers). Just as with rRBCs, rRATs in orthogonal orientations can cross each other with no consequence; and the late-time instability due to numerical noise has been eliminated [19].

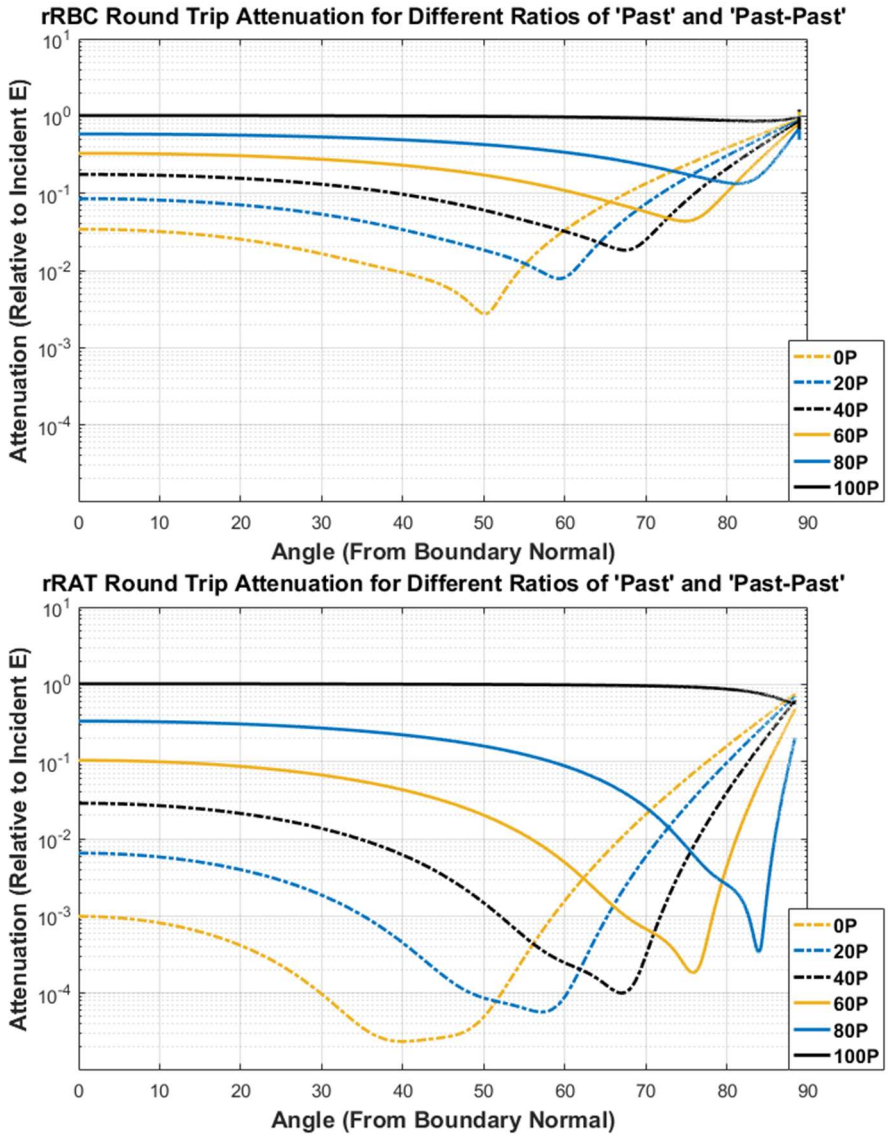

Figure 2.5: Double rRBC roundtrip attenuation vs. angle (top) and double rRAT roundtrip attenuation versus angle (bottom) for different proportions of Past and Past-Past stored fields. (e.g. 20P is 20% 'past' and 80% 'Past-Past')

Although the rRATs do not provide a one-directional attenuation as great as that of an equivalent number of rRBCs, in the vast majority of problems it doesn't matter what passes through the boundary, rather it matters what reflects back from the total boundary termination into the simulation space. Since the rRATs attenuate equally from either side, the total round-trip attenuation of rRATs compared to standard rRBCs is up to 100 times better when two layers of each are used (Figure 2.5).

In the standard termination of the FDTD domain with rRBCs, the last row (or plane in 3D) of the domain is terminated in a simple absorber such as a Huygens termination or a 377Ohm impedance surface. However, to highlight the difference between the rRBC and the new rRAT termination, we have removed this termination (therefore the end of the boundary behind the absorbing layers has a reflection coefficient of 100%) and compare in Figure 2.5 the round-trip attenuation versus angle of incidence presented by two layers of the original rRBC and two layers of the rRAT layer. As pointed out earlier, the proportion of the "past" information (one time-step behind) and the "past-past" (two time-steps behind) is used to tune the angular response.

The standard rRBC recipe is to use 100% of the "past-past", that is, 0% "past" or 0P, the yellow dashed curve with the strongest attenuation at normal incidence (0 degrees) and a deep null near 40-50 degrees off normal in both sets. By increasing the proportion of "past" used, the attenuation at normal incidence is reduced and that off normal increased.

Figure 2.5 shows that in addition to the rRATs being more attenuating overall, the off-normal null is much deeper and that there is significantly more attenuation at 90 degrees, true grazing incidence. This performance, combined with the ability to stack as many rRATs as desired (recall that the computational burden of a rRAT is one dimension

less than the space – a plane in 3D, a line in 2D) means that the rRAT termination can be used to eliminate artificial reflections from the sky boundary even after a ray has traversed 4km in a 36m tall domain, or an angle of incidence of 0.2 degrees from grazing.

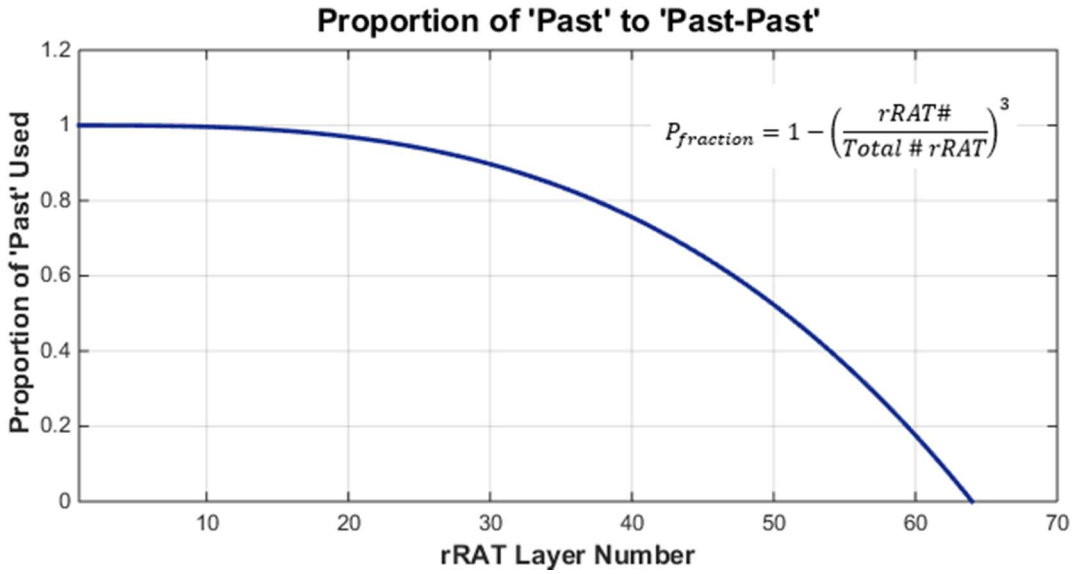

Figure 2.6: Proportion of past field used for each layer of rRATs for a very long sea simulation. Layer 0 is on the interior, 64 is on the exterior.

The winning strategy is to vary the angular performance of the stack so that some layers provide the high near-normal incidence attenuation and others take care of the shallow angle of incidence attenuation. Figure 2.6 shows the typical proportions of "past" field used for a stack of 64 rRATs when simulating a very long domain with a high proportion of waves with shallow angles of incidence upon the top boundary. The values near 1 at the outermost layers are optimized for near grazing incidence whereas the values near 0, closest to the interior of the domain, are optimized for near normal incidence.

Programming a stack of rRATs with a functional form as shown in Figure 2.6 is easily accomplished by a for-loop which iterates over each boundary when performing the calculations. For the GPU implementation this is more efficiently accomplished by adding an additional dimension to the work item division scheme for the boundary update kernel.

2.4 Adapting the Sub-Domain to the Problem of Interest

The typical FDTD sub-domain used to analyze a long sea surface with inhomogeneous atmosphere is illustrated in Figure 2.7. Since the result of interest is the frequency domain total field in the last sub-domain, there is no need for a near-field to far-field transformation. The total fields in the last sub-domain are Fourier transformed to select the center frequency of the pulse, to allow one-for-one comparison with the frequency domain methods (VTRPE, MoM, and BIE.)

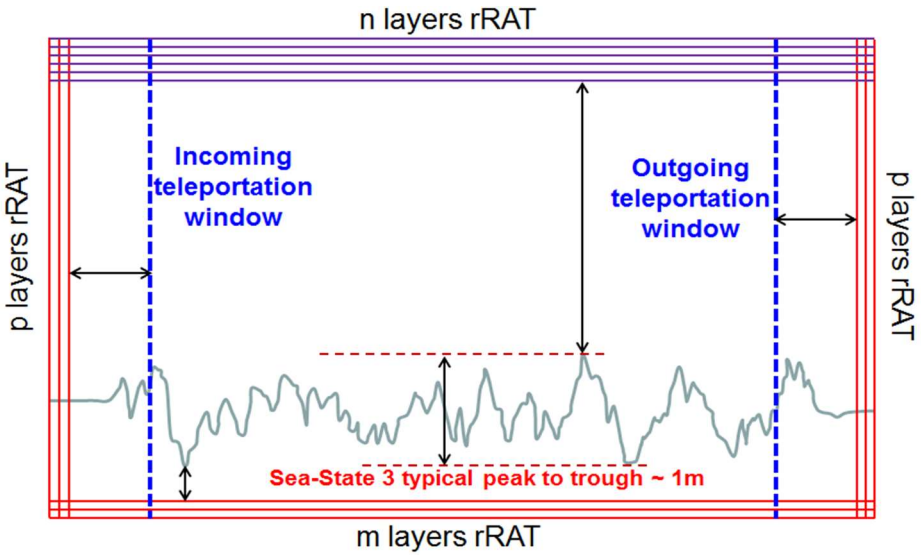

Figure 2.7: Typical domain layout containing a large amount of atmospheric space above the sea.

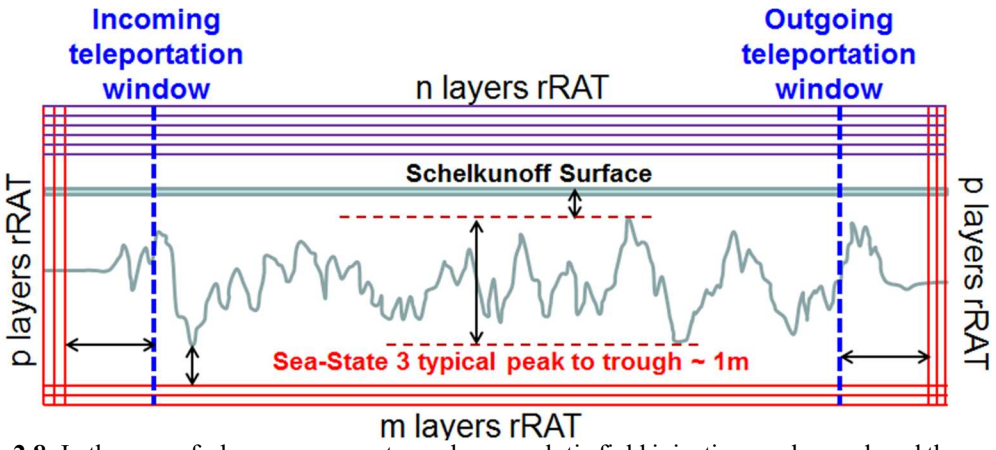

Figure 2.8: In the case of a homogeneous atmosphere, analytic field injection can be used, and the scattered field can be reproduced in post-processing by storing Schelkunoff currents on a bounding surface just above the sea.

When the atmosphere is inhomogeneous we are forced to use the total-field formulation of FDTD, with a physical source inside the domain; in the case where the atmosphere is everywhere homogeneous (free space), we can use a scattered-field formulation (Figure 2.8). The most economical of these formulations is analytic field injection, where the source field is impressed directly into the scattering media in the FDTD domain.

By invoking the Volume Equivalence Principle (VEP), the electromagnetic field incident on a material scatterer ($\mu_r, \varepsilon_r \neq 1$) induces volumetric electric and magnetic currents in the materials which then radiate into the FDTD grid [39].The derivation proceeds as follows: First write Maxwell's curl equations for the problem of interest (in this case a space with *dielectric* inhomogeneities, i.e. $\mu_r = \mu_0$) with the incident field ($\vec{E}_0$, $\vec{H}_0$) and the scattered field ($\vec{E}_s$, $\vec{H}_s$) explicitly separated.

$$\nabla \times \left[\vec{E}_0 + \vec{E}_s\right] = -\mu_0 \frac{\partial}{\partial t}\left[\vec{H}_0 + \vec{H}_s\right] \tag{2.13}$$

$$\nabla \times \left[\vec{H}_0 + \vec{H}_s\right] = \varepsilon_0 \frac{\partial}{\partial t}\left[\varepsilon'\left(\vec{E}_0 + \vec{E}_s\right)\right] + \sigma\left[\vec{E}_0 + \vec{E}_s\right] \tag{2.14}$$

Next assume the solution for the free space case in the presence of the source field of interest is known:

$$\nabla \times \vec{E}_0 = -\mu_0 \frac{\partial}{\partial t}\vec{H}_0 \tag{2.15}$$

$$\nabla \times \vec{H}_0 = \varepsilon_0 \frac{\partial}{\partial t}\vec{E}_0 \tag{2.16}$$

Subtracting equations 2.15 and 2.16 from equations 2.13 and 2.14, respectively, yields,

$$\nabla \times \vec{E}_s = \mu_0 \frac{\partial}{\partial t}\vec{H}_s \tag{2.17}$$

$$\nabla \times \vec{H}_s = \varepsilon_0 \frac{\partial}{\partial t}\left[(\varepsilon' - 1)\vec{E}_0 + \varepsilon' \vec{E}_s\right] + \sigma\left[\vec{E}_0 + \vec{E}_s\right] \quad 2.18$$

The first curl equation for the scattered field, Faraday's Law, remains unchanged, but rearranging the second it can be seen that the known source field, in the free space case, appears as an additional source term (eq. 2.19). (In FDTD the curl terms are the source terms for the time integration.)

$$\nabla \times \vec{H}_s - \left\{\varepsilon_0 \frac{\partial}{\partial t}\left[(\varepsilon' - 1)\vec{E}_0\right] + \sigma \vec{E}_0\right\} = \varepsilon_0 \frac{\partial}{\partial t}\left[\varepsilon' \vec{E}_s\right] + \sigma \vec{E}_s \quad 2.19$$

This means the incident field, which appears as the direct ray at the distant target, does not need to be carried by the FDTD grid, and can, in fact, be replaced by an analytic source expression (Section 3.4). It therefore never interacts at shallow angles with the upper truncation boundary of the sky and causes no artificial reflection. Only the locally scattered field is carried by the grid and teleported forward from subsection to subsection. Of course, for very long domains (kilometers long and meters high), the forward-travelling scattered fields themselves will also strike the sky boundary at near grazing incidence. And thus, the rRAT termination is still required.

The major advantage of this scattered-field formulation is that all the "dead space" above the sea surface is superfluous. The upper truncation boundary can be brought below the height of the source, as close as desired to the sea surface (Figure 2.8). The only requirement is that this boundary clear the top peaks of the highest wave in the full sea being modeled.

It follows that if we store the time-history, or frequency-domain representation, of the propagating fields (in the form of Schelkunoff equivalent currents) at every such surface in every sub-domain, then when the entire problem is solved this total bounding

surface constitutes a single surface closed at infinity ($x = \pm\infty$) for the purpose of the equivalence theorem. Such a surface, with its surface normal equal to $\hat{z}$, can only radiate upwards [40]. We do not need to use the equivalent currents on this surface to couple from one sub-domain to the next as was done in [12].

Although it is true that if we only considered a given sub-domain at a time, and closed the Schelkunoff surface on all sides of that sub-domain, we would expect the currents on all the walls to radiate outside the resulting box, and therefore into the adjacent sub-domains, a moment's consideration is enough to realize the equivalent currents existing on the rightmost wall of the n^{th} domain are exact negative copies of the currents existing on the left wall of the $n+1^{st}$ domain and therefore the latter will exactly cancel the former.

In our formulation, the coupling from the n^{th} to the $n+1^{st}$ domain occurs exclusively through the remaining, albeit small, vertical teleportation walls at the two ends of the sub-domain. Figure 2.8 shows this arrangement, as used for the case where the atmosphere is homogeneous. The scattered-field formulation can be readily verified as correct, and missing no information, by solving the same 2D problem in the total-field and in the scattered-field formulation and comparing the results to each other. This is examined in Chapter 3. Solving the 4km sea space as shown in Figure 2.8 leads to scatter angles as low as 0.03 degrees incident upon the rRATs.

2.5 Conclusion

To solve FDTD problems of extreme size, spatial decomposition becomes an absolute necessity for the management of computational burden. For the case of radar

propagation in a maritime environment, we break down the sea space into many sub-domains along the direction of travel, record the time-history of the fields in the form of discrete surface currents as they exit the outgoing boundary of each sub-domain, then impress the currents on the incoming boundary of the current sub-domain. This process is repeated until all sub-domains in the total domain have been simulated. Connecting the sub-domains by means of the teleportation algorithm allows for a numerically seamless interconnect with no overlap. The sub-domains need to be large enough to include sufficient back-and-forth interactions between the scatter from large sea features such as to accurately preserve the forward-travelling scatter. The tests to determine the required sub-domain length are discussed in Chapter 5.

The other equally critical requirement for problems of size are high-performance truncation boundaries which function down to extremely low grazing angles of incidence. This is accomplished via the novel re-Radiating Absorbing Termination (rRAT), which is a computationally efficient, innumerably stackable, bi-directionally attenuating, and angle-tunable boundary condition. The rRAT has been used to terminate spaces containing rays as low as 0.03 degrees off grazing.

In an effort to further reduce computational burden, we introduce the Volume Equivalence Principle formulation of FDTD. This permits the elimination of all the "dead space" of a homogeneous atmosphere by allowing the source fields to be directly injected into the FDTD scattering material. The source fields can be any fields: FDTD-generated or analytic. We later introduce the use of a dispersion-matched analytic source in Chapter 3. This approach is critical to allow for sufficiently large 3D sea spaces.

CHAPTER 3

CROSS-VALIDATION

3.1 Cross-Validation against VTRPE and MoM

Neither VTRPE nor MoM suffer from grid dispersion. But in an FDTD domain there will always be a *combination* effect of dispersion and anisotropy, *intrinsic* to propagation within the finite space and time discretization of FDTD. As it is intrinsic to FDTD, it is consistent, source independent, and can be easily determined by comparing FDTD propagation in an empty space against the theoretical propagation in the continuum of free space. For our typical grid discretization of $\lambda/17.8$ (dx = dy = dz) and Courant number of 0.707, the apparent speed of light in the dominant x-direction is 98.13% of c_0. This factor was used for all the comparisons with VTRPE. (Doubling the spatial resolution would improve the factor to 99.24%.)

Therefore, for propagation over a long sea surface the approach, to first order, is to "pre-compress" the sea surface in the propagation direction by this 98.13% factor, run the FDTD program, and at the end expand the space back to its original dimensions. With this approach the FDTD and VTRPE results are extremely close, specifically close to the sea surface. However, the results deviate somewhat the higher we look above the sea surface. This is precisely what would be expected due to the anisotropy in FDTD, because the rough sea surface continually creates high angle radiation by scatter. (And we know in FDTD the propagation velocity is much less dispersive close to the diagonal.) However, we must not be too quick to attribute all differences in results to dispersion. As will be noted in Section 3.2, the representation of the sea itself is critical. Nevertheless, the agreement is very good.

The standard result of interest in these problems is the Pattern Propagation Factor (PPF), a contour plot of the total field strength in the space normalized to the incident field, in dB, as given in equation 3.1. Figure 3.1 shows the results for vertical polarization in

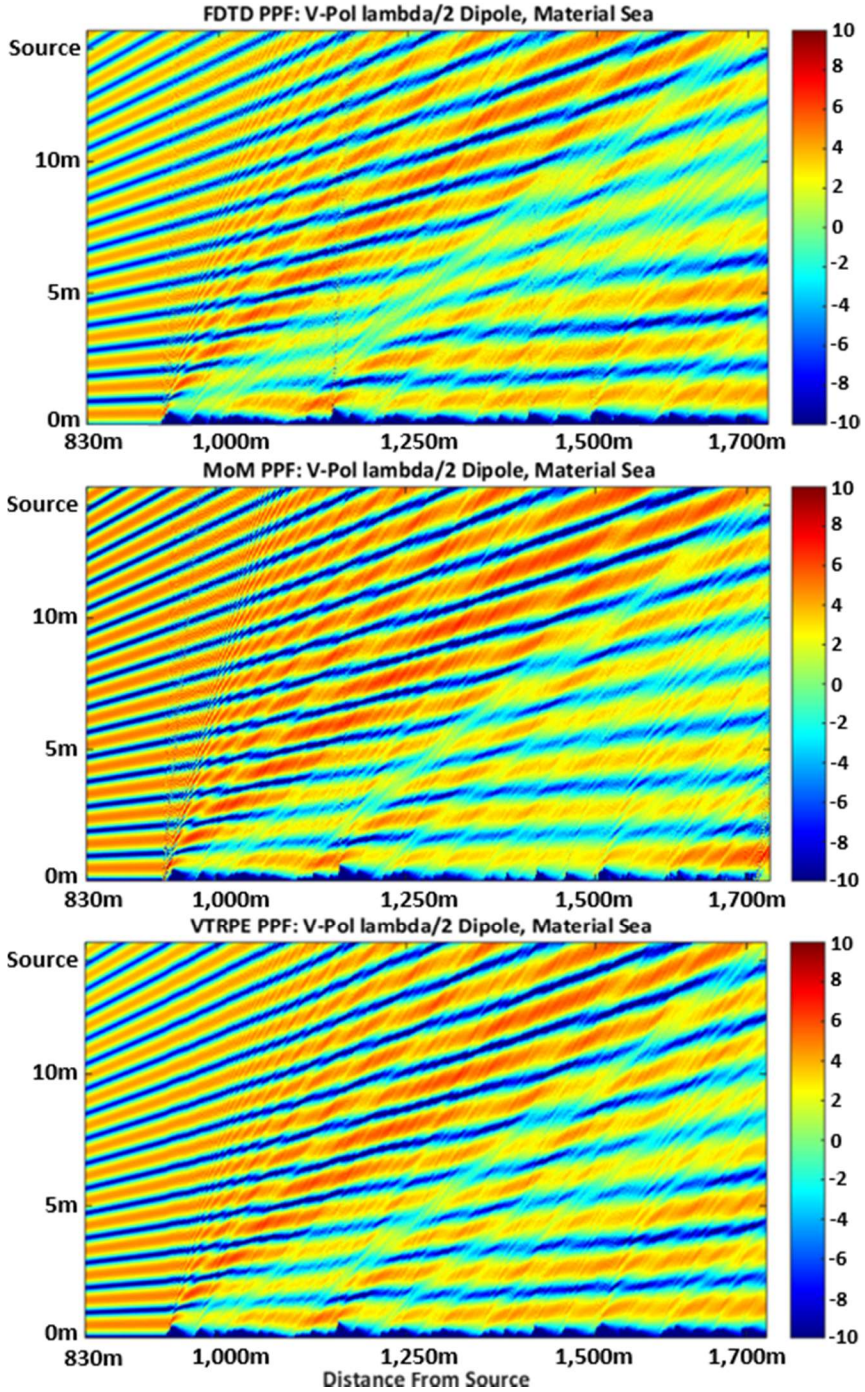

Figure 3.1: Contour plots of PPF (normalized Total Field) in a common simulation space for FDTD (top), MoM (middle), and VTRPE (bottom).

the 2D space of Figure 1.1 for FDTD, Forward-Backward MoM [41] and VTRPE [42]. The source is a half-wave dipole at 9.2GHz, suspended 14.3m above mean sea level.

$$PPF \equiv 20\log_{10}\left(\frac{Total\ Field}{Free\ Space\ Source\ Field}\right) \qquad 3.1$$

A more direct quantitative comparison is obtained by plotting the PPF along a slice of the space. Figure 3.2 shows such a cut at the plane 1,650m from the source for both vertical (TM, V-Pol, H field out of the plane, and TE, H-Pol, E field out of the plane).

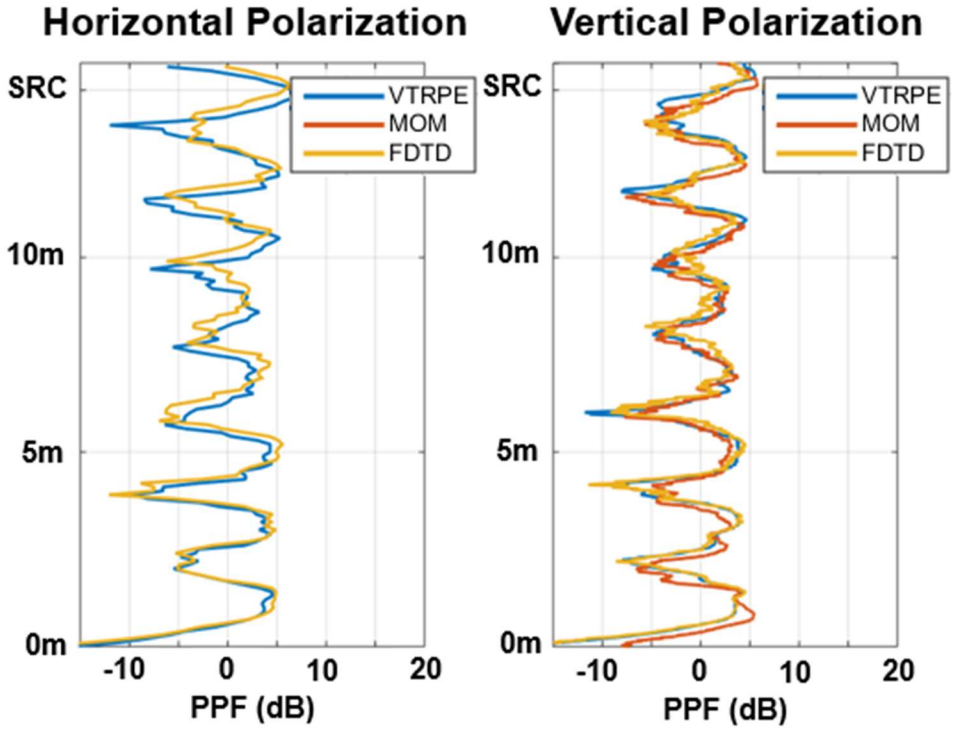

Figure 3.2: Vertical cut of PPF 1,650m down range comparing FDTD, MoM, and VTRPE with H-Pol (left) and V-Pol (right) sources.

MoM uses an impedance boundary condition to model the sea surface while VTRPE uses the local Fresnel coefficients. For the TM problem the sea is treated as a lossy dielectric of $\varepsilon = 59.2 + j35.94$ (which includes both the dielectric loss and the saline conductivity). To enable easy comparison with MoM, the sea surface for TE was treated as a PEC boundary. Figure 3.2 shows excellent agreement between FDTD and VTRPE for

the TM case while MoM deviates somewhat. In the TE case, PEC surface, MoM and VTRPE agree exactly while FDTD matches the results close to the surface and deviates slightly with increasing height. The deviation with height in this problem is a combination of the anisotropy of FDTD and potential discretization differences in the sea. However, the issue of anisotropy is of minimal importance for the ultimate application because as the propagation distance increases, this error gets smaller and smaller.

3.2 Cross-Validation in an Inhomogeneous Atmosphere

Since the inhomogeneities of the atmosphere are due to variations in the permittivity of air of the order of 0.0005, it takes propagation through tens of kilometers to see significant refraction of the propagating wave (Figure 3.15). Such domains are too large for the BIE method. Therefore, in [8] the authors artificially enhanced the atmospheric variations by a factor of 10 to get the phenomena to manifest within a 3.6km distance. They then compared their Forward/Backward (FB) BIE results to a Parabolic Wave Equation (PWE) approach using the Split-Step Fourier (SSF) method.

To enable solution with the BIE method the authors used an index of refraction ($n^2=\varepsilon_r$) profile that admits an accurate Green function [43]:

$$n^2(z) = n_1^2 + \varepsilon(h - z) \qquad 3.2$$

Figure 3.3 shows the authors' problem of concern. For their comparison with a parabolic equation method they assume a duct height, h = 50m, a "duct parameter", ε=0.001, and a flat sea of dielectric constant 70.4 - j40.6. We model the same space in FDTD, simulated up to an altitude of 60m, and with a very fine spatial discretization of

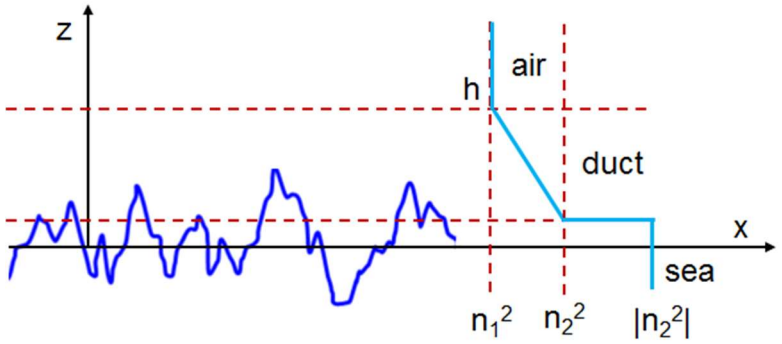

Figure 3.3: The inhomogeneous atmosphere problem solved by the BIE method in [8].

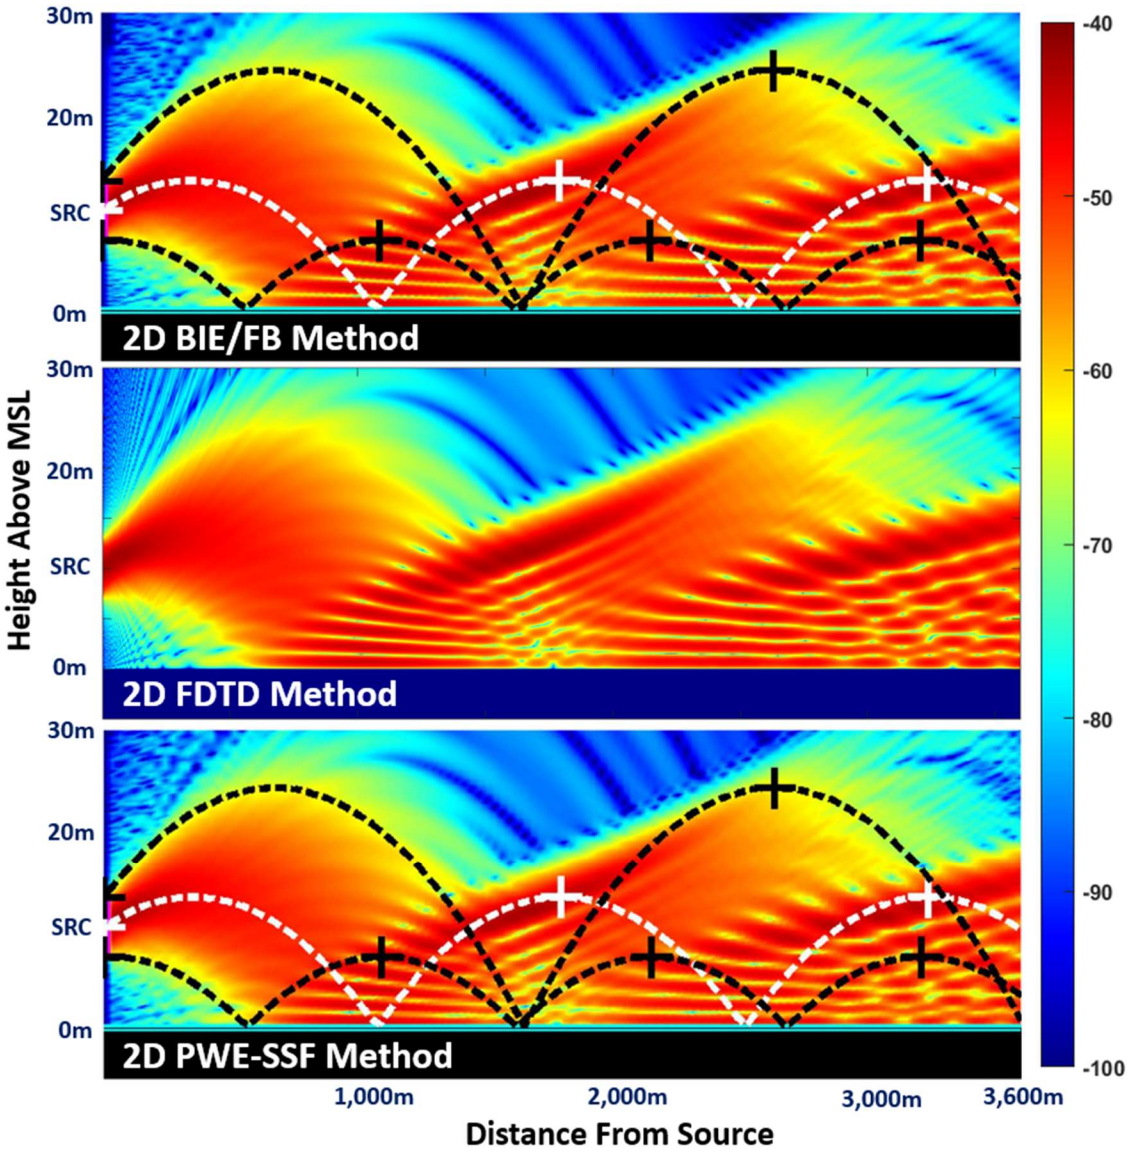

Figure 3.4: Total field contour plots of a high directivity 3GHz source in the presence of an exaggerated atmospheric duct. Top: BIE from [8], Middle: FDTD, Bottom: PWE-SSF from [8].

53.4 points per wavelength to allow for minimal dispersion and anisotropy. The authors' source is a 6m-wide tapered vertical planar aperture, suspended at its center 10m above the surface and radiating coherently at 89 degrees incidence towards the 3.6km-long surface at 3GHz.

Given the parameters the authors used for duct height of 50m and $\varepsilon = 0.0001\text{m}^{-1}$, their peak susceptibility ($\varepsilon - 1$) is 0.005, which is about ten times higher than in a typical duct found in nature. The resulting large rate of change of permittivity with height readily bends upward traveling rays down toward the plane below to be reflected and create a "bouncing" trajectory, over and over again.

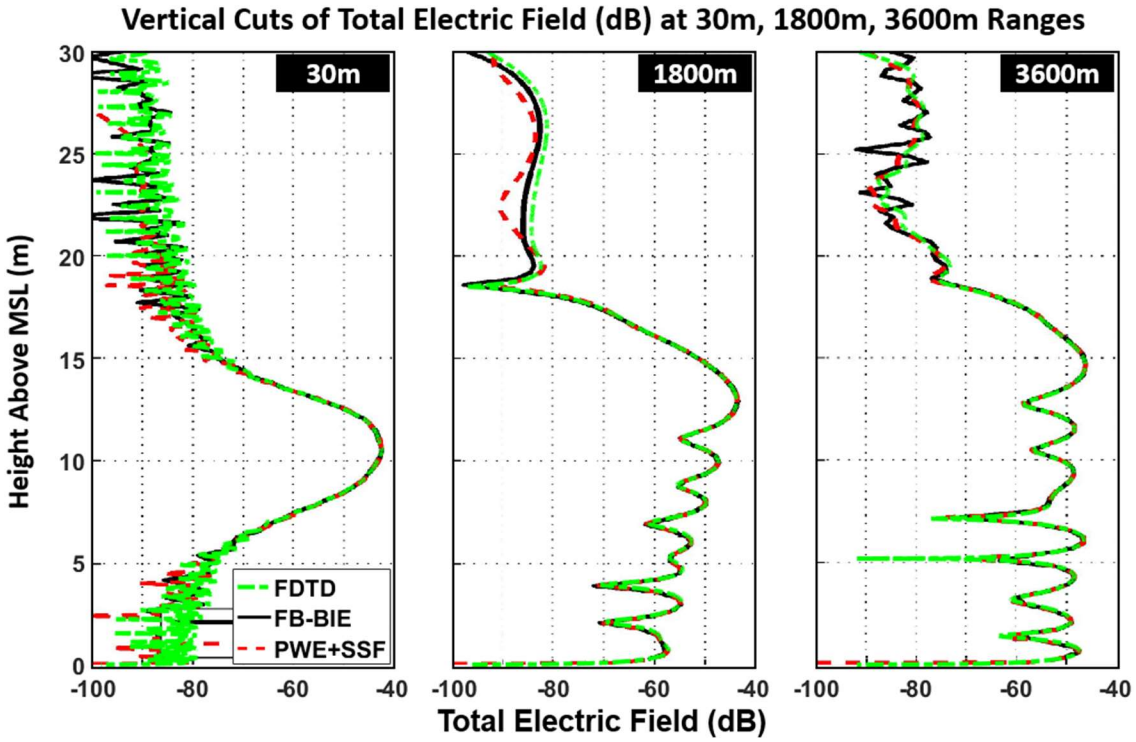

Figure 3.5: Vertical cuts of total E-Field from FDTD (green), BIE (black), and PE (red) methods taken at 30m, 1800m, and 3600m ranges.

Figure 3.4 shows the result of the forward-backward BIE method from [8] on top, that from our FDTD simulation in the middle, and the PWE-SSF method from [8] on the bottom. Figure 3.5 shows vertical cuts of the total E-field, taken at several ranges. The

agreement is excellent, with some very slight variations occurring between all three cases in the areas of low field strength. The slight differences observed between FDTD and the BIE method may be slightly in part due to FDTD's inherent dispersion and anisotropy, but with such fine spatial discretization this is unlikely. The more likely culprit is FDTD's discretized representation of the source aperture.

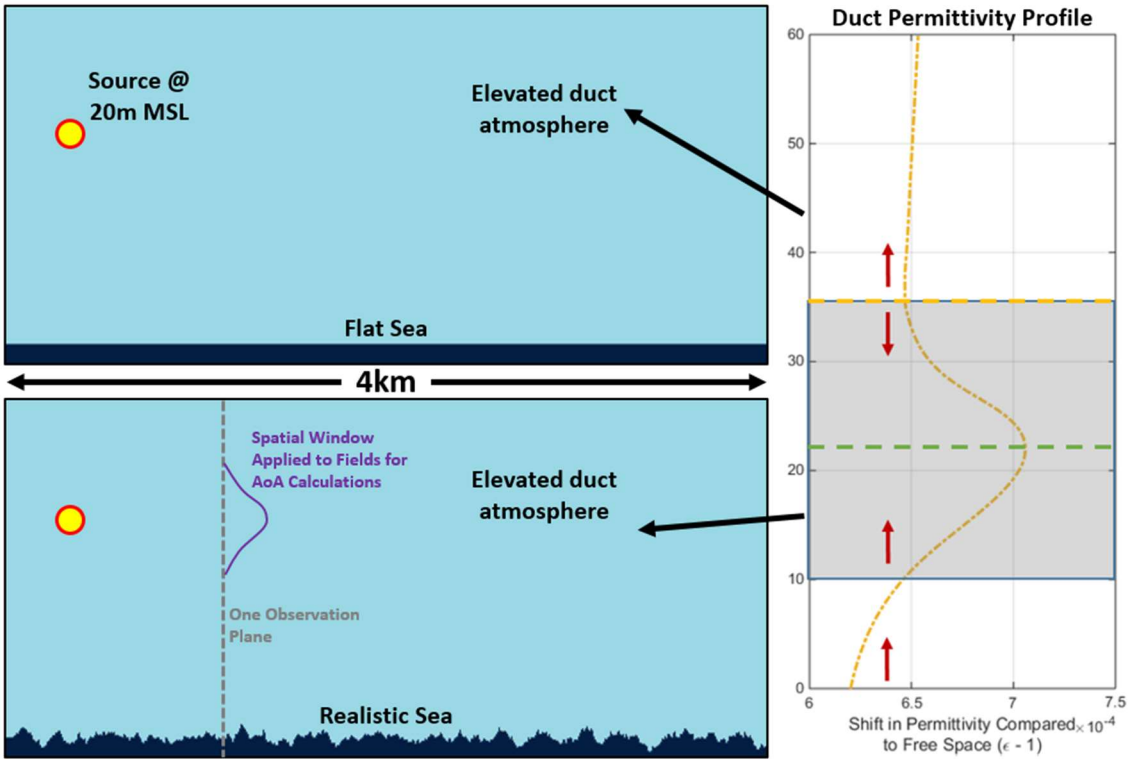

Figure 3.6: Diagram of two test spaces – a flat sea (top), realistic sea (bottom) – containing an elevated duct profile (right) for comparing FDTD and VTRPE. Vertical field cuts are taken on several observation planes (gray), AoA is calculated from the spatially windowed fields on the observation planes (purple).

A more realistic case of an inhomogeneous atmosphere, an elevated duct, in a 4km space is now considered. The profile is given in Figure 3.6 (right). The same space was solved using FDTD and VTRPE [42] for a horizontally polarized λ/2 tall (half-sine) planar source located 20m above mean sea level, at 9GHz with a sea permittivity of 56.05 - j38.95. The quantities of interest are the pattern propagation factor and Angle of Arrival (AoA). The AoA is a very practical quantity, it is comprised of a Plane Wave Spectrum (PWS)

decomposition of the fields on a vertical observation plane (gray line in Figure 3.6). Thus, the AoA essentially indicates how much energy is travelling in a given direction across an observation plane at a given range, and it does so by breaking the fields down into plane waves with a given amplitude and direction. The AoA is calculated by application of a Fourier transform into k-space. The propagation angles (eq. 3.3) are derived from the wave vector component relations.

$$\theta(k_z) = \operatorname{acos}\left(\frac{k_z}{k_0}\right) \tag{3.3}$$

Furthermore, the AoA can be spatially windowed by applying a spatial taper to the fields used in the calculations. In the test cases considered, we apply a Gaussian window centered on the source height (20m above mean sea level) with a 5m-wide Full-Width Half-Maximum (FWHM). Spatial windowing serves a highly practical purpose for radar scattering. Suppose a radar target, such as a ship, is sitting just beyond, but very close to, the observation plane, and we are interested in the scatter from a specific feature of the ship. A spatial window can be applied at the same height and of the same size as the scattering feature and, in doing so, we can limit the fields to only those which will interact with the scatterer. The fields which are not at the correct height to interact with the scatterer will be removed by the window.

The AoA results are presented first for the case of the flat conducting plane of sea material (Figure 3.7 and Figure 3.8). Figure 3.7 contains the AoA calculated at all observation planes within the flat sea simulation space. The direct ray from the source is clearly visible around 0 degrees off horizontal. The specular ray reflected from the flat, conducting plane is clearly defined at higher angles early down range, but blends together

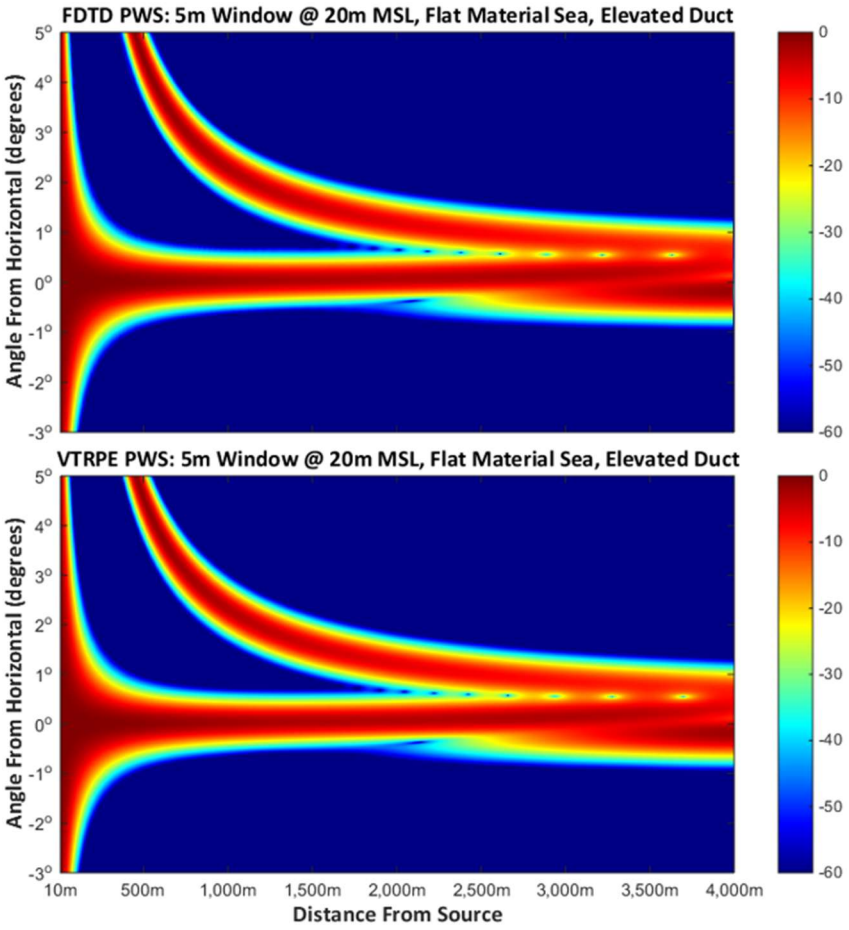

Figure 3.7: AoA of FDTD (top) and VTRPE (bottom) calculated at all observation planes for the flat sea.

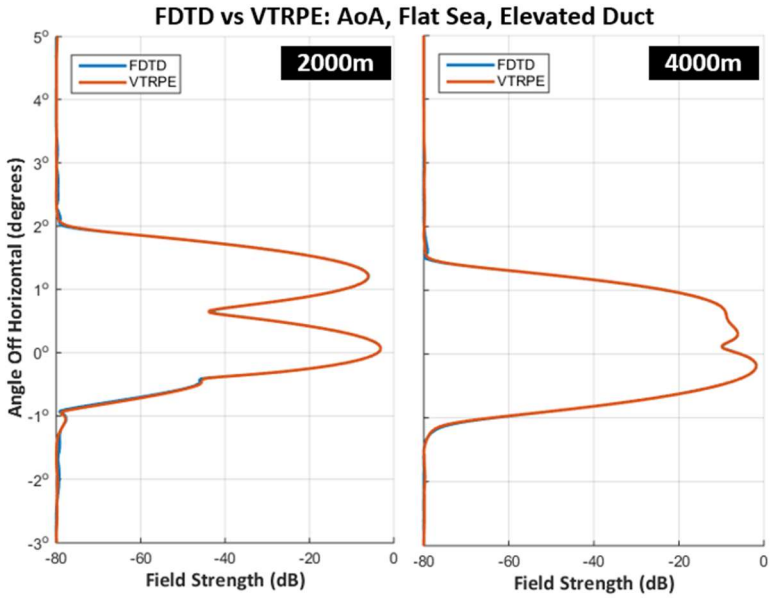

Figure 3.8: AoA of FDTD (blue) and VTRPE (orange) calculated from fields along observation planes at 2km (left) and 4km (right) down range for the case of the flat sea.

with the direct ray further down range where the reflected ray starts propagating at progressively shallower angles. The influence of the duct can on both direct and specular rays can start to be seen around 2km down range, where the results appear to distort, and information starts to be refracted toward angles below the horizon. Figure 3.8 shows vertical cuts of the data in Figure 3.7 taken at 2km (left) and 4km (right). The cuts from FDTD and VTRPE are overlaid on top of one another, and they exactly match.

Next, we perform the same AoA comparisons for the case of the realistic sea (Figure 3.9 and Figure 3.10). It's immediately apparent from Figure 3.9 that the realistic sea scatters an appreciable amount toward higher angles, and the scatter is much more

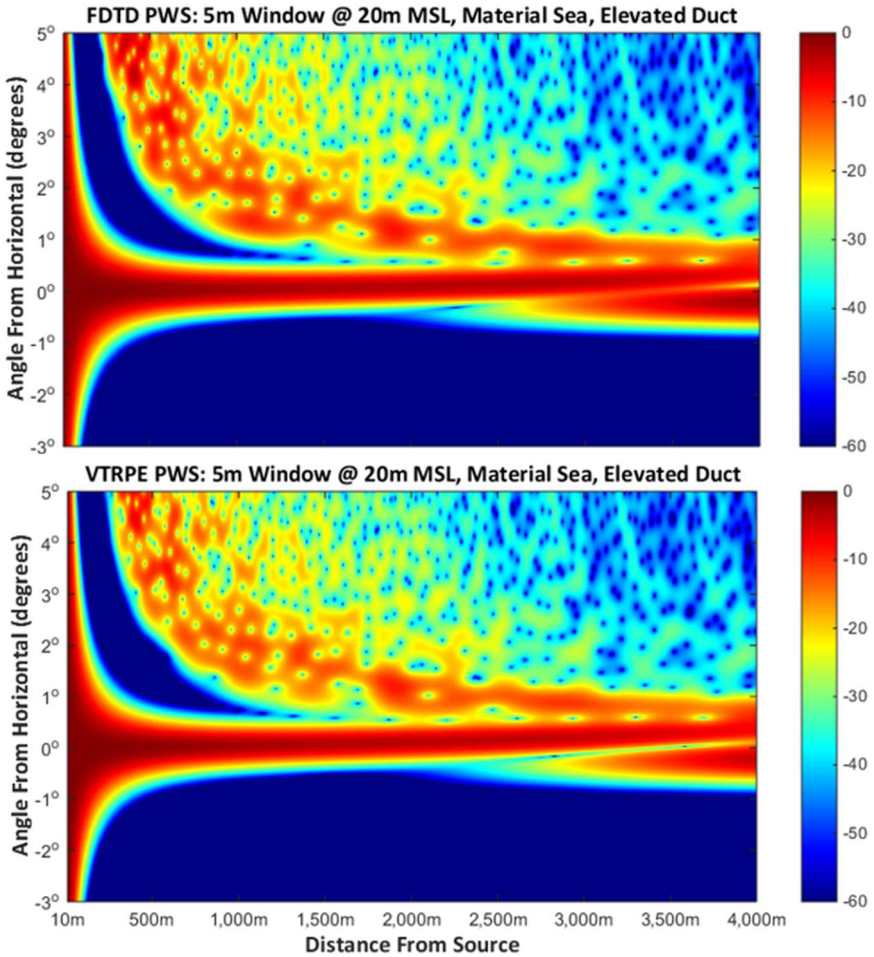

Figure 3.9: AoA of FDTD (top) and VTRPE (bottom) calculated at all observation planes for the realistic sea-state 3 sea.

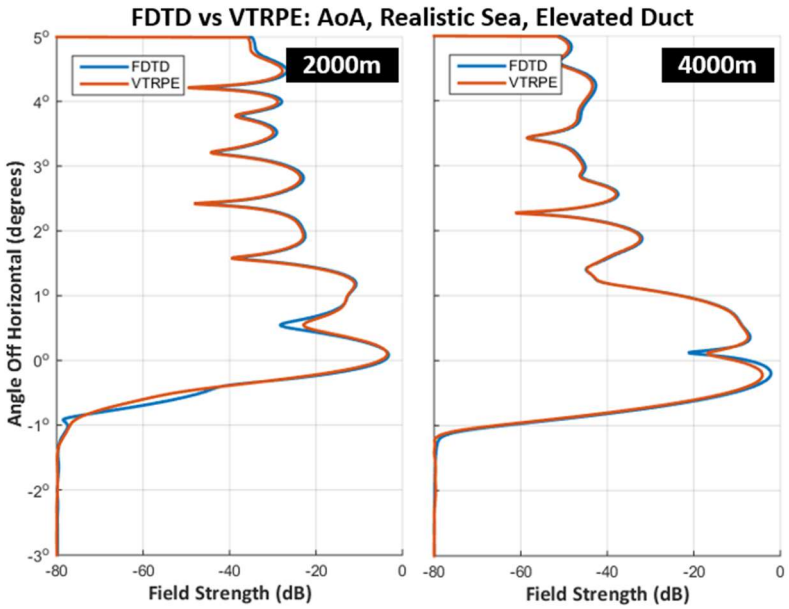

Figure 3.10: AoA of FDTD (blue) and VTRPE (orange) calculated from fields along observation planes at 2km (left) and 4km (right) down range for the case of the realistic sea-state 3 sea.

random, which is expected given the stochastic nature of the sea surface. Examination of the same vertical AoA cuts at 2km and 4km for the case of the realistic sea demonstrates excellent agreement between FDTD and VTRPE, however it's not exact agreement like the case of the flat sea. After the presentation of the PPF results, a hypothesis will be drawn.

Next, the PPF is considered. The full-domain colormap of the PPF for the case of the flat sea is shown in Figure 3.11, while that for the realistic sea is shown in Figure 3.12. Figure 3.13 and Figure 3.14 contain vertical PPF cuts at ranges of 3km and 4km, respectively. Note, the post-simulation lengthwise "decompression" (as discussed in Section 3.1) has not been performed for Figure 3.11 and Figure 3.12, however it has been performed for Figure 3.13 and Figure 3.14.

The perfect PPF agreement in the case of the flat sea for Figure 3.13 and Figure 3.14, in addition to the perfect AoA agreement, indicates FDTD and VTRPE are indeed capturing the same physics. The PPF agreement in the case of the realistic sea is less than

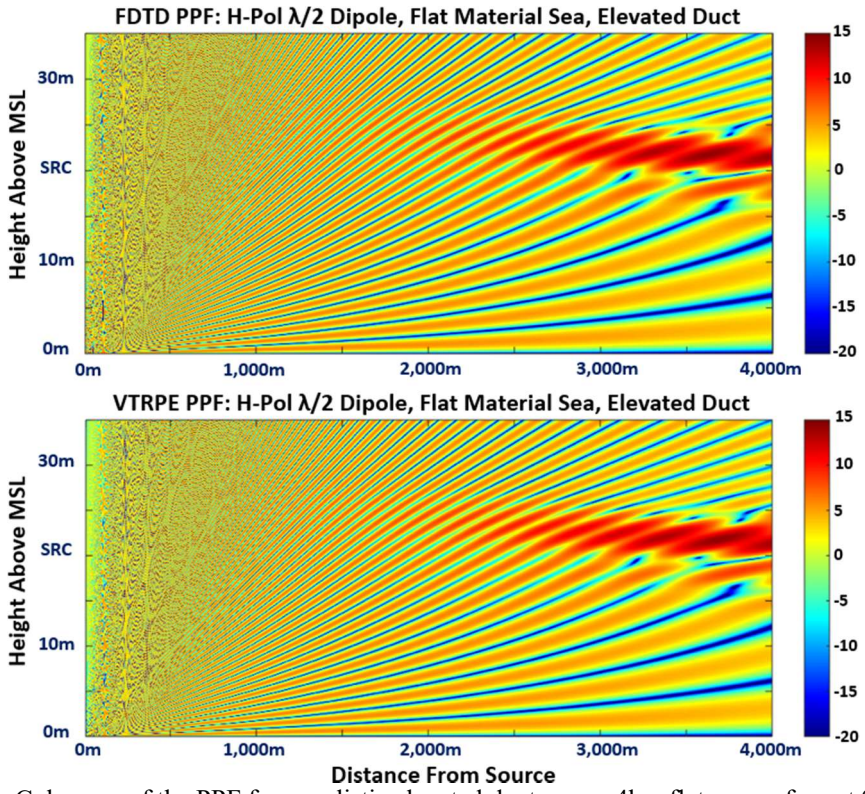

Figure 3.11: Colormap of the PPF for a realistic elevated duct over a 4km flat sea surface at 9GHz.

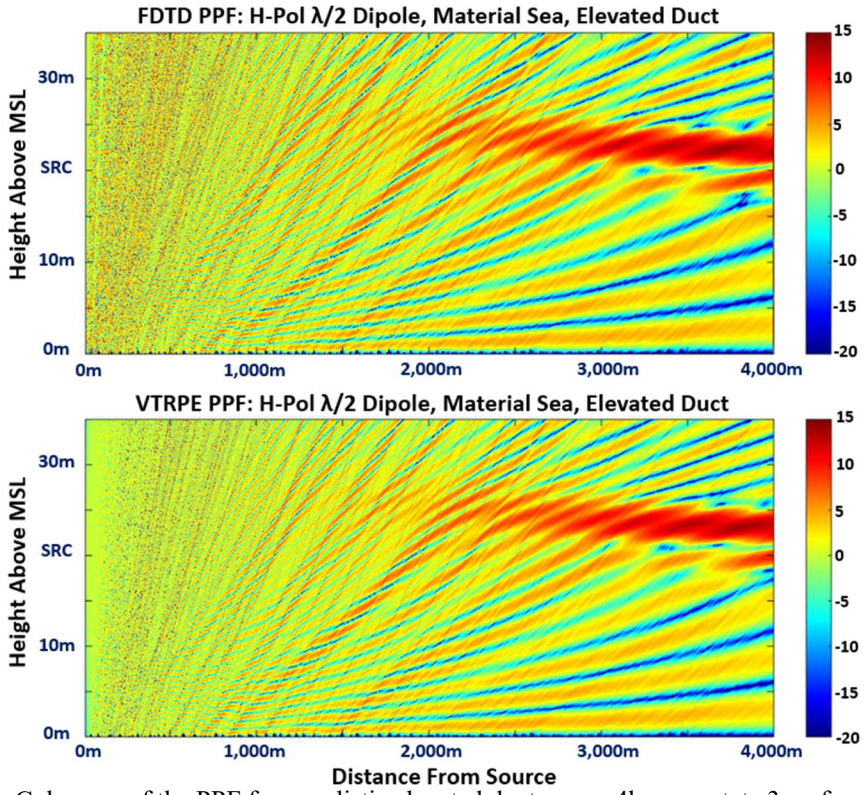

Figure 3.12: Colormap of the PPF for a realistic elevated duct over a 4km sea-state 3 surface at 9GHz.

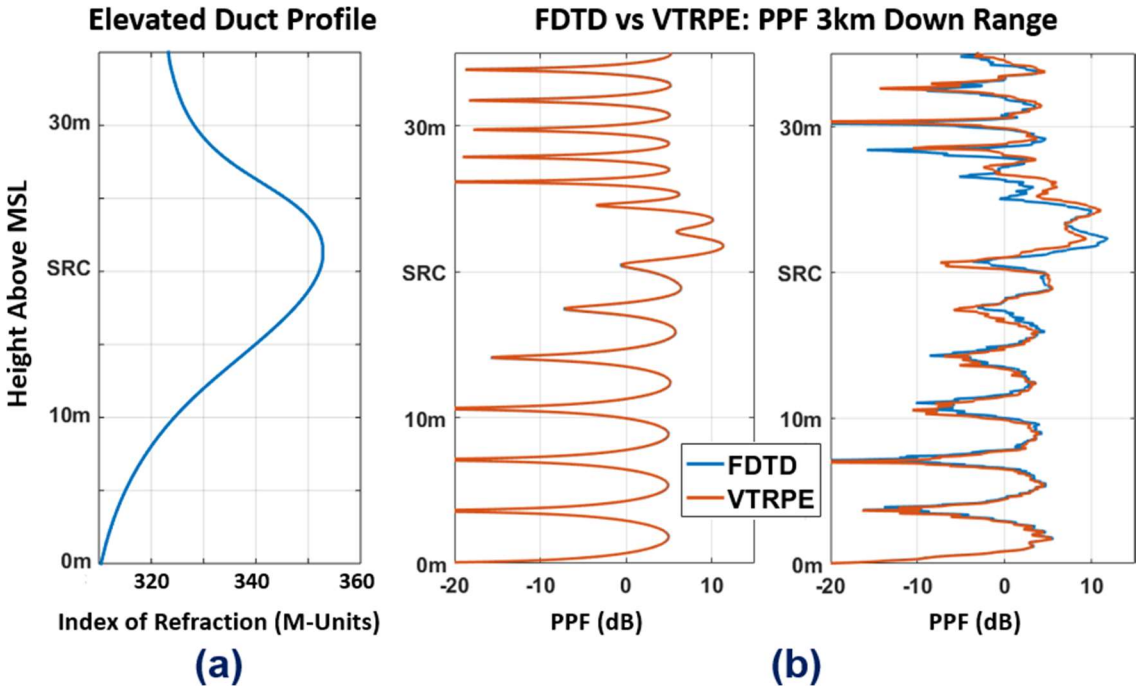

Figure 3.13: Duct profile (a), and PPF of FDTD (blue) and VTRPE (orange) along an observation plane at 3km down range for the case of the flat sea (b – left) and realistic sea (b – right).

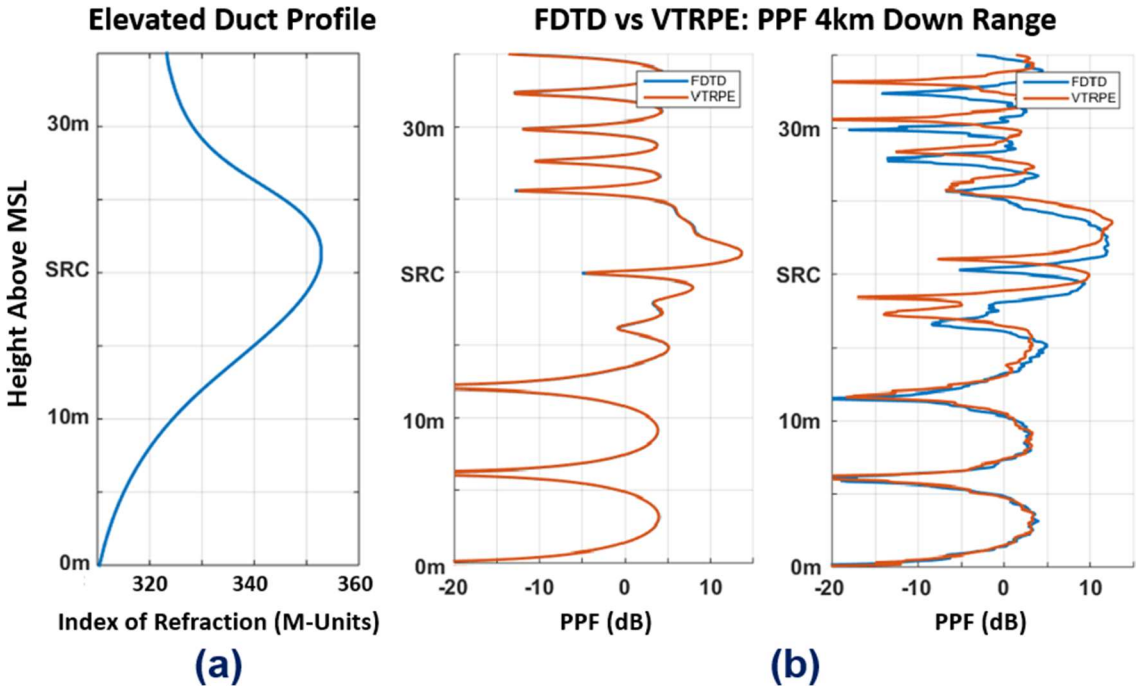

Figure 3.14: Duct profile (a), and PPF of FDTD (blue) and VTRPE (orange) along an observation plane at 4km down range for the case of the flat sea (b – left) and realistic sea (b – right).

exact, but the plots show essentially the same positions of most nulls and peaks. The fact that the PPF has better agreement at 3km than at 4km suggests the dispersion and anisotropy of the FDTD grid may not be the primary contributor of error – if it was, agreement would be better further down range where the angular spread of the dominating source and scattered fields would be narrow, and the grid would appear to be more isotropic. This, in combination with the perfect PPF agreement for the flat plane case, suggests that the disagreement at higher angles with the realistic sea may be due in part to the different resolution of the surface model in both methods. The VTRPE implementation has a variable range step from 5.301mm up to 5cm and a vertical resolution of 0.1359mm. The FDTD simulation used a square cell of size dx = dz = 1.875mm. The spatial discretization difference between FDTD and VTRPE is the most likely explanation. In future work, a potential remedy might be to implement an effective media algorithm to minimize FDTD "staircasing" effects on the sea surface.

3.3 Demonstration of the Scalability of the Applied Techniques

As a demonstration of the scalability of the teleportation-based domain decomposition and the performance of the new rRAT boundaries, the case of an enormous 2D domain 15km (450,000λ) in length and 41m (1,240λ) tall is solved. The domain contains a realistic sea-state 3 sea surface, a surface duct profile, and a 9GHz horizontally-polarized electric (E_z) point source placed 20m above mean sea level. The surface duct is constantly pulling fields toward the sea, trapping them near the sea surface, which creates the appearance of the fields bouncing along the sea. This space contains 178.3 billion FDTD cells (558 million square wavelengths) and requires 11.4 million time steps to solve,

appearing to be the largest 2D FDTD simulation in published literature to date. The propagation distance of 15km was chosen arbitrarily, and there are no reasons the space could not have been made larger.

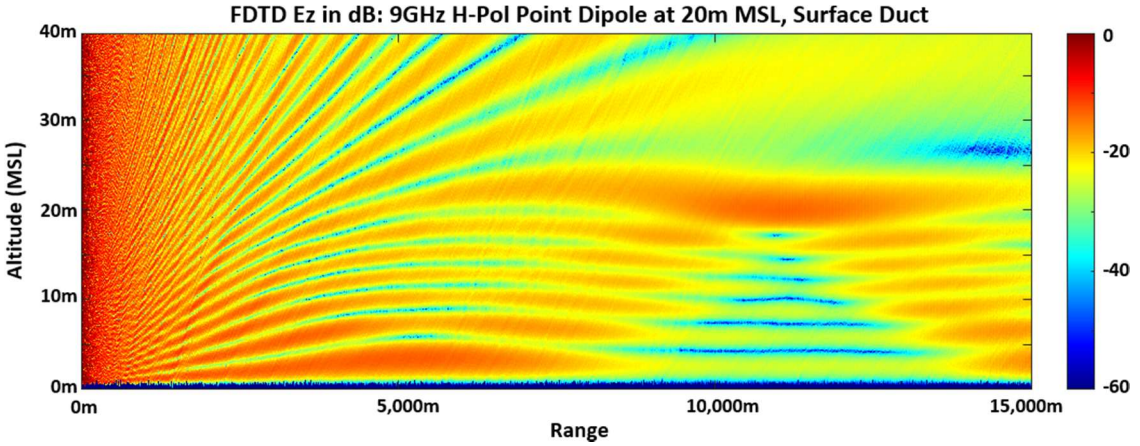

Figure 3.15: An extremely large 15km sea-atmosphere space containing a realistic surface duct profile which tends to trap the fields near the sea surface. (E_z shown here, in dB)

The only concerns introduced with a longer domain are increased simulation run time (obviously), but also spreading of the FDTD source pulse due to the grid's inherent dispersion in conjunction with the time-enveloped pulse's finite bandwidth. The total spreading of the source pulse can be calculated a priori, and the sub-domain to sub-domain time-history windowing (Section 2.2) can be adjusted accordingly to accommodate the final width of the source pulse in the final domains of the simulation space. First, the Fourier transform of the time domain source pulse function is performed to obtain the frequency content. Next, we set a signal strength threshold (usually within -80dB of the center frequency) and determine the frequencies at the low and high frequency edges of the threshold band. The FDTD space is set to run for a given number of time steps, thus the amount the pulse spreads by the end of the simulation is simply the difference between

the final propagation distances of the frequency components at the edge of the threshold band, which are determined by the group velocities (eq. 2.3) of each frequency.

3.4 Cross-Validation of the Total-Field and Scattered-Field VEP Formulation

To ensure a scattered-field FDTD problem solved with an impressed source (Section 2.4) agrees with the same FDTD problem in which the source is inside the domain, the fields of the impressed source must possess the same anisotropy and dispersion as the FDTD grid. Much work has been done in this respect in the literature, ranging from field injection using auxiliary FDTD grids to the creation of discrete Green's functions that exactly mimic FDTD [44, 45, 46, 47, 48].

In this section we briefly show the results of a very simple implementation. Since our source of interest is a linear λ/2 dipole located many wavelengths away from the scatterer, the propagating fields of interest are all outside the radian sphere and known; they are all "far fields". In that case the normalized fields, F_y (E_y for H-pol, H_y for V-pol, 2D space), transverse to the plane of incidence a distance $r(x,z)$ from the source, at time step N, are given as:

$$F_y = sin\left(\tilde{k}(\theta)\, r(x,z) - \omega_0 N\Delta t\right) \cdot \frac{Env}{\sqrt{r(x,z)}} \tag{3.4}$$

$$Env = exp\left(-\left[\left(r_0(t) - r(x,z) - \frac{1}{2}W\right) \cdot \frac{2}{gW}\right]^2\right) \tag{3.5}$$

$$g = \frac{1}{\sqrt{-\ln(Env_{min})}} \tag{3.6}$$

In the case of the sea, the source is placed in free space outside of the scatterer (the sea), thus the singularity is avoided in equation 3.4 as the source's field is never calculated

at r = 0. In principle, the analytic source could be placed within a scattering medium, but an exact Green's function for such a source embedded inside the given scatterer, which includes the source term, would be required.

To inject these into the FDTD space we create a traveling wave packet given by equations 3.4-3.6. Here g is a parameter to ensure the envelope reaches a small enough value, Env_{min}, at the outside edges to create a smooth taper. The parameter W is the full spatial width of envelope. The propagation constant is made angle-dependent at the frequency of interest to correspond with the anisotropy of FDTD. From the standard FDTD dispersion relation, the minimum and maximum phase velocities in the grid can be found.

$$\frac{1}{S^2} sin^2\left(\frac{ds\,\pi S}{\lambda_0}\right) = sin^2\left(\frac{ds}{2}\tilde{k}\,cos(\theta)\right) + sin^2\left(\frac{ds}{2}\tilde{k}\,sin(\theta)\right) \quad 3.7$$

Equation 3.7 is solved for the numerical wavenumber, $\tilde{k}$, at the angles of minimum and maximum phase velocity, 0- and 45-degrees, respectively (for square cells).

$$\tilde{v}_{pmin2D} = \pi c\left[\frac{\lambda_0}{\Delta} sin^{-1}\left(\frac{1}{S} sin\left(\frac{\Delta\pi S}{\lambda_0}\right)\right)\right]^{-1} \quad 3.8$$

$$\tilde{v}_{pmax2D} = \pi c\left[\frac{\lambda_0}{\Delta}\sqrt{2}\, sin^{-1}\left(\frac{1}{S\sqrt{2}} sin\left(\frac{\Delta\pi S}{\lambda_0}\right)\right)\right]^{-1} \quad 3.9$$

$$\Delta v_{p2D} \equiv \tilde{v}_{pmax2D} - \tilde{v}_{pmin2D} \quad 3.10$$

The angular-dependent phase velocity at the frequency of interest is then:

$$v_p(\theta) = \tilde{v}_{pmin2D} + [2\,cos(\theta)\,sin(\theta)]^2 \cdot \Delta v_{p2D} \quad 3.11$$

Thus, the numerical wavenumber used in equation 3.4 is simply:

$$\tilde{k}(\theta) = \frac{\omega}{v_{p2D}(\theta)} \quad 3.12$$

These fields are injected according to the VEP recipe into the sea dielectric. Because the skin depth in the sea is less than one cell, about 1.25mm at 9GHz, there is no need in our case to inject fields below about 5 or 6 cells under the sea surface. To avoid scattering from discontinuity of the truncation of the sea surface at edges of each domain, the injected field is tapered to zero after the teleportation window, over a distance of 5 wavelengths. The injection in the neighborhood of the teleportation window must be done carefully to avoid "double dipping", since the injected field appears again in the subsequent sub-domain. The key step is to begin the taper one cell past the time history collection boundary. Anything scattered backward from the taper will in fact appear in the subsequent domain when the time history is injected, but it will be travelling strictly backward toward the absorbing boundary after injection and will not scatter forward. Sufficient tapering ensures this artificially scattered field is very weak compared to the relevant fields and does not interfere with the simulation results.

Following this procedure, the same 4km sea surface from the VTRPE comparison Section 3.2 with a free space atmosphere was solved using the total field formalism (a space 36m tall, 43.3 billion total FDTD cells, in conventional FDTD) and the scattered field approach (a space 2.25m tall, 2.7 billion total FDTD cells.) The comparison between the two methods is then done at the scattered field level. That is, an FDTD free space of the same dimensions is run and its fields are subtracted from the total field solution to yield the scattered field only. To obtain the fields at 9GHz the results of the total field formalism are Fourier transformed at runtime. In the scattered field VEP approach, the Schelkunoff currents on the entire collection surface above the sea are Fourier transformed and then the field radiated by those currents into the space are calculated by using the conventional free

space Green's function corrected for FDTD's dispersion and anisotropy. Details of the reconstruction are given in Chapter 8 and the specific 2D implementation in APPENDIX A. Figure 3.16 and Figure 3.17 show the results.

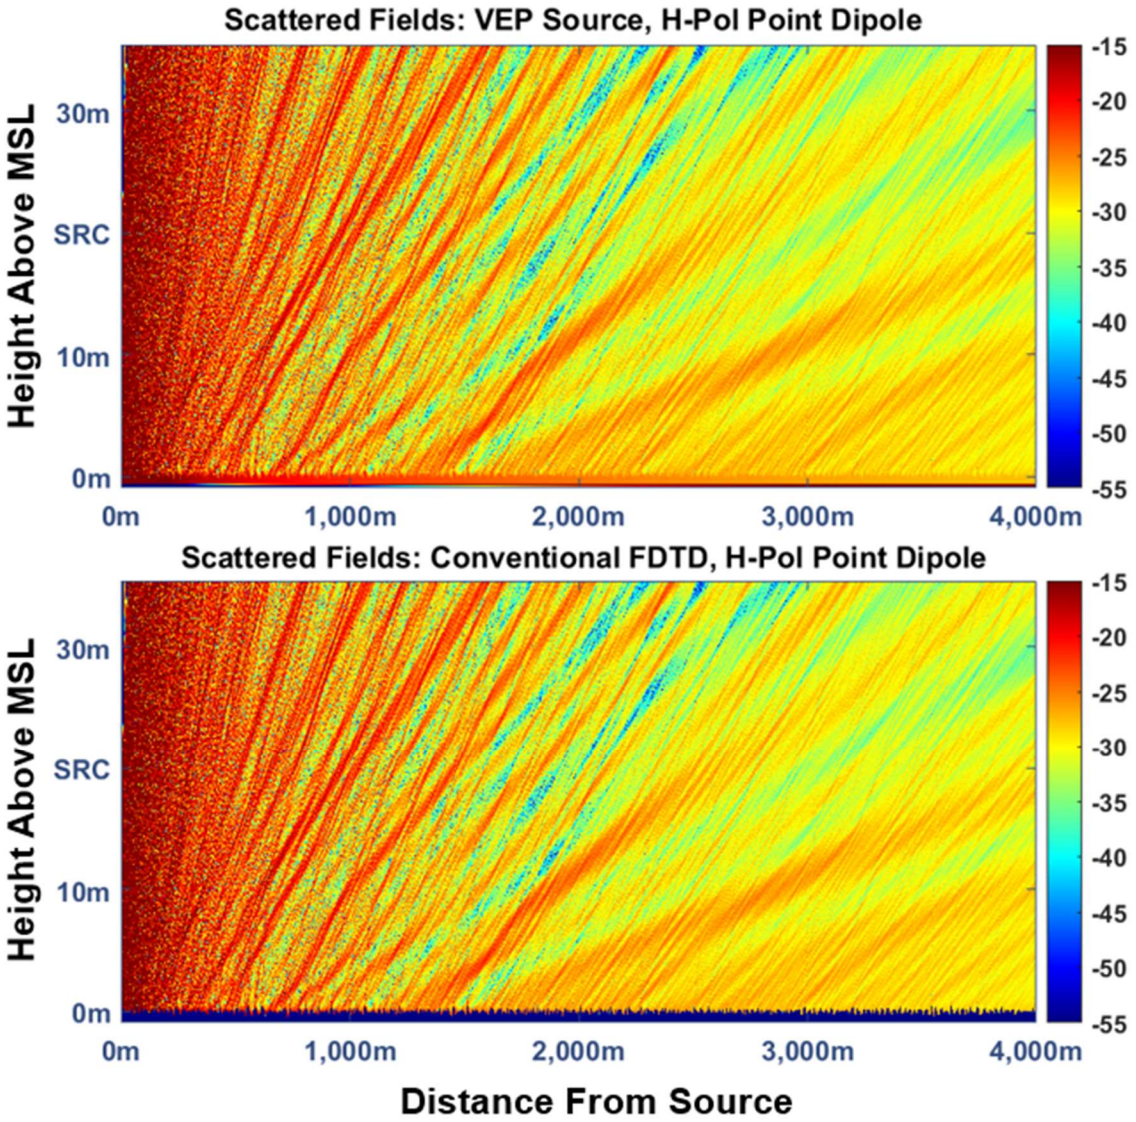

Figure 3.16: Contour plots of the scattered field from a realistic sea calculated using the total field FDTD formalism (top) and the scattered field approach using VEP sources (bottom).

The excellent agreement shows (i) that the very simple anisotropy and dispersion correction of equation 3.4 is good enough for this problem and (ii) that the short height teleportation windows at the ends of the sub-domains (of Figure 2.8) are indeed all that is

required to transfer the relevant (coupled) fields from sub-domain to sub-domain and that the equivalent currents on the collection surface only contribute to the field radiated upwards in the space.

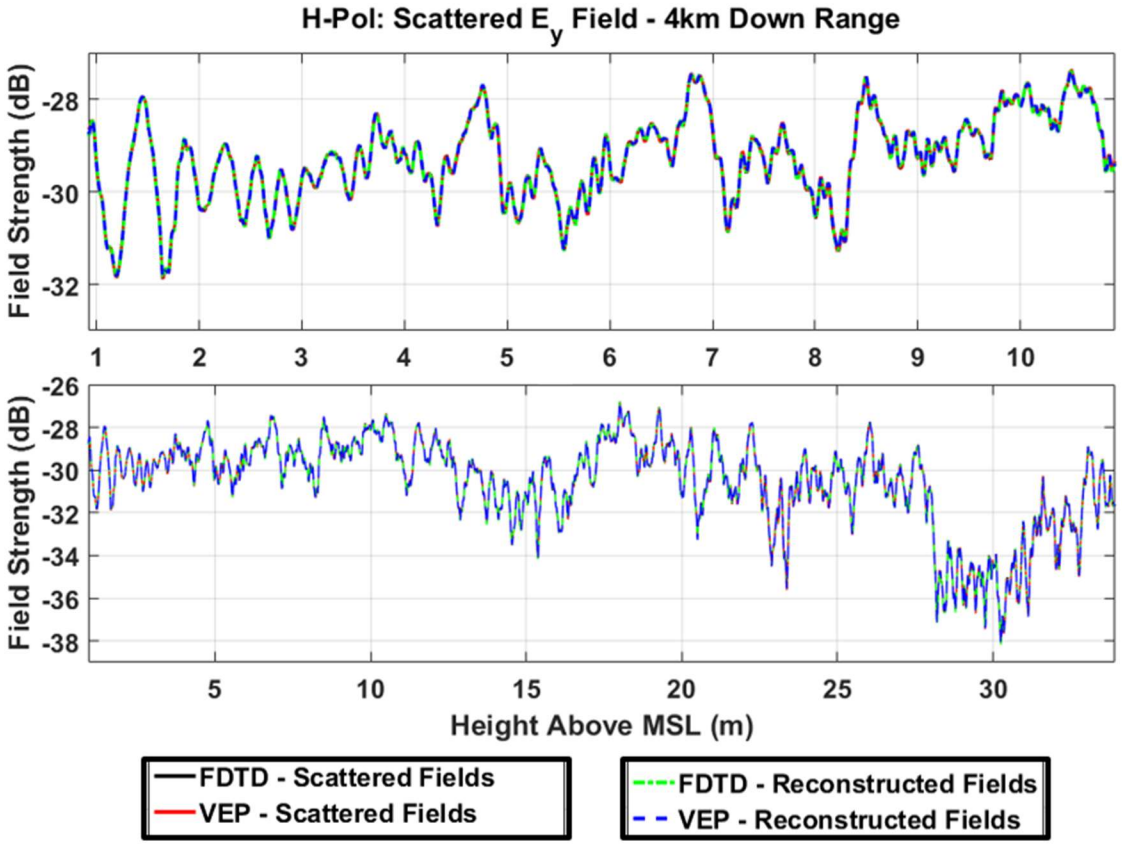

Figure 3.17: Scattered field along a vertical cut 4km down range. Conventional FDTD (black), VEP analytic source injection (red), dispersion-corrected Green function reconstruction from Schelkunoff surface in conventional FDTD space (green) and VEP scattered field (blue).

3.5 Conclusion

Implementation of the techniques described in 0 have allowed for the cross-validation of enormous test spaces against commonly used integral and parabolic equation methods. Due to FDTD's inherent and unavoidable grid dispersion and anisotropy, we implement a first-order correction to the FDTD domains which involves first pre-compressing the length of the domain by an empirically derived factor, running the simulation, and decompressing the results back out to full length. This technique was

applied to the line plot comparisons against MoM and VTRPE in Sections 3.1 and 3.2. The comparisons against the published FB-BIE and SSF-PWE methods showed negligible effects of dispersion and anisotropy due to much finer grid discretization ($\lambda/53.4$), and the aforementioned "scaling" was not applied. The agreement is excellent, with only minor differences in regions of very low signal strength.

Another, perhaps more significant source of error than dispersion and anisotropy, arises as a result of the spatial representation of the sea surface between FDTD, MoM and, particularly, VTRPE. The comparisons against VTRPE in Section 3.2 indicate due to the exact agreement in the case of the flat sea, and for the case of the realistic sea's better agreement closer down range, that FDFD's anisotropy may not be the culprit of the differences. Rather, differences in spatial representation of the sea due to differences in FDTD's discretization (a 1.875mm square cell) compared to VTRPE's coarser horizontal (5.301mm) and much final vertical (0.1359mm) discretization sizes are much more plausible. A possible solution may involve implementation of an effective media model for the sea to minimize any "staircasing" effects introduced into the material surface by FDTD's grid.

Finally, we implemented the volume equivalence principle source injection method described in Section 2.4 with an analytic source pulse. To exactly reproduce the conventional FDTD scattered fields, the analytic source's Green's function must be dispersion-matched to the FDTD grid, which is accomplished by modification of the numerical wavenumber based on the direction-dependent phase velocity. Dispersion correction allows for an exact replication of conventional FDTD scattered fields.

CHAPTER 4

DEMONSTRATION THE TECHNIQUES SCALE DIRECTLY FROM 2D TO 3D

4.1 Direct Scalability to a 3D Corridor of Sea

The techniques applied to this point transfer directly from the 2D problem to the 3D problem. The only change is the 3D problem involves more field terms. For instance, the application of the teleportation algorithm to create rRATs on the upper domain truncation boundary uses essentially the same equations for the TEy case (eqs. 2.4-2.12) but combines them with what would be the boundary case for TMy. Thus, in 3D, there are four field components to update rather than the two in 2D. Furthermore, the boundaries form a plane in 3D rather than a line, so one additional index (j in this case) must be iterated over. Again, consider the case of a rRAT boundary placed at z-position K_{rRAT}, and let $K_E \equiv K_{rRAT}$, $K_H \equiv K_{rRAT} + \frac{1}{2}$, and for simplicity $I \equiv i + \frac{1}{2}$, $J \equiv j + \frac{1}{2}$.

1. Standard H update
2. Update H in rRAT boundaries:
   a. Recall stored 'past' field and store new 'past' field:

Forward wall:

$$\begin{aligned} Ex_{past}(i,J,K_E) &= Ex_{stored}(i,J,K_E) \\ Ex_{present}(i,J,K_E) &= S_c \cdot E_x(i,J,K_E) \\ Ex_{stored}(i,J,K_E) &= Ex_{present}(i,J,K_E) \end{aligned} \quad 4.1$$

$$\begin{aligned} Ey_{past}(I,j,K_E) &= Ey_{stored}(I,j,K_E) \\ Ey_{present}(I,j,K_E) &= S_c \cdot E_y(I,j,K_E) \\ Ey_{stored}(I,j,K_E) &= Ey_{present}(I,j,K_E) \end{aligned} \quad 4.2$$

Reverse wall:

$$Ex_{past}(i,J,K_E+1) = Ex_{stored}(i,J,K_E+1)$$
$$Ex_{present}(i,J,K_E+1) = S_c \cdot E_x(i,J,K_E+1) \quad (4.3)$$
$$Ex_{stored}(i,J,K_E+1) = Ex_{present}(i,J,K_E+1)$$

$$Ey_{past}(I,j,K_E+1) = Ey_{stored}(I,j,K_E+1)$$
$$Ey_{present}(I,j,K_E+1) = S_c \cdot E_y(I,j,K_E+1) \quad (4.4)$$
$$Ey_{stored}(I,j,K_E+1) = Ey_{present}(I,j,K_E+1)$$

b. Update boundary fields

Forward wall:

$$H_x(I,j,K_H+2) = H_x(I,j,K_H+2) - B \cdot Ey_{past}(I,j,K_E) \quad (4.5)$$
$$H_y(i,J,K_H+2) = H_y(i,J,K_H+2) + B \cdot Ex_{past}(i,J,K_E) \quad (4.6)$$

Reverse wall:

$$H_x(I,j,K_H) = H_x(I,j,K_H) + B \cdot Ey_{past}(I,j,K_E+1) \quad (4.7)$$
$$H_y(i,J,K_H) = H_y(i,J,K_H) - B \cdot Ex_{past}(i,J,K_E+1) \quad (4.8)$$

3. Standard E update

4. Update E in rRAT boundaries:

a. Recall stored ‘past’ field and store new ‘past’ field:

Forward wall:

$$Hx_{past}(I,j,K_H+1) = Hx_{stored}(I,j,K_H+1)$$
$$Hx_{present}(I,j,K_H+1) = S_c \cdot H_x(I,j,K_H+1) \quad (4.9)$$
$$Hx_{stored}(I,j,K_H+1) = Hx_{present}(I,j,K_H+1)$$

$$Hy_{past}(i,J,K_H+1) = Hy_{stored}(i,J,K_H+1)$$
$$Hy_{present}(i,J,K_H+1) = S_c \cdot H_y(i,J,K_H+1) \quad (4.10)$$
$$Hy_{stored}(i,J,K_H+1) = Hy_{present}(i,J,K_H+1)$$

Reverse wall:

$$Hx_{past}(I,j,K_H+1) = Hx_{stored}(I,j,K_H+1)$$
$$Hx_{present}(I,j,K_H+1) = S_c \cdot H_x(I,j,K_H+1) \quad (4.11)$$
$$Hx_{stored}(I,j,K_H+1) = Hx_{present}(I,j,K_H+1)$$

$$Hy_{past}(i,J,K_H+1) = Hy_{stored}(i,J,K_H+1)$$
$$Hy_{present}(i,J,K_H+1) = S_c \cdot H_y(i,J,K_H+1) \quad (4.12)$$
$$Hy_{stored}(i,J,K_H+1) = Hy_{present}(i,J,K_H+1)$$

b. Update boundary fields

Forward wall:

$$E_x(i,J,K_E+1) = E_x(i,J,K_E+1) + A \cdot Hy_{past}(i,J,K_H+1) \quad (4.13)$$
$$E_y(I,j,K_E+1) = E_y(I,j,K_E+1) - A \cdot Hx_{past}(I,j,K_H+1) \quad (4.14)$$

Reverse wall:

$$E_x(i,J,K_E) = E_x(i,J,K_E) - A \cdot Hy_{past}(i,J,K_H+1) \quad (4.15)$$
$$E_y(I,j,K_E) = E_y(I,j,K_E) + A \cdot Hx_{past}(I,j,K_H+1) \quad (4.16)$$

In accordance with the discussion in Chapter 1, the reason a 3D fullwave method such as FDTD needs to be applied to the problem of the realistic maritime radar environment is to enable the discovery of out-of-plane or other physical phenomena that might be missed by the strictly 2D implementation of VTRPE. Therefore, the transition to 3D proceeds in two steps: First, we must verify that the algorithms have been correctly programmed. This is accomplished by solving the exact same problem in 2D and 3D. This means, as Figure 4.1 suggests, that a horizontal y-directed line source stretching from a PEC plane at y = -0.65m to a PEC plane at y = +0.65m, over an "extruded" version of the 2D sea must give the same results as the 2D problem when solved using the same space-step and Courant number. Sure enough, the results exactly match (Figure 4.2).

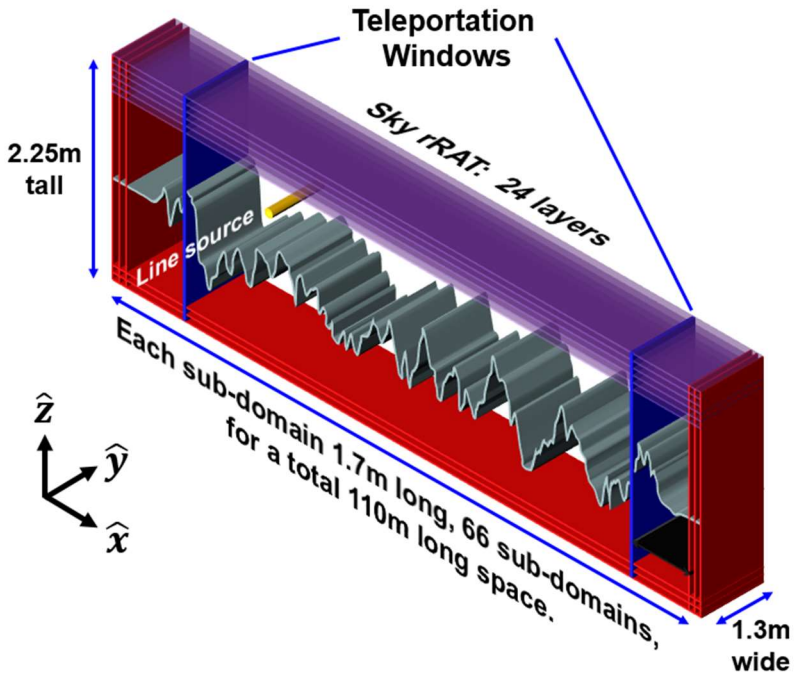

Figure 4.1: Extruded 3D sea and source used to verify that all algorithms are functional in 3D.

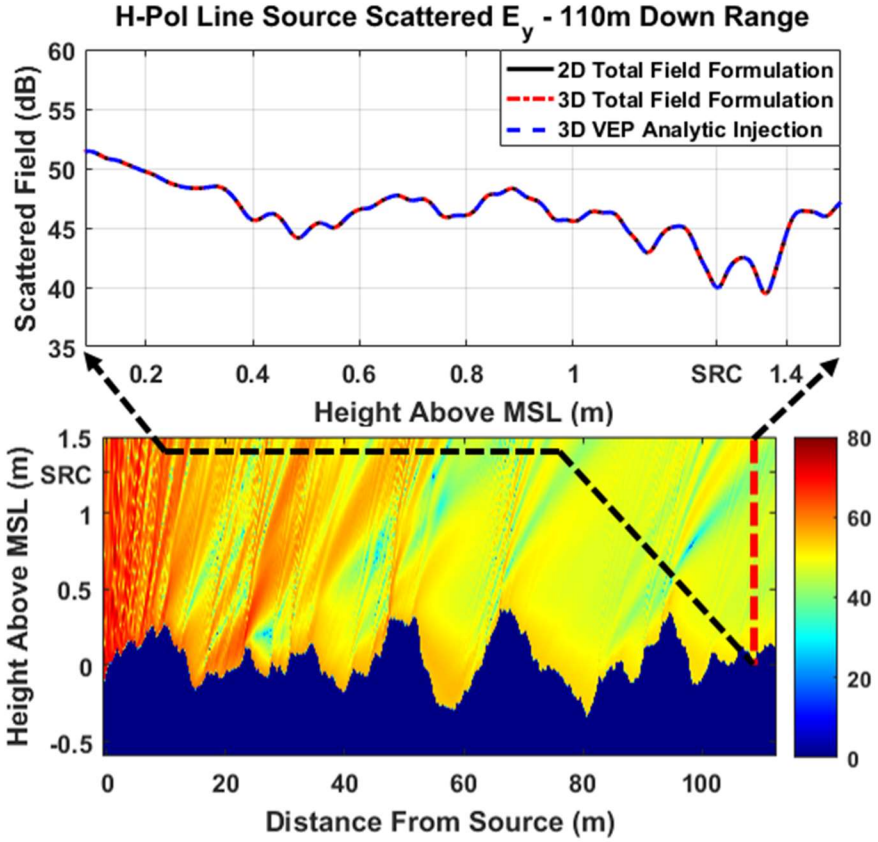

Figure 4.2: Scattered fields for the 2D to 3D comparison on extruded sea with source inside of FDTD domain: (top) scattered fields on vertical line 110m down range, (bottom) contour plot of scattered fields.

To maximize the total amount of sea which can be simulated in 3D, we must use the scattered field VEP approach. Therefore, the same scenario as suggested in Figure 4.1 is used but the line source is now moved to a realistic height above the sea surface (and above the upper rRAT truncation boundary) and its fields are injected as discussed earlier. Figure 4.3 shows the results of the comparison of the extruded space to the 2D space for both the total field formulation and the scattered field VEP approach. In 3D, this 110m-long space is over 3,300 wavelengths long and contains 50.5 billion FDTD cells.

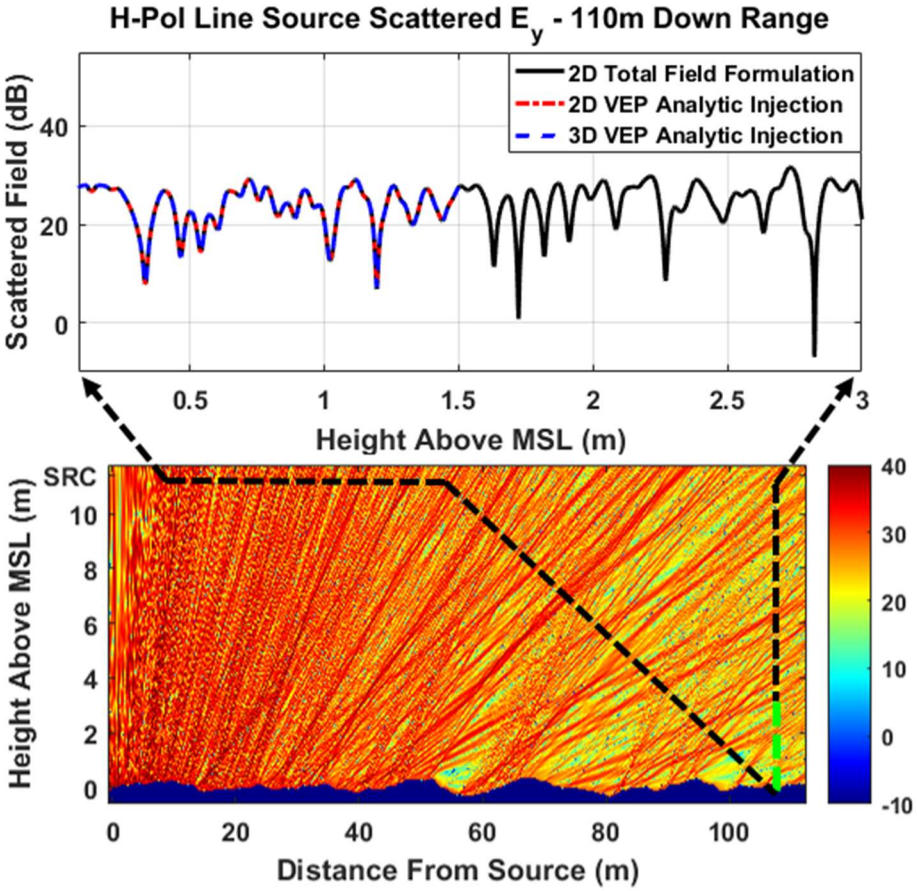

Figure 4.3: Scattered fields for the 2D to 3D comparison sea with source outside of FDTD domain: (top) Fields on vertical line 110m down range, (bottom) contour plot of equivalent FDTD Total Field Formulation space.

Before presenting the first case of a true 3D source using VEP analytic injection, the appropriate 3D dispersion relationship must be considered. The 3D dispersion relation

differs from the 2D relation (equation 3.7) through inclusion of an extra term on the right-hand side, as well as another angular dependence (φ).

$$\frac{1}{S^2} sin^2\left(\frac{\pi S}{\lambda_0}\right) = sin^2\left(\frac{\mathrm{ds}}{2}\tilde{k}_x\right) + sin^2\left(\frac{\mathrm{ds}}{2}\tilde{k}_y\right) + sin^2\left(\frac{\mathrm{ds}}{2}\tilde{k}_z\right) \tag{4.17}$$

Where, in physics notation,

$$\begin{aligned} \tilde{k}_x &= \tilde{k}\sin\theta\cos\phi \\ \tilde{k}_y &= \tilde{k}\sin\theta\sin\phi \\ \tilde{k}_z &= \tilde{k}\cos\theta \end{aligned} \tag{4.18}$$

Equation 4.17 leads to a generally more complicated relation to solve. In the case of the line source described in Figure 4.1-Figure 4.3, the dispersion relation between 2D and 3D is identical since ϕ is simply zero in the x-z plane. In the case of the true 3D source used to produce Figure 4.5 and Figure 4.6, the full 3D directional dependence must be considered. To obtain the 3D phase velocity, a numerical fit was performed on equation 4.17, and an analytic expression was constructed to match it exactly (for cubic cells):

$$\begin{aligned} v_{p3D}(\theta,\phi) = \; & v_{pmin2D} + \Delta v_{p2D}\cdot[2\sin(\theta)\cos(\theta)]^2 + \\ & +\Delta v_{p2D}\cdot[2\sin(\phi)\cos(\phi)]^2\cdot\sin^4\theta \end{aligned} \tag{4.19}$$

In equation 4.19, Δv_{p2D} is the same definition as in equation 3.10. To obtain the 3D numerical wavenumber, $v_{p3D}(\theta,\phi)$ from equation 4.19 is substituted into equation 3.12 in place of $v_{p2D}(\theta)$. Figure 4.4 shows the contour plot of the directional dependence of the 3D numerical wavenumber relative to that of free space.

Finally, the source in the 3D extruded space is changed to a 3D horizontal dipole point source. Given that both the line source and point source electromagnetic fields approach plane waves in the far field, we expect the scattered field distribution in the mid-

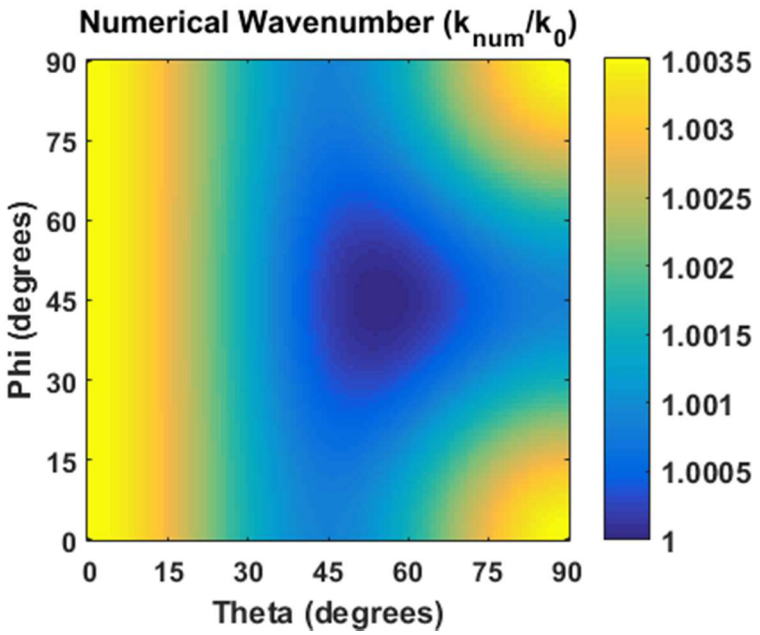

Figure 4.4: Contour map of the directional scaling of the numerical wavenumber relative to the free space wavenumber in 3D for ds = λ/17.8 and S = 0.577.

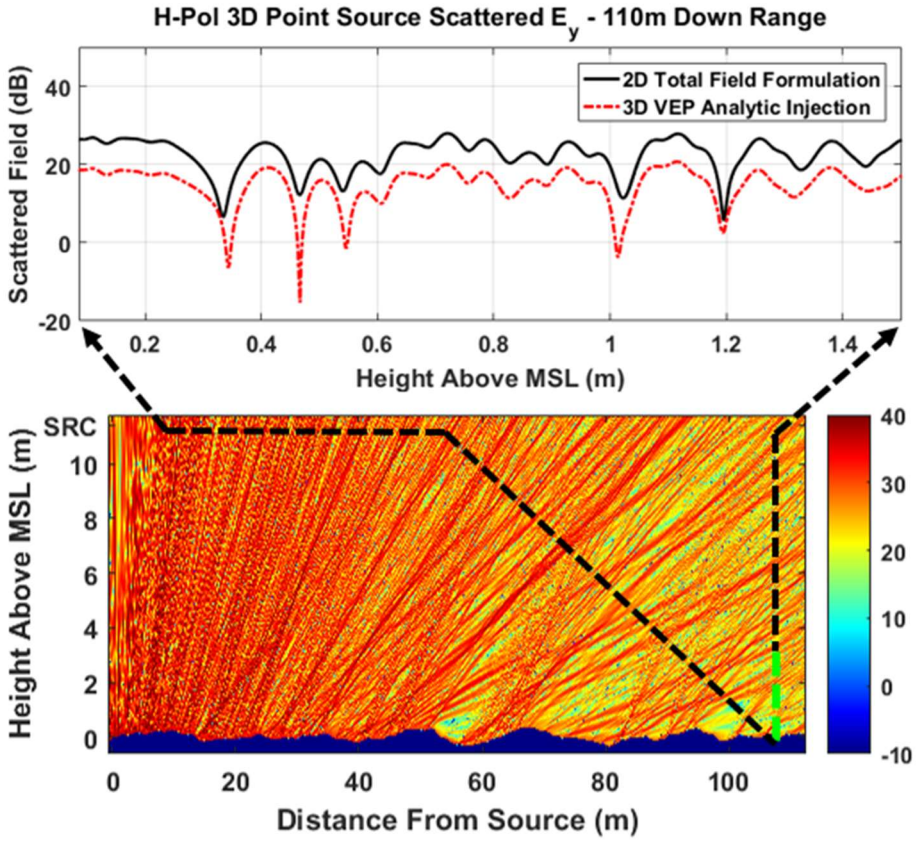

Figure 4.5: The scattered field, along a vertical observation plane 110m down range, produced by a true 3D source over the extruded sea compared to the 2D total field formulation result.

plane of the 3D space, far from the source, to show approximately the same behavior as the 2D result, except the scaling factor of the radial dependence ($1/\sqrt{r}$ for the line, and $1/r$ for the point source). The results of Figure 4.5 and Figure 4.6 agree with this expectation.

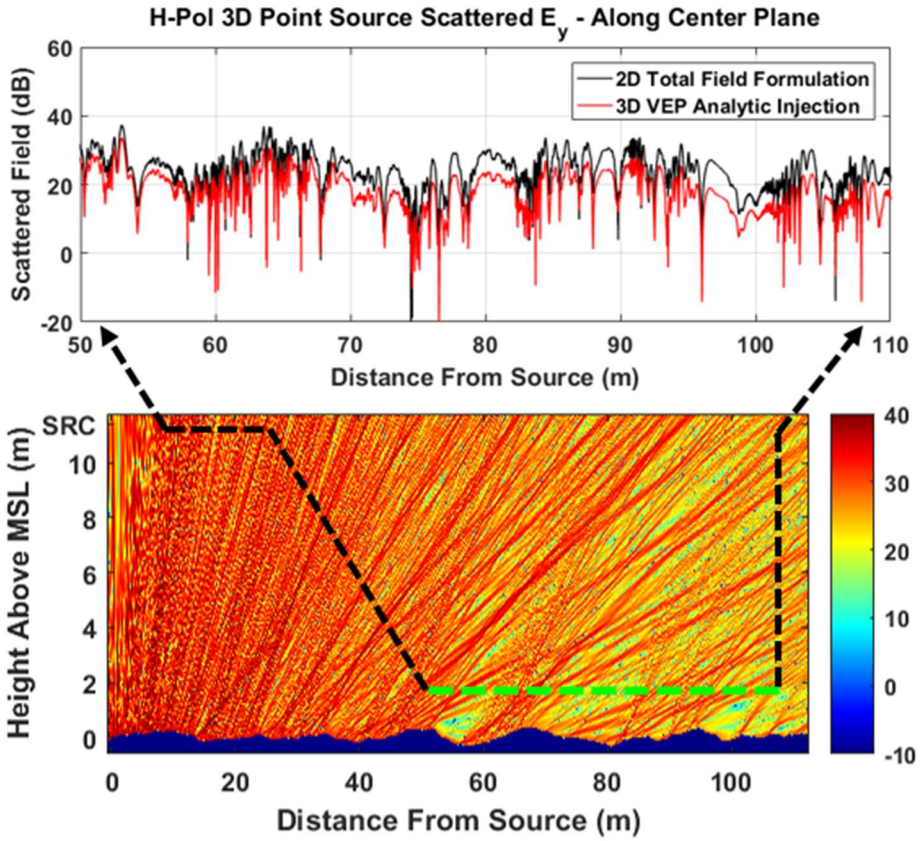

Figure 4.6: The scattered field, along a horizontal observation plane just above the sea surface, produced by a true 3D source over the extruded sea compared to the 2D total field formulation result.

4.2 Conclusion

The teleportation-based domain decomposition and the rRAT boundaries can be directly applied in 3D FDTD in virtually the same manner as in 2D, the only extra consideration are the two extra field components in the teleportation algorithm. Similarly, the analytic source VEP injection is directly applicable to 3D as well, however 3D FDTD has a different, more complicated dispersion relation than in 2D, but it does possess an analytic expression, which allows for 3D analytic sources to be dispersion-matched just as

well as in 2D. Tests of a 3D space containing a corridor of extruded 2D sea, with PEC walls and a line source, exactly match the 2D solution for both conventional 3D FDTD and 3D VEP source injection, which indicates proper 3D functionality of our methods. The same sea corridor, but with absorbing terminations on all walls and a 3D analytic point source, was solved, and the scattered fields along the end-plane and mid-plane of the corridor were compared to the 2D solution. Since the sea contains no transverse variations, the basic shape of the scattered fields should be very nearly the same between the 2D and 3D cases, and they are, the only major differences arise from the $1/\sqrt{r}$ radial dependence of the 2D source field amplitude compared to the $1/r$ dependence in 3D. We can now move forward in confidence of the validity of the new techniques.

CHAPTER 5

ENSURING 3D FDTD CAN MODEL A RELEVANT CORRIDOR OF SEA SURFACE

5.1 Demonstration that rRATs do not Introduce Artefacts

The FDTD computational domain must be terminated in absorbing boundary conditions that do not scatter significant energy back into the domain. In 2D the most demanding termination is the sky boundary because after traveling thousands of wavelengths downrange the angle of incidence of the waves onto that boundary is extremely shallow. The re-Radiating Absorbing Termination (rRAT) is the only termination reported to date that can absorb rays traveling at arbitrarily shallow angles (down to hundredths of a degree from grazing). This boundary enabled our 2D FDTD results to agree with VTRPE and BIE simulations of both sea and inhomogeneous atmosphere in domains 4km long.

In 3D the same problem now arises at the sidewalls of the domain. To faithfully reproduce the relevant out-of-plane physics, expected to include side to sideways-scattering of energy between the side surfaces and crests of the waves, the transverse extent of the 3D domain must be wide enough to contain more than one wavelength of the sea surface gravity waves. Since these are of the order of 1m the goal is to make a sea corridor at least 3m wide (about 100 wavelengths across.) Once the corridor is wide enough to contain the important phenomena it must be terminated on the sides with absorbing boundaries good enough to guarantee that any reflection from them does not falsify the out-of-plane phenomena we are looking for.

In addition to reflection artefacts, imperfectly absorbing boundaries could also artificially remove energy from a traveling wave in the domain. Again, in our case this is not acceptable because our input signal must travel hundreds of meters to develop the relevant downrange scattered field. To verify that no significant reflection is being returned into the domain and no excess energy is being drained off we simulate a corridor of pure free space with a point radiator, and compare the field propagated along that space to the theoretical free space propagation from a point source

Three 3D test spaces were constructed. Each consists of a 215m-long corridor of varied cross-sectional dimensions (Figure 5.1) and discretized with cubic cells of side length λ/17.8 (1.875mm) at 9GHz. Each sub-domain is 800 cells in length (1.5m long) and the domain decomposition method of Section 2.2 is used to propagate the pulse the full 215m. The cross sections of the test domains were 800x800, 400x400, and 200x200 cells. In every case the source is placed in the center of the corridor, and there are 16 rRAT layers on the walls of the corridor. Each rRAT consumes 3 cells.

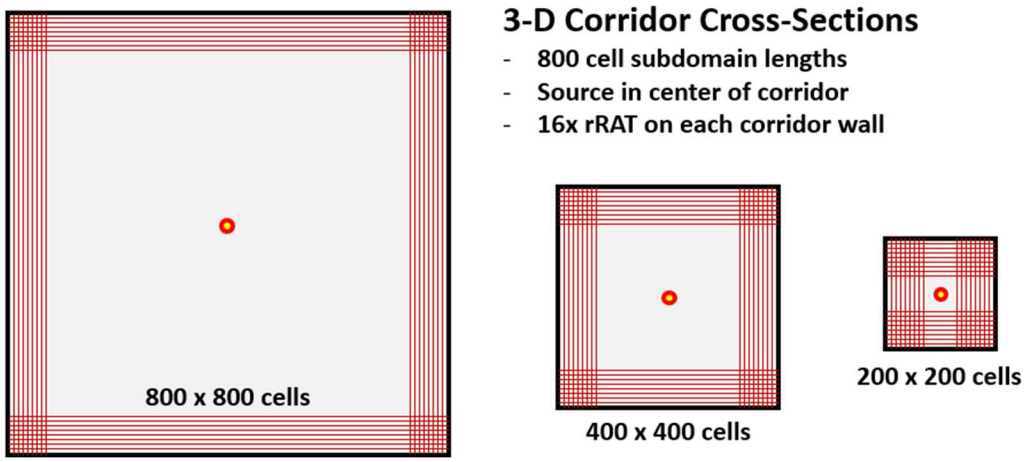

Figure 5.1: Corridor cross-sections for rRAT energy leakage tests.

One way to assess the fidelity of the propagation is to determine whether the fields along the center of the corridor decay in strength as the inverse of the distance from the

source. As a qualitative test, plotting the fields along the center of the corridor down its length should have overlapping results for all three test cases, which is the case (Figure 5.2).

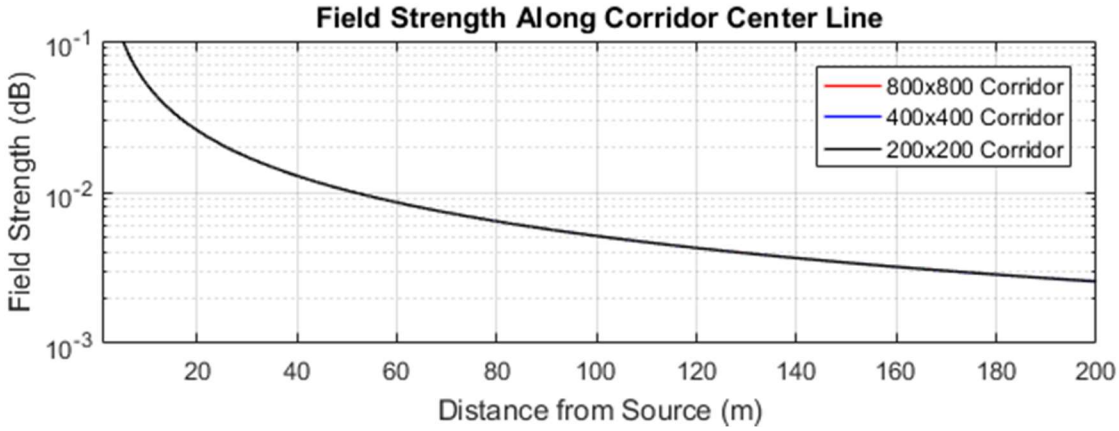

Figure 5.2: Field strength vs. distance along center of corridor.

A more quantitative test is to check that the logarithmic slope of the absolute field amplitude is -1 over the length of the corridor since the Green function solution follows an r^{-1} dependence in the far field. However, we must keep in mind that the -1 power of the Green's function solution arises from the propagation of spherical waves in Maxwell's equations and we know that all FDTD domains exhibit anisotropy. Table 5-1 shows the slopes for each case for two points taken at different ranges, the logarithmic slope will be called α – the field drop with distance is then r^{α}. The α are very close to -1, but not exactly.

TABLE 5-1: LOGARITHMIC SLOPE OF THE FIELD'S "SPACE LOSS" ALONG THE CENTER OF THE CORRIDOR

Range of Slope	α (800x800)	α (400x400)	α (200x200)
10m – 100m	-0.999821	-0.999820	-0.999823
100m – 200m	-0.999961	-0.999674	-0.999966
10m – 200m	-0.999871	-0.999854	-0.999426

The results of Figure 5.2 indicate that the size of the domain has no significant effect on the exponent of the $1/r$ space loss. No excess energy appears to be drained form

the traveling wave. At our chosen spatial discretization, propagation of a spherical wave along a principal Cartesian coordinate of FDTD appears to have an exponent within 4 hundredths of a percent of 1 (that is 0.9996+/- 0.0003). The variation induced by random interference of the faint scattering from the rRAT sidewalls is of the same order, 3 hundredths of a percent. This interference is overstated because we have evaluated the field along the center line of the domain where all echoes will add coherently. This worst-case echo translates to about -70dB down. This means we could discern out-of-plane phenomena anywhere along the 3D corridor at least -50dB down from the main signal and still have 20dB clearance from the "noise floor". We conclude that the rRATs developed to eliminate the "sky-bounce" artefact in 2D are also the ideal terminations to truncate the transverse extent of the 3D domains.

5.2 Effects of Sub-domain Length on Forward-Travelling Scatter

In the sea problem using the extreme domain decomposition method, the 2D FDTD sub-domains can be made tens of thousands of wavelengths long due to the relatively low memory and processing requirements compared to 3D. For such long sub-domains in 2D there is no doubt that all relevant back-and-forth interactions of the sea surface are accounted for. (And this was confirmed previously by comparison with the forwards-backwards MoM solution of the entire sea surface by Ohio State in Section 3.1.)

In 3D FDTD, the computational burden increases rapidly, which in turn limits the length of each sub-domain in the 3D corridor to be about 1,000 cells long (about 50-100 wavelengths long). The question must be asked, "How short can a sub-domain be and still capture the forward-travelling information relevant for scattering off of a target many

kilometers down range?" Figure 5.3 below illustrates how a very short sub-domain (black rectangle) might miss capturing a multipath reflection from a large sea crest down range that a much wider domain (blue box) would not miss.

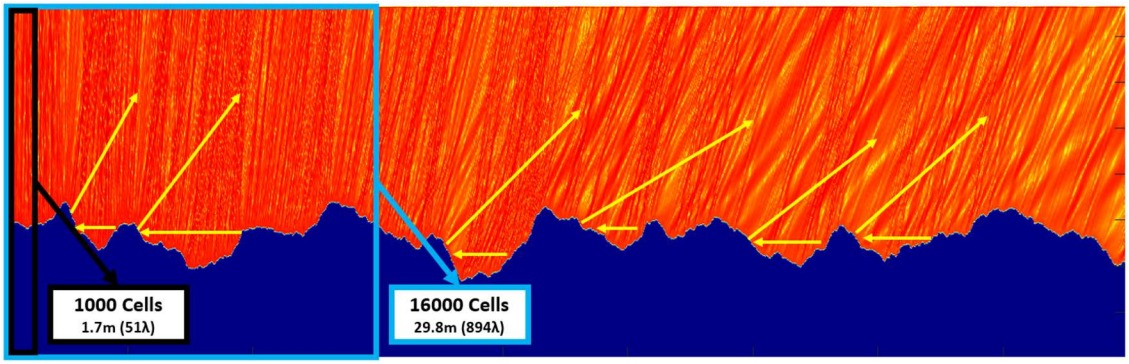

Figure 5.3: Potential multiple-bounce rays may not be captured by short sub-domains.

To answer this question, a sea-state 3 sea space 1km in length with free space atmosphere was run in 2D FDTD and divided into sub-domains of lengths varying from 400 cells up to 144,000 cells. Firstly, the results will be compared in the form of the scattered fields which arrive at a 35m-tall vertical observation plane 1km down range. Secondly, the angle of arrival (AoA) will be compared for the fields at the same vertical plane. The FDTD space in all cases is run with a spatial discretization of λ/17.8, a CFL of 0.707, with dispersion-matched analytic source field injection. The sub-domain lengths tested are 400, 800, 1000, 1200, 1600, 2000, 4000, 8000, 16000, 32000, 48000, 64000, 96000, 144000 cells. For each sub-domain, 100 of the total cells are consumed for absorbing boundaries and analytic field injection tapering – i.e. the 400 cell sub-domains contain 300 cells of sea. A few relative sub-domain sizes are shown in Figure 5.4 below.

For the following figures, the 64000, 96000, and 144000 cell configurations were not included. They are identical to the 48,000-cell configuration in all cases. The scattered fields 1km down range versus height are shown in Figure 5.5 below. For clarity the results

near the surface (1-5m) and near the top of the observation plane (30-35m) are shown separately. Across the 35m height there are only minor differences in the peaks and nulls of the field plots for all cases.

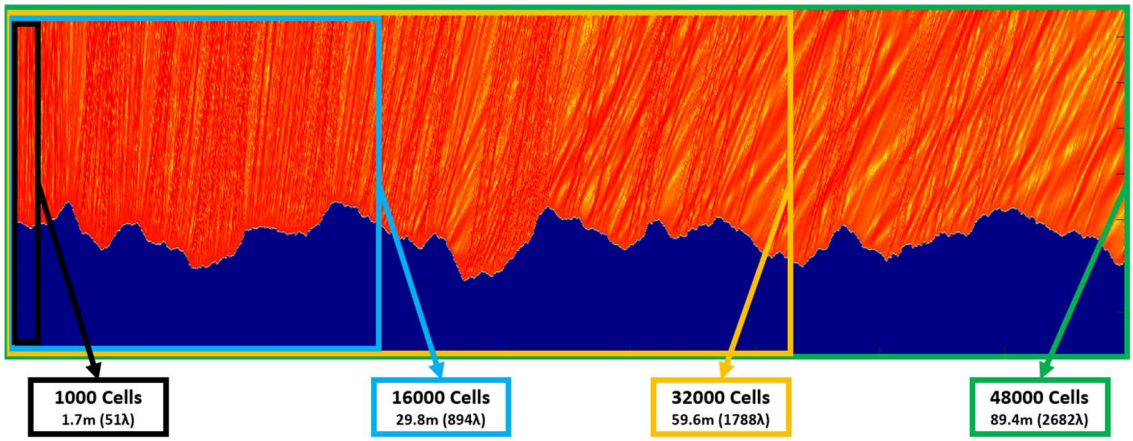

Figure 5.4: Size comparison for several sub-domains used in testing.

The fields are almost indistinguishable between all cases at 30m – 35m above mean sea level (MSL). This is not surprising, the strongest fields reaching higher altitudes are those affected less by the local sea and more by the longer history of the sea it has passed. This also has implications for the AoA, especially since the AoA calculated by the spatially-limited fields on the 35m vertical plane. As discussed earlier, of the rays which cross that plane, the higher the angle of scatter, the closer the section of sea is which gave rise to those rays.

If only single-bounce forward reflections occurred from the sea surface, then the length of the sub-domain would not matter, and all fields and AoA would look the same. The sea surface, however, like most rough surface scattering, generates a significant amount of multi-bounce reflection. Reference back to the ray paths of Figure 5.3 shows how a short sub-domain cannot capture a multi-bounce between two peaks of adjacent large gravity waves, therefore this information will be missed. Luckily, since the path

length between gravity wave peaks is of the order of 1m, the $1/r$ (or $1/\sqrt{r}$ in 2D) field strength drop is significant and thus in this kind of multi-path the energy decreases with every bounce and its contribution to the forward-travelling scatter of interest is generally weak. Furthermore, the multi-bounce scatter typically sends most of its energy into higher angles, so there is only a small portion travelling at the low angles which reach far enough down range to interact with the target. The AoA plots of Figure 5.6 show that very short domains tend to miss some of the details of high angle scatter. For clarity we separate the low angle results (0-10 degrees) from the high angle results 40-45 degrees.

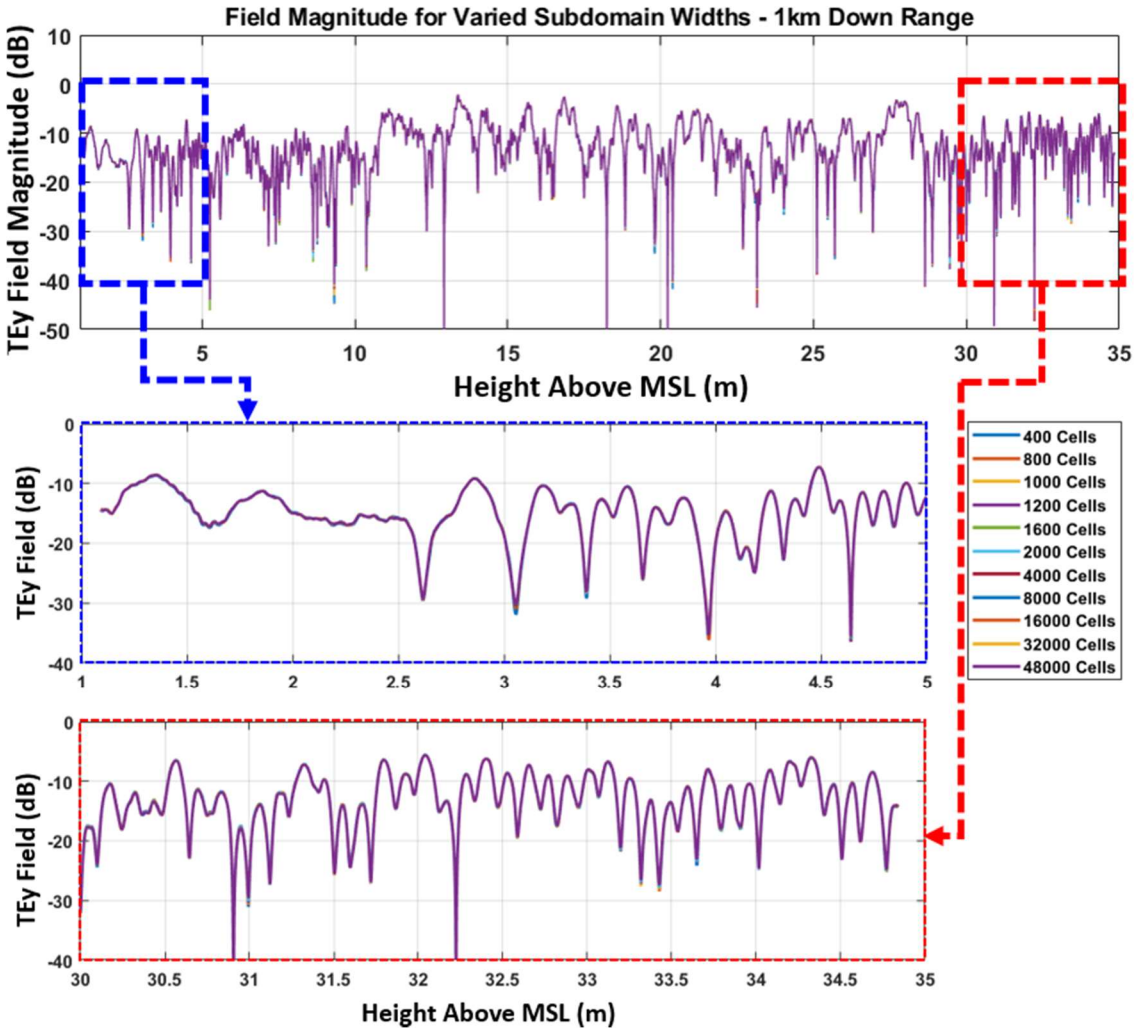

Figure 5.5: Scattered field magnitude vs. height at an observation plane 1km down range.

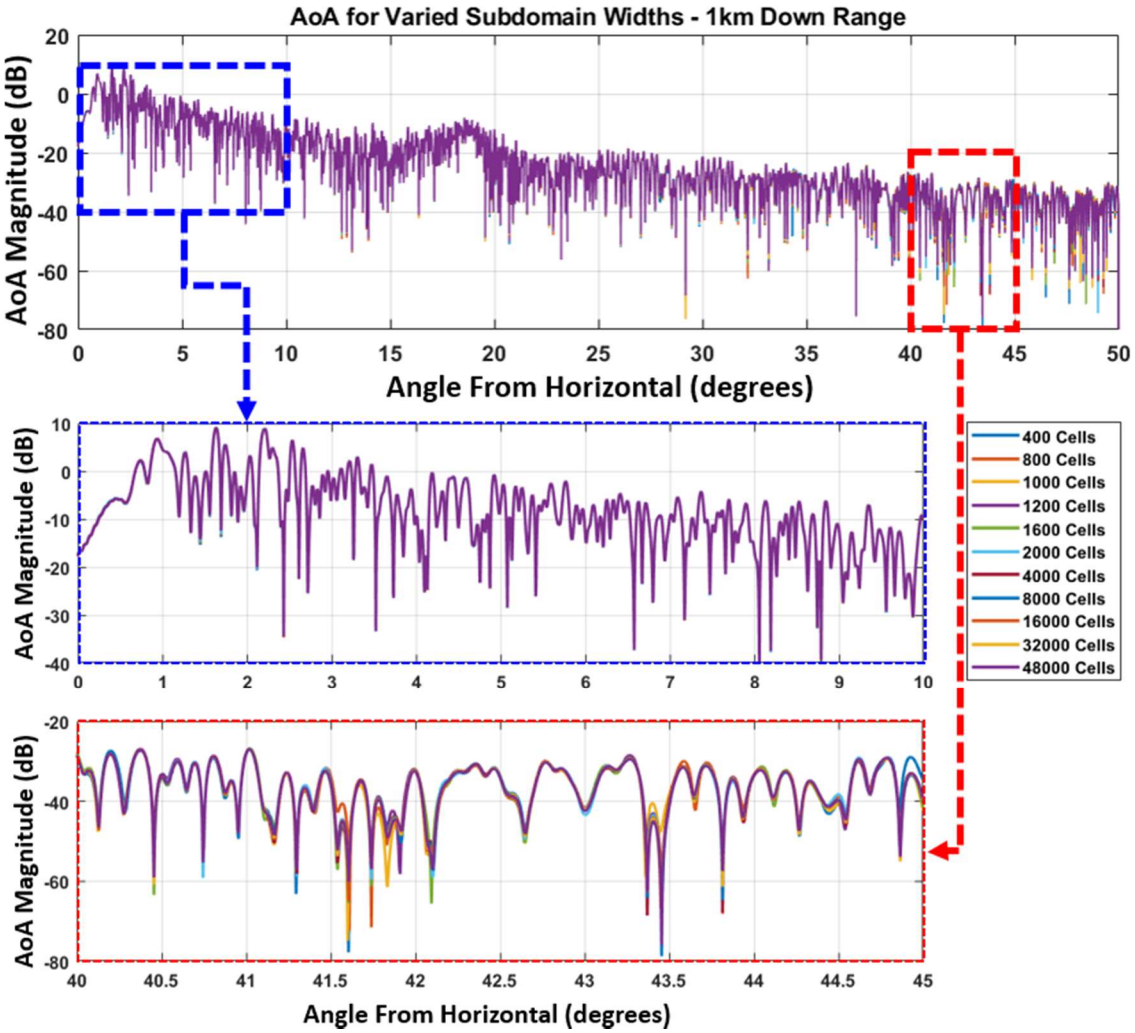

Figure 5.6: AoA of scattered fields 1km down range (top) 0 to 50 degrees, (middle) 0 to 10 degrees agrees extremely well, (bottom) 40 to 45 degrees show some disagreement in localized angular ranges occurring with shorter sub-domains.

The previous analysis shows that for the scattered fields which cross the 35m-tall observation plane – a problem very relevant to naval platforms – the fields and their AoA are affected to some degree by the size of the sub-domain. The main disagreement appears to occur over localized angular regions at the higher angles, does not exceed more than half of a degree in extent, and becomes small at domain lengths of 1,000 cells and above. Therefore 1,000 cells in length is our target minimum domain size for the 3D simulations for sea-state 3 or lower seas.

5.3 Conclusion

The fact that the rRATs are suitable for the truncation of the constant-contact fields along the multiple walls of the 3D FDTD corridor is demonstrated by placing a 3D point source in several free space corridors of varied cross-sectional width, surrounded on the outside by rRAT boundaries. The near $1/r$ dependence of the source fields in all cases indicates the rRATs are sufficiently absorptive and do not sap energy away from the simulation space. Prior to venturing into full 3D realistic sea spaces, we need to be certain the domain decomposition sub-domain sizes do not eliminate any multiple-bounce phenomena occurring between large sea features which is relevant to the forward-propagating scattered fields. Variation of 2D sub-domain lengths from a very short 400 cells, to an extremely large 144,000 cells, shows that to faithfully capture the forward scatter the sub-domain lengths should be at least 1,000 cells in length (more than 50λ) for a sea-state 3 or lower sea surface.

CHAPTER 6

SEA SURFACE CREATION AND STATISTICAL ANGLE OF ARRIVAL PROPERTIES

6.1 Sea Surface Generation from an Energy Spectrum

Through the first two years of our investigation, the only sea surfaces available to test in FDTD were two-dimensional slices provided to ASU by outside sources – all originally generated by VTRPE. To meet the goals of this project we needed to be able to generate sea surfaces in-house. We have developed the capability to produce our own sea realizations in 2D and 3D from any spectrum, as well as the ability to replicate commonly used sea spectra generated by VTRPE. Additionally, in this section a brief statistical analysis is performed on six instances each of VTRPE sea and ASU sea using the VTRPE spectrum. We will compare the angle of arrival (AoA) of scatter reaching a target 1km down range, as well as the rough surface scattering characteristics from the full sea.

The two most common spectra used for the creation of sea surfaces [50] are the Pierson-Moskowitz (PM) and the Joint North Sea Wave Observation Project (JONSWAP) spectra. The two are similar, however the PM spectrum assumes a fully-developed sea – that is, a very large, deep body of water with a steady wind blowing for a long time – whereas the JONSWAP accounts for the distance over the water the wind has been blowing (known as fetch). The JONSWAP spectrum takes the same form as the PM spectrum, however it includes one additional term which multiplies the PM spectrum part and uses a different definition for the α and ω_p parameters. The PM spectrum is given in equation 6.1. The JONSWAP spectrum is given in equation 6.2.

$$S_{PM}(\omega) = \left(\frac{\alpha g^2}{\omega^5}\right) \exp\left[-\frac{5}{4}\left(\frac{\omega_p}{\omega}\right)^4\right] \qquad 6.1$$

$$\omega_p \equiv \frac{0.855g}{U_{10}} \qquad \alpha \equiv 0.0081$$

$$S_J(\omega) = \gamma^r \left(\frac{\alpha g^2}{\omega^5}\right) \exp\left[-\frac{5}{4}\left(\frac{\omega_p}{\omega}\right)^4\right] \qquad 6.2$$

$$\omega_p = 22\left(\frac{g^2}{U_{10}F}\right)^{\frac{1}{3}} \qquad \alpha = 0.076\left(\frac{U_{10}^2}{Fg}\right)^{0.22}$$

$$1 \le \gamma \le 7 \qquad r = \exp\left[-\frac{(\omega-\omega_p)^2}{2\sigma^2\omega_p^2}\right] \qquad \sigma = 0.07\ for\ \omega \le \omega_p\ ,\ \ 0.09\ otherwise$$

Descriptions of each parameter are shown in Table 6-1.

TABLE 6-1: PARAMETER DESCRIPTIONS FOR PM AND JONSWAP SPECTRA

PARAMETER	DESCRIPTION
$\boldsymbol{\alpha}$	Phillips constant
$\boldsymbol{\omega}$	Wave frequency
$\boldsymbol{\omega_p}$	Frequency of spectrum peak
g	Acceleration due to gravity (g = 9.80665 m/s^2)
$\boldsymbol{U_{10}}$	Wind speed 10m above sea level (m/s)
$\boldsymbol{\gamma}$	Peak enhancement factor, usually 3.3
F	Sea fetch in meters

The following discussion covers the limitations of the base PM and JONSWAP spectra, the process for generating 2D and 3D seas, the addition of directionalized dispersion to the sea, and finally verification of the ability to match the VTPRE sea spectra. As a point of clarification, the 2D sea surface is generated from a 1D spectrum, but the spectrum will be called the "2D spectrum" such as to relate it to the 2D sea. Similarly, a 2D spectrum is used to generate the 3D sea, but that spectrum will be called the "3D spectrum" for the same reasons.

The typical PM and JONSWAP spectra are limited in scope. They account well for gravity and large wind waves, but they do not capture higher-frequency phenomena such as capillary waves, which are relevant for scattering at X-band frequencies. In the surfaces generated for use in FDTD, these higher-frequency effects are incorporated (Figure 6.1). In addition, the ASU sea surfaces can incorporate wind from any direction as well as dispersion effects, in both 2D and 3D.

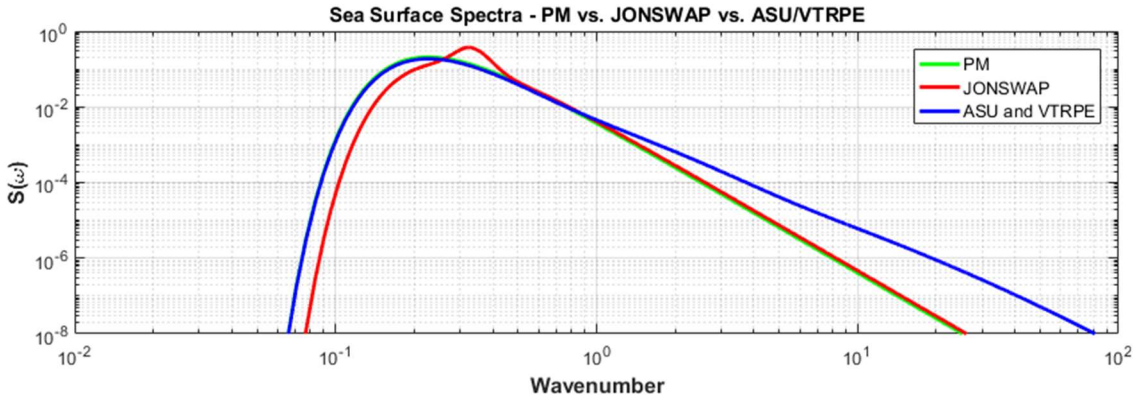

Figure 6.1: Sea surface spectra, (green) PM, (red) JONSWAP, (blue) ASU/VTRPE (which contain more high-frequency energy).

The sea wave spectrum is a purely real function. To generate a randomized sea surface each element of the spectrum must be multiplied by a complex random number, generated from a normal distribution (Gaussian) of random numbers. This assigns a random amplitude and phase to each frequency component, after which the sea surface is simply the Fourier transform of the randomized spectrum.

Comparison of seas generated by the baseline PM and JONSWAP spectra to the VTRPE and ASU seas (Figure 6.2) shows the effects of the higher-frequency components in the surface spectrum. The PM and JONSWAP surfaces are very smooth, with the smallest surface features still much larger than X-band wavelengths.

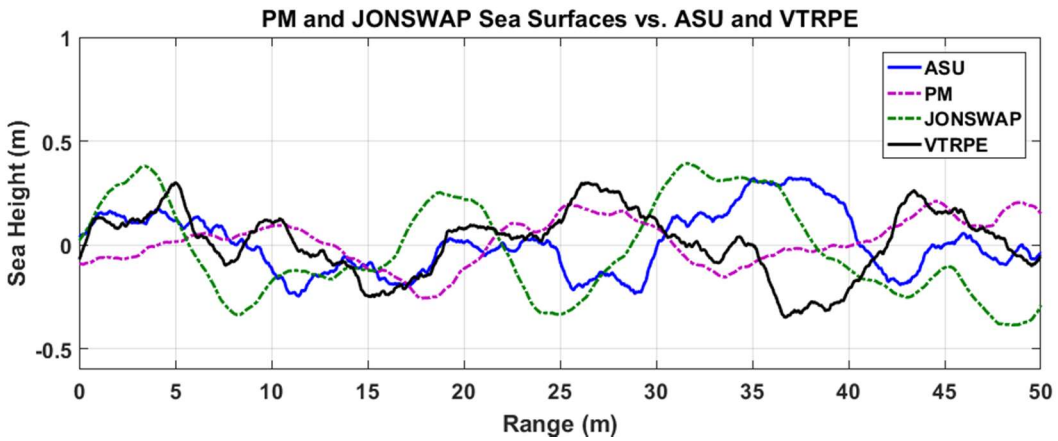

Figure 6.2: Sea surfaces, (blue) ASU, (purple) PM spectrum, (green) JONSWAP spectrum, (black) VTRPE. The ASU and VTRPE spectra show more high-frequency roughness.

In order to extend any of the spectra to create a 3D sea, the wavenumber is simply broken down into k_x and k_y components and the spectrum now has two dimensions instead of one. The spectrum equations don't change, with the exception of a couple of energy scaling factors in the Fourier transform. The randomization of the phase and amplitude remains the same, and the 1D Fourier transform becomes a 2D transform. A symmetric spectrum in line with this straight-forward extension into 3D is shown in Figure 6.3 left.

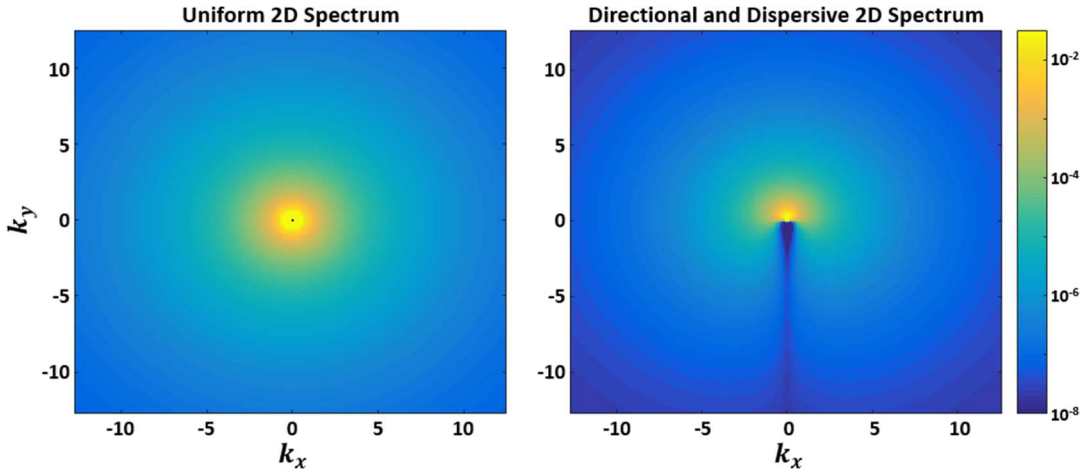

Figure 6.3: Spectra for 3D sea generation, (left) Symmetric of a 2D sea spectrum directly extended into 3D, (right) Dispersive spectrum due to defined and prolonged wind flow over the sea surface.

The symmetric spectrum physically means that if a 2D cut is taken of the 3D surface from anywhere within its plane and in any direction, the spectrum of the 2D cut of sea will

be the same regardless of where it was cut from. This would imply no defined wind direction over the entire sea. In most real cases, the wind blows for reasonably long periods of time over the sea in a more or less constant direction. This will lead naturally to sea waves travelling favorably in the direction of the wind, and results in a dispersive, directionalized, and asymmetric spectrum (Figure 6.3 right).

The dispersion factor [49] used to create the 3D spectrum, such as in Figure 6.3 right, is incorporated as an additional multiplicative factor $D(\omega, \theta)$ to the base spectrum, $S(\omega)$, as shown in eq. 6.3. The resulting sea surface (in 3D) of a dispersive spectrum is shown in Figure 6.4 (bottom), which clearly shows a favored wave direction.

$$S(\omega, \theta) = D(\omega, \theta) \cdot S(\omega) \tag{6.3}$$

$$D(\omega, \theta) = N\big(s(\omega)\big) \cos\left(\frac{\theta_w - \theta}{2}\right)^{2s(\omega)} \tag{6.4}$$

$$N\big(s(\omega)\big) = \frac{1}{2\sqrt{\pi}} \frac{\Gamma(\mathrm{s}(\omega) + 1)}{\Gamma\left(\mathrm{s}(\omega) + \frac{1}{2}\right)} \qquad s(\omega) = 11.5 \left(\frac{g}{\omega_p U_{10}}\right)^{2.5} \left(\frac{\omega}{\omega_p}\right)^{\mu}$$

$$\mu = 5 \;\; for\; \omega \leq \omega_p \;, \;\; -2.5 \; otherwise$$

TABLE 6-2: PARAMETER DESCRIPTIONS FOR THE DISPERSION FACTOR

Parameter	Description
μ	Variance
θ_w	Wind direction
θ	Local direction in the spectrum
$\Gamma(x)$	Gamma mathematical function

As discussed previously, VTRPE uses a modified JONSWAP spectrum in order to generate seas, and then modifies it again to account for the higher-frequency surface waves relevant at X-band. To enable direct cross-validation between our FDTD simulations and VTRPE results (including the relevant high frequency surface texture) it was necessary to

match as closely as possible the VTRPE spectrum. The resulting ASU spectrum shown in Figure 6.1 was obtained by simply taking the inverse Fourier transform of several instances of 2D and 3D ASU seas and modifying the baseline spectra until agreement with the VTRPE spectrum was obtained. Figure 6.5 shows that the inverse FFT of a single instance of VTRPE sea and that of a single instance of the ASU-generated sea are essentially the same, both in 2D and 3D.

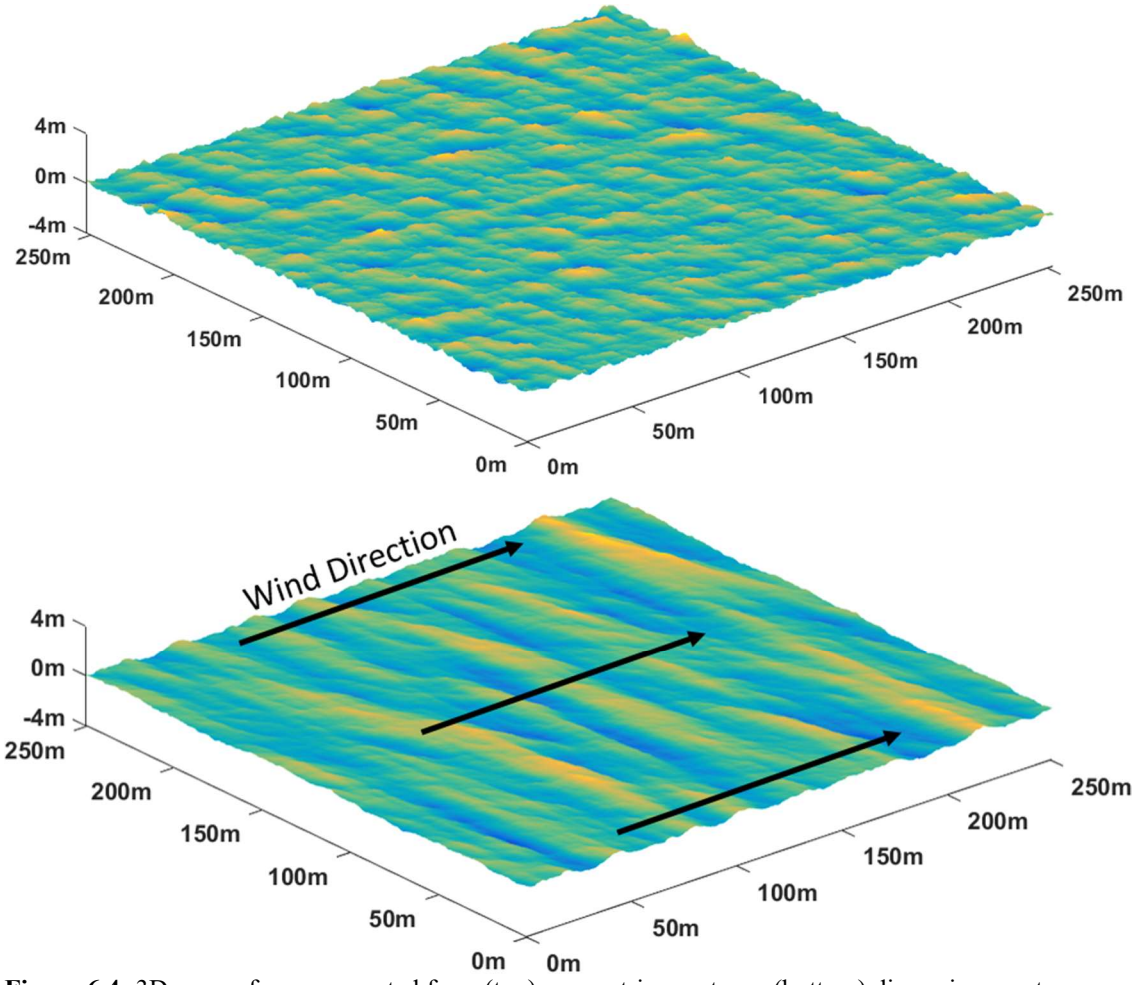

Figure 6.4: 3D sea surfaces, generated from (top) symmetric spectrum, (bottom) dispersive spectrum.

These results demonstrate the capability to generate 2D and 3D sea surfaces of arbitrary dimension, sea state, and wind properties without the need to outsource. This

allows for quicker statistical studies of the sea scatter, such as those with angle of arrival shown later. Additionally, 2D and 3D seas spectrally-identical to VTRPE's 2D sea can be created, allowing for easier cross-comparisons in the future.

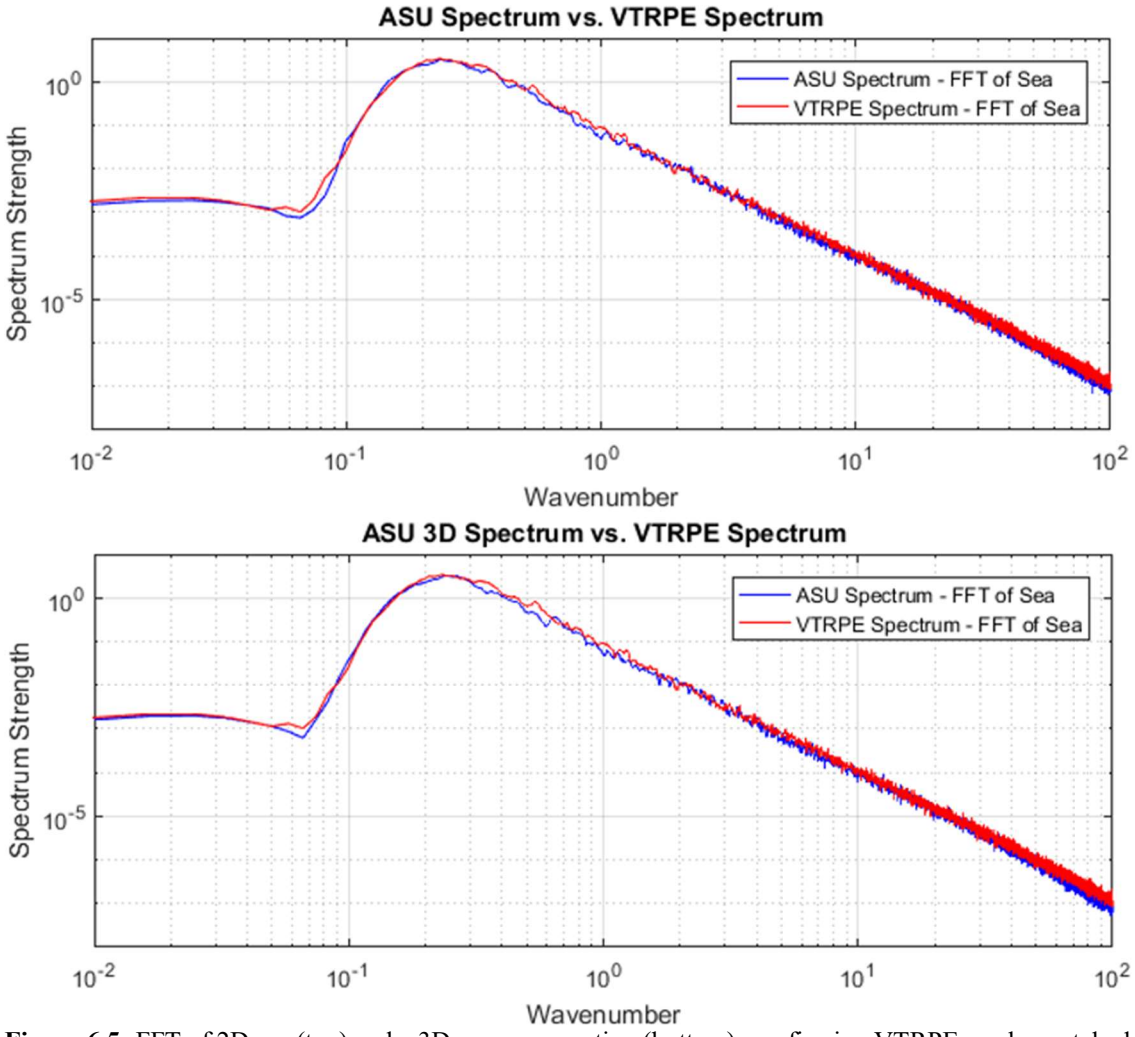

Figure 6.5: FFT of 2D sea (top) and a 3D sea cross-section (bottom), confirming VTRPE can be matched if need be.

6.2 Angle of Arrival and Rough Surface Scatter Comparison of Several VTRPE and ASU 2D Sea Surfaces

More important for practical purposes than the surface spectrum of the sea, is the angular spectrum of scattered electromagnetic energy from the sea. The angular spread of the scatter provides a quantitative measure of the proportion of the scattered energy that is

travelling in a given direction toward the far field region. This angular spread calculation is performed on two surfaces, illustrated in Figure 6.6.

The first calculation is commonly referred to as Angle of Arrival (AoA) and is most relevant to radar scattering from targets in the sea. The AoA is simply a Plane Wave Spectrum (PWS) representation (a Fourier transform into k-space) of the fields along the relatively short, vertical observation (target) plane (Figure 6.6a). Because the target plane is positioned directly in front of a potential sea-based radar target, essentially all the fields crossing the target plane will inevitably interact with the target. A unique practical feature of the AoA is how the vertical plane can be spatially windowed (Figure 6.7) in order to limit the calculation to include only fields between a specific set of heights. This is of critical practical importance in the scattering assessment of smaller structures on a ship which would only be illuminated by the fields within said height range.

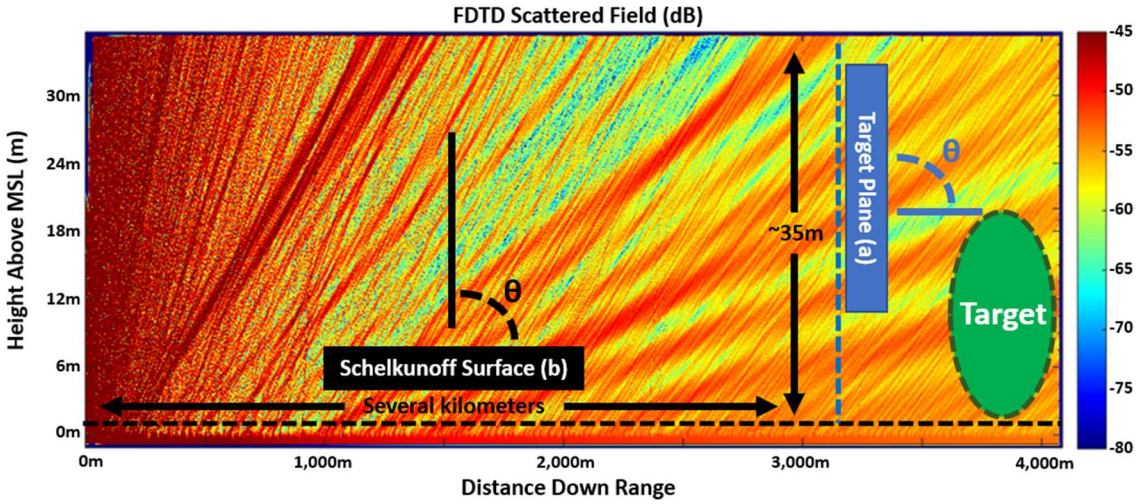

Figure 6.6: Planes to calculate AoA from (a) Vertical observation (target) plane, relevant for the scattering from a target. (b) Schelkunoff surface bounding the entire sea, which captures all scatter from the simulation.

The second method, what we'll call the rough surface scattering (RSS) profile, is another PWS representation of the scattered fields, but is a more complete characterization

of the surface itself because it accounts for the angular spectrum of the far-field scatter from the entire sea surface at once. To calculate the RSS profile, the scattered fields are collected along a bounding surface just above the peak height of the sea (the Schelkunoff surface) and converted to Schelkunoff equivalent currents (Figure 6.6b). The Schelkunoff currents contain all the scattering information for the entire sea at once, they are Fourier transformed into k-space to produce the angular spectrum of the scatter in the far field.

In the following sections, we compare six ASU-generated seas to six VTRPE-generated seas, each 1km long. In every case, a source pulse originating from 20m above mean sea level (MSL) is propagated for 1km. We will show that averaging the AoA and RSS profiles of the six seas produces very similar results between the ASU and VTRPE seas. In order to perform a proper statistical analysis, six sea instances are too few, and should ideally be closer to 50. Unfortunately, at present we have only six sea instances generated by VTRPE. We first look at the AoA results, followed then by the RSS profile.

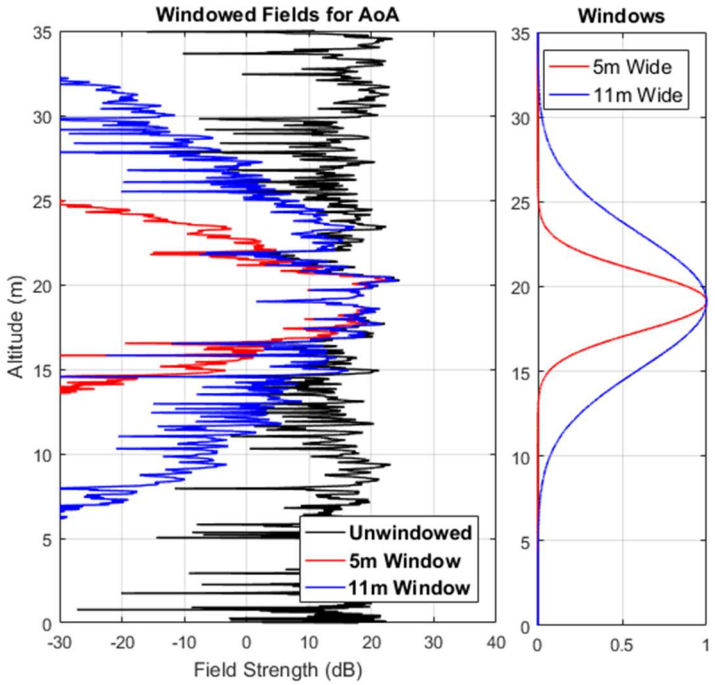

Figure 6.7: Windowed fields used for AoA calculations. Both windows are centered on the source height.

6.3 Angle of Arrival

As previously stated, the AoA is calculated from the scattered fields which physically traverse the target plane in Figure 6.6a. Compared to the Schelkunoff surface, very different information will cross the target plane. Since the target plane extends from the sea up to a low altitude (10's of meters), yet is potentially kilometers down range, the portion of the sea which contributes to scatter crossing this plane decreases in size with higher angles. Scatter travelling at low angles will have interacted with a much longer stretch of ocean before crossing the target plane than scatter at higher angles. High-angle scatter generated far from the observation plane will simply overshoot the target plane and not be captured. Thus, AoA contains high-angle information that is less dependent upon the statistics of the whole sea and more dependent on the local sea surface near the observation plane.

Following this logic in the AoA comparison between different instances of the sea surface, low-angle rays should agree more closely than high-angle rays on the target plane due to the increased dependence on the local sea of high-angle rays, especially when only six VTRPE seas are available for averaged comparison. Conversely, the AoA data near zenith (near 90 degrees) contains data that is extremely dependent upon the local sea features. In Figure 6.8-Figure **6.10**, this is indeed the case for every instance.

Since the sea surface scatters randomly and over a wide range of angles, the unwindowed AoA is littered with peaks and nulls. Additionally, the unwindowed AoA experiences ringing caused by the abrupt discontinuities of the surface currents at the edges of the target plane, which are removed when the spatial window is applied in the other cases. To help visually clarify the plots, Figure 6.10 contains a smoothed version of the un-

windowed and 11m windowed plots from Figure 6.8 and Figure 6.9, respectively. The principal scatter at lower angles agrees very closely between the sea sets and deviates slightly at higher angles.

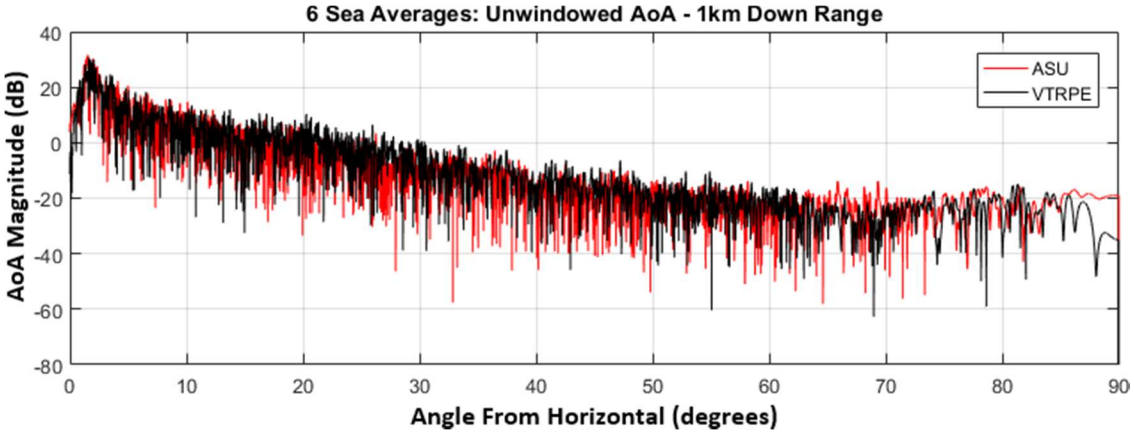

Figure 6.8: AoA 1km down range over target plane with no windowing.

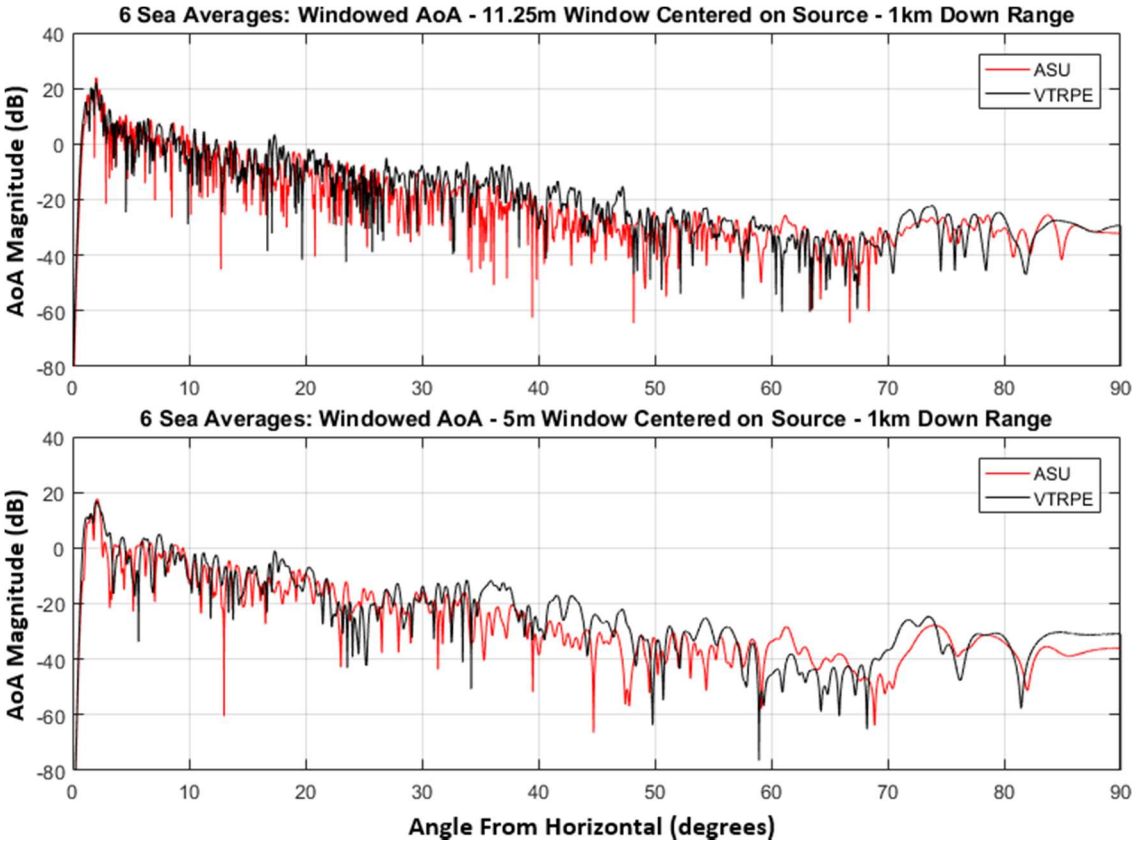

Figure 6.9: AoA 1km down range (top) 11m window, (bottom) 5m window.

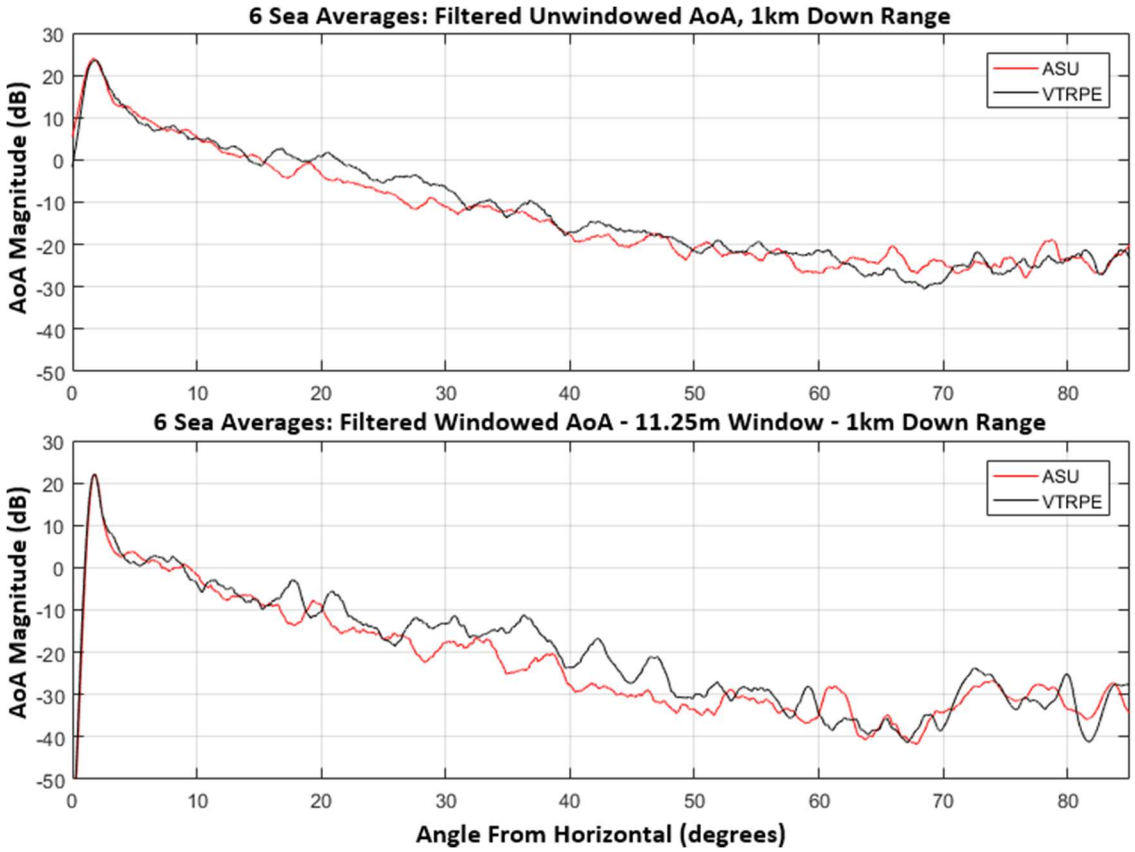

Figure 6.10: AoA 1km down range over target plane, filtered (top) no window, (bottom) 11m window.

6.4 Rough Surface Scattering Profile

In much the same way as the AoA, the RSS profile is calculated from the Fourier transform of the currents along the Schelkunoff surface (Figure 6.6b) to produce the PWS representation of the far-field radiation. In the AoA, only a relatively small portion of sea is responsible for scatter at angles higher than grazing because only the local sea contributes high-angle rays. In the RSS profile on the other hand, scatter at all angles is equally dependent upon the sea as a whole because the PWS representation of the Schelkunoff surface currents removes any locality distinctions in the fields on the surface, and only considers angular directions, summed over the entire surface. Since the full sea surface is represented equally, then comparatively to the AoA, the RSS profile should

require fewer sea averages to yield the same agreement between the ASU and VTRPE seas. This is indeed the case, as is demonstrated in the following figures.

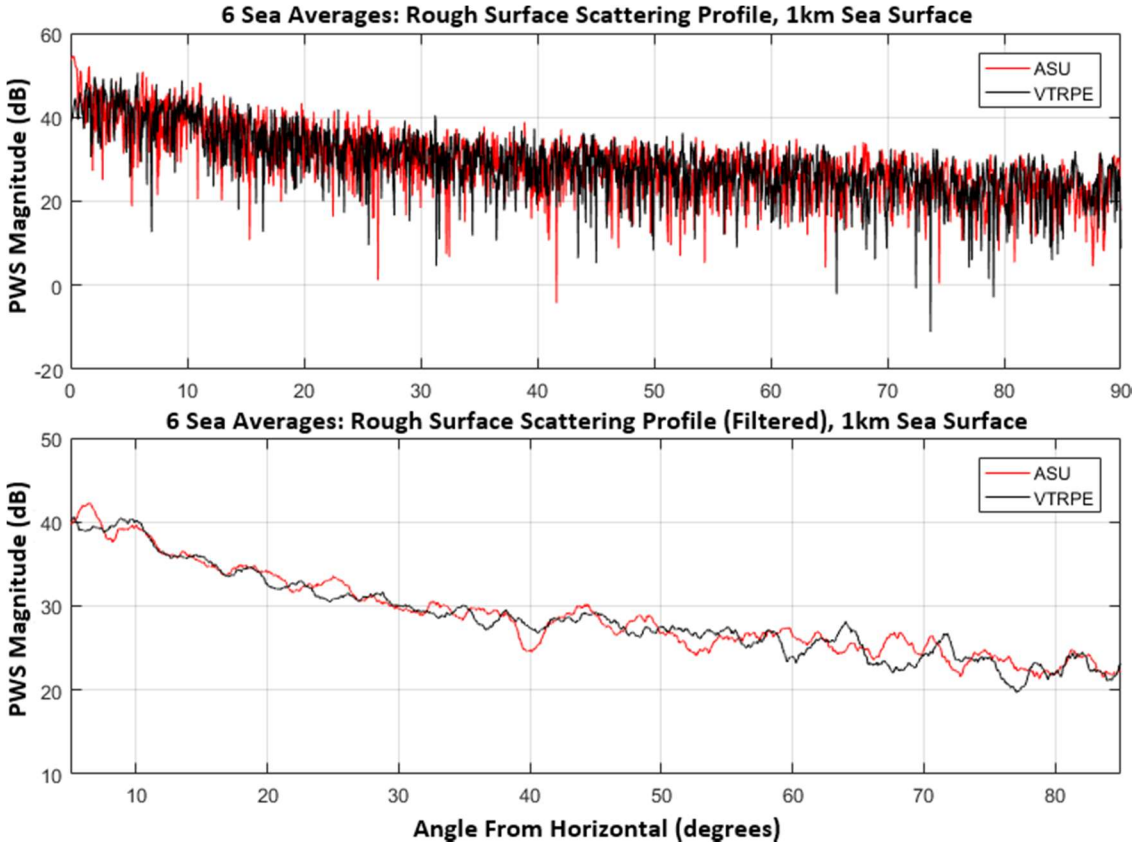

Figure 6.11: RSS profile over 1km of sea calculated from the Schelkunoff surface currents, (top) unsmoothed, (bottom) smoothed for better visualization. Note, the distorted ends of the smoothed data are a result of the smoothing.

For the presentation of the RSS profile, only the forward-travelling scatter propagating between the forward horizon (0 degrees) and zenith (90 degrees from horizontal) will be considered. The extreme domain decomposition method used to simulate the sea spaces in FDTD limits backward scattering information. In Figure 6.11 (bottom), the upper and lower 5 degrees of the RSS profile have been distorted (and omitted) by the smoothing algorithm (5th order Savitzky-Golay filter) which was required in order to eliminate the noisy nature of the data seen in Figure 6.11 (top). The Savitzky-

Golay filter was chosen because it is designed not to average out data but to preserve a signal amidst noise, with the downside of potentially distorting the data endpoints.

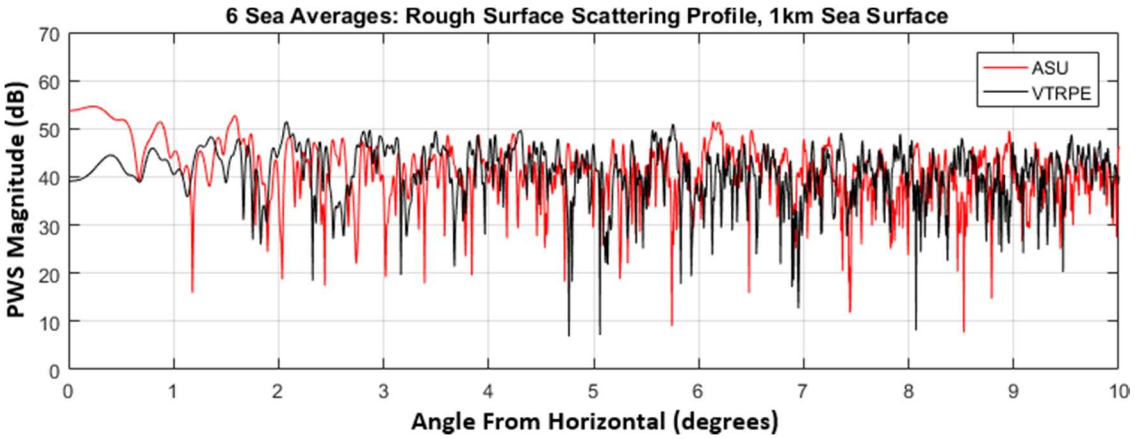

Figure 6.12: Deviation below 1 degree potentially due to grazing incidence.

6.5 Conclusion

The capability to create realistic sea spectra which simulate a wide variety of sea conditions, as well as to generate realistic and arbitrarily-sized 2D and 3D sea realizations from them was implemented in order to avoid the necessity to outsource the creation of the sea surfaces required for 2D and 3D simulation spaces. Creation of VTRPE-like sea surfaces in either 2D or 3D, has also been demonstrated via seas which mimic the VTRPE spectrum, thereby further eliminating the need to outsource seas to perform cross-validations. The AoA and RSS profile averages over the set of 1km seas further demonstrates the comparable characteristics of the in-house surfaces with VTRPE's seas. This similarity is important to obtain because it allows us to create 3D sea spaces (for our subsequent 3D FDTD analyses) which can be characteristically analogous to the 2D spaces used by current tools of the trade, particularly VTRPE, to better highlight the relevance of the 3D phenomenology introduced by the 3D source and sea.

CHAPTER 7

ASSESSING THE PRESENCE OF OUT-OF-PLANE PHENOMENA

7.1 Realistic 3D Sea vs. Extruded 2D Sea

The out-of-plane scattering phenomena in a realistic maritime environment cannot be captured in a 2D or 2.5D simulation space, the physics are simply not present. It's inevitable that the scatter from a 3D sea surface will send energy in all directions. Specifically, when we consider the plane of incidence to be the 2D slice at the center plane of a 3D space, it is statistically conceivable (unavoidable, in fact) that sea features on either side of such a plane will scatter energy back through the center plane, which should lead to "hot" and "cold" spots in the fields. These would not be predicted in the 2D simulation. In this chapter we measure some of the phenomenological differences between 2D and 3D sea surface simulation via the Plane Wave Spectrum (PWS), which breaks down the far-field scatter from the sea into plane waves of different strengths and travelling in different directions. Additionally, we compare the scattered field strengths in the space just above the sea, as well as spatially window the PWS. Two simulation spaces will be considered:

TABLE 7-1: 3D SIMULATION SPACES TO BE COMPARED

Simulation Type	Space Length/Width	Source Type/Height	Sea Type
3D	225m / 5m	3D point dipole, 9GHz y = 0m, z = 20m MSL	Extruded center contour of realistic 3D sea
3D	225m / 5m	3D point dipole, 9GHz y = 0m, z = 20m MSL	Realistic 3D Sea

The first simulation in Table 7-1 is the relevant 2D simulation because, as was shown in Chapter 4, the scattered fields from a 2D sea simulation with a 2D point source

very closely match the scattered fields of a 3D extruded sea simulation with a 3D point source. The shape of the scattered fields is nearly identical and differs in intensity only due to the difference in the Green's function radial dependence of the 2D point source ($1/\sqrt{r}$) compared to that of the 3D point source ($1/r$), as is shown in Figure 7.1. We can thus perform the out-of-plane phenomena evaluation with the two test spaces in Table 7-1. Additionally, the two 3D simulations allow for a full hemispherical mapping of the PWS, rather than simply the 1D line of the 2D simulation. Figure 7.2 and Figure 7.3 illustrate the computational domain and the test sea surface, respectively.

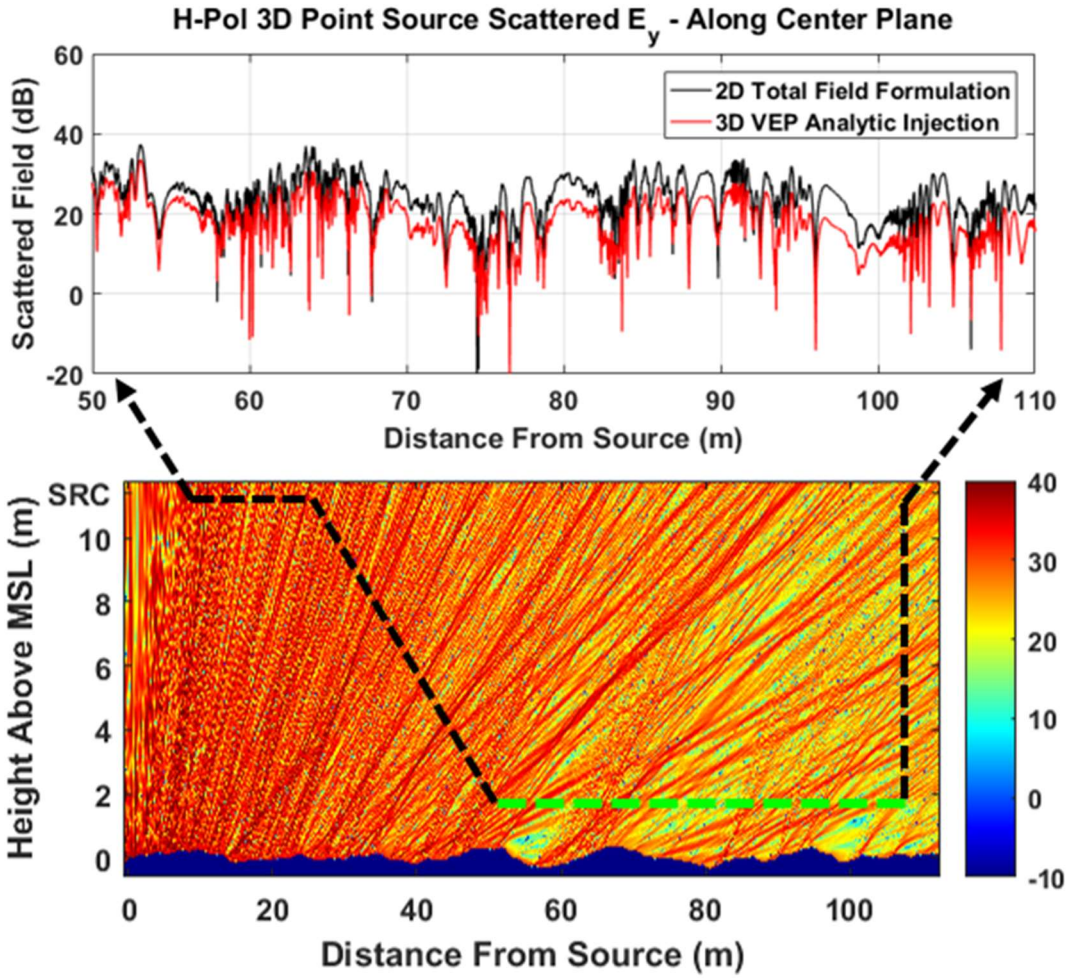

Figure 7.1: (repeated Figure 4.6) The scattered fields, along a horizontal observation plane just above the sea surface, produced by a true 3D source over the extruded sea compared to the 2D total field formulation result. Characteristically similar to the 2D result, varying mostly by the radial dependence.

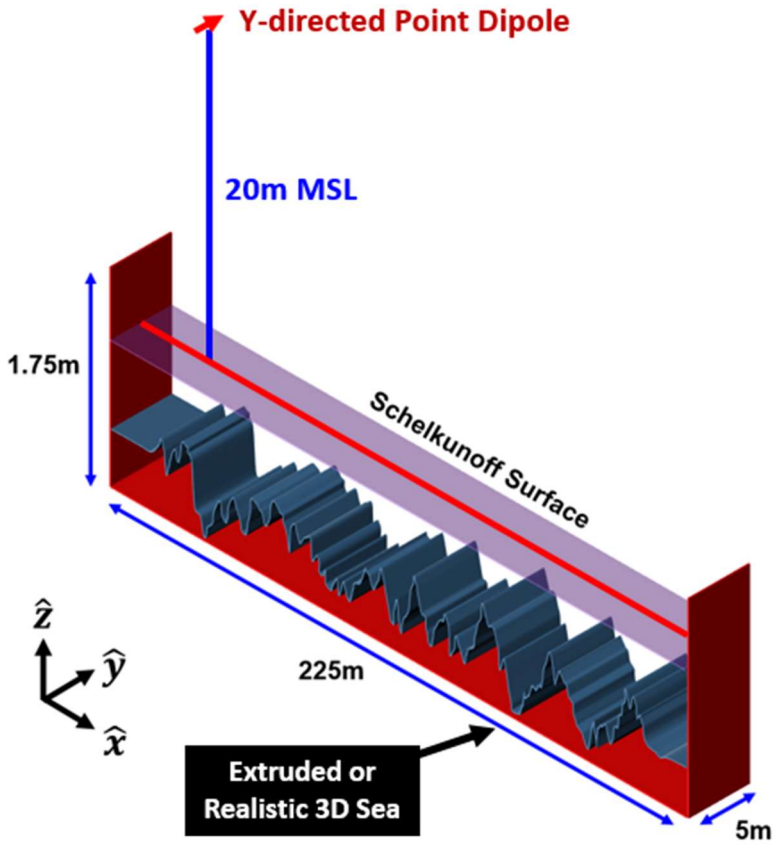

Figure 7.2: Basic layout of 3D domains to be tested, consisting of a corridor of sea and propagation in +x.

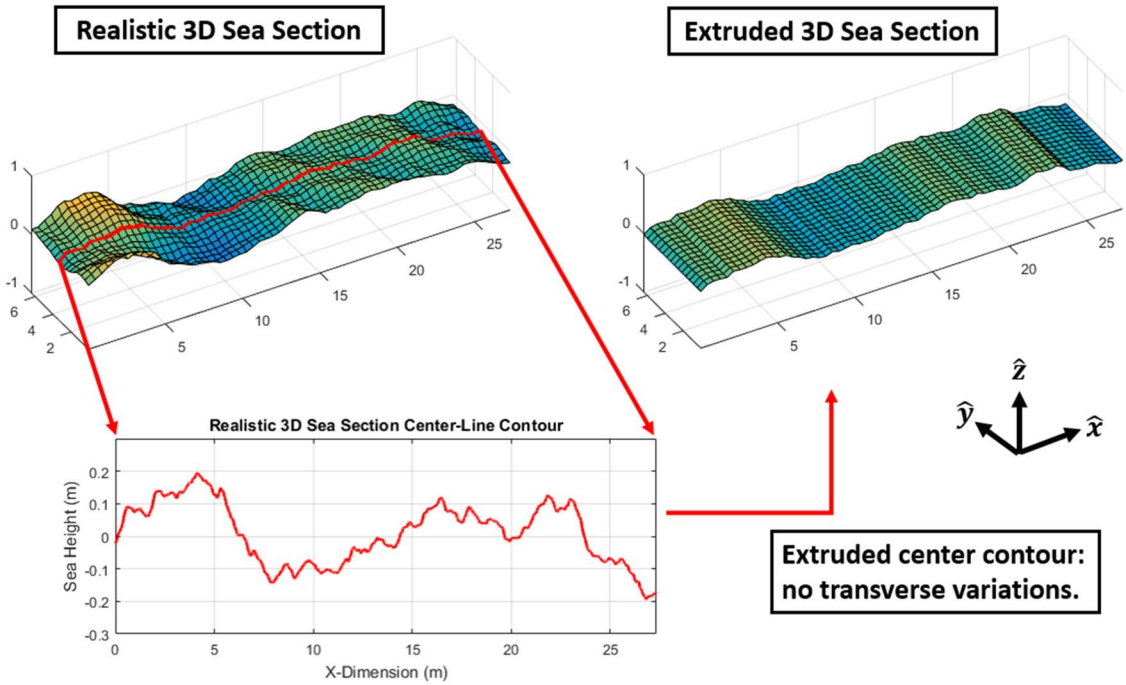

Figure 7.3: The center contour of the 3D sea is extruded to the walls of the 3D corridor to produce the extruded 3D sea, which has no transverse variations.

A y-directed point dipole is placed 20m directly above the center line of the 3D sea surface (Figure 7.2). This center line is the surface contour which will be extruded to make up the surface for the 3D extruded sea case (the "2D sea" case, Figure 7.3). Each simulation space is 225m in length, and each contains a 5m-wide section of sea, with the principal propagation direction along the x-axis.

The surface is excited in both cases by a dispersion-matched 3D analytic source impressed onto the sea using the volume equivalence principle (VEP). With the analytic VEP source, the only FDTD fields produced are the scattered fields. The scattered fields are then collected on a bounding plane above the sea surface (Schelkunoff surface of Figure 7.2) and converted into Schelkunoff equivalent surface currents. These currents are used to produce the PWS. In the next section, we present a brief discussion on the generation of the PWS, followed then by a far field comparison of the TE and TM scatter for both cases presented in Table 7-1.

7.2 Creation of the Plane Wave Spectrum

The output of the FDTD program contains the frequency-domain fields collected over the entire sea on the Schelkunoff surface shown in Figure 7.2. Only the tangential field components are required on an aperture to determine what radiates into space, and the Schelkunoff surface exists in the x-y plane, thus we record the E_x, E_y, H_x, and H_y field components. The first step is to convert the electric and magnetic fields to surface currents. Electric fields are converted to magnetic surface currents (M_s):

$$M_s(x,y) = -\hat{n} \times \vec{E}(x,y) \;\rightarrow\; -E_x(x,y)\,, E_y(x,y) \;\rightarrow\; M_y(x,y)\,, M_x(x,y) \qquad 7.1$$

Here $\hat{n}$ is the normal vector to the Schelkunoff surface, $\hat{z}$. Magnetic fields are converted to electric surface currents (J_s) in equation 7.2.

$$J_s(x,y) = \hat{n} \times \vec{H}(x,y) \rightarrow H_x(x,y), -H_y(x,y) \rightarrow J_y(x,y), J_x(x,y) \qquad 7.2$$

As an additional note, in the FDTD temporal update loop, the E and H fields are separated by half of a time-step in time. If we are to use the Schelkunoff surface fields in conjunction with one another, the H-fields must first be synchronized with the E-fields. Since the Schelkunoff surface contains the frequency-domain currents, this is achieved by simply multiplying the H-fields by the appropriate complex exponential, as shown in equation 7.3.

$$J_s = \left(\hat{n} \times \vec{H}\right) \cdot e^{-\frac{j\omega t}{2}} \qquad 7.3$$

Next, a 2D FFT is performed on the surface currents. This creates spectral electric (Ke) and spectral magnetic (Km) surface currents in terms of the wavenumber components k_x and k_y (note, lowercase k represents wavenumber).

$$Km(k_x, k_y) = FFT[M_s(x,y)] \qquad Ke(k_x, k_y) = FFT[J_s(x,y)] \qquad 7.4$$

In order to radiate these currents into space as spherical TE ($\hat{\phi}$-polarized E) and TM ($\hat{\theta}$-polarized E) waves, it is helpful to convert the spectral currents to spherical coordinates. Since the spectral currents lie entirely within the x-y plane, the relevant radial coordinate on that plane is ρ. By the same token there is no $\hat{\theta}$ component of the spectral currents since $\hat{\theta}$ is normal to the x-y plane. The spectral currents in spherical coordinates are simply decompositions of the x-y currents into longitudinal and transverse traveling currents. (That is, current waves polarized along or perpendicular to their direction of propagation on the plane).

$$Ke_{\phi}(k_x, k_y) = -Ke_x(k_x, k_y)\frac{k_y}{k_\rho} + Ke_y(k_x, k_y)\frac{k_x}{k_\rho} \quad 7.5$$

$$Ke_{\rho}(k_x, k_y) = Ke_x(k_x, k_y)\frac{k_x}{k_\rho} + Ke_y(k_x, k_y)\frac{k_y}{k_\rho} \quad 7.6$$

Here k_ρ is obtained by the typical wavenumber constraint equation $k_\rho = \sqrt{k_0^2 - k_z^2}$. The spectral magnetic surface currents follow the exact same form:

$$Km_{\phi}(k_x, k_y) = -Km(k_x, k_y)\frac{k_y}{k_\rho} + Km_y(k_x, k_y)\frac{k_x}{k_\rho} \quad 7.7$$

$$Km_{\rho}(k_x, k_y) = Km_x(k_x, k_y)\frac{k_x}{k_\rho} + Km(k_x, k_y)\frac{k_y}{k_\rho} \quad 7.8$$

Now that the spectral electric and magnetic currents are known in spherical coordinates, the radiated spherical TE and TM electric fields are obtained through the following expressions.

$$TE: \quad E_{\phi}(k_x, k_y) = -\frac{\omega\mu_0}{2k_z} \cdot Ke_{\phi}(k_x, k_y) + \frac{1}{2}Km_{\rho}(k_x, k_y) \quad 7.9$$

$$TM: \quad E_{\theta}(k_x, k_y) = -\frac{\omega\eta_0\varepsilon_0}{2k_z} \cdot Km_{\phi}(k_x, k_y) - \frac{\eta_0}{2}Ke_{\rho}(k_x, k_y) \quad 7.10$$

It's important to note that the "slow" waves ($k_\rho > k_0$) are not radiated into the far field. The TE and TM far fields are only radiated when $k_\rho \leq k_0$, which limits the range of observable k_x and k_y to match this condition. We can only observe radiation within the unit circle of so-called "visible space".

A brief note on the plotting of far-field contour plots in k-space, the spatial directions θ, ϕ of a far field ray are directly connected to the k-vector components k_x, k_y, 1 to 1. To every point inside the unit circle $k_x^2 + k_y^2 \leq k_0^2$, corresponds the direction of propagation.

$$\phi = \text{atan}\left(\frac{k_y}{k_x}\right) \qquad \theta = \text{asin}\left(\frac{\sqrt{k_x^2 + k_y^2}}{k_0}\right) = \text{asin}\left(\frac{k_\rho}{k_0}\right) \tag{7.11}$$

Therefore a contour plot in the k_x-k_y plane is a circular region truncated at the $k_\rho = k_0$ radius where the conventional angle ϕ in azimuth is measured from the k_x axis and the radial distance from the origin k_ρ/k_0 is the sine of the angle θ, measured from the vertical; such that at the center of the circle we have zenith and at the perimeter of the circle we have the horizon. The symmetry plane that is usually modeled in 2D solutions is the $k_y = 0$ plane. A contour plot in k_x-k_y space is what would be seen if we were to plot the strength of the given function (say, amplitude of $E_\phi(\theta, \phi)$) on the surface of a hemisphere of unit radius and then observed that hemisphere directly from above. This is illustrated in Figure 7.4 below.

A more intuitive picture is demonstrated in Figure 7.5, in which the Schelkunoff surface fields are shown to be radiated onto the upper far-field hemisphere. The upper

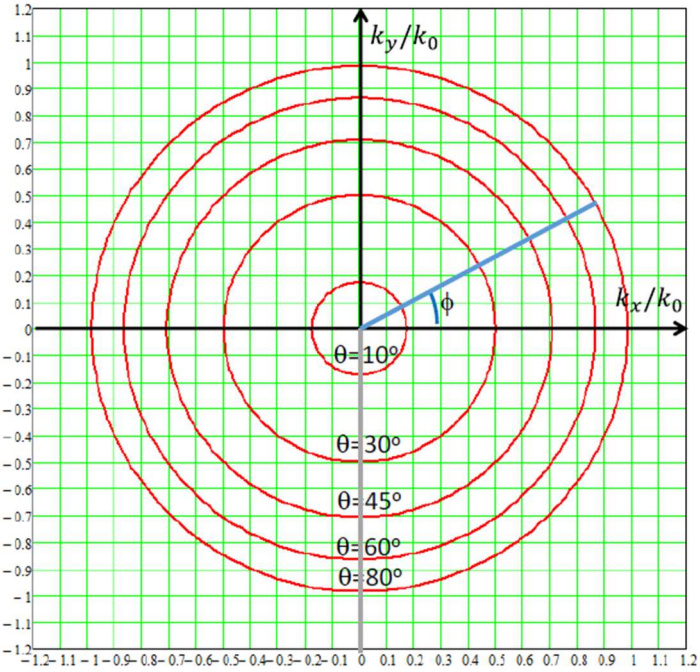

Figure 7.4: Definition of k-space for the far field plots.

hemisphere is then collapsed downward onto a flat plane with k_x and k_y as axes (normalized to the free space wavenumber, k_0). Every (k_x, k_y) pair represents a unique propagation direction in the far-field hemisphere, with x as the direction of forward propagation down the length of the corridor, and y as the transverse propagation direction.

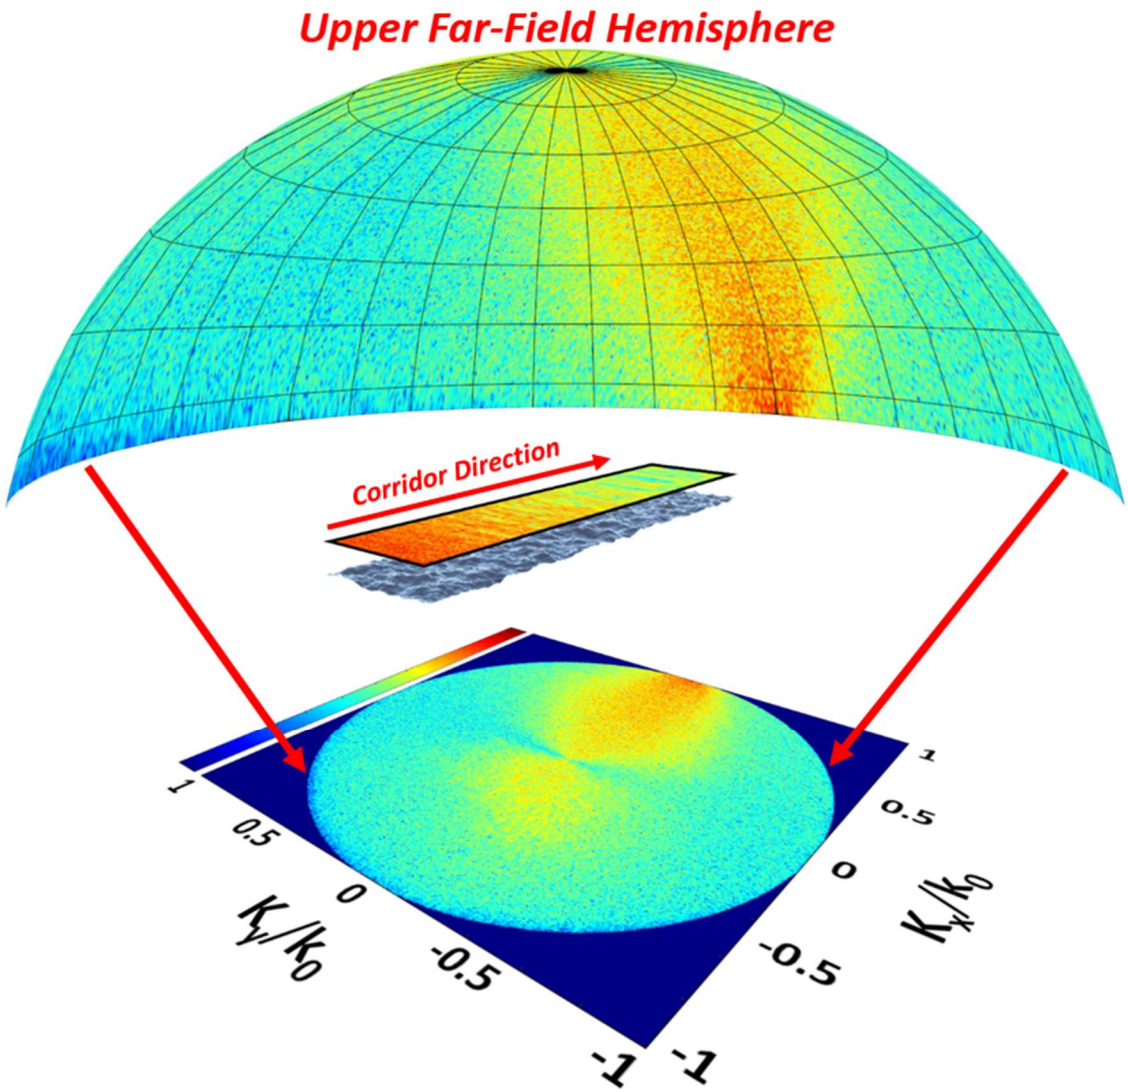

Figure 7.5: The upper far-field hemisphere is collapsed onto a flat plane in k-space.

7.3 Plane Wave Spectrum: Realistic Sea vs. Extruded Sea

In a 2D simulation space (and by extension very nearly for the 3D extruded sea space) the center plane contains the full extent of the scattering information. However, it

has been reported in the literature [9] that 3D features on the sea surface may be responsible for strong flashes or field nulls being pointed toward a target in the sea. In order to see such 3D phenomenology, we will explore two facets: a) how the two simulation cases differ significantly in the information contained on the center plane, and b) show what information is not captured by 2D outside of the center plane. The second facet is of particular interest to the radar scattering problem because the information just outside of the center plane is immediately relevant to scattering from a realistic 3D distant target.

Before diving into the PWS results, it is important to articulate what is to be expected. Since, as shown in Figure 7.2, the source radiator is a y-directed dipole which radiates zero $\hat{\theta}$-directed (TM) E-field in the x-z plane, the fields are purely $\hat{\phi}$-directed (TE) E-fields in this plane. Next consider that the 3D corridor of the simulation extends in the x-direction and follows the x-z plane very closely. Combining those facts indicates the sea surface in the simulation corridor will be very strongly excited by TE fields and very weakly excited by TM fields, so the TE fields should (and do) dominate the far field.

The contour plot of Figure 7.6 shows this result in the polar plot described in Figure 7.4 and Figure 7.5. Note the k_x axis is on the vertical axis, and k_y is represented on the horizontal axis. Thus, the forward horizon, the direction of radiation towards the distant target is the bottom edge of the circles shown.

Even though the scattering of Figure 7.6 represents a "2D sea" we see energy scattered over a band of angles to either side of the symmetry plane. This is because the source is a 3D source of spherical waves, thus the extruded sea surface is illuminated not only along the center plane but also at a variety of angles ϕ from the center planes. With an extruded sea, invariant in the y-direction, we would expect the principal angle of scatter

off every feature to be equal to the angle of incidence. (This is the rule from diffraction theory: the interior angle of the diffraction cone off an edge is equal to the angle of incidence.) Thus, radiation impinging on the sea at high ϕ angles should excite scattering also at high ϕ angles.

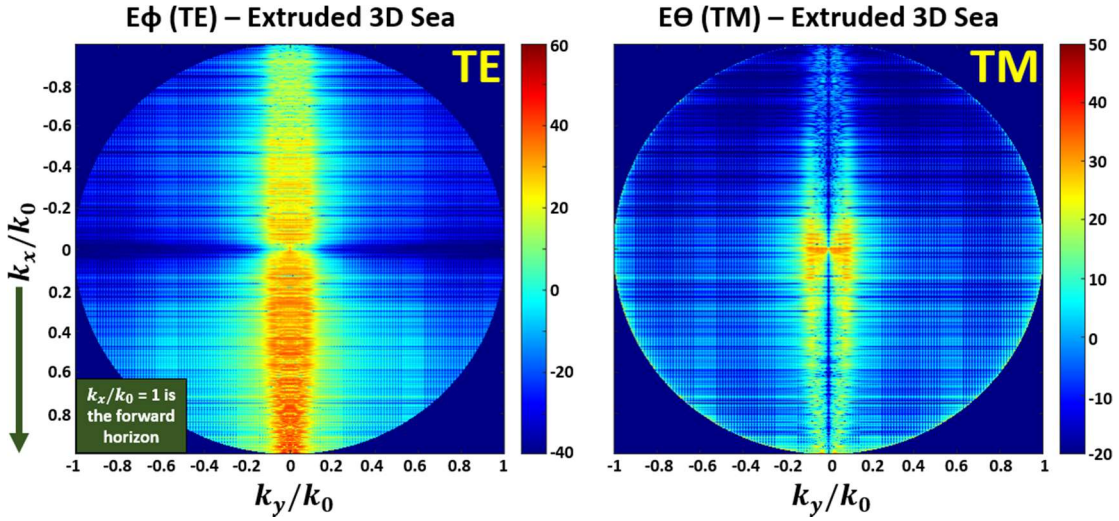

Figure 7.6: TE and TM scattered fields from the extruded 3D sea. NOTE: TM color scale is weaker.

The reason the scattering plot appears restricted to a finite band in k_y about the symmetry plane, $k_y = 0$, (a narrow band of angles in ϕ that is wider near zenith and narrower near the horizon) is a result of the finite transverse extent of the FDTD domain. The portions of the source radiation and its scatter traveling at high ϕ angles simply exit the domain at the truncation boundaries. Near the source, that scattering occurs near the $k_y = 0$ plane and some of those rays are captured, but as the energy travels down range there are less and less source and scattered rays traveling at large ϕ angles to excite the corresponding scattering. Thus, the appearance of the plot in Figure 7.6 is as expected. The case of a true 3D sea is altogether different. The surface is not invariant in the y-direction but rather is rough. Therefore, every ray impinging on a point of the surface can induce radiation in all directions (and all polarizations.)

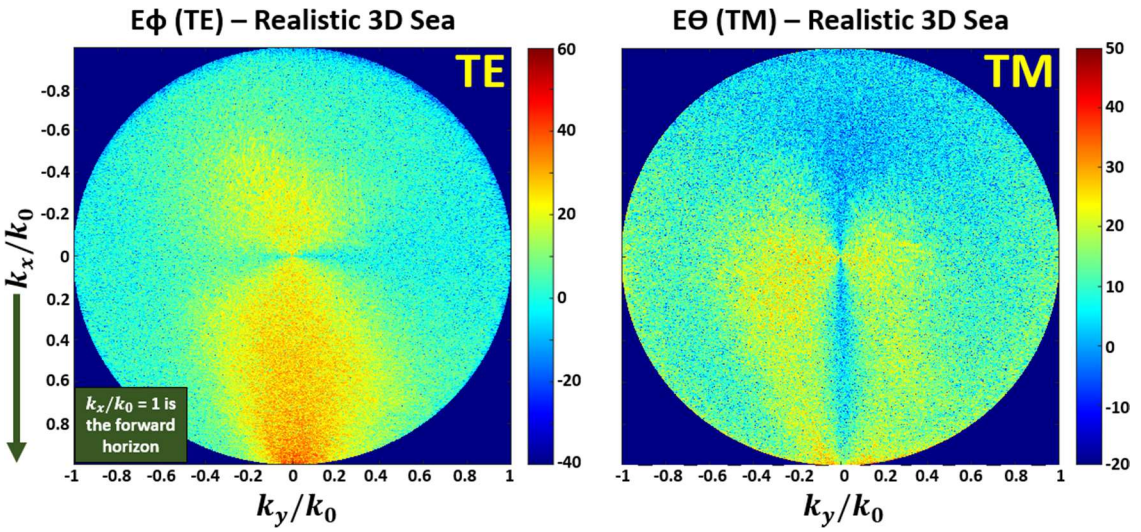

Figure 7.7: TE and TM scattered fields from the realistic 3D sea. NOTE: TM color scale is weaker.

Since the extruded sea is symmetric about the x-z plane, and since there is no TM source in the problem, there can be no TM scatter at all above the extruded sea at any point where $k_y = 0$. (The symmetry plane simulated in 2D solutions.) Conversely, the TE source fields are very strong in the x-z plane, it is no surprise the TE scatter is strongest near $k_y = 0$. The realistic 3D sea does not possess such symmetries, and scatters energy in a much wider angular range (Figure 7.7).

For both the realistic and extruded seas, the strongest direction of scatter is towards the forward horizon (the edge of the unit circle) because the length of the sea is much greater than the height of the source, so the majority of the total sea surface is illuminated at very shallow angles by the source, and thus the scatter inherits the shallow angle characteristics as well. We saw this in the Rough Surface Scatter plots in Chapter 6, where the scattering toward the horizon is 20dB higher than scattering to zenith. The reason the plots show more energy towards the forward horizon than the backward horizon is simply because the FDTD domain extends forward from the location of the source. We are not interested in energy radiated by the source away from the target. (It is for this same reason

that the domain decomposition method used in our work focuses on tracking the forward traveling scatter.)

As expected, Figure 7.7 shows that the realistic 3D sea's transverse variations lead to a much more random k-space spectrum with a significant portion of energy travelling at a much wider range of angles. Although the TE scatter still dominates, the diffraction and re-polarization by off-plane tilted surfaces of the sea produces significantly more TM radiation than the extruded sea. Next, the TE and TM cases are presented individually.

7.4 Plane Wave Spectrum – TE

In this section we take a closer look at the dominant TE fields. In Figure 7.9 the 2D plot of the whole scattering hemisphere is successively zoomed in closer toward the forward horizon. As the scale is zoomed toward the horizon in, the forward-travelling scattering features of each sea surface appear more clearly. Both seas have their scatter concentrated around $k_y = 0$. The scatter from the realistic sea is much more random than the extruded sea. The extruded sea scatter is completely symmetric about $k_y = 0$. In the contour plots, it's possible to tell that the scatter from both seas has approximately the same

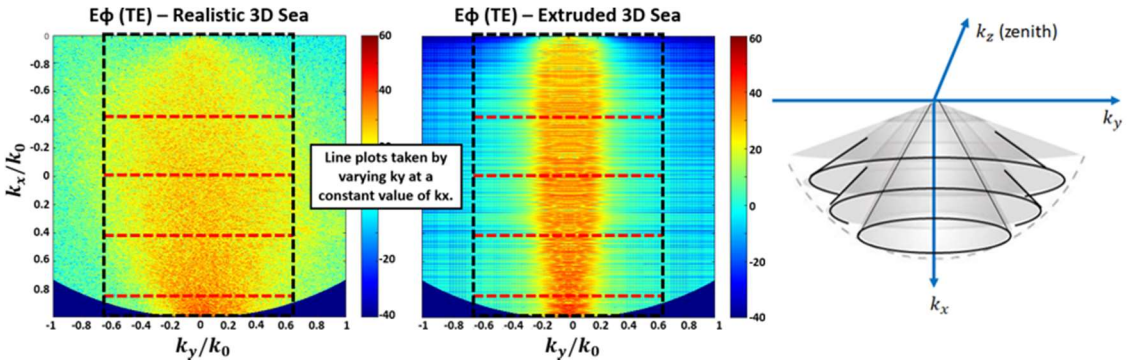

Figure 7.8: Example cuts of the 2D colormaps used for line plots. NOTE: not the actual cuts used in the following figures.

peak magnitude, but it's hard to precisely quantify. To better observe this, cuts are taken at a constant k_x and are plotted against k_y. Such cuts are equivalent to far field "conical cuts" centered on the x-axis. Examples of how the cuts will be made are shown in Figure 7.8.

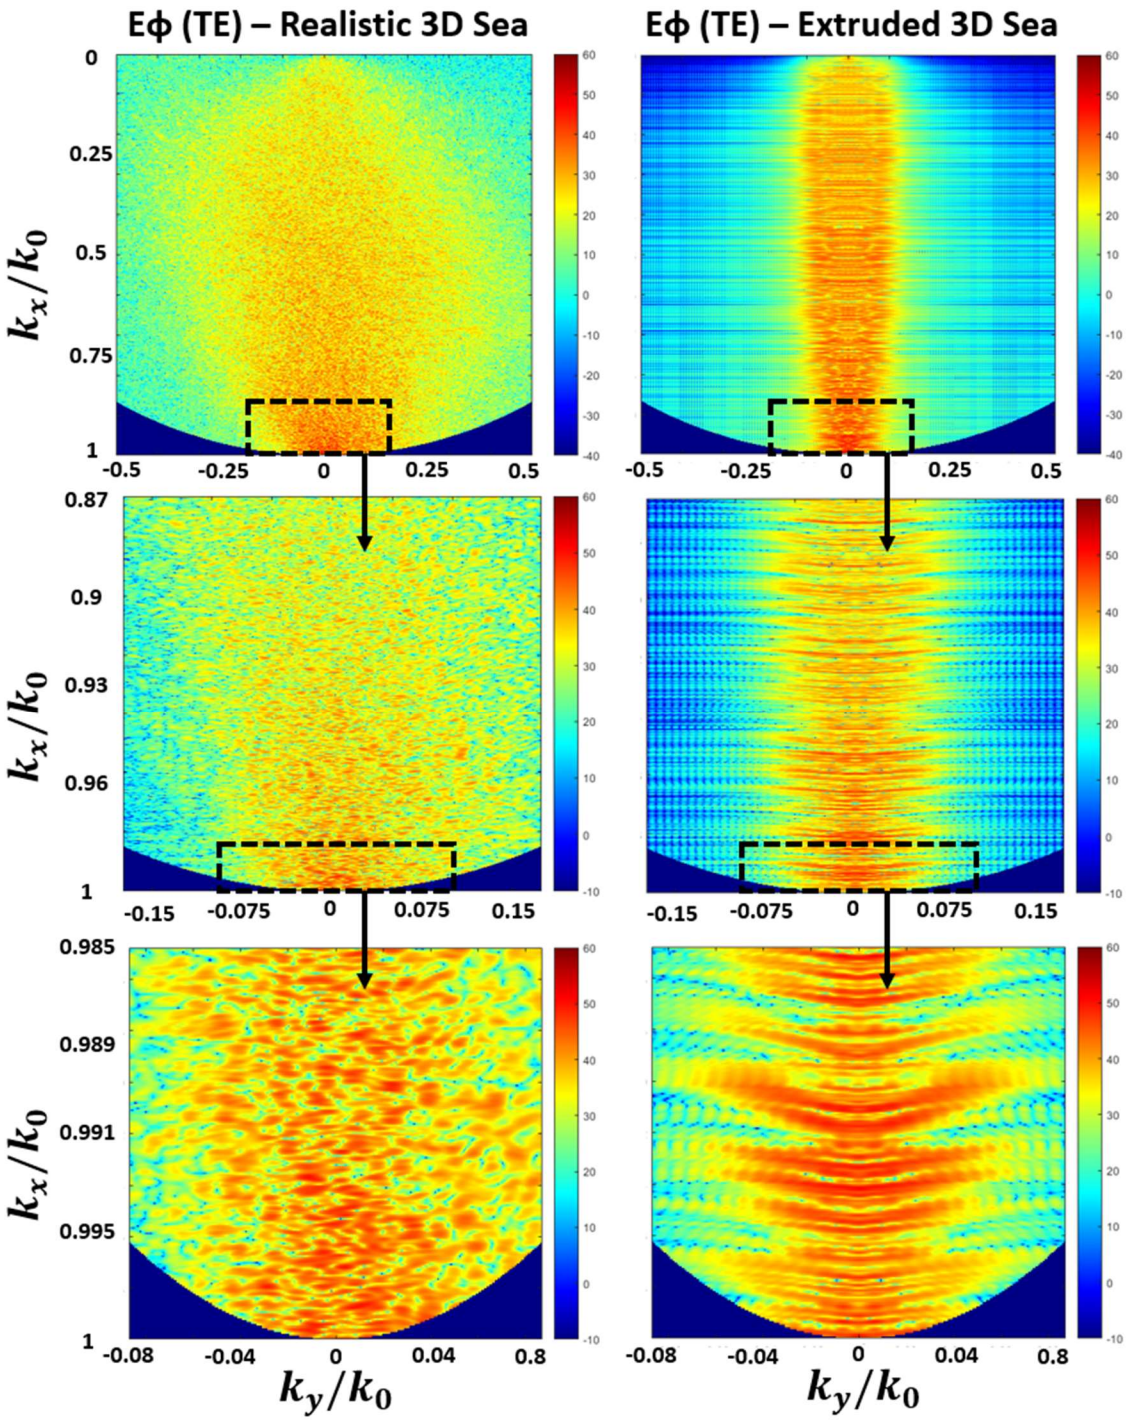

Figure 7.9: TE Scattered fields for realistic and extruded seas, successively zooming in toward the horizon.

Figure 7.10 shows a comparison of the conical cuts with interior cone angles of 70 and 50 degrees. That is, on the x-z plane this cut crosses at $\theta = 20^o, 40^o$ from zenith, respectively, and they are so labelled.

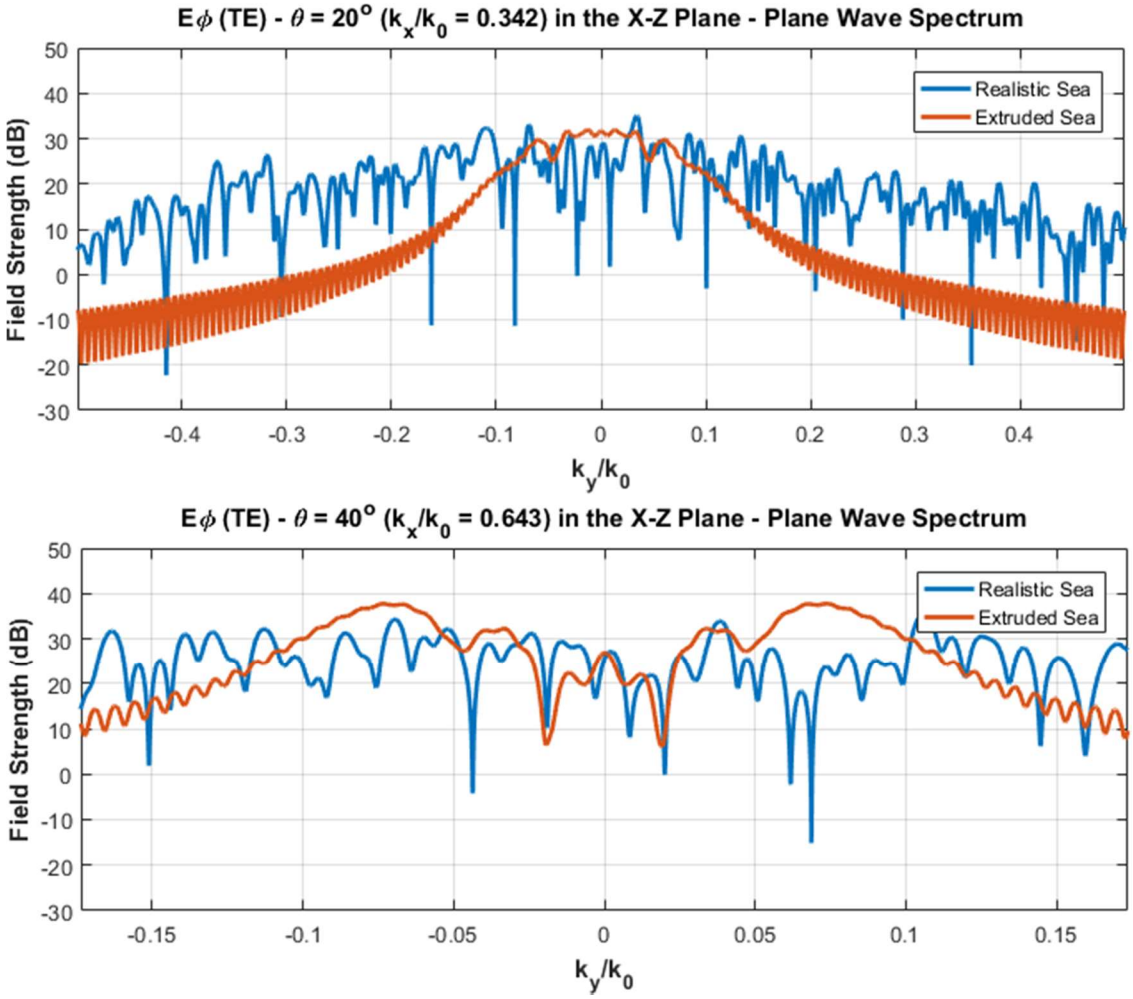

Figure 7.10: Top: Conical cut of TE at 70° off the x-axis ($\theta = 20^o$), Bottom: Conical cut of TE at 50° off the x-axis ($\theta = 40^o$).

These particular angles aren't important for scattering from a target far down range since a ray generated at such an angle will overshoot a target far away. The importance of the data in Figure 7.10 is three-fold. First, it shows that the realistic and extruded seas differ in strength along the center (x-z) plane. Secondly, there are both high-intensity flashes and very deep nulls along directions very near to the center plane. Lastly, there is a very large

amount of energy scattered by the realistic sea outward along the y-direction, transverse to the direction of propagation along the 3D simulation corridor.

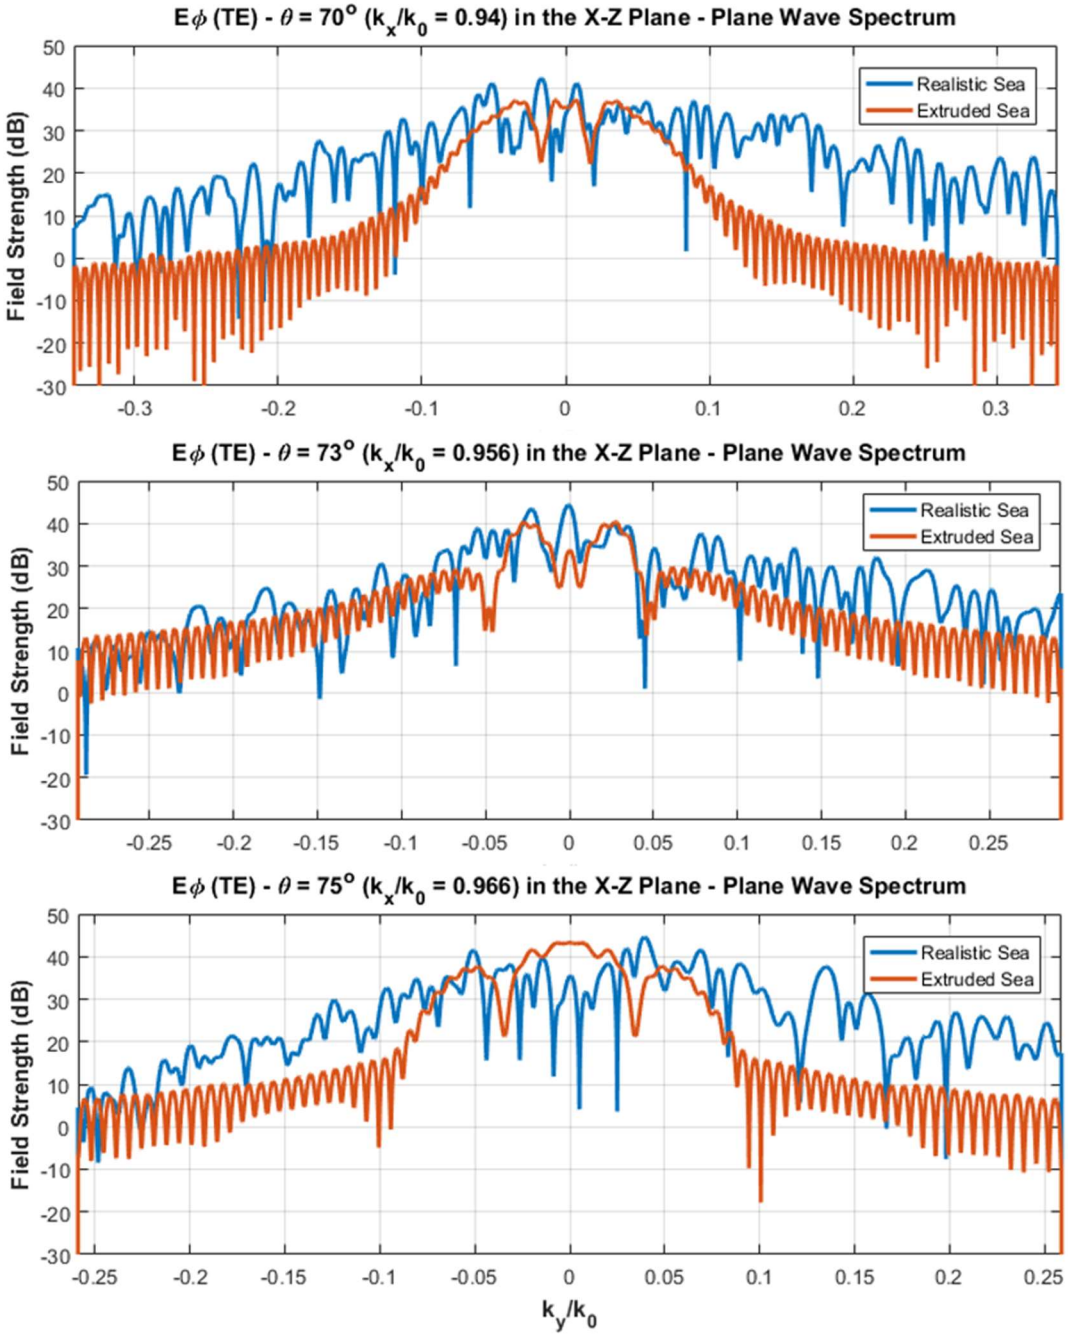

Figure 7.11: Conical Cuts of TE at 70-75° in the x-z plane. NOTE: Changing ky scale.

Figure 7.11 includes three cuts, crossing the x-z plane at θ angles of 70, 73, and 75 degrees. The same effects seen at $\theta = 20^o, 40^o$ continue to occur as the forward horizon

is approached. A look at the center plane shows the scatter is nearly equal at 70 degrees, but nearby at 73 degrees the realistic sea generates a flash 10dB stronger than the extruded sea, and nearby again at 75 degrees the realistic sea scatter is instead 10dB weaker.

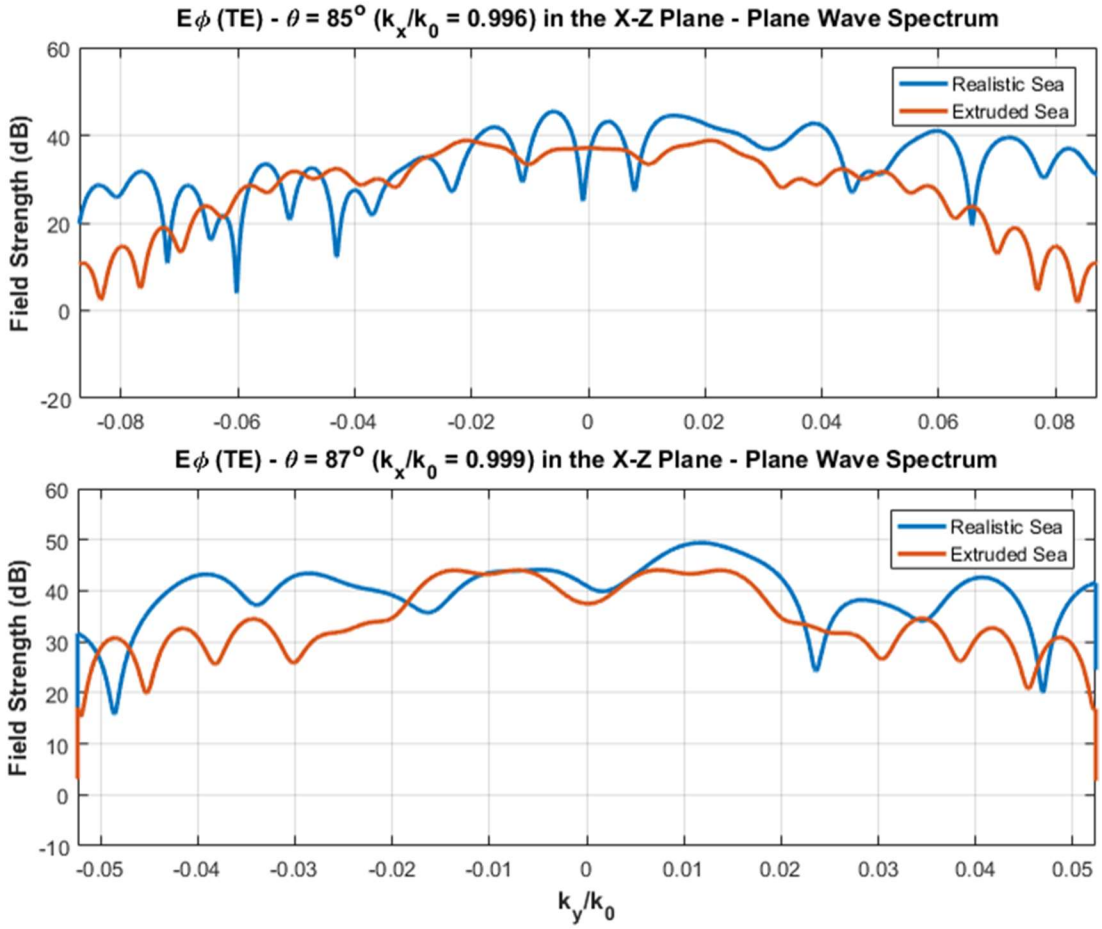

Figure 7.12: Conical Cuts of TE at 85° and 87° in the x-z plane, approaching grazing. NOTE: Changing ky scale.

Figure 7.12 contains cuts approaching very close to the horizon, which continue to show the same types of variations between the two seas. First, it shows that the realistic and extruded seas differ in strength along the center (x-z) plane. Secondly, there are both high-intensity flashes and 20dB deep nulls along directions very near to the center plane. Lastly, there is a very large amount of energy scattered by the realistic sea outward along the y-direction, transverse to the direction of propagation along the 3D simulation corridor.

7.5 Plane Wave Spectrum – TM

The TM fields are unique to the 3D simulation space since the y-directed electric dipole source of our test spaces produces zero TM field in 2D. This is echoed in the extruded sea case by the fact that along the center ($k_y = 0$) plane the TM fields are exactly zero. Since an actual 3D point source is used in both 3D simulations, very weak TM source

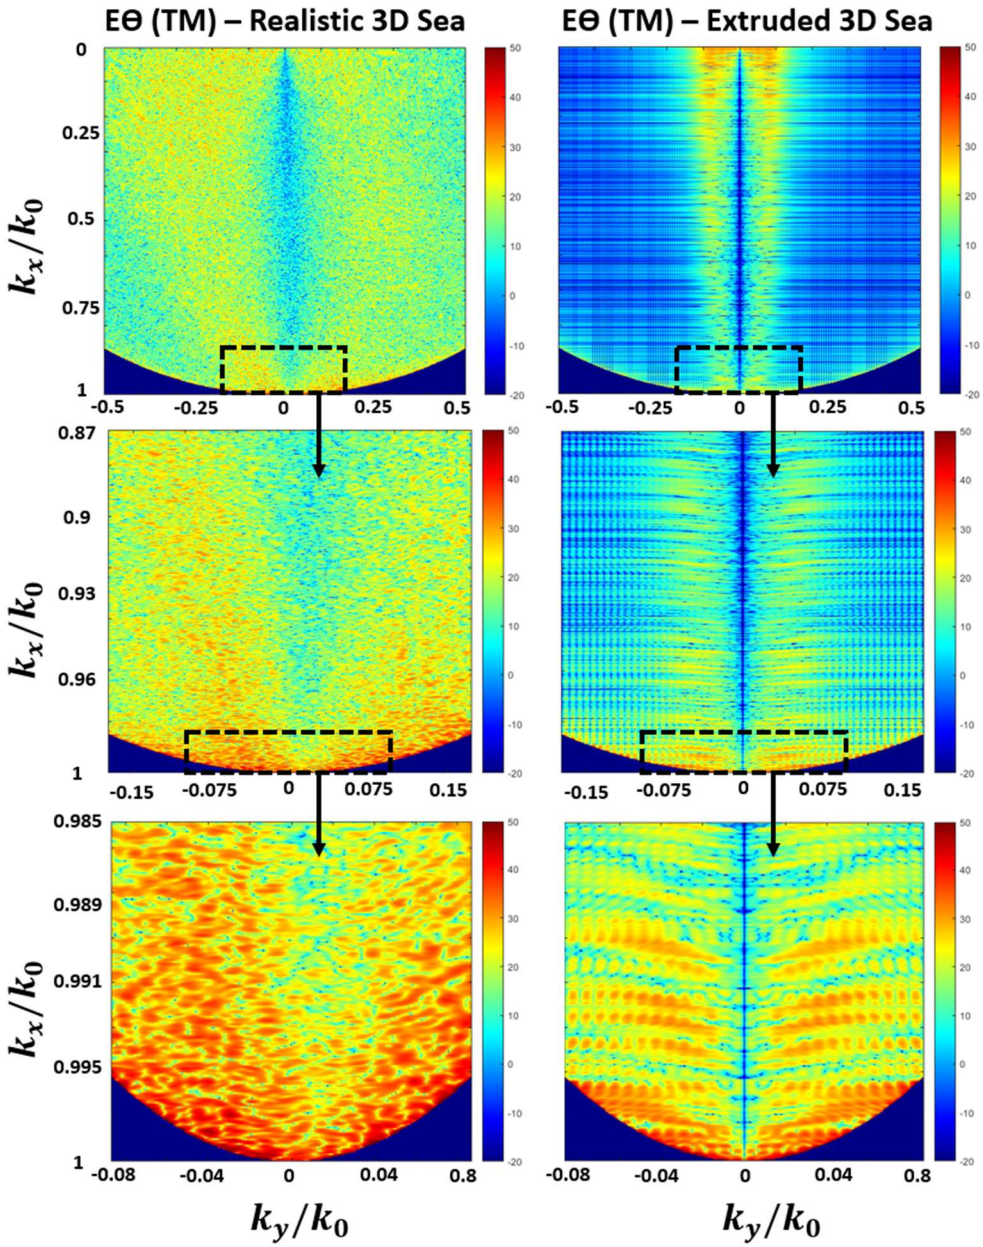

Figure 7.13: TM Scattered fields for realistic and extruded seas, successively zooming in toward the horizon.

fields exist outside of the center plane for both the realistic and extruded sea cases. Figure 7.13-Figure 7.15 showcase one of the most drastic differences between the realistic and extruded sea cases, which is the relevantly strong presence of TM scatter from the realistic sea, most crucially toward the forward horizon.

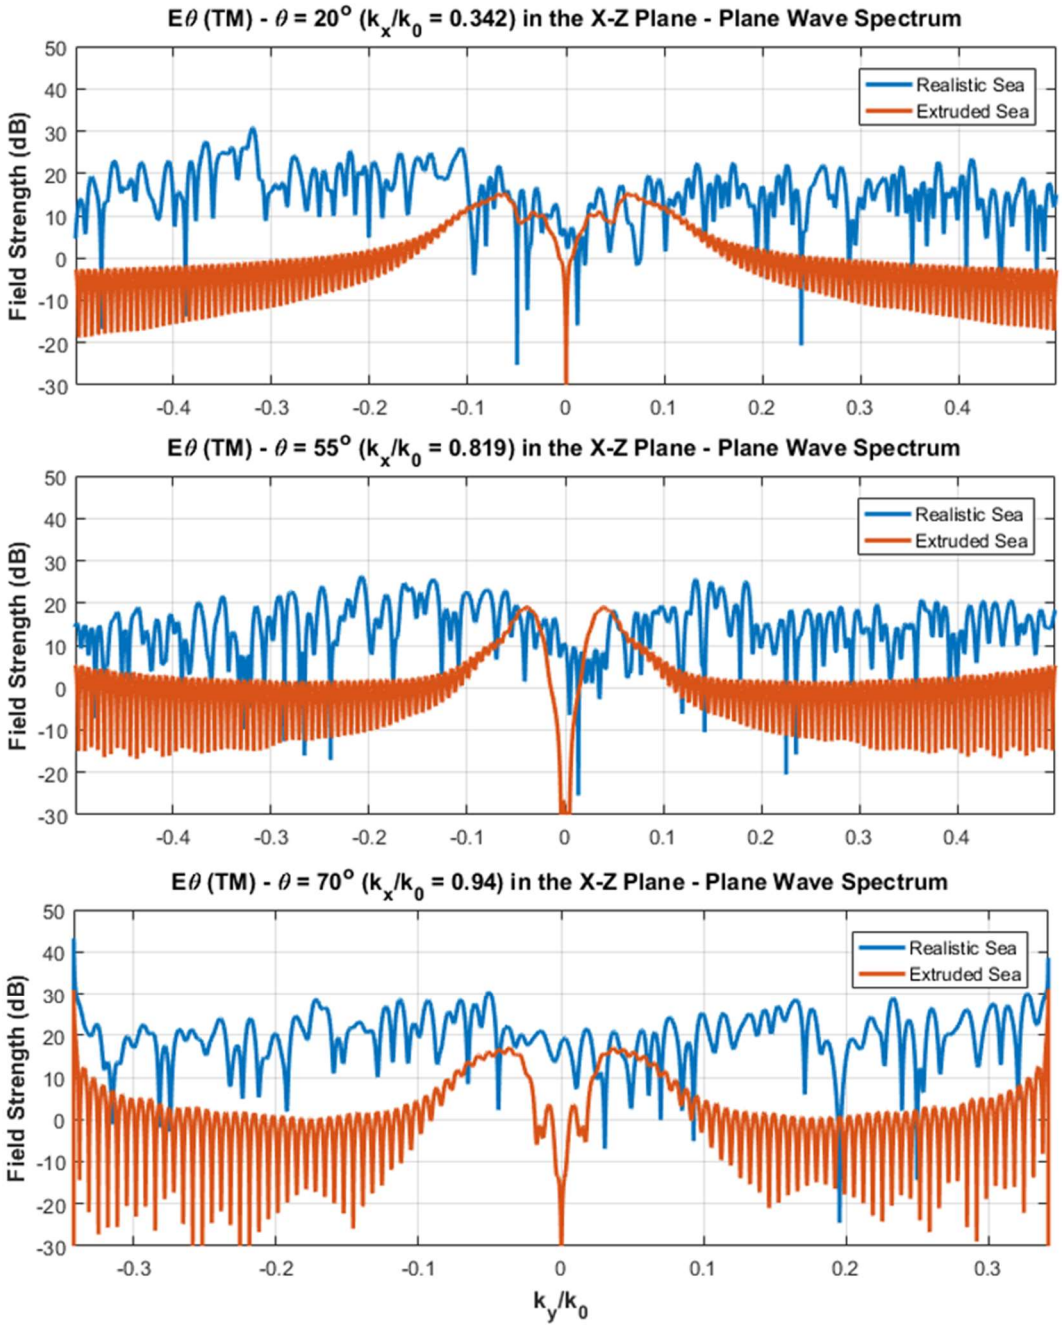

Figure 7.14: Conical Cuts of TM at θ angles between 20-70° in the x-z plane. NOTE: Changing ky scale.

The bottom two pictures in Figure 7.13 show the dramatic presence of TM scatter above the realistic sea in the direction of the forward horizon – the direction of a radar target. Figure 7.14 shows a deep null for the extruded sea around $k_y = 0$ for all angles in the x-z plane. However, there is increasingly strong TM scatter for the realistic sea as angles in the x-z plane approach the horizon.

Close to the forward horizon (Figure 7.15), the TM scatter near $k_y = 0$ is within 10dB of the TE scatter. Additionally, analogous to the TE case, the realistic 3D sea scatters TM field transversely in the y-direction much more strongly than the extruded sea.

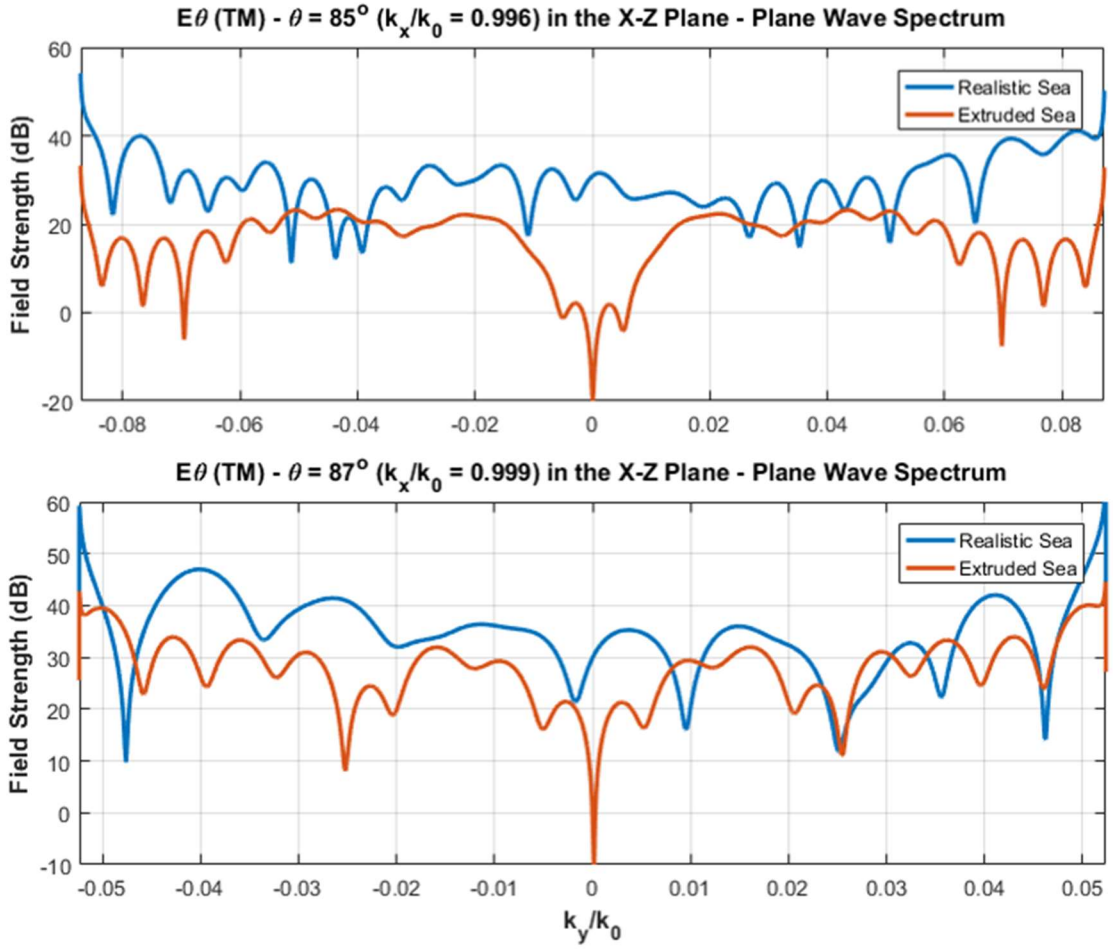

Figure 7.15: Conical Cuts of TM at 85° and 87° in the x-z plane, approaching grazing. NOTE: Changing ky scale.

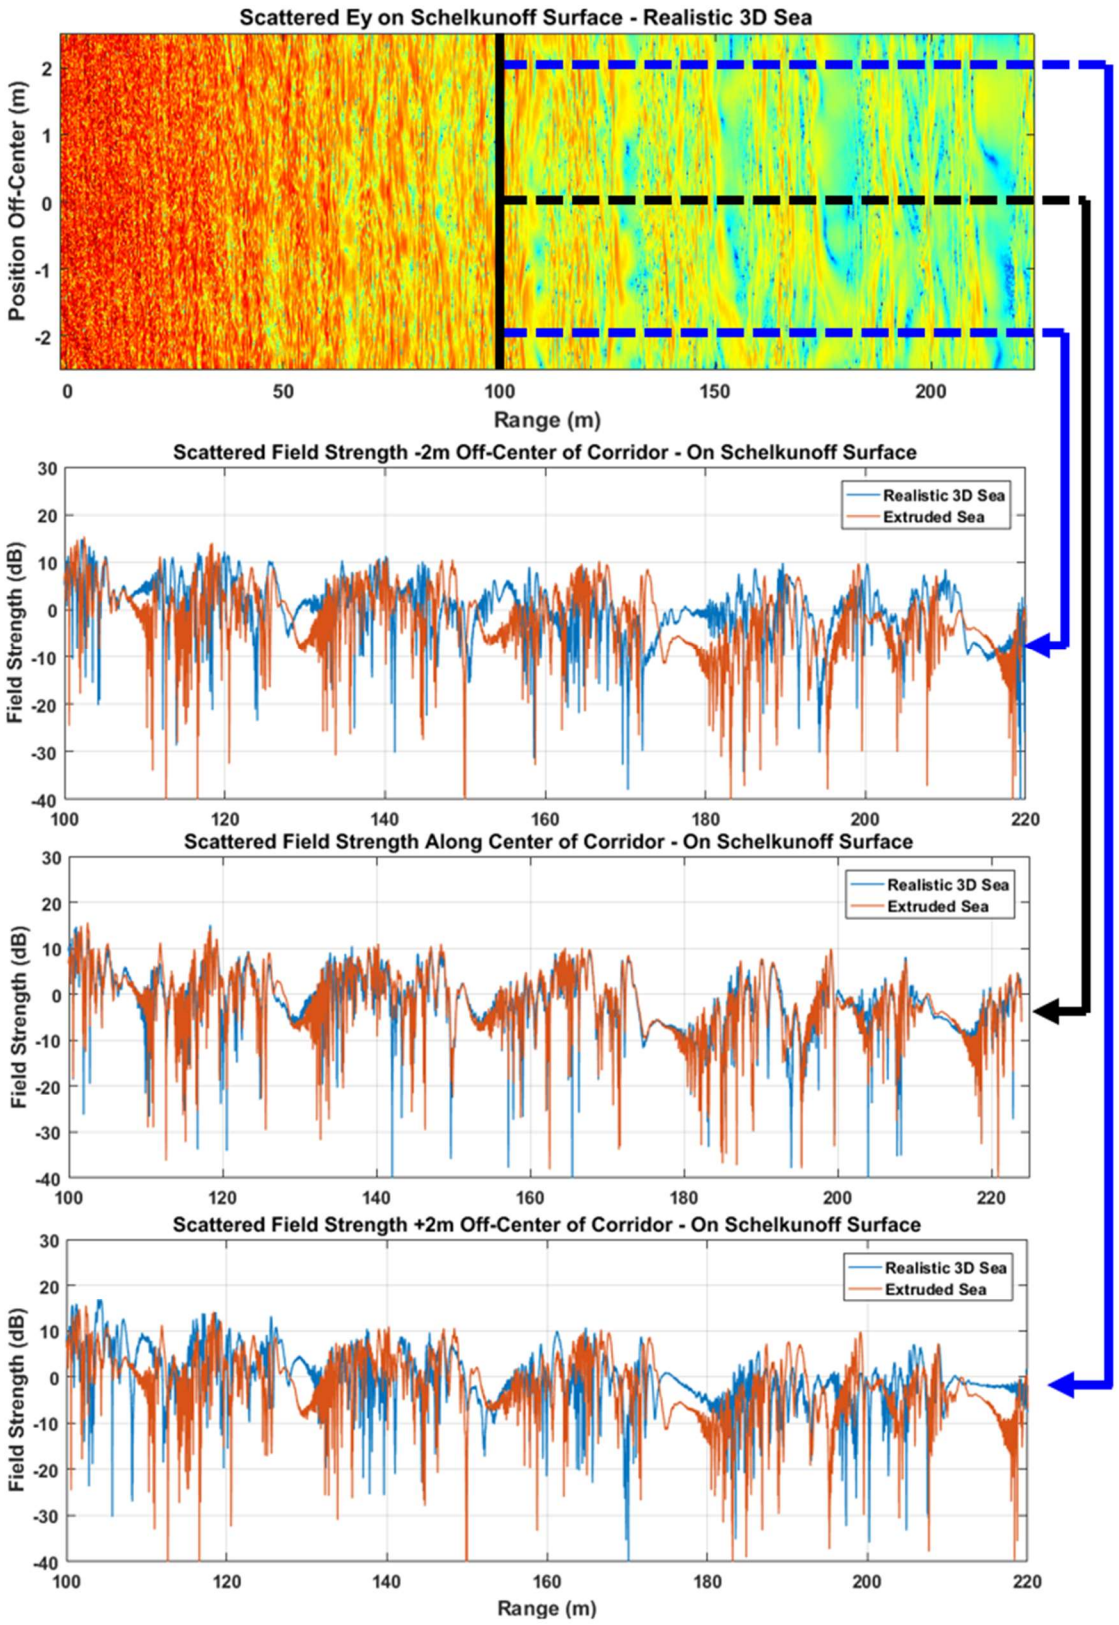

Figure 7.16: Slices of the Ey field on the Schelkunoff surface plane above the sea.

7.6 Off-Center Field Magnitude Variations on the Schelkunoff Surface

Thus far only the plane wave spectrum data has been presented, which takes into account the energy travelling out to the far field in a given direction; it doesn't take into account spatial distributions of the scattered fields above the sea. Figure 60 presents the y-directed (principal polarization) electric field collected on the Schelkunoff surface compared between realistic and extruded sea surfaces. The Schelkunoff surface sits very close to the sea surface at about six wavelengths above it, which is why the scattered fields of both seas nearly match up along the center line. Moving just two meters along the y-dimension in either direction produces up to 10dB variations in the envelope, variations which would not be present in a 2D simulation space.

7.7 Windowed Plane Wave Spectrum Demonstrates How Local Patches of Sea Contribute Out-of-Plane Rays

It's inevitable that the scatter from a realistic 3D sea surface will send energy in all directions, and this was clearly demonstrated in the plane wave spectrum of the scattered field from the 3D sea (Figure 7.7). The PWS operates in the frequency domain and automatically treats the entire set of scattered field data at once, and hence it includes the whole area of the sea in one go, with all spatial distinctions removed. The E-fields collected on the Schelkunoff surface, shown in Figure 7.17 (top), clearly demonstrate the presence of significant and abrupt transverse spatial variations in the field – influenced by the local sea. Suppose now we spatially window the Schelkunoff surface to include only a small section in the PWS calculations (Figure 7.17 bottom). The result would be the collection of rays scattered from only the patch of sea near the window, and demonstrate how

localized patches of sea contribute to out-of-plane scattering phenomena, which is particularly important in the vicinity of the target.

The Schelkunoff surface fields are spatially windowed at locations of our choosing by choosing a circular Gaussian taper with a FWHM of 1 meter, and tailored to taper down to a scaling factor of 10-7 by the edge of the window – which is essentially 0 since the scattered fields are calculated in 32-bit floating point. For this demonstration, the windows are placed at various positions within the corridor with transverse offsets of y = -1.5m, 0m, and +1.5m, at ranges of x = 0m, 10m, 100m, and 200m. An example of the scattered fields and the Gaussian windowing is shown in Figure 7.17. NOTE: The windows shown, and their apparent scales and aspect ratios, are not the windows used in the upcoming PWS calculations, they are for visualization purposes. Additionally, for the contour plots in Figure 7.19-Figure 7.22, the x- and y-axes are swapped compared to previous PWS plots, which places the forward horizon now on the right side.

Expectations for the extruded sea are fairly simple, for both the TE and TM cases the same assumptions can be made. Since the point source is located above the y = 0m line of the 3D corridor the PWS results for the window also placed at a y = 0m transverse offset should be symmetric about $k_y = 0$. By a similar token, moving the window in the positive or negative y-direction away from the center of the corridor should show the majority of radiation travelling with a positive or negative k_y component, respectively, which indicates it's moving outward from the center. This behavior results from the extruded sea's lack of transverse variation.

Expectations for the realistic sea are simpler, since the surface is stochastic and random in both the x- and y-directions, the scatter direction doesn't depend explicitly on

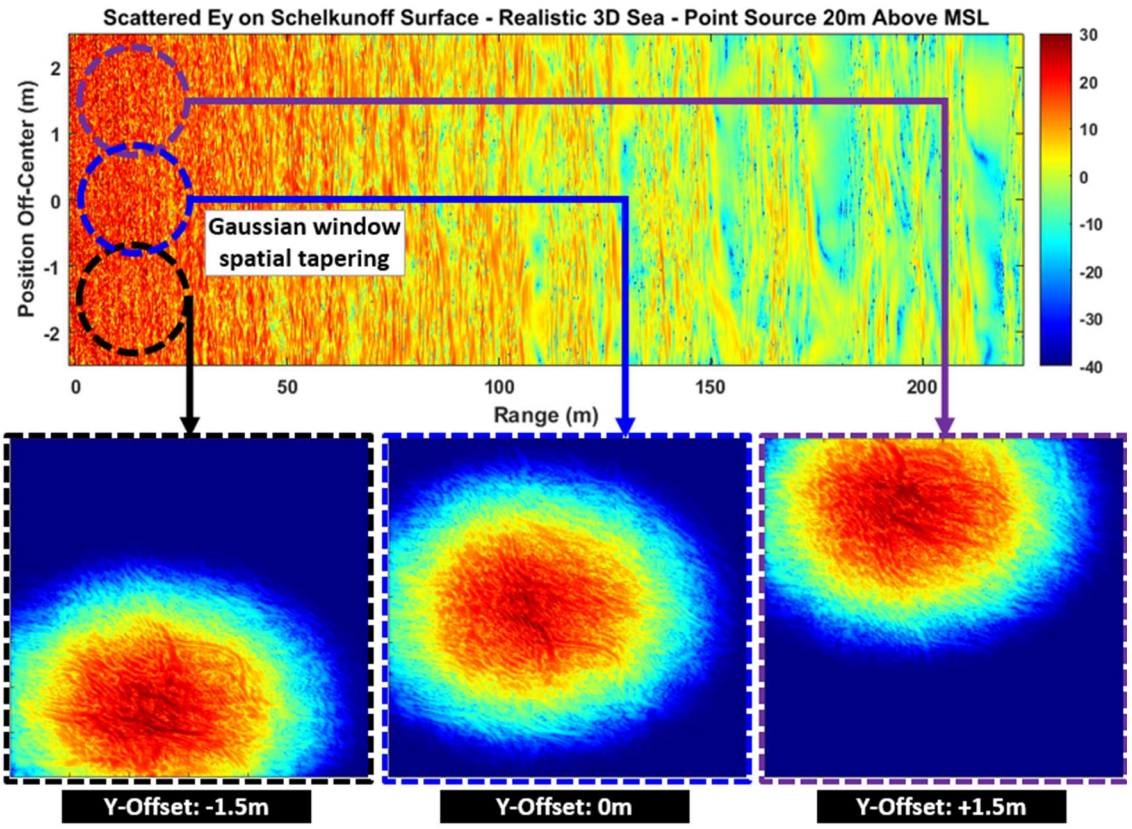

Figure 7.17: Circular Gaussian spatial windowing of the scattered fields. The windows used in the PWS calculations have a FWHM of 1m and are placed at y = -1.5m, 0m, and +1.5m.

the transverse offset of the spatial window, but rather the local scattering features of the sea. In both sea cases the common expectation is that with increased range, more and more energy will be concentrated at an increasingly narrower angular spread for both TE and TM. The 5m-wide simulation corridor will accommodate a narrow range of propagation angles down the length of the corridor. Energy travelling within such angles will accumulate, whereas wider-angle scatter will simply exit the domain (be absorbed at the terminations) and be unable accumulate over long distances.

In Figure 7.19 and Figure 7.20, at transverse offsets of y = -1.5m and y = 1.5m there appear to be plane waves propagating back toward the center, which based on previous reasoning regarding the lack of transverse variation in the sea should be impossible. These

rays are produced as a direct result of the analytic source injection VEP method used. The method artificially tapers the injected field strength, albeit softly over many wavelengths, along the walls of the corridor - diagramed Figure 7.18. This produces a minor artificial scatter, which is very weak (at least -25dB down) compared to the relevant scatter.

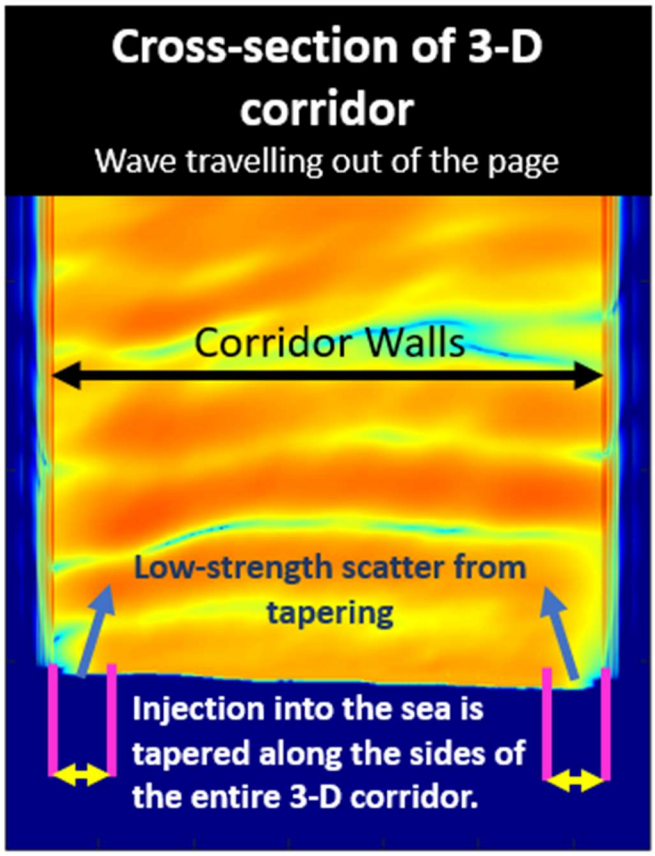

Figure 7.18: Analytic source injection tapering along sides of 3D corridor produces a small amount of artificial scattering.

In Figure 7.19 and Figure 7.20 confirm the assertions of symmetry about the center plane of the corridor, and all of Figure 7.19-Figure 7.22 show the concentration of propagating rays near the horizon as distance increases. Figure 7.21 and Figure 7.22 show how over the realistic 3D sea such drastically different angular spreads can result from patches of sea in close proximity, thus confirming the strongly present 3D local scattering phenomena unable to be captured in 2D.

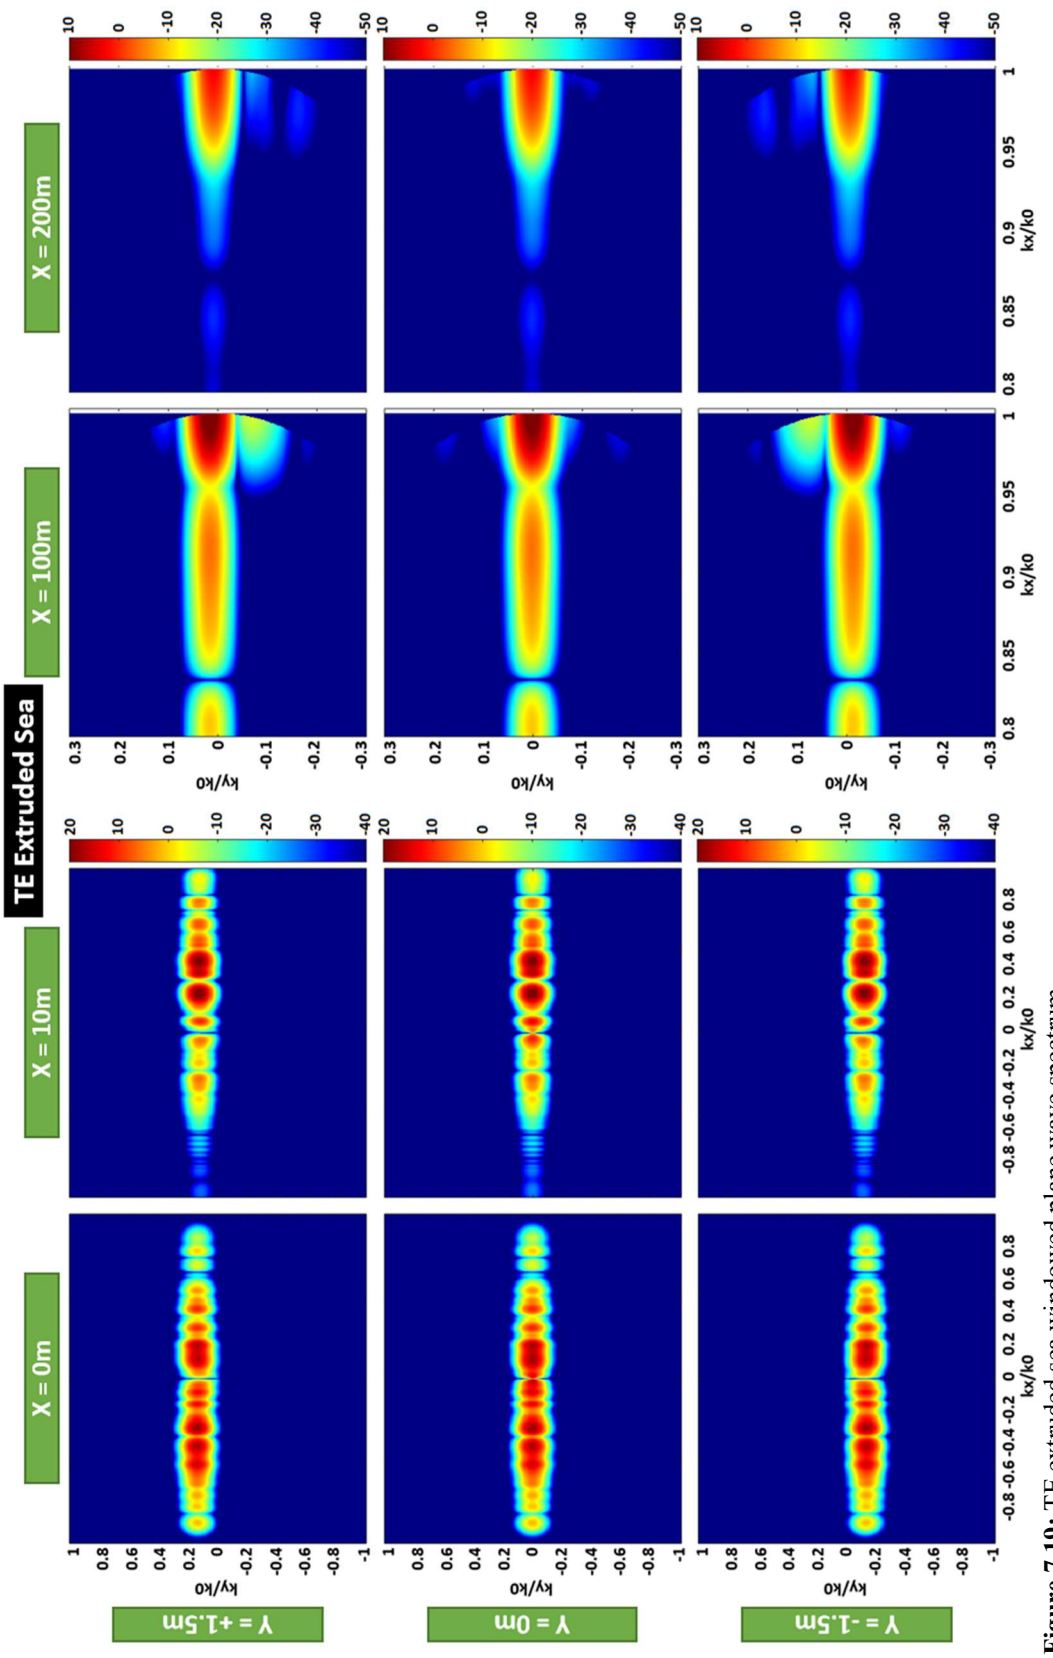

Figure 7.19: TE extruded sea windowed plane wave spectrum.

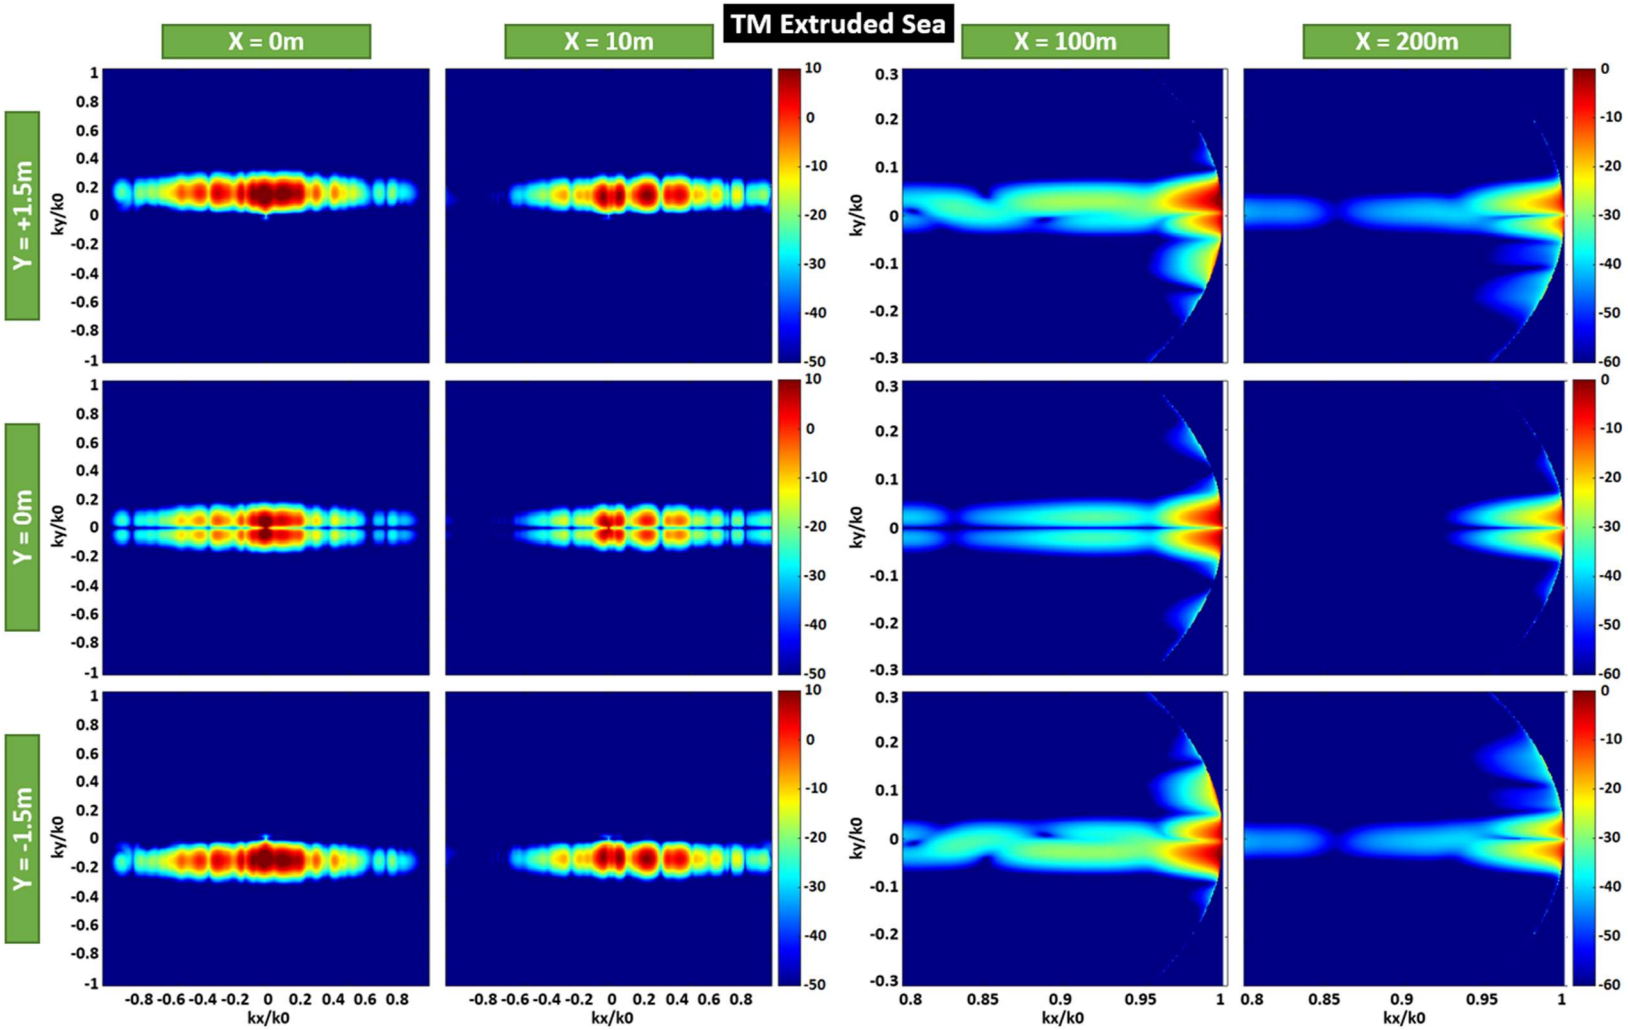

Figure 7.20: TM extruded sea windowed plane wave spectrum.

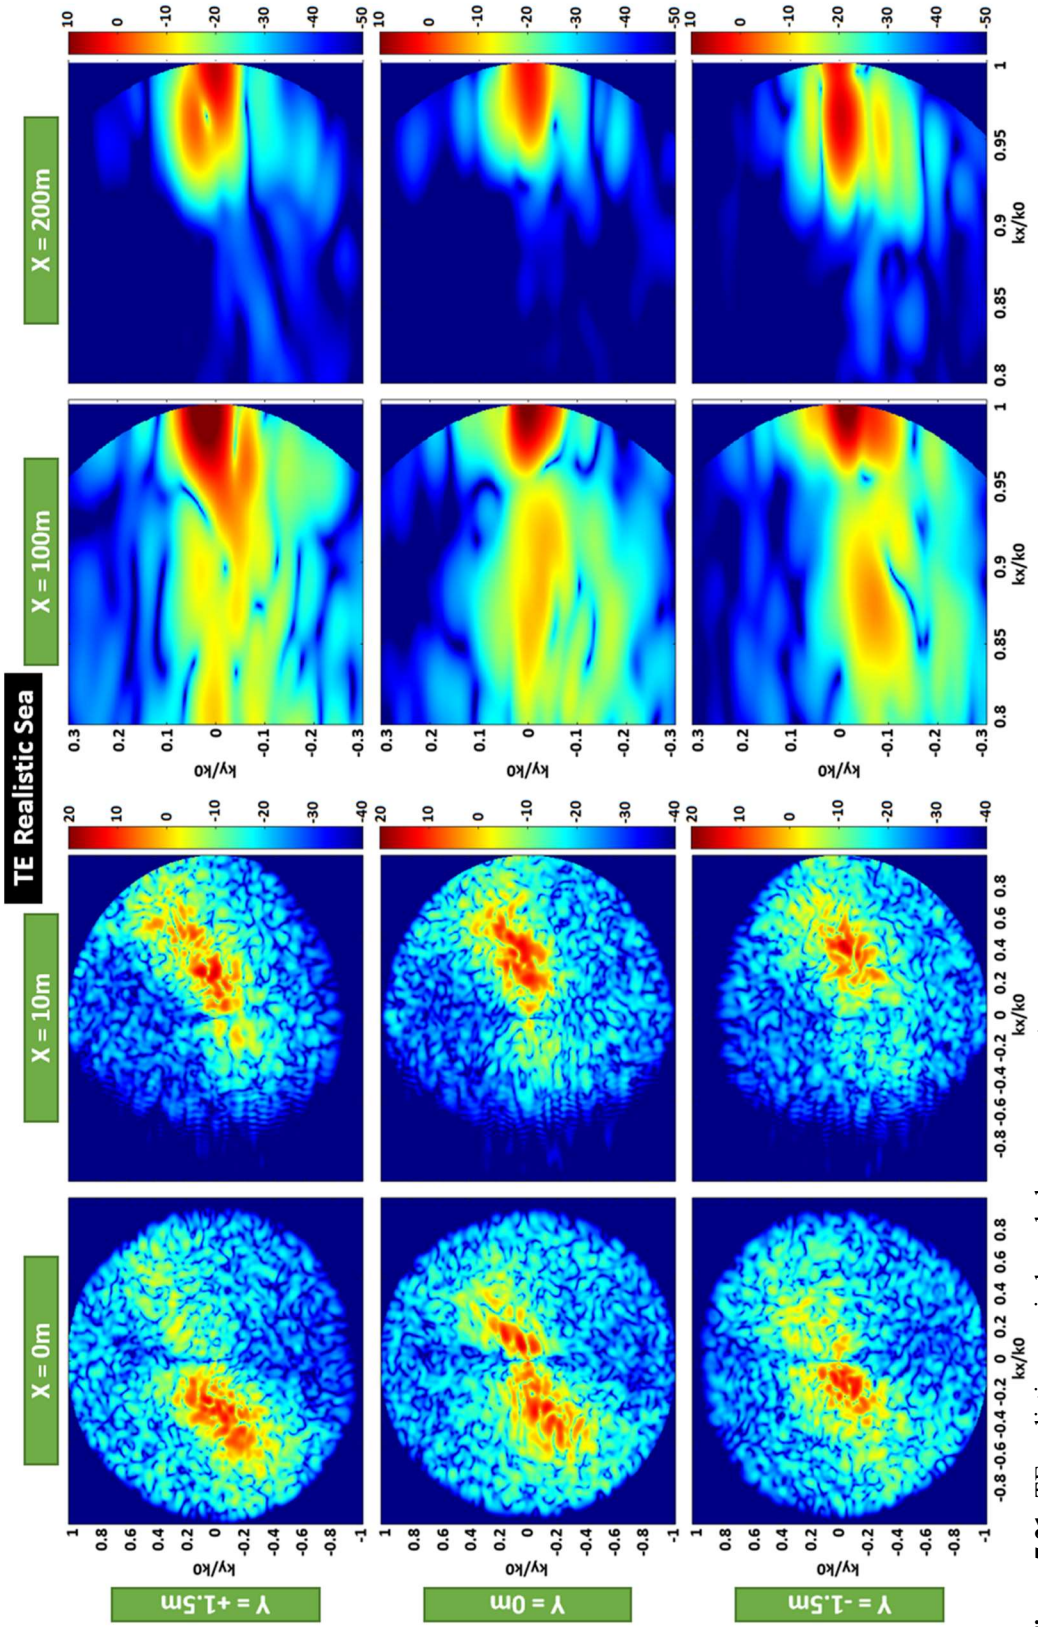

Figure 7.21: TE realistic sea windowed plane wave spectrum.

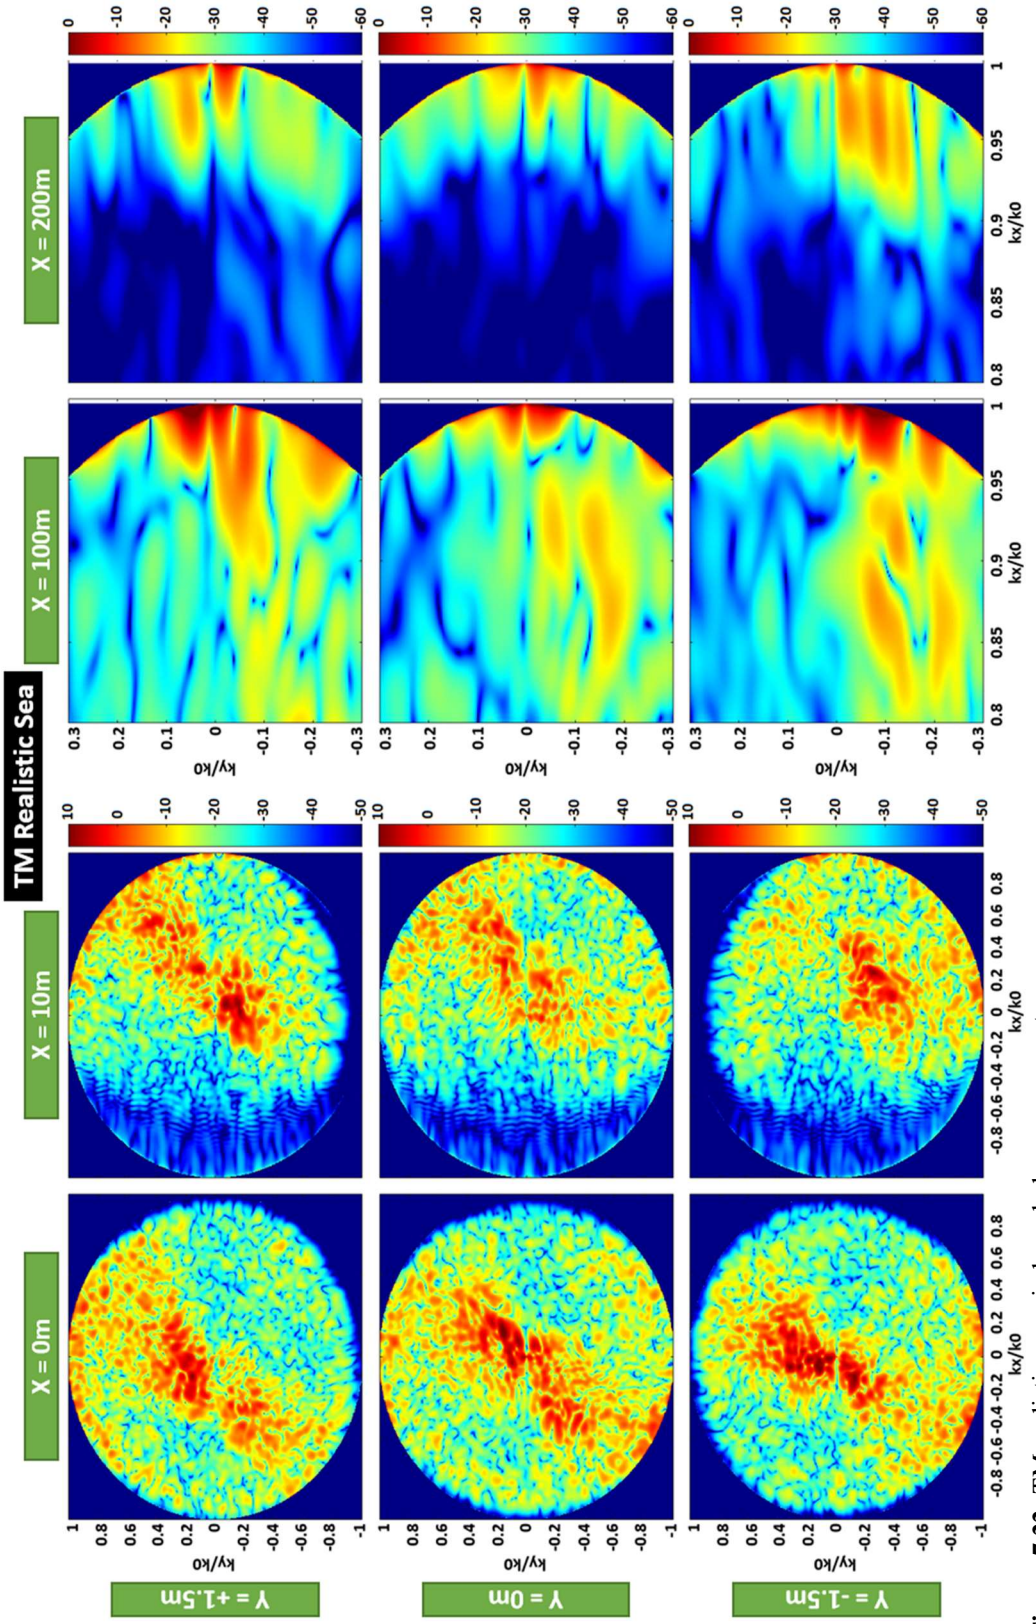

Figure 7.22: TM realistic sea windowed plane wave spectrum.

7.8 Conclusion

A 225m-long, 5m-wide corridor of realistic 3D sea as well as an extruded 3D sea excited by a horizontal electric dipole has been simulated in FDTD and yielded important scattering characteristics not able to be captured in 2D simulation spaces. A plane wave spectrum analysis of the scattered fields collected on the Schelkunoff surface just above the sea produced distinguishing characteristics between the case of the realistic sea and that of the extruded sea. First, the realistic sea scatters energy over a much wider range of angles. Second, the transverse texture of the realistic sea geometry leads to scattered fields which contain a relevantly large TM component radiating toward the forward horizon in the principal propagation plane (x-z). These TM components are physically non-existent in the 2D scenario. Third, the realistic sea also generates strong flashes and nulls in different directions which are unobserved in the extruded or 2D cases.

Regarding spatial variations of the scattered fields near the surface of the sea, the fields collected along the Schelkunoff surface show strong variations above a realistic 3D sea when observed off the center axis. A spatially-windowed plane wave spectrum analysis of the scattered fields collected on the Schelkunoff surface provides useful insight into how rapidly the 3D scattering characteristics can change as a result these strong off-center scattered field variations. The spatially-windowed PWS shows how minor transverse offsets of only a few meters outside the principle direction of propagation can lead to notably varied angular spreads in the scattered field, further indicating the relevance of the 3D scattering characteristics of the realistic sea.

There is an important distinction between the conclusions that can be drawn from spatial field variation near the sea and the plane wave spectrum analysis of the radiation

into the far field. If we focused only on the fact that the field strength varies as the sampling plane is moved transversely (as in Figure 7.16) it could be argued that the 2D modeling philosophy of VTRPE might capture such a variation also. After all, the way VTRPE and other methods are used with a 3D sea is that the 2D slice of the sea is taken along the given principal plane direction of interest. It follows than that moving the principal plane a few milliradians to the left or right (in ϕ) (as suggested in Figure 7.23a, below) would produce a different slice of sea that, then, would yield different scattered fields along that plane.

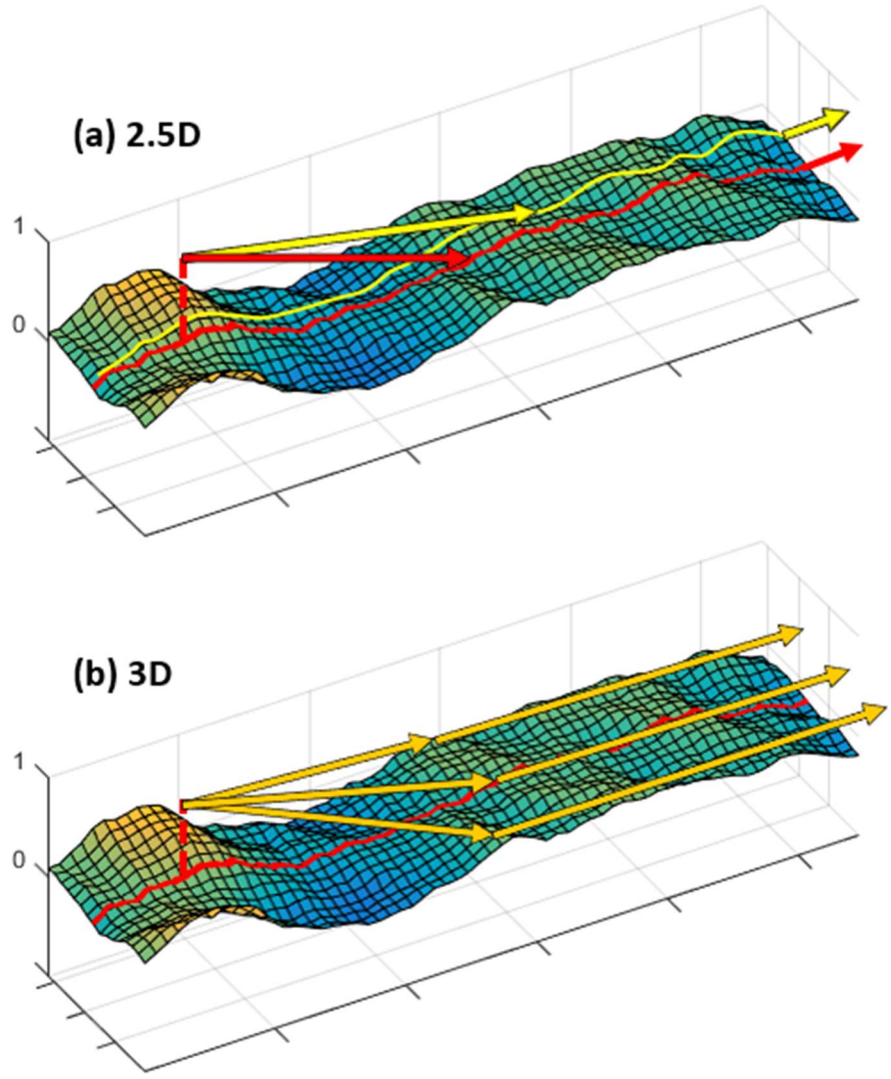

Figure 7.23: Key differences in the downrange results between 2.5D (and 2D) simulations (a) and realistic 3D simulations (b) arise because the 3D sea surface redirects energy towards the distant target.

However, that is not the meaning of the variations we observe here. The clearest way to see the difference is to return to the plane wave spectrum and conical cuts near the forward direction as shown in Figure 7.11, Figure 7.12, and Figure 7.15. In those plots we see that energy gets scattered into the $k_y = 0$ direction by the 3D sea that is not present in the 2D extruded sea simulation. In other words, it is not a matter of the field being different along adjacent planes but rather that rays that were emitted in other directions in ϕ away from the principal plane are redirected by the 3D texture into traveling in the $\phi = 0$ direction. Energy that would never reach the distant target using the modeling philosophy of Figure 7.23a, reaches it in the realistic model of Figure 7.23b. This is most evident when we consider the TM cross-polarized signal completely absent from the principal plane in the 2D model but very much present in the 3D realistic model.

CHAPTER 8

SCATTERED FIELD RECONSTRUCTION BASED ON SPATIAL INTEGRATION

8.1 Spatial Integration of FDTD Sources

The plane wave spectrum analysis of the Schelkunoff surface currents yields a broad insight into the physics of the sea space, and readily points out many vital differences between 2D and 3D simulation environments. However, it is only a part of the full picture. In converting to spectral domain, we are immediately taken to the far field, and important spatial features near the scattering surface (the sea) are lost. To some degree the PWS of the spatially windowed surface currents recovers some of the locality of the problem but at the same time it eliminates important contributions from other parts of the sea. Furthermore, near the horizon, the PWS may experience artefacts due to the proximity to the so-called invisible region of the spectrum, requiring extreme angular resolutions to mitigate.

In reality, we are not interested in the plane waves reaching the far field, but rather the local plane waves propagating within meters of the sea surface which would scatter off individual features on a radar target. That is, we are interested in plane waves traveling within the “Fresnel Zone” of the scattering surface. They must be reconstructed from the FDTD simulation in order to calculate useful angle of arrival data (AoA) for rays scattering off radar targets.

The ideal scenario for the generation of AoA would be to know the fields on a large enough target plane facing the intended target (a y-z plane in our case). Such a plane should

be at least 30m tall. However, given the computational burden of a 3D simulation, the FDTD solution uses the analytic field injection VEP method. Thus, the FDTD-generated fields in the simulation space extend barely a meter above MSL – which is scarcely above the sea swell size. To produce the fields on an appropriately-sized target plane, we will perform in 3D the same technique as previously demonstrated in 2D: Spatial-integration-based reconstruction.

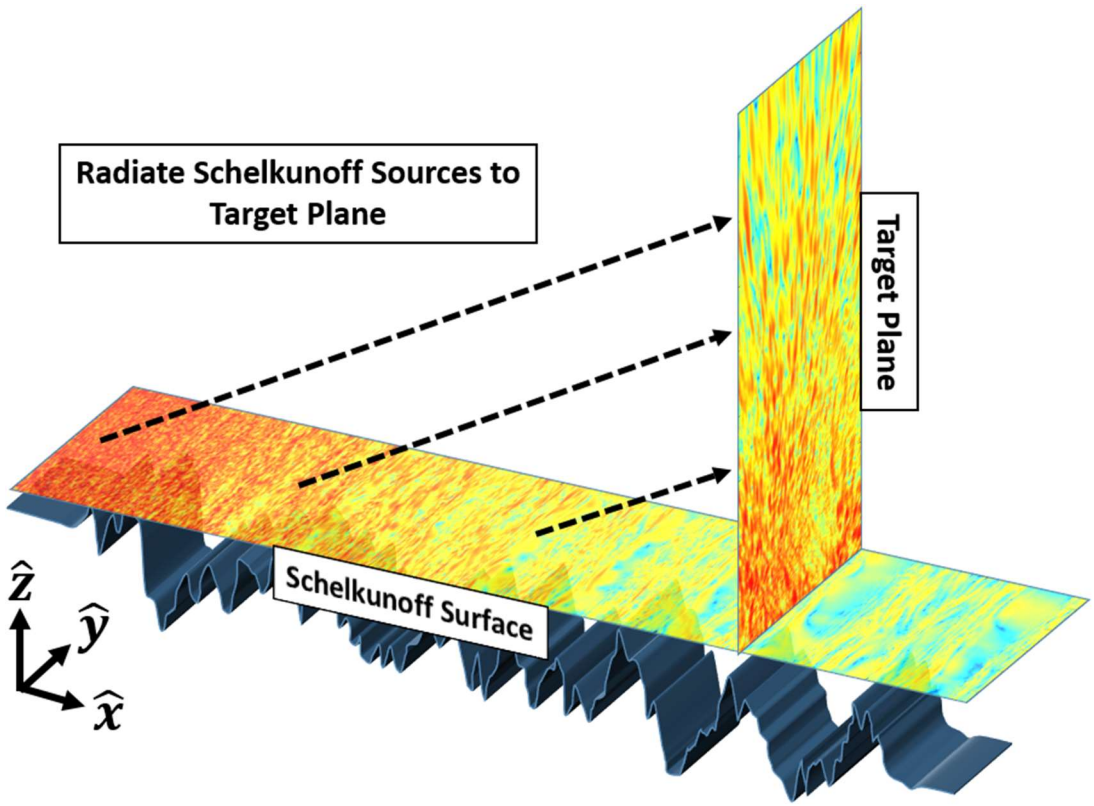

Figure 8.1: The goal of the spatial integration reconstruction is to radiate the Schelkunoff sources to produce the fields on a vertical target plane.

With the spatial integration reconstruction approach, we can readily reproduce the fields at any point in space, particularly the points in space beyond the scope of the FDTD domain, since the 3D FDTD space is only a couple of meters in height. Figure 8.1 displays the general setup for the reconstruction process, which is analogous to 2D. The Schelkunoff current elements are considered as individual sources, and their radiated contributions

summed over all points at which the scattered fields are to be. Reconstruction in 3D carries with it a few extra considerations and complications. Most of the reconstruction demonstrations here were performed with a 9GHz y-directed point dipole, 20m above a 1,050m (31,500λ) long, and 5m (150λ) wide corridor of realistic 3D sea.

In 2D, it was possible to perform extremely accurate reconstruction using only one field component – the field transverse to the plane of propagation. In 3D, all four field components ($E_x,\ E_y, H_x, H_y$) stored on the Schelkunoff surface are required. The electric and magnetic surface currents, J_s and M_s, respectively, are responsible for radiating into the space above the surface. The currents are calculated from the conventional surface current relations:

$$J_s = \hat{n} \times \vec{H} \qquad 8.1$$

$$M_s = -\hat{n} \times \vec{E} \qquad 8.2$$

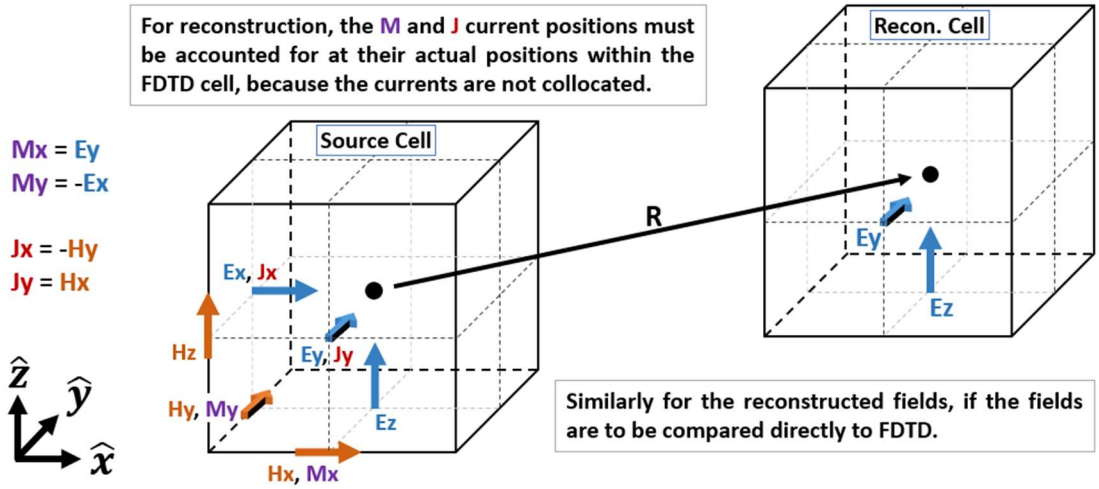

Figure 8.2: The half-cell field component offsets within the Yee cell are critical in the calculation of the radius between the source and reconstruction points.

Since the Schelkunoff surface resides in the x-y plane, the normal, $\hat{n}$, is in the $\hat{z}$ direction, and we get the relations for the electric and magnetic currents shown in Figure 8.2. Note the locations of each current element, they are not at the same point in space as

the field element which produced them, they are offset by half a cell due to the way the discrete curl operator works. Additionally, since both E and H are used to compute the surface currents, and in the FDTD update loop E and H are separated by half a time step, the surface currents need to be synchronized before they are used for reconstruction (just as in Section 7.2). In our case, we calculate H first in our update loop, so we must add a phase delay back into H before calculating the electric surface currents (eq. 8.3).

$$J_s = \left(\hat{n} \times \vec{H}\right) \cdot e^{-\frac{j\omega t}{2}} \tag{8.3}$$

Each current element is treated as an infinitesimal dipole (treatment as a small dipole yields no discernable benefit.) As shown in Figure 8.2, because of the staggered nature of the Yee grid, none of the current or field elements share the same location in the unit cell, and the slight offsets in location must be taken into account when calculating the source current's location relative to the reconstructed field element's location. For each current element dipole, the fields are calculated from well-known expressions in spherical coordinates and then converted to Cartesian coordinates, the math is simple, but the bookkeeping is tedious. For an infinitesimal z-directed electric dipole:

$$H_r = H_\theta = 0 \tag{8.4}$$

$$H_\phi = j\frac{k_0 I_0 \Delta s \sin\theta}{4\pi r}\left[1 + \frac{1}{jk_0 r}\right] e^{-j\tilde{k}r} \tag{8.5}$$

$$E_r = \frac{\eta_0 I_0 \Delta s \cos\theta}{2\pi r^2}\left[1 + \frac{1}{jk_0 r}\right] e^{-j\tilde{k}r} \tag{8.6}$$

$$E_\theta = j\frac{\eta_0 I_0 \Delta s \sin\theta}{4\pi r}\left[1 + \frac{1}{jk_0 r} - \frac{1}{(k_0 r)^2}\right] e^{-j\tilde{k}r} \tag{8.7}$$

$$E_\phi = 0 \tag{8.8}$$

For a small z-directed magnetic dipole:

$$H_r = \frac{I_m \Delta s \cos\theta}{2\pi\eta_0 r^2}\left[1 + \frac{1}{jk_0 r}\right] e^{-j\tilde{k}r} \tag{8.9}$$

$$H_\theta = j\frac{I_m \Delta s \sin\theta}{4\pi\eta_0 r}\left[1 + \frac{1}{jk_0 r} - \frac{1}{(k_0 r)^2}\right] e^{-j\tilde{k}r} \tag{8.10}$$

$$H_\phi = 0 \tag{8.11}$$

$$E_r = E_\theta = 0 \tag{8.12}$$

$$E_\phi = -j\frac{k_0 I_m \Delta s \sin\theta}{4\pi r}\left[1 + \frac{1}{jk_0 r}\right] e^{-j\tilde{k}r} \tag{8.13}$$

Here again $\tilde{k}$ is the numerical wavenumber, dispersionHere again $\tilde{k}$ is the numerical wavenumber, dispersion-corrected by the 3D FDTD phase velocity definition, $v_{p3D}(\theta, \phi)$, which is shown in equation 4.19, producing the following definition:

$$\tilde{k}(\theta) = \frac{\omega}{v_{p3D}(\theta, \phi)} \tag{8.14}$$

Following this recipe, the Green's function reconstruction is a mostly straightforward, brute-force computation. To verify the reconstruction works, we compute the fields on the target plane within 1m of the sea surface since the FDTD fields also exist in the region. For simplicity, only the y- and z-directed E field reconstructions will be shown. The H_z and H_y results are analogous to E_y and E_z, respectively, and an E-field analysis essentially doubles as an H-field analysis for our purposes. The E_y and E_z field components, reconstructed at a range of 13m, are shown in dB in Figure 8.3. Included in the plots are three scattered fields: those produced by conventional VEP FDTD, those produced by application of special pre-processing techniques to the Schelkunoff surface fields (discussed shortly), and those reconstructed with no pre-processing.

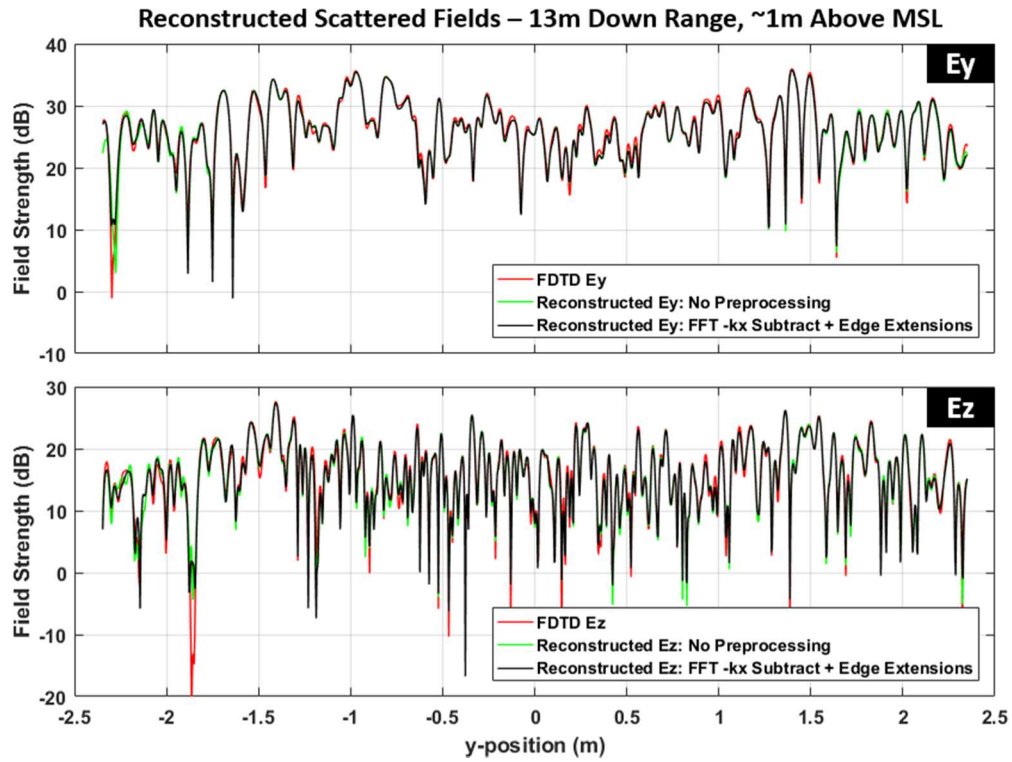

Figure 8.3: Reconstruction results compared to FDTD at 13m down range.

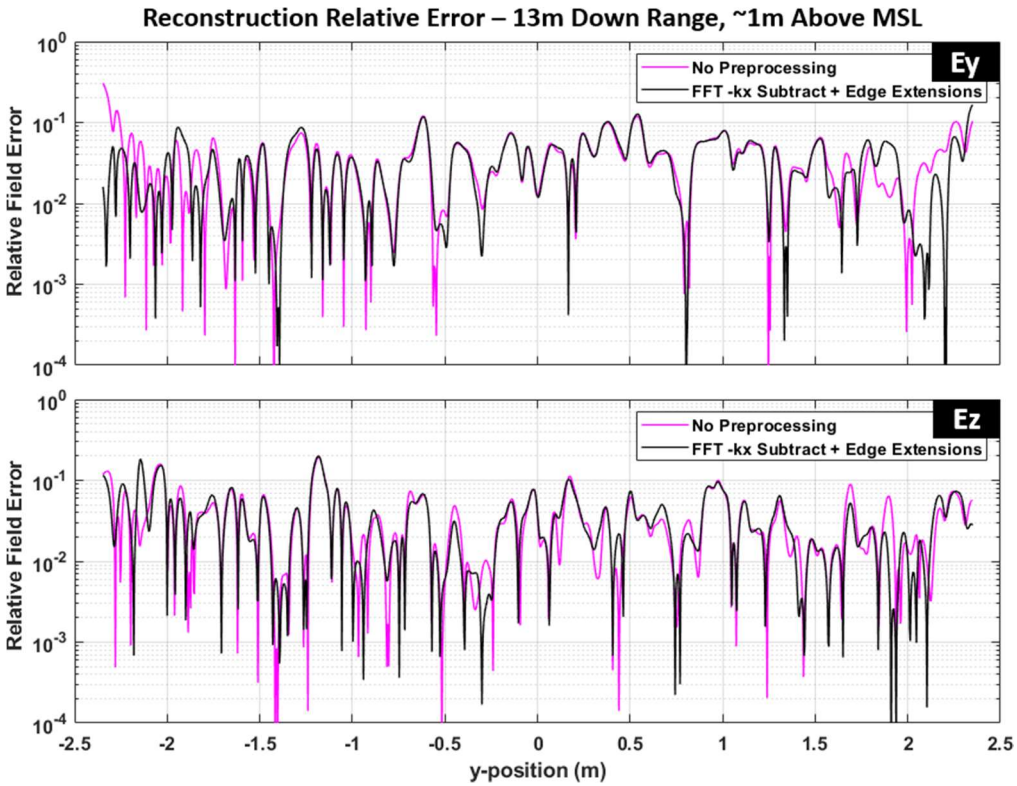

Figure 8.4: Reconstruction relative field error to FDTD, compared at 13m down range.

Figure 8.4 shows the relative field error of the reconstructed fields compared to the FDTD fields (A field error of 0.1 means the error is -20dB down relative to the field itself.) Computation of the relative error for the electromagnetic fields, particularly those which have picked up the stochastic nature of the sea, is not necessarily a straightforward consideration. The field data consists of a densely-packed collection of peaks and nulls, and very slight differences in null positions can lead to enormous errors. However, those enormous errors typically occur in regions of very weak fields, fields so weak such that an order of magnitude difference would introduce no effects for any practical purpose. Thus, a desirable relative error recipe is one which eliminates the heavy skewing caused by the deep nulls but preserves the behavior of the curves outside the nulls.

This is accomplished by first applying a Gaussian-weighted moving average filter to the squared modulus of the field data and taking the square root of the results. The filter is shown in equation 8.15, the smoothing kernel in equation 8.16, and the relative error formula in equation 8.17. Here F represents the field component to be smoothed, F_S represents the smoothed field, and x is the position along the dimension of smoothing.

$$F_{s_i} = \left[\frac{\sum_j |F_j|^2 \cdot K\left(\frac{x_i - x_j}{w}\right)}{\sum_j K\left(\frac{x_i - x_j}{w}\right)} \right]^{\frac{1}{2}} \tag{8.15}$$

$$K(\alpha) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{\alpha^2}{2\sigma^2}} \quad , \quad \sigma = \frac{1}{e} \tag{8.16}$$

$$Error = \frac{|F_{s_{FDTD}} - F_{s_{Recon}}|}{|F_{s_{FDTD}}|} \tag{8.17}$$

The rule we define for the parameter w, the width of the smoothing window, is to choose a value such as to not reduce the strength of any peaks in the field data by more

than 3dB. By smoothing the fields in this manner prior to performing the relative error computations, the skewing caused by deep nulls is minimized, and an accurate representation of the error is maintained elsewhere.

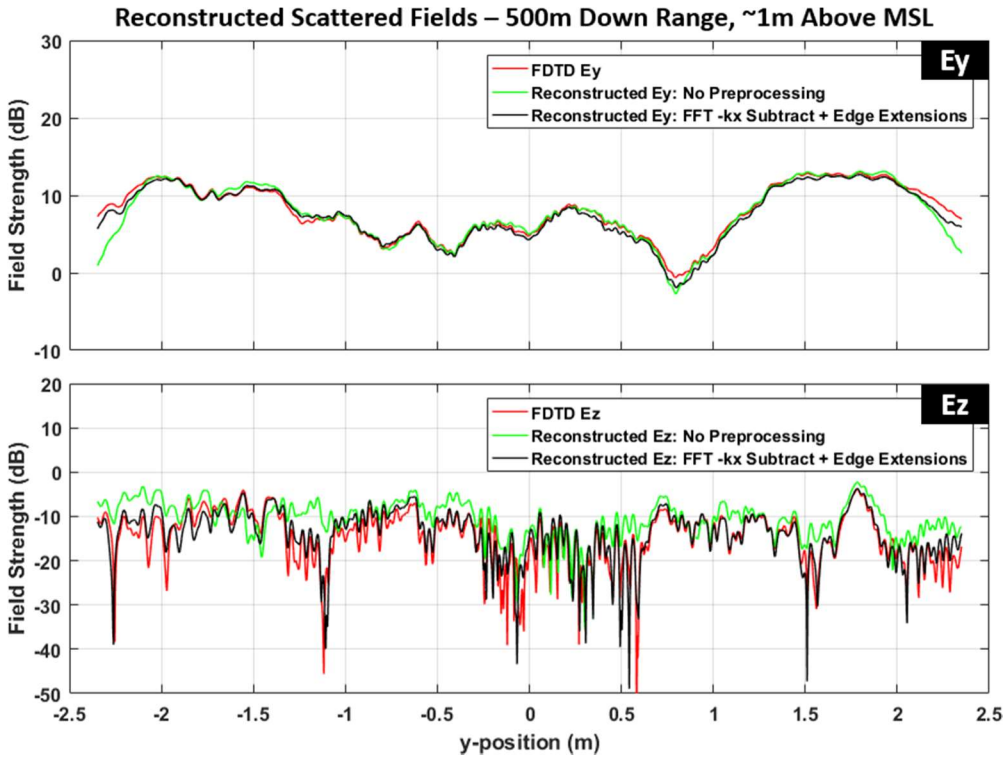

Figure 8.5: Reconstruction results compared to FDTD at 500m down range.

Very close down range, the reconstructed fields agree with the conventional FDTD fields quite well. An important aspect of Figure 8.3 occurs near the edges of the space, at which both E_y (the co-pol field) and E_z (the cross-pol field) tend to show the worst disagreement. This arises because the Schelkunoff surface forms a rectangular aperture of surface currents roughly 150 wavelengths wide and many thousands of wavelengths long (pictured in Figure 8.1 and Figure 8.7). Since there is no data outside the simulation domain, there is an apparent persistent surface current discontinuity along the edges of this "aperture". Although in the time domain the FDTD termination boundary simply allows the energy to exit with no further adverse effects, in the frequency domain this absence of

data introduces diffraction emanating from the discontinuity, which grows in distance down the corridor.

The diffraction from the edges results in the distortion of the reconstructed fields when compared directly to the FDTD fields. The cumulative effects of the diffraction increase with range since the discontinuity at the edges becomes larger with increased distance due to the requirement to include more field sources in the reconstruction (APPENDIX A). With increased reconstruction range, the E_y (and H_z) fields fare quite well, with diffraction distortion mostly isolated to the edges (Figure 8.5 top). The cross-pol E_z (and H_y) fields are much more susceptible to the edge diffraction, and without corrective pre-processing can become heavily distorted, as demonstrated in the bottom of Figure 8.5.

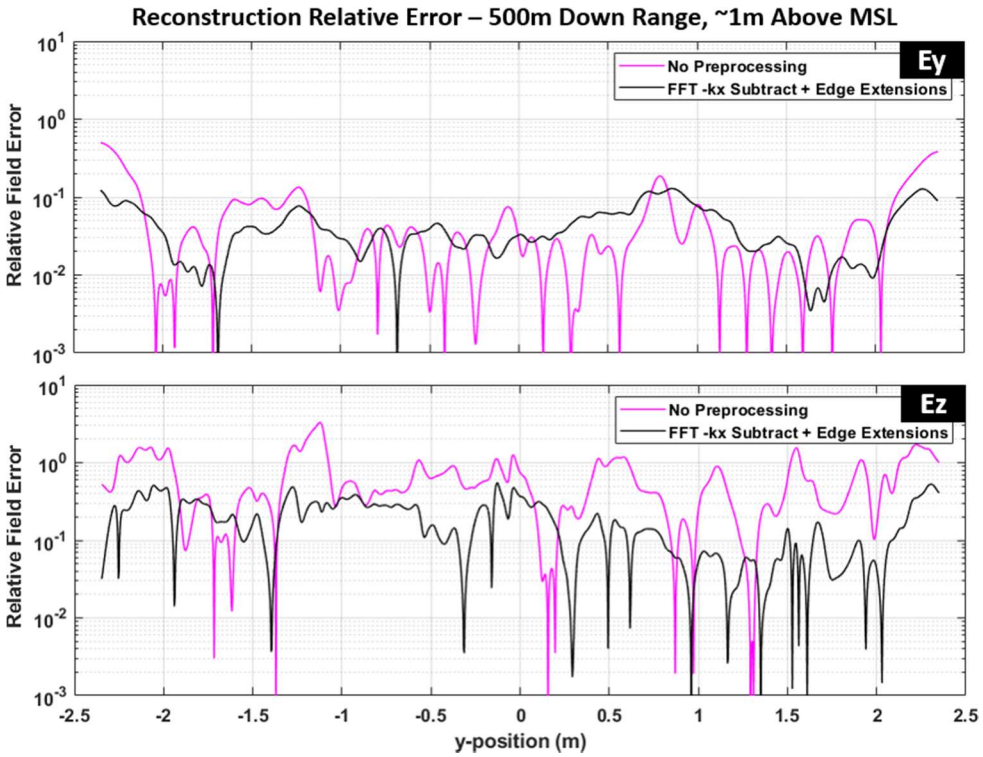

Figure 8.6: Reconstruction relative field error to FDTD, compared at 500m down range.

In order to minimize the effects of the edge diffraction, we introduce a technique to reduce the effects of the abrupt current discontinuities created by the edges of the Schelkunoff surface's aperture. We could simply apply a spatial taper to the Schelkunoff currents, however a sufficiently large taper removes too much field information, which creates more distortion than the edge diffraction itself. Instead we opt to artificially extend the surface currents in the $\pm y$ directions. In doing so, we must taper off the amplitude, but preserve the phase progression. The procedure is simple:

1. Compute the natural logarithm of the complex Schelkunoff surface currents.
2. Extend the real part as a constant using its edge value, and the imaginary part (phase) linearly by using its edge slope.
3. Convert everything back to electric and magnetic surface currents by computing the complex exponential over the whole surface to undo Step 1.
4. Apply a spatial taper to the extension regions.

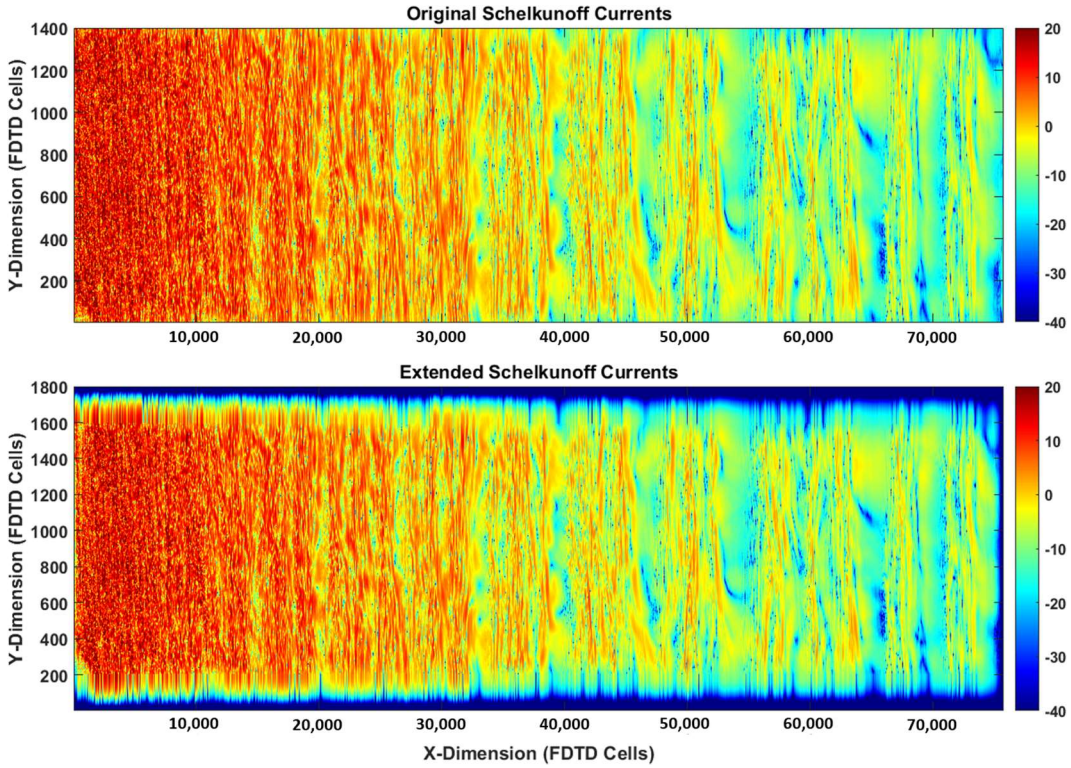

Figure 8.7: Artificial extension and tapering of the Schelkunoff surface currents to reduce diffraction.

The extended Schelkunoff surface currents are shown in Figure 8.7. The extension lengths on either side are typically on the order of 50 or 100 wavelengths per side for domains 30,000+ wavelengths long, which increases the computational size of the reconstruction by about 67-133%, but it is an effective and extremely straightforward method to minimize edge diffraction.

In addition to the edge diffraction, another issue emerges due to our FDTD domain decomposition process which, as it stands, carries the forward-travelling information successively from sub-domain to sub-domain (it's only the forward-travelling scattered fields which are of interest at a target.) If there exists a large sea feature in sub-domain (N+1) which scatters strongly back toward sub-domain N, the backward scatter will not re-enter sub-domain N. As a result, local backward scatter may create a large accumulation of currents at the leading edge of sub-domain (N+1), which generates a discontinuity on the Schelkunoff surface at the boundary with sub-domain N, as Figure 8.8 shows.

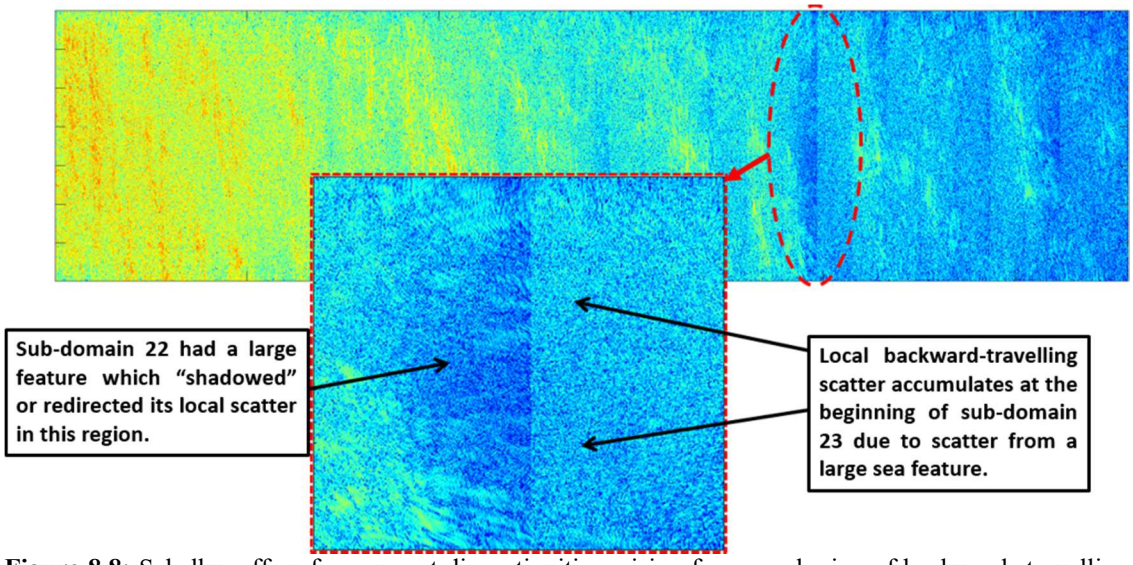

Figure 8.8: Schelkunoff surface current discontinuities arising from no sharing of backwards-travelling scattering information between sub-domains.

These abrupt discontinuities, which can occur at every sub-domain transition, skew the reconstructed fields by strongly scattering. The solution to this problem is simply to apply an FFT along the direction of propagation (in this case a 1D FFT along the x-direction) to each of the Schelkunoff current components in order to transform them from frequency domain currents into k-space currents. Then, we can simply remove any currents propagating backward on the surface (i.e. currents with $k_x < 0$). The currents are removed with a taper applied to the regions of $-k_x$ and $k_x > k_0$ (Figure 8.9.)

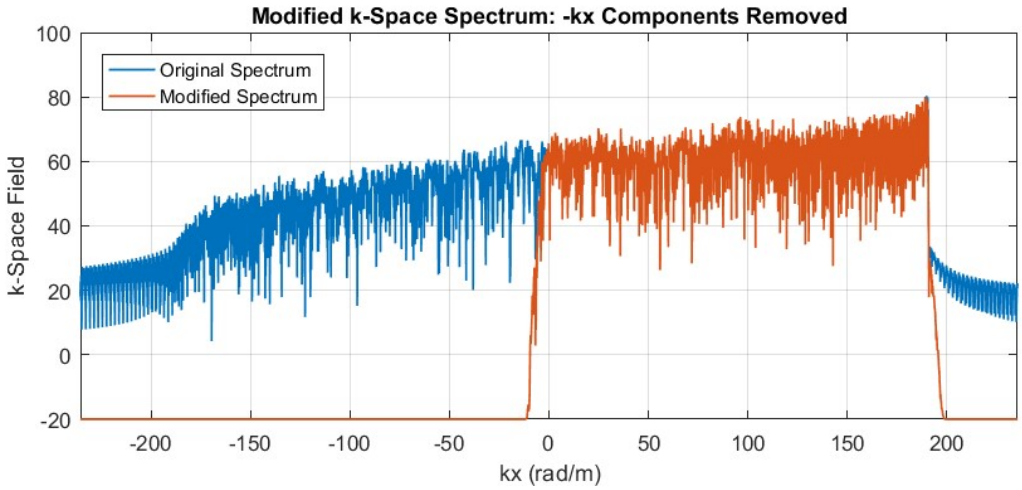

Figure 8.9: Modification of the Schelkunoff currents' k-space spectrum to remove backward-travelling waves.

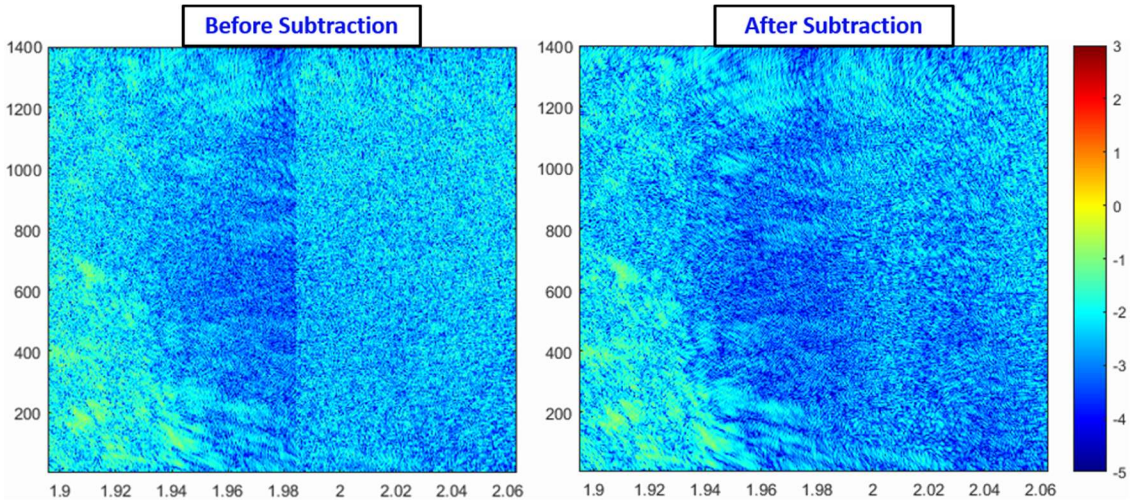

Figure 8.10: Removal of the Schelkunoff surface current discontinuities by way of modifying the k-space spectrum.

An example of the removal in 1D is shown in Figure 8.9. Once the backward-travelling portion of the spectrum is removed, the inverse transform to return to frequency-domain is performed, and the discontinuities vanish (Figure 8.10). Removing the backward-travelling waves on the Schelkunoff surface has no bearing on what reaches a target plane down range, since all that matters is the forward-travelling information.

8.2 Angle of Arrival: FDTD Fields vs. Reconstructed Fields

The surface current extensions and the filtering of the $-k_x$ components produce results quite close in line with what is the only base truth available: the scattered fields predicted by FDTD. The most relevant information to be gleaned from the reconstructed fields will be the angle of arrival calculations performed on the target plane. Thus, it's pertinent to investigate how minor differences in the reconstructed fields transition into the plane wave spectrum analysis which produces the target plane AoA. In order to assess this, a comparison of AoA calculations will be performed between a plane of FDTD scattered fields and the reconstructed scattered fields on the same plane.

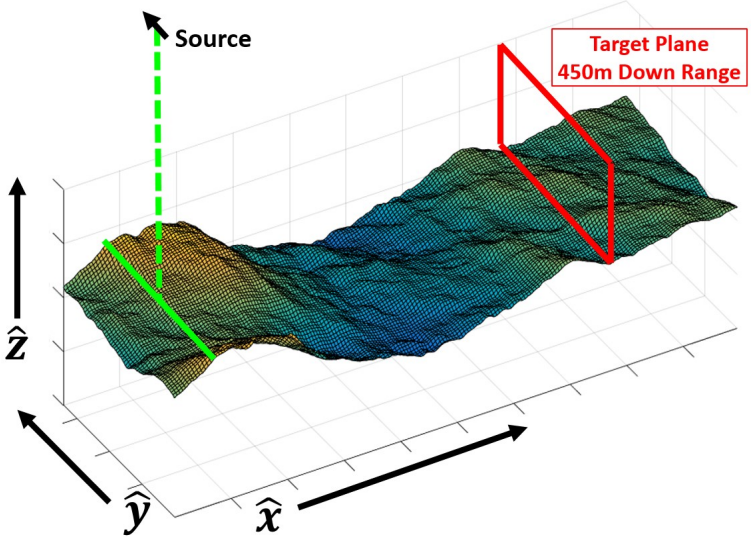

Figure 8.11: Not to scale: test domain setup to analyze proliferation of reconstruction error into AoA.

The basic test space is diagramed in Figure 8.11, and consists of a 9GHz y-directed point source placed 20m above MSL in a 485m corridor of sea, with the target plane placed 450m down range. The region of overlap of the conventional FDTD space and the reconstructed scattered fields is quite small – only about 10 wavelengths in the z-dimension, and 150 wavelengths in the y-dimension. Figure 8.12 demonstrates the size of the comparison region to what would be considered a practically-sized (900 wavelengths, or 30m tall) target plane. The relative error of the reconstructed scattered fields compared to the FDTD scattered fields is shown in Figure 8.13, using a power-of-10 color scale.

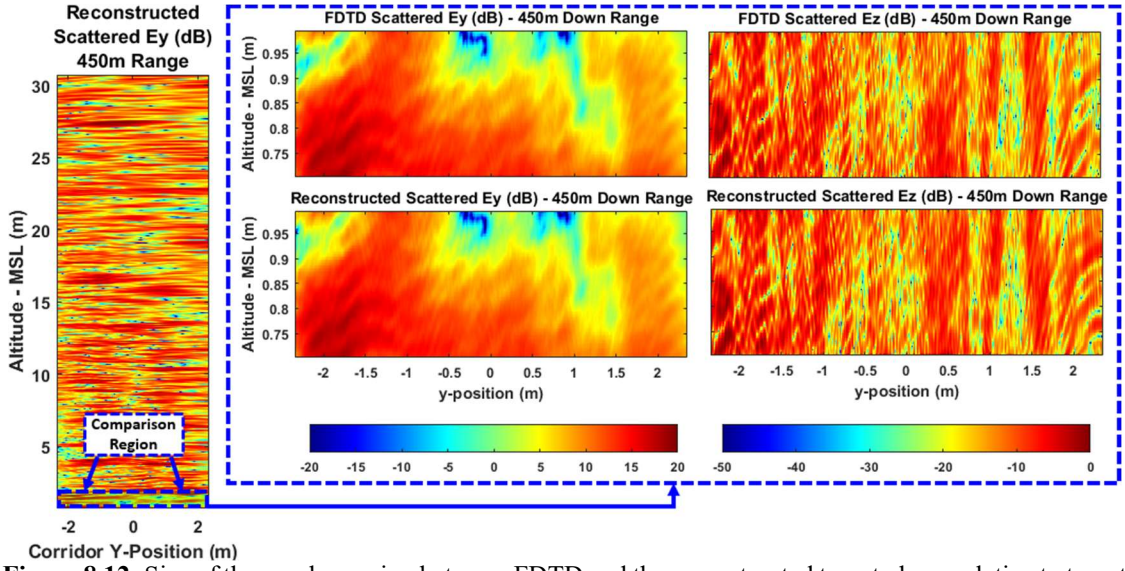

Figure 8.12: Size of the overlap region between FDTD and the reconstructed target plane, relative to target plane of practical size. (Not to exact scale)

Note the cross-polarized field (E_z) scale in Figure 8.12 starts -20dB lower than the co-polarized field (E_y). The largest relative errors for the co-polarized field (Figure 8.13 left) are of the order of 10^{-1}, or -20dB down. However, for the much weaker cross-polarized field there are regions in which the relative error is as high as $10^{+0.2}$. That is, there is a region on the right edge of the scattered E_z field (Figure 8.12 right) where the FDTD computation yields a broad null surrounding a local peak, but the reconstruction sees a

much narrower null region around the same peak. Thus, peak values get compared to null values resulting in a large computed local error. Overall, the typical relative error on the reconstructed E_z field is on the order of -10dB down. Clearly, the diffraction artefact from the edges of the Schelkunoff surface more strongly affects weak cross-polarized field. The key question is: Are these errors on the cross-polarized field relevant to the angle of arrival results?

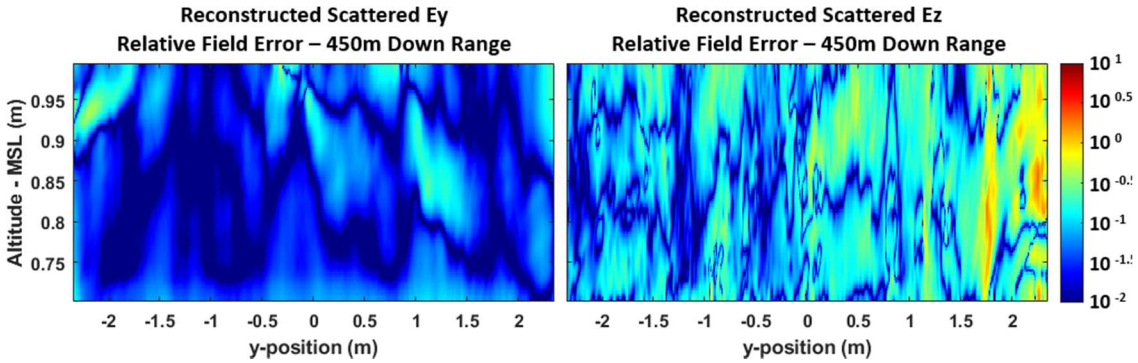

Figure 8.13: Relative field error of the reconstructed scattered fields compared to the FDTD scattered fields within the overlap region.

In 3D, the AoA is determined via a PWS analysis of the target plane fields, similar to the PWS analysis in the previous chapter. The 3D AoA calculations are performed analogously to those in 2D space – instead of a 1D Fourier transform on a line of currents, a 2D Fourier transform on the target plane is applied. Due to the spatial orientation of the target plane, the AoA will be considered in the upper-forward quadrant of the far-field sphere, or rather the region in k-space where $0 \leq k_z \leq k_0$ and $-k_0 \leq k_y \leq k_0$. A visualization of this space is demonstrated in Figure 8.14. The results are presented as a 2D plot in which the upper-forward quadrant is essentially collapsed and pressed flat onto the k_y-k_z plane. The field components are broken down into TE and TM, or E_ϕ and E_θ, respectively. Figure 8.15 shows the AoA of the FDTD and reconstructed scattered fields

of the overlap region from Figure 8.12. The AoA agreement is very good between both cases.

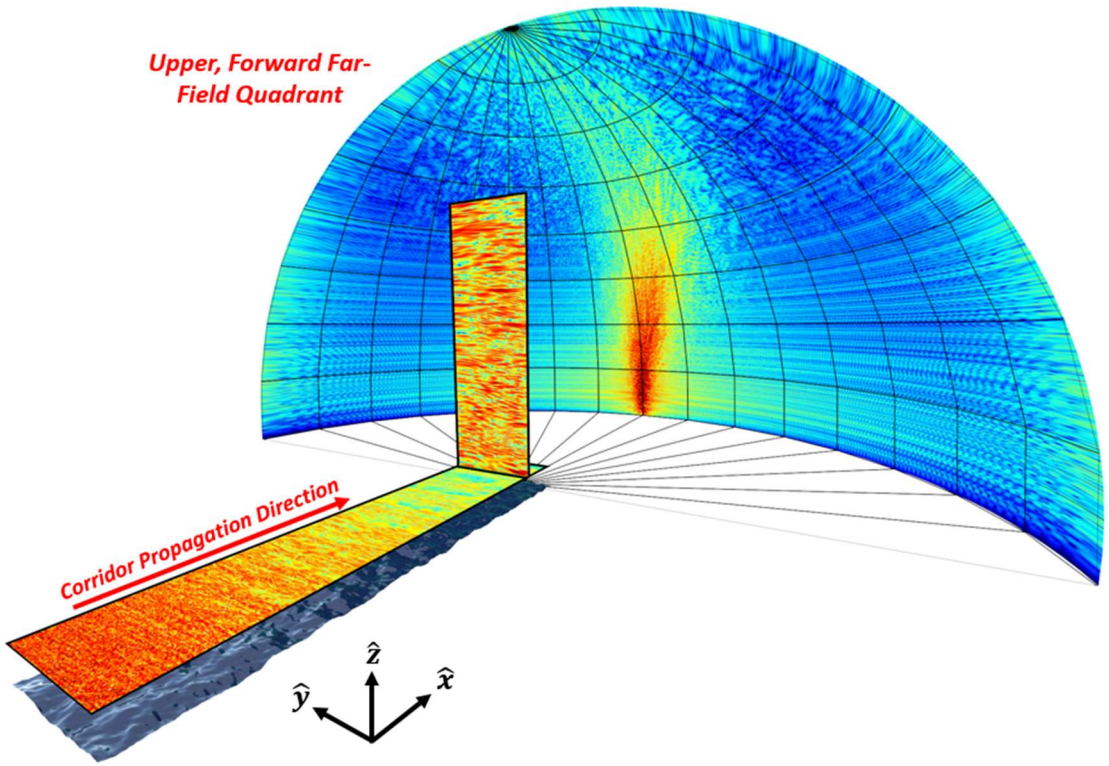

Figure 8.14: AoA calculations are performed on the target plane fields, and determine what fields are radiated to the upper-forward far-field quadrant.

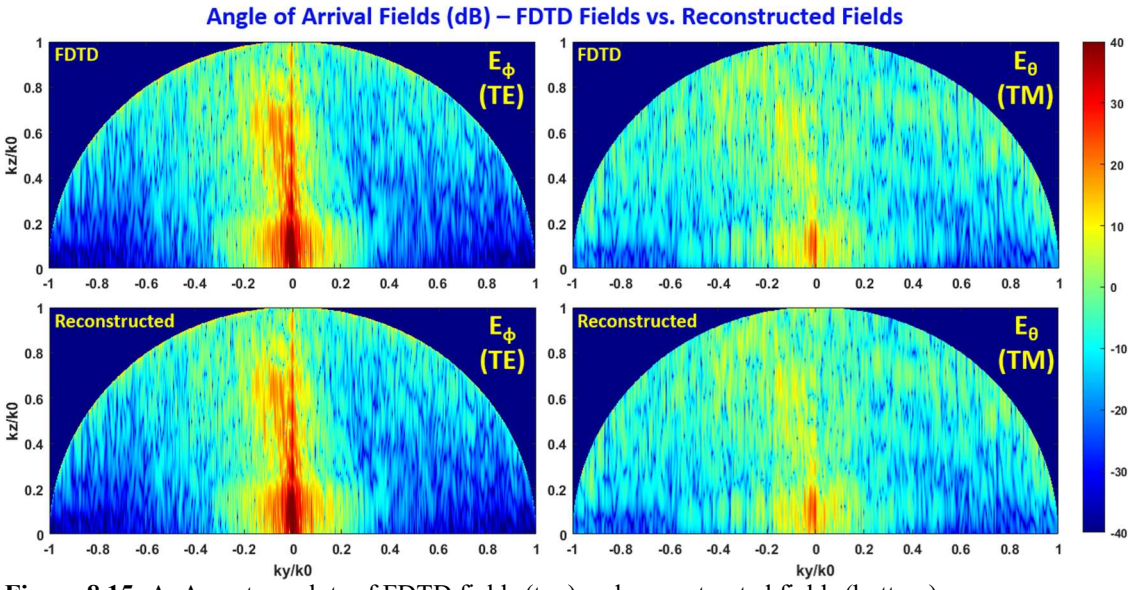

Figure 8.15: AoA contour plots of FDTD fields (top) and reconstructed fields (bottom).

The AoA tells a story similar story to that of the Schelkunoff surface PWS in Figure 7.7 of Chapter 7. Most of the energy at 450m down range is propagating over a fairly narrow band of angles centered about the middle (y = 0, k_y = 0) xz-plane of the corridor, and mostly concentrated toward the forward horizon, with the realistic sea sending strong flashes up toward higher elevations. Much like with the PWS of the Schelkunoff surface, a significant presence of cross-polarized (TM) field is observed to be generated by the realistic sea, which otherwise wouldn't exist in a 2D simulation space. The cross-pol fields in this case are roughly 20dB below the co-polarized fields. Keep in mind, this comparison is performed over a plane of very small height (about 10 wavelengths) sitting close to the sea surface, thus this plane contains mostly very low-angle fields accumulated over a large length of the sea, while only a minimal portion of high-angle fields produced by the local sea is captured. In Section 8.3, much larger target planes are considered, ranging from 500 up to 900 wavelengths in height. In order to take a closer look at the AoA results, Figure 8.16-Figure 8.19 contain line plots cut from constant angles of inclination and azimuth, θ and ϕ, respectively. Under the physics definition of spherical coordinates, θ = 90o and ϕ = 0o represents the forward horizon.

While the E_z scattered fields of the comparison region are essentially the sole sources of the cross-polarized (TM) AoA fields, the co-polarized (TE) AoA fields contain contributions from both E_y and E_z scatter, depending on the direction of propagation. For propagation near the forward horizon, the E_y scattered fields are the dominant contributor. However, outside of the $k_y = 0$ plane and as k_z increases, the TE AoA fields gain an increasing dependence on E_z. A reference back to the relative field error plots between the FDTD and reconstructed fields of the comparison region in Figure 8.13, which shows

closer agreement in E_y, would lead to expectations of better AoA agreement within the regions in which E_y fields are the primary contributor, and worse agreement where contributions from E_z increase.

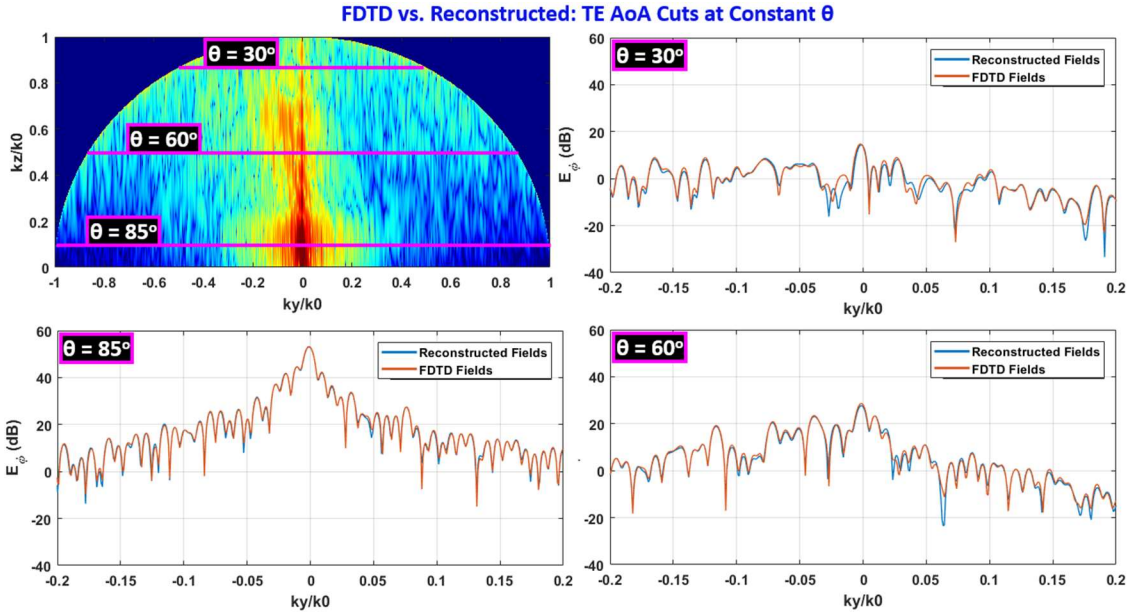

Figure 8.16: AoA E_ϕ (TE) line cuts at constant values of θ, over a range of $-0.2k_0 \leq k_y \leq 0.2k_0$.

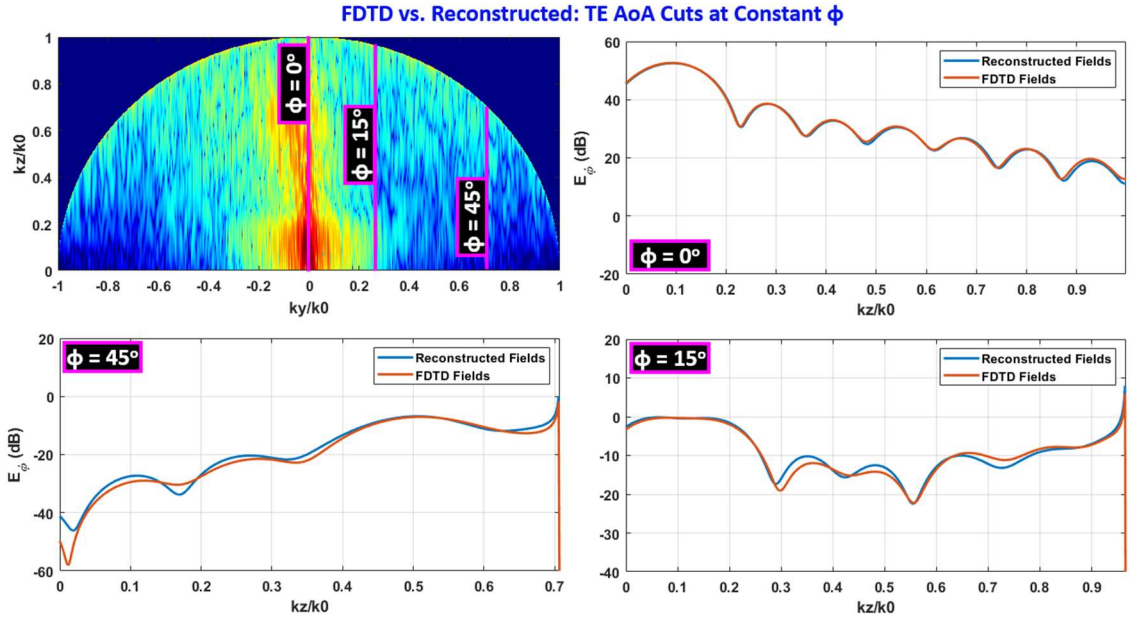

Figure 8.17: AoA E_ϕ (TE) line cuts at constant values of ϕ, over all k_z in the visible region.

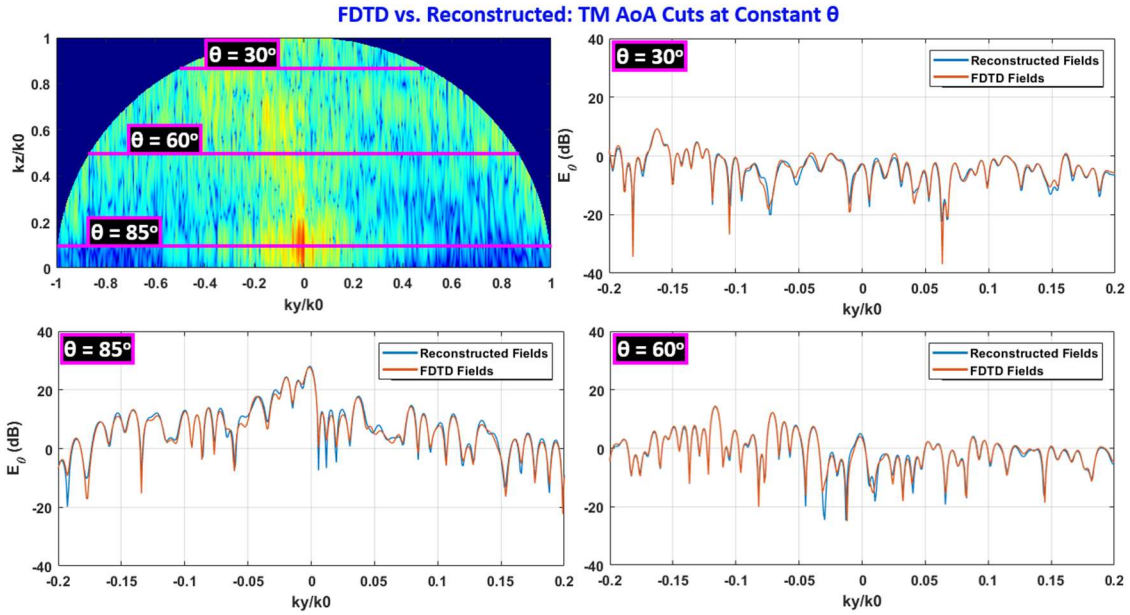

Figure 8.18: AoA E_θ (TM) line cuts at constant values of θ, over a range of $-0.2k_0 \leq k_y \leq 0.2k_0$.

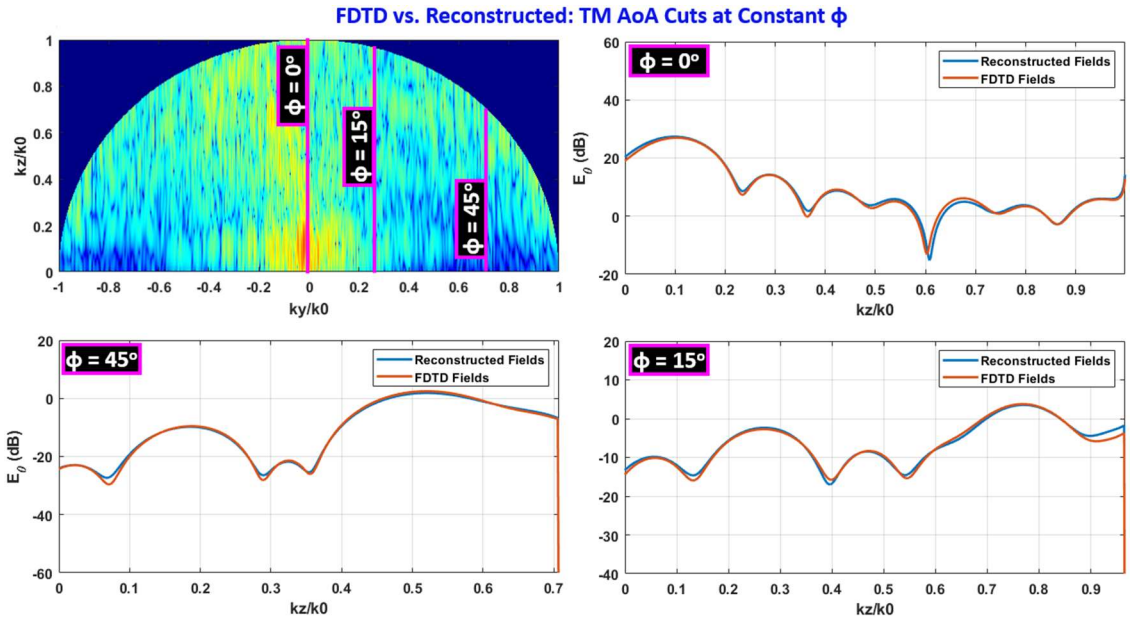

Figure 8.19: AoA E_θ (TM) line cuts at constant values of ϕ, over all k_z in the visible region.

The constant-θ TE plots (Figure 8.16) show very strong AoA agreement between the FDTD and reconstructed scattered fields. This agreement is particularly strong in the regions within about 30dB of the peak of the TE fields (53dB), which corresponds closely to the regions most heavily dependent upon the E_y sources. The constant-ϕ plots (Figure

8.17) demonstrate the same behavior, but it's clear the regions of weaker fields are more susceptible to the target plane field variations. The TM fields in Figure 8.18 and Figure **8.19** show good AoA agreement between FDTD and the reconstructed fields as well, however not as strong of an agreement as most of the TE fields. This can be attributed to the TM AoA fields' sole dependence upon the E_z scattered field sources.

Before analysis of relative error, it's important to discuss how the AoA is applied practically in order to determine the relevance of any given error. A common practice for the investigation of radar scattering off of specific sea-based targets is to divide up the problem into two regions. The first region propagates the radar pulse from the source to the target plane, where the output of region 1 consists of the AoA on the target plane. Propagation in region 1 is computed by methods similar to VTRPE or our FDTD model. The second region consists of a detailed model of the scattering target, typically within a ray-tracing environment. The input for the ray-tracer in region 2 consists of the AoA output on the target plane from region 1, where each input "ray" corresponds to a given peak in the AoA above a given field strength threshold. Each peak in the AoA is converted to a plane-wave ray with the corresponding amplitude, phase, and direction, with typical thresholds set at 25dB below the peak field strength of the AoA. That is, all field information more than 25dB below the AoA peak is completely disregarded. Therefore, for practical purposes, weaker AoA fields are not as useful, and at very low signal strengths large errors don't in fact appreciably matter. Relative error plots along with their regions of relevance for the TE and TM cuts at $\theta = 85^o$ are shown in Figure 8.20 and Figure 8.21, respectively.

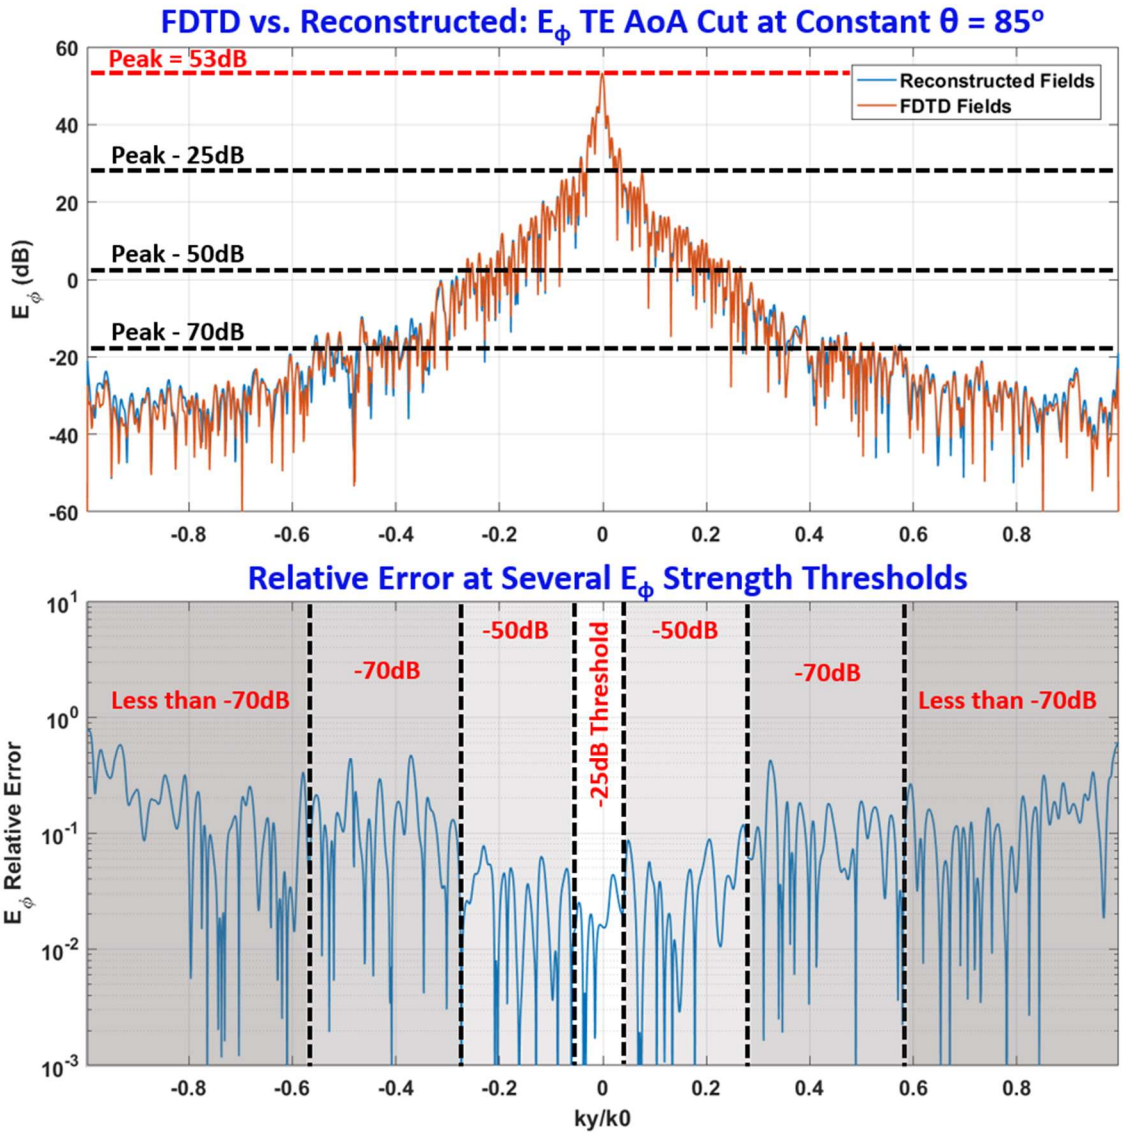

Figure 8.20: AoA relative field error in E_ϕ (TE), showing signal strength thresholds below the peak value.

The relative error for the TE AoA at $\theta = 85^o$ is excellent in the vicinity of the field peak. Within a threshold of -25dB below the peak, a maximum relative field error of 0.041 (-28dB down) is observed, and across almost the entirety of the -50dB threshold band the relative field error is well below 0.1 (-20dB). As the signal becomes weaker, toward the -70dB threshold and below, the relative differences in the FDTD and reconstructed field AoA results begin to rise above 0.1, but this is of little concern as the signal is exceedingly

weak and not relevant to most practical applications. The TM fields at $\theta = 85^o$ have a lower dynamic range than the TE fields, with a somewhat higher average relative error, which peaks in a few places at 0.2 (-14dB), but mostly remaining roughly constantly below 0.1 (-20dB) across the spectrum.

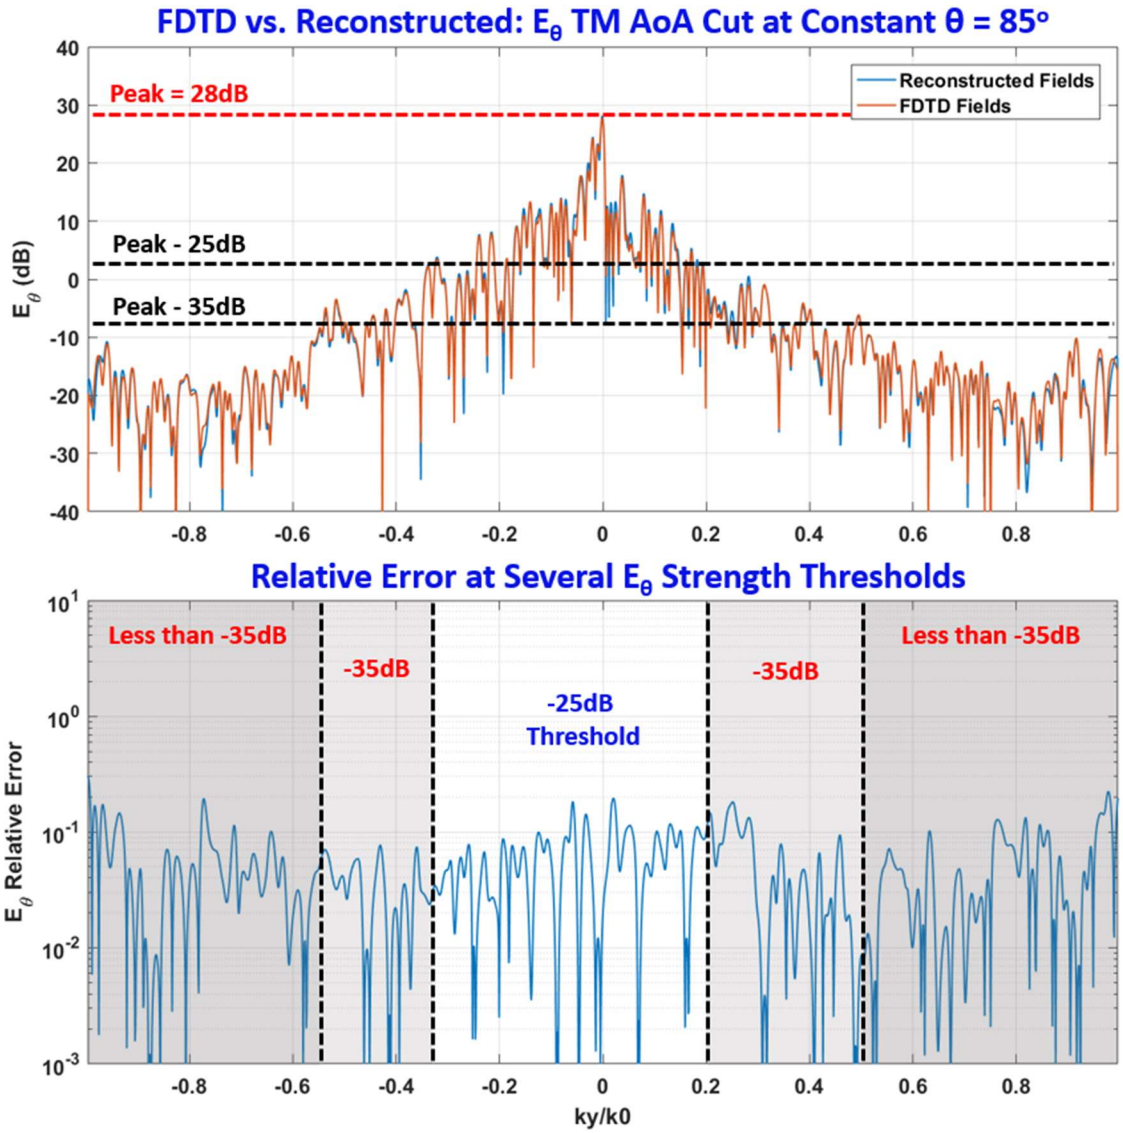

Figure 8.21: AoA relative field error in E_θ (TM), showing signal strength thresholds below the peak value.

The relative field error in the -50dB threshold region of the TE case roughly corresponds to the E_y field errors in Figure 8.13 (left), which can be expected since E_y is

essentially the only contributor to E_ϕ in this region. The TM case is more interesting, however. Despite the regions of high error seen in the cross-polarized field in Figure 8.13 (right), the TM AoA error is largely unaffected. Appendix B contains relative field error plots for the remainder of the line cuts shown in Figures 91-94, the results of the collective are on par with or better than Figure 96, for both TE and TM.

Colormaps of the relative error for the entire AoA far-field quarter-sphere are shown in Figure 8.22 and Figure 8.23 for TE and TM, respectively. To aid the eye, a dashed red boundary has been added to indicate the region in the map which lies within 50dB of the peak co-polarized field (i.e. all fields above +3dB in strength). The plots outside the region have been de-emphasized to remind the reader they are not relevant to the target plane rays which are to be scattered from the target. We conclude that the relative errors in the reconstructed fields compared to the FDTD fields do not adversely affect the AoA calculations and we are free to use the spatial integration techniques for 3D field reconstruction everywhere over the target plane.

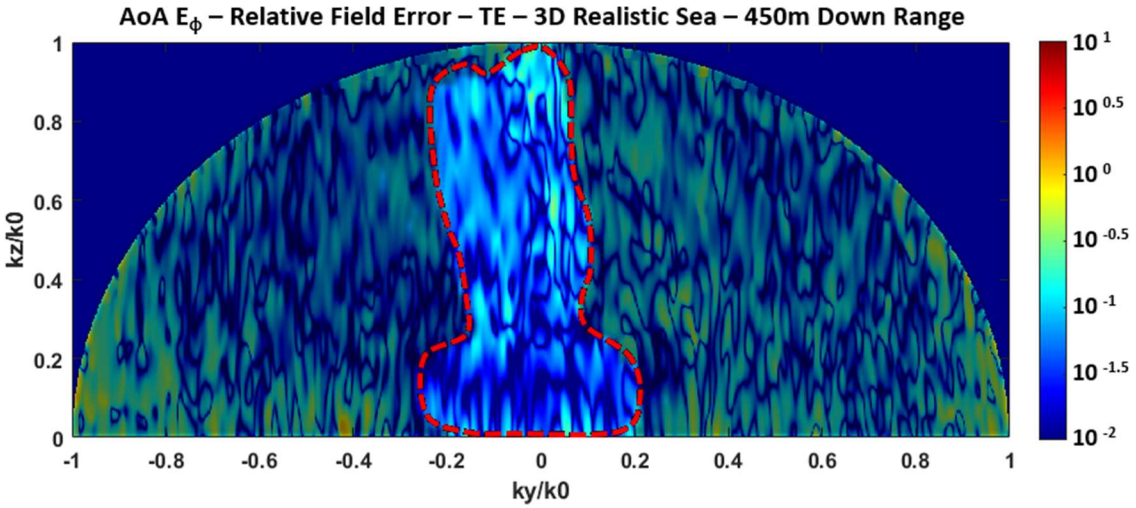

Figure 8.22: FDTD vs. Reconstruction: TE AoA (E_ϕ) relative error colormap, highlighted region surrounded by the red dashed line represents fields within 50dB of the peak of the TE (E_ϕ) field.

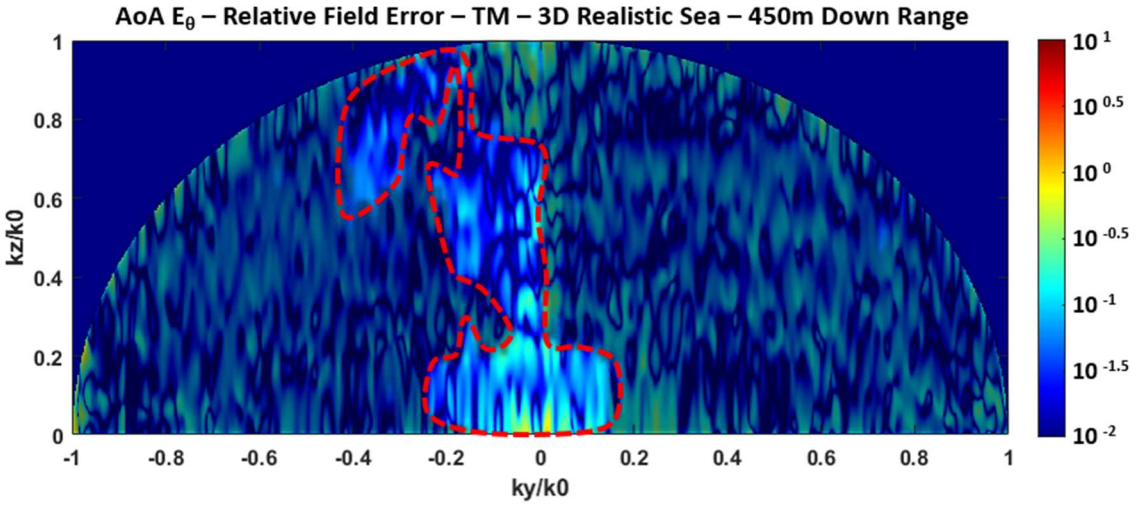

Figure 8.23: FDTD vs. Reconstruction: TM AoA (E_θ) relative error colormap, highlighted region surrounded by the red dashed line represents fields within 50dB of the peak of the TE (E_ϕ) field.

8.3 Angle of Arrival Results of Practically-Sized Target Planes

In practical scenarios, almost all sea-based radar targets extend to altitudes well beyond the roughly 1m above MSL space covered by the 10 by 150 wavelength overlap region in the previous section. This fact precisely represents the reason why analytic VEP injection in combination with field reconstruction is required. While the FDTD space is much more limited by computational resources, the field reconstruction is far less bound, particularly in terms of computer memory. In principle, since the primary limiting factor is run time, it's possible to reconstruct a target plane up to any altitude. Two FDTD sea simulation domains are presented in this section with practically-sized reconstructed target planes. A 9GHz y-directed point dipole source is used, placed 20m above MSL in each domain. Both domains are comprised of 5m-wide (150λ) corridors of realistic 3D sea surfaces, each corridor's sea is unique, and each contains a homogeneous atmosphere. Domain 1 is simulated to a range of 240m (7,200λ), while domain 2 extends to 485m (14,550λ) down range.

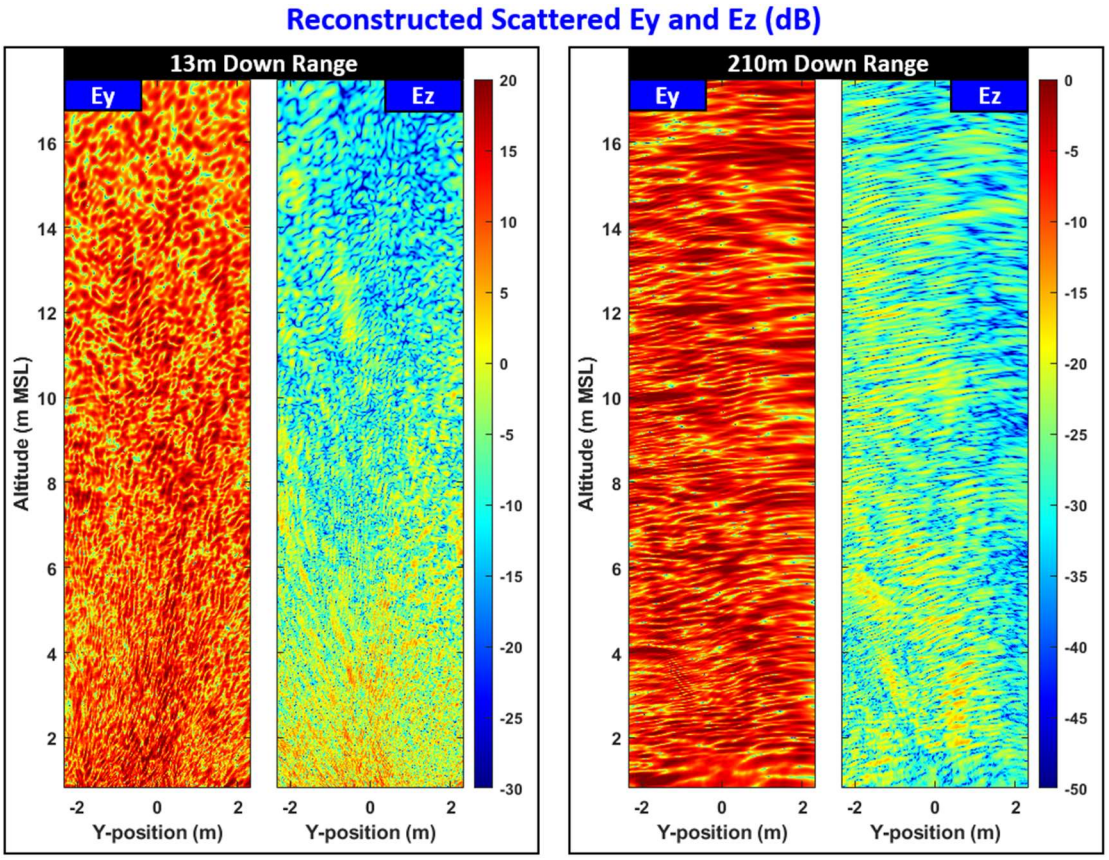

Figure 8.24: Domain 1 reconstructed scattered E_y and E_z on a 5m-wide, 17m-tall target planes placed down range at 13m (left) and 210m (right).

Two target planes have been reconstructed for domain 1: one 13m down range, and the second at 210m down range. The fields are reconstructed up to an altitude of 17m (500λ) and are shown for each target plane in Figure 8.24. The target plane fields 13m down range are very near to the source, so a very broad range of scattering angles exists at this range, which creates a very chaotic field picture. Therefore, we expect the resulting AoA calculations at 13m should include a large amount of scatter over a broad range of k_y and concentrated toward higher valued of k_z. Further down range at 210m most of the energy of the forward scatter consists of rays travelling over a narrower band of angles which is concentrated toward the forward horizon (lower values of k_z) and centered mostly about the middle of the corridor ($k_y = 0$).

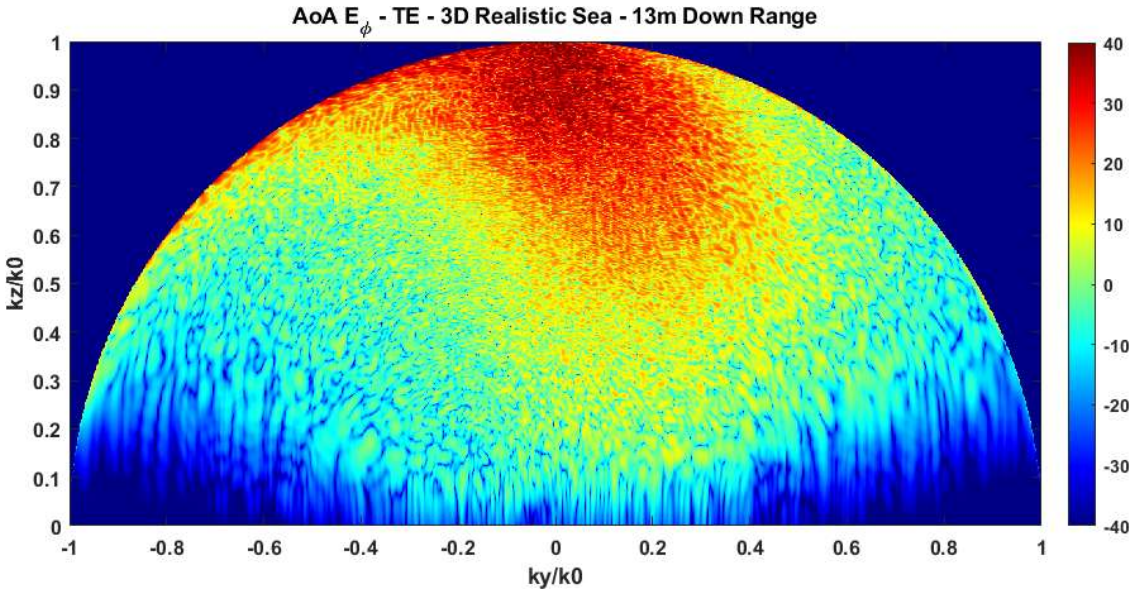

Figure 8.25: Domain 1 TE AoA (E_ϕ) on a 5m-wide, 17m-tall target plane 13m down range.

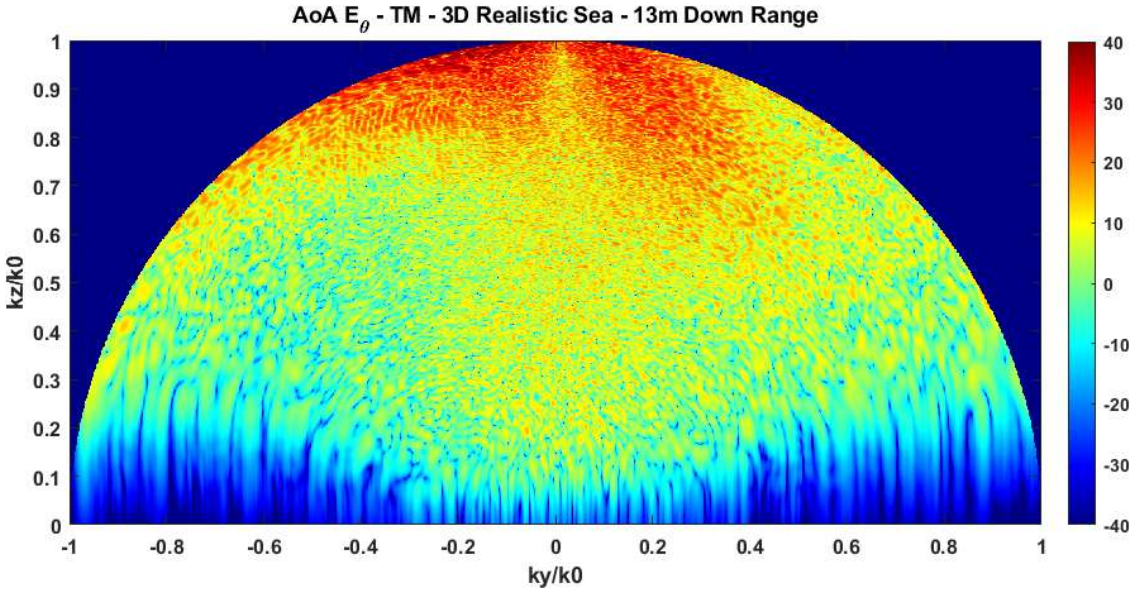

Figure 8.26: Domain 1 TM AoA (E_θ) on a 5m-wide, 17m-tall target plane 13m down range.

Figure 8.25 and Figure 8.26 show the TE and TM AoA results at 13m down range, which clearly demonstrates scatter over a very broad angular spectrum, with an appreciable amount of energy scattering out toward the horizons of visible k-space. The peak of the TE AoA field is concentrated around the $k_y = 0$ plane, which represents the center plane of the corridor. This is expected for the TE fields since the y-directed point dipole source

produces purely TE fields in the $k_y = 0$ plane. The TM AoA fields on the other hand demonstrate a dip in amplitude around the $k_y = 0$ plane since the source produces no TM fields on that plane. The 3D nature of the sea, however, generates cross-polarized fields in the center plane of the corridor, which is a purely 3D effect unable to be captured by 2D.

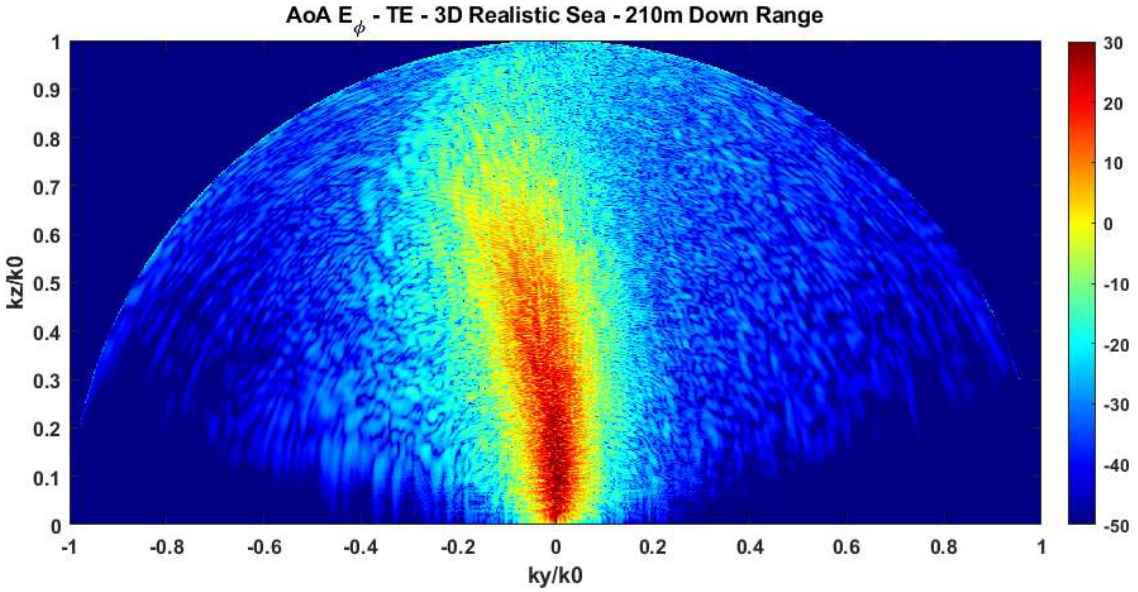

Figure 8.27: Domain 1 TE AoA (E_ϕ) on a 5m-wide, 17m-tall target plane 210m down range.

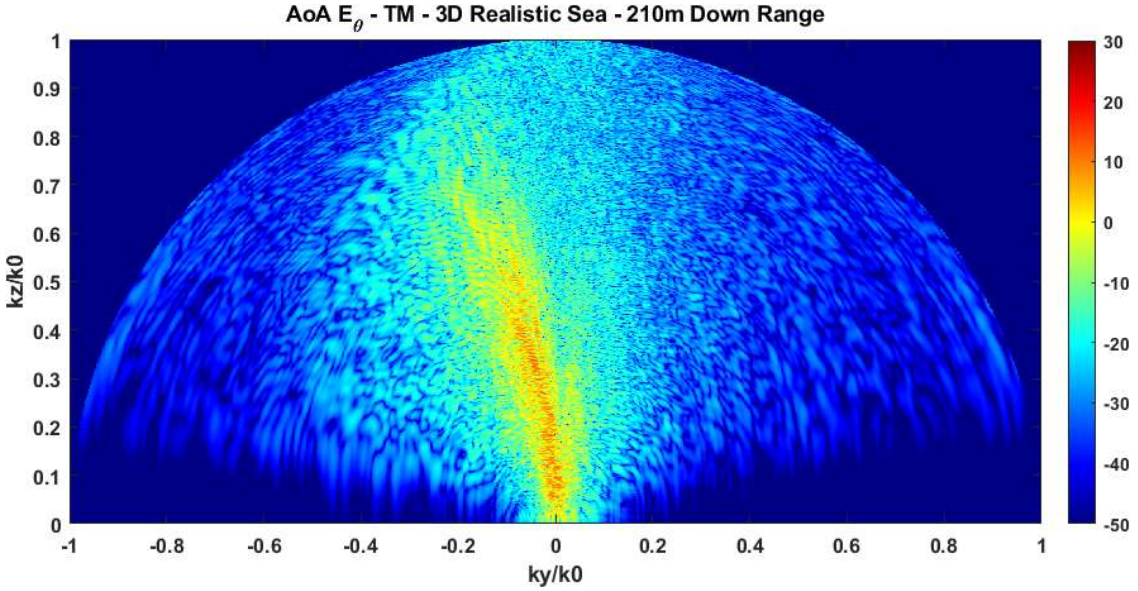

Figure 8.28: Domain 1 TM AoA (E_θ) on a 5m-wide, 17m-tall target plane 210m down range.

Figure 8.27 and Figure 8.28 demonstrate the narrowing angular spread of the scatter within the corridor further down range at 210m. The forward-travelling scatter in the vicinity of the sea surface naturally tends to accumulate at shallow angles approaching the forward horizon, since only low-angle rays remain close to the sea as propagation range increases. The corridor's finite width guarantees any scatter with a large transverse propagation angle (large k_y, that therefore would not hit the target plane) to exit the domain (absorbed at the walls of the corridor) after a relatively short distance, thus the fields tend to accumulate toward propagation angles down the center of the corridor around the plane where $k_y = 0$. This phenomenon explains the success of 2D methods (which work only on the $k_y = 0$ plane) in providing reasonable AoA information for the co-pol scatter.

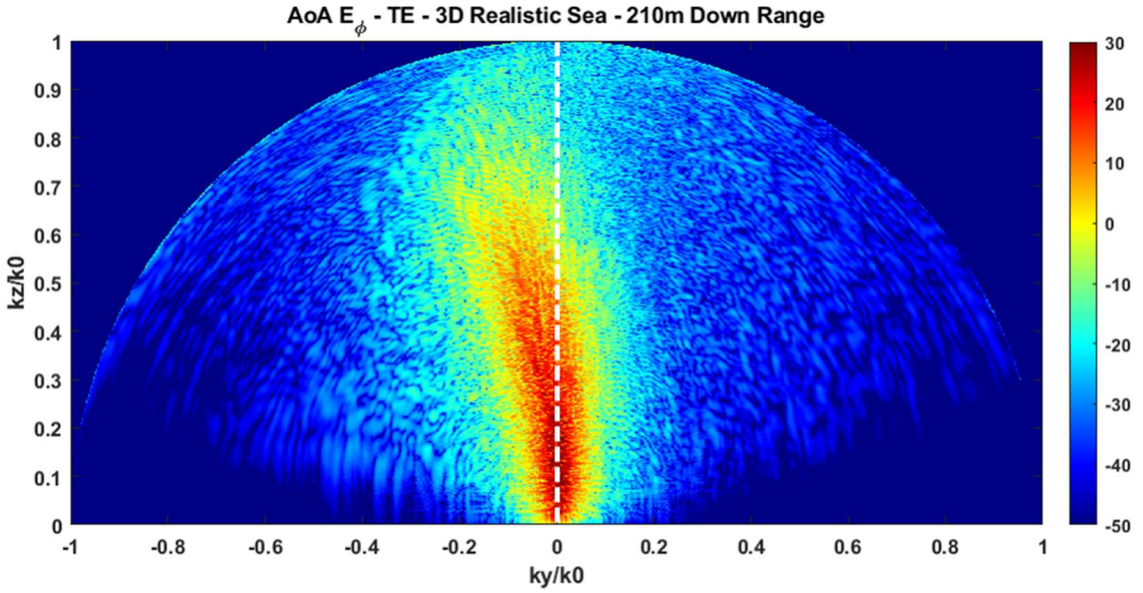

Figure 8.29: The AoA at 210m down range in domain 1 clearly shows a shift toward the -k_z direction.

The AoA of domain 1 demonstrates a particularly interesting trend. Drawing a vertical line at $k_y = 0$ shows a significant amount of energy on the target plane is directed toward the -y direction (Figure 8.29). Some large feature of the sea surface leading up to the target plane appears to be redirecting a significant amount of energy. Sure enough,

examination of the 40m of domain 1's sea leading up to the target plane shows a clear downward slope toward the -y direction (Figure 8.30), which causes the energy redirection witnessed in the AoA. This energy redirection due to transverse variations in the realistic 3D sea comprises another phenomenon unique to the 3D sea environment.

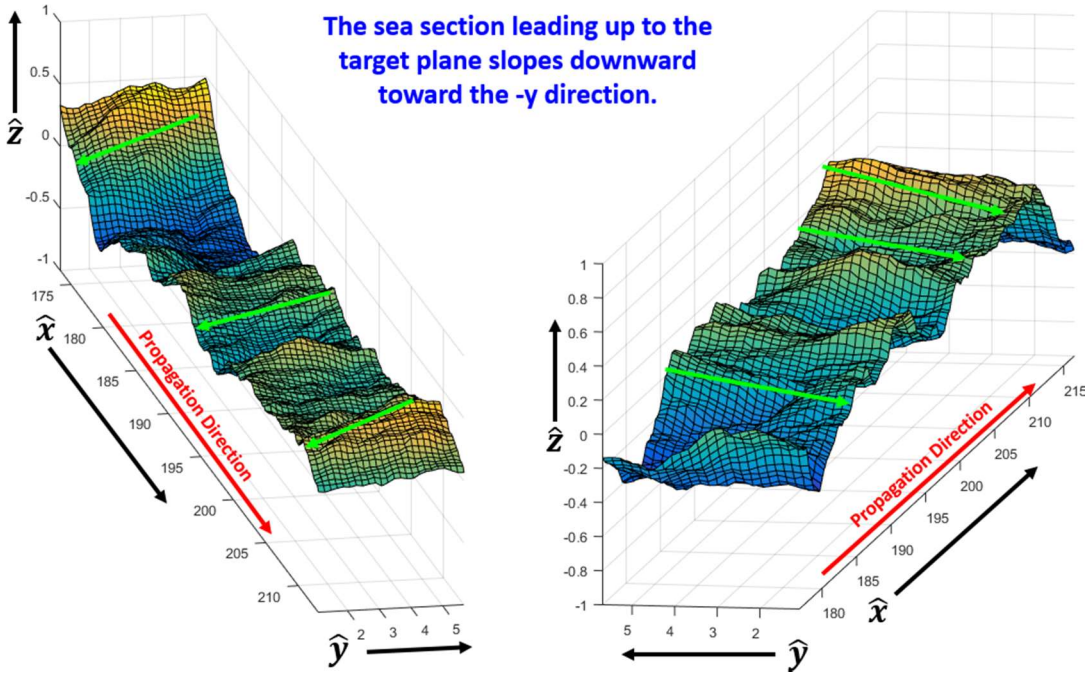

Figure 8.30: The AoA at 210m down range in domain 1 clearly shows a shift toward the $-k_z$ direction.

Domain 2 contains a different sea instance than domain 1 and is more than twice the length. Additionally, the reconstructed target plane, while it's similarly 5m in width, is 30m tall (900λ). Since the target plane for domain 2 is placed 450m down range, its fields are propagating at angles focused toward the forward and about the center plane of the corridor (Figure 8.31 and Figure 8.32). The sea leading up to the target plane in domain 2 does not exhibit such well-defined features as in domain 1, so the fields are not as dramatically shifted toward a preferred direction in this case.

In both simulation domains, the influence of the transverse variations of a realistic 3D sea are evident. In fact, everything which exists outside of the $k_y = 0$ line on any one

of these AoA plots is a direct result of the influence of 3D phenomenology and is unable to be captured within a 2D space. Furthermore, is the entire existence of the cross-polarized fields is a direct result of 3D physics. The magnitude by which such phenomena change the present 2D results of the radar scattering properties of sea-based targets has yet to be quantified, but these new phenomena are indeed present and must be investigated.

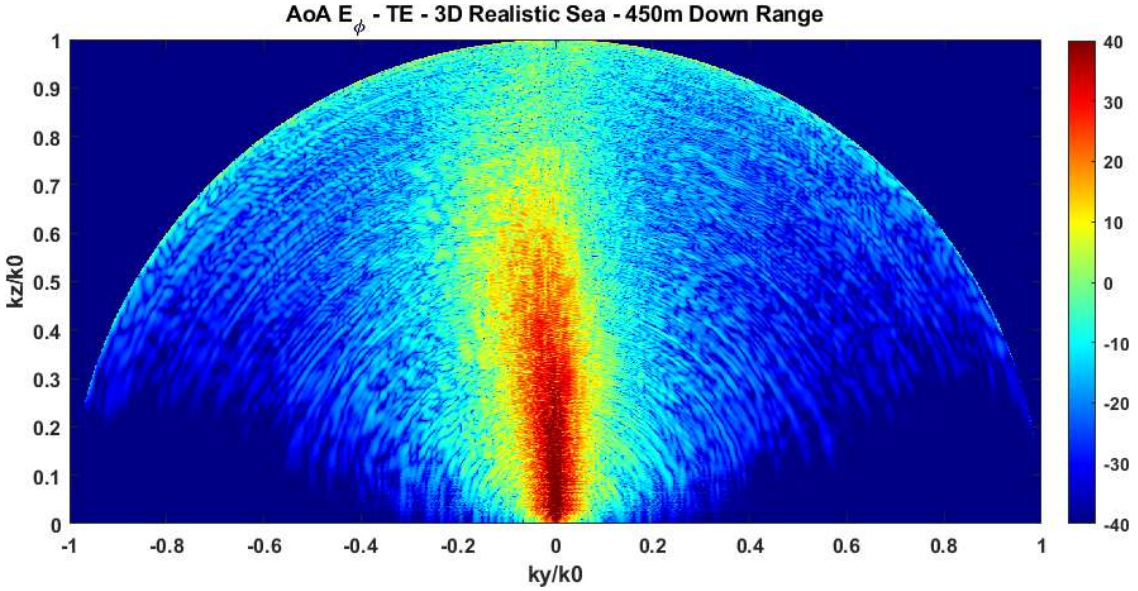

Figure 8.31: Domain 2 TE AoA (E_ϕ) on a 5m-wide, 30m-tall target plane 450m down range.

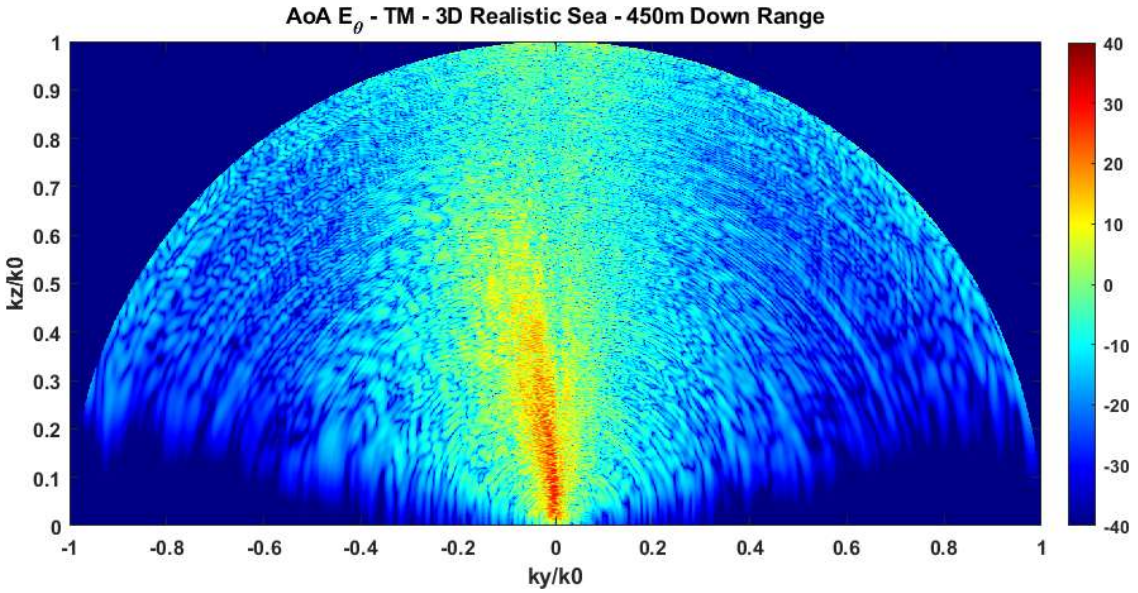

Figure 8.32: Domain 2 TM AoA (E_θ) on a 5m-wide, 30m-tall target plane 450m down range.

8.4 Conclusion

To obtain the scattered fields which interact with a sea-based radar target from the VEP 3D FDTD space, we need to extend beyond the plane wave spectrum analysis of Chapter 7. The PWS produces the far-field radiation, but the plane waves propagating within the Fresnel zone of the sea are the fields which impinge upon the target, and which are input into subsequent RCS analysis applications. To provide this scattered field information, the fields from the Schelkunoff surface need to be propagated to a target plane near the end of the sea corridor. Since the target plane can be 30m tall or more, whereas the compact FDTD computational domain only extends to about 1m above MSL, these fields are calculated via spatial integration-based reconstruction. To verify the accuracy of the integration, a direct comparison was performed with the FDTD scattered fields and the reconstructed fields on a plane inside the FDTD corridor.

For the reconstruction to be accurate, two significant differences between frequency domain electromagnetics and time domain electromagnetics must be remedied. First, in the time domain, FDTD allows rays exiting the computational domain to simply leave without causing any non-physical reflection from the absorbing boundary. However, in the frequency domain, since no information exists outside the computational domain, there is an apparent large discontinuity in the surface currents of the Schelkunoff surface along the walls of the corridor. We apply a computationally inexpensive analytic extension of the complex currents to the edges which minimizes diffraction from the discontinuity. Secondly, in our FDTD domain-decomposition method we carry the forward-propagating scatter from sub-domain to sub-domain since it's the only scatter of interest to a radar target. Local backward scatter within sub-domains is not relevant to the final solution,

however, backward scatter from large sea features can lead to discontinuities on the Schelkunoff surface created by accumulation of energy at the boundary of the current sub-domain with the previous one. Again, such a discontinuity constitutes a diffracting edge, sending false information forward into the target plane. This artefact is eliminated by filtering the plane wave spectrum of the Schelkunoff currents to eliminate the (irrelevant) backward-travelling waves prior to performing the spatial integration.

Pre-processing of the Schelkunoff surface fields prior to spatial integration helps account for the aforementioned discontinuities, and the resulting reconstructed scattered fields at the target plane agree well with the conventionally-produced 3D VEP FDTD scattered fields. The relative error in the E_y (co-polarized) fields is very low, and peaks near -20dB. For the weaker cross-polarized field, E_z, which is overall about 20dB weaker than the E_y field, the typical errors are around -10dB, except for some regions of very faint fields where large errors can appear. We can verify the unimportance of large errors in faint field regions by calculating the plane wave spectrum of the fields at the target plane, which demonstrates that the angle-of-arrival information (AoA) is essentially the same for the 3D VEP FDTD scattered fields and the reconstructed fields. Field reconstructions on very large target planes (up to 30m tall) and their AoA are demonstrated on two corridors of 3D seas, one of which extending in length to almost 500m, the results of which continue to further demonstrate the presence of relevant 3D phenomena warranting further investigation.

CHAPTER 9

A POSSIBLE 2D-3D HYBRID APPROACH TO INCLUDE INHOMOGENEOUS ATMOSPHERES

9.1 Iterative Coupling of a Sea Space to an Inhomogeneous Atmosphere

In order to make the 3D computational burden manageable, the analytic field injection VEP method must be used. The computational domain size is drastically reduced (Figure 2.7 and Figure 2.8) for homogeneous atmospheres under this method since the non-scattering space between the source and the sea can be completely eliminated. After the elimination of most of the atmosphere space via analytic source VEP injection, the total computational space dimensions of the corridor of realistic 3D sea used in Chapter 7 are 2m (50λ) in height, 5.75m (170λ) in width, with sub-domains of length 3m (100λ). This is still an enormous computational burden, requiring more than 32GB of GPU RAM, and over 100GB of system RAM, with a 1km corridor requiring about a week of runtime. To include the same 35m+ atmosphere space in 3D as in 2D would require a supercomputer.

A potential solution involves combining 2D and 3D approaches. The atmosphere's variation is much gentler than the variation in the topography of the sea surface. The vertical variations in atmospheric ducting are of principal concern, while the transverse variations above the corridor of sea are negligible. Therefore, the atmosphere region could be modeled with a coarser grid, or an entirely different computational method, and perhaps even using 2D FDTD. The 3D and 2D spaces can be coupled iteratively across the Schelkunoff surface.

In order to prove the feasibility of such an iterative approach, we prove its efficacy in 2D. The procedure is outlined in Figure 9.1 below. The first iteration begins by propagating the source fields in the "duct space" which contains only the source and the atmosphere, no sea, and collect and store the time-history of the fields as they traverse the Schelkunoff surface at the bottom of the domain – these are all of the downward-travelling fields which would normally impinge upon the sea surface. There are absorbing terminations below the Schelkunoff surface, nothing is scattered back upward through the it at this stage. After the duct space is finished, the recorded time-history is then injected into the sea space, allowed to scatter off the sea, and the upward-travelling scatter is once again collected at the Schelkunoff surface and stored as a time history. This time-history is now re-injected into the duct space, a new time-history is once again collected on the Schelkunoff surface – this time containing what gets scattered back downward by the duct – and the process keeps repeating. The procedure is iterated for as many passes as desired, and a single "pass" will be defined as the sequence of simulations: duct space to sea space and back again to duct space (Figure 9.1).

In principle, the iterative procedure should yield exactly the same results as the full conventional 2D spaces given enough iterations. In order to test the effectiveness of the iterative method, we will examine two test cases. The first test case will be a 2D TE (to the plane of incidence) 9GHz y-directed point source placed 20m above MSL, in the presence of a realistic elevated duct, and propagated for 4km (the same exact parameters as the VTRPE comparison space shown in in Figure 3.12). The second scenario involves a 6m-wide 9GHz TEy aperture source, aimed 1^{o} above the forward horizon, centered 10m above MSL above a 4km flat plane of sea material in the presence of a highly exaggerated linear

duct profile (Figure 3.4), the same profile used in the FB-BIE and SSF-PWE comparisons of Section 3.2. In both cases the sea space contains the inhomogeneous atmosphere as well.

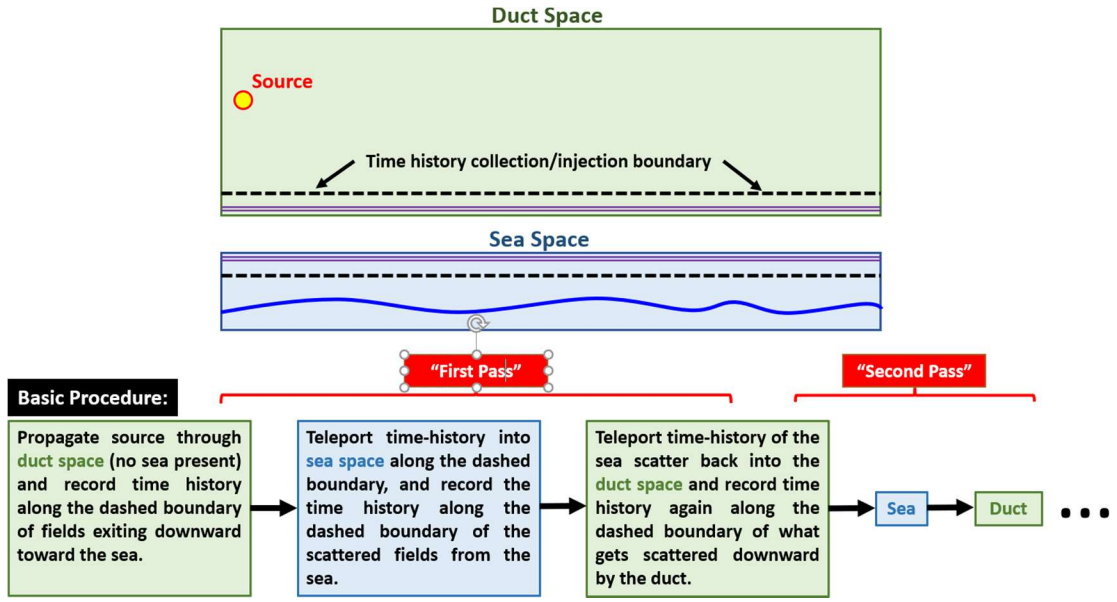

Figure 9.1: Division of the FDTD space for the iterative injection technique, and outline of the procedure.

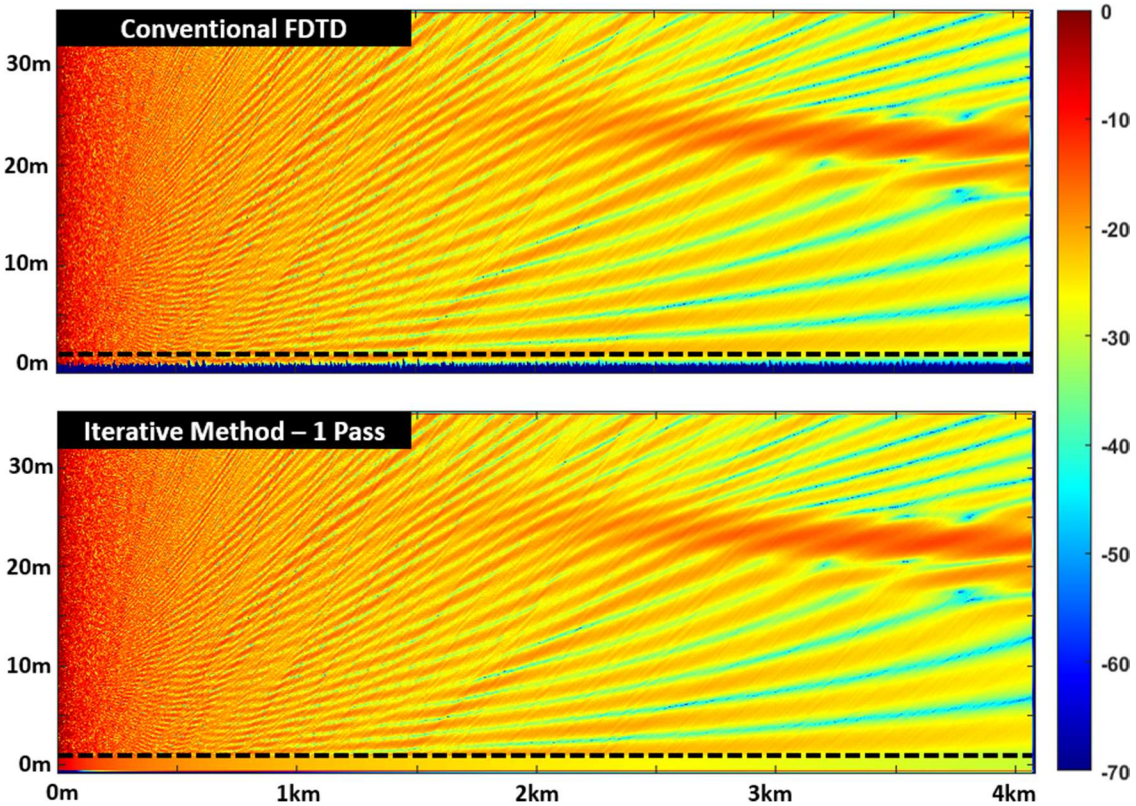

Figure 9.2: Contour plots of the conventional FDTD results (top) and after the first pass of the iterative injection method (bottom) for the first 2D test case containing a sea and elevated duct atmosphere.

The results of the first test case of the point source above the sea in the presence of an elevated duct are shown in Figure 9.2. The results are virtually identical after only a single pass. Figure 9.3 shows a contour map of the fields during the last part of the first pass, where the sea scatter is re-injected back into the duct space. The fields seen below the Schelkunoff surface are those scattered back downward toward the sea by the duct. They are very weak compared to the principal scatter off the sea but increase in strength with distance as the duct diffracts more of the energy downward.

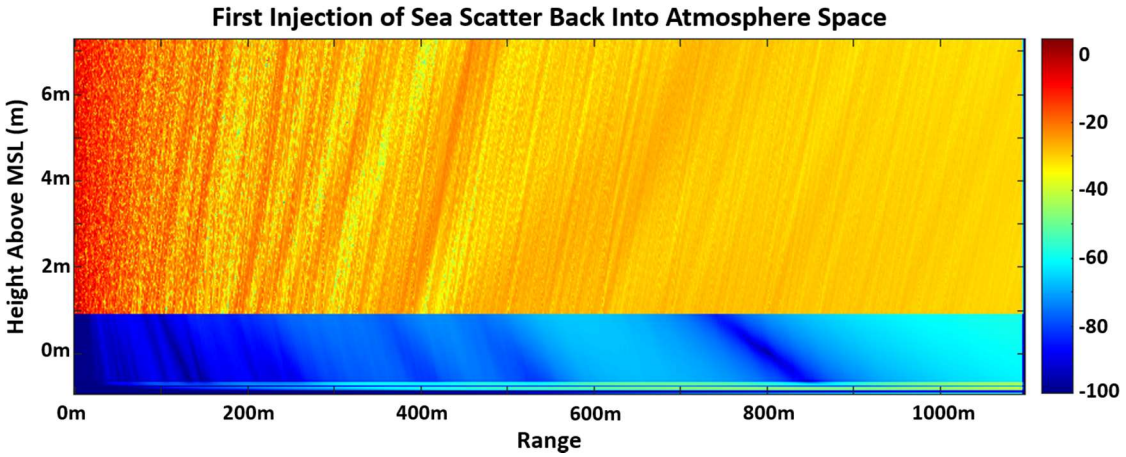

Figure 9.3: Colormap of the fields at the end of the first pass, when the sea scatter is injected back into the duct space.

In the following Figure 9.4 and Figure 9.5, cuts of the iterative injection fields are compared to the conventional FDTD case at 500m and 4,000m down range, respectively. Nearly the same total field picture occurs 500m down range for both the first and second passes. Further down range, by 4,000m, multiple reflections between the sea and duct become more relevant. The absence of scatter beyond the first "bounce" off the sea is noticeable. The second pass is even closer than the first, lending credence to the iterative approach's validity. The relative error in the E_y field with respect to conventional FDTD (Figure 9.6) falls below -30dB at 500m and mostly below -20dB for 4,000m down range.

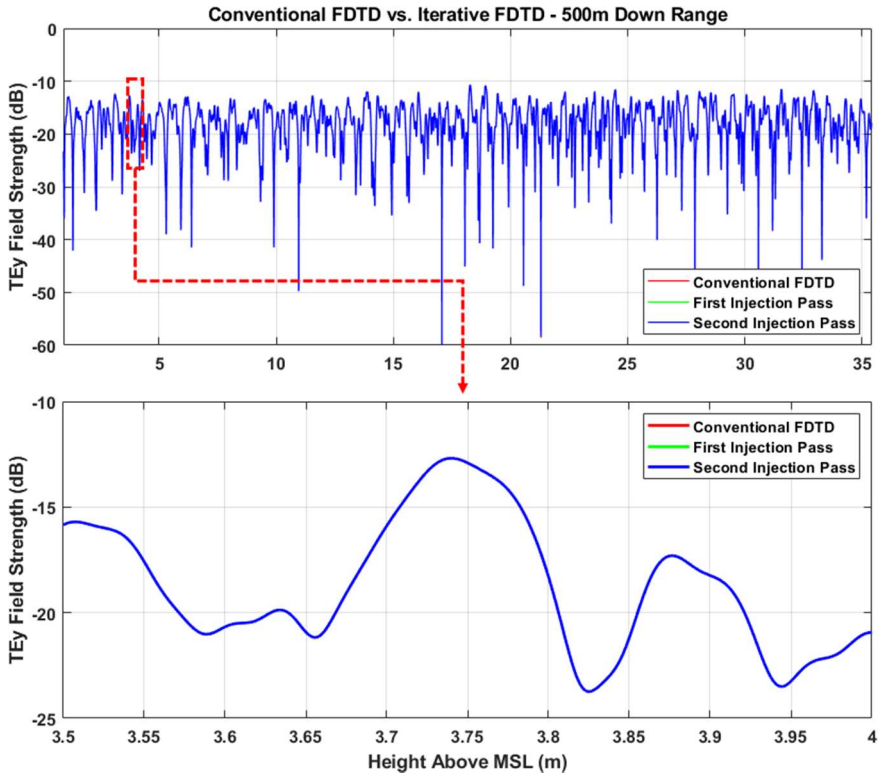

Figure 9.4: Line plot comparisons of Ey taken 500m down range.

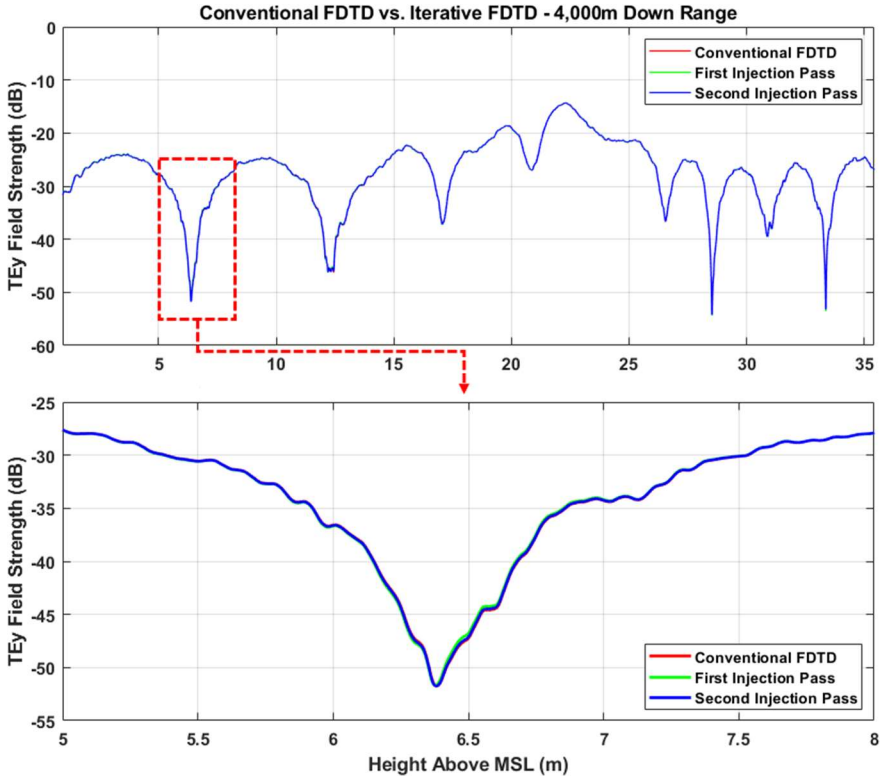

Figure 9.5: Line plot comparisons of Ey taken 4,000m down range.

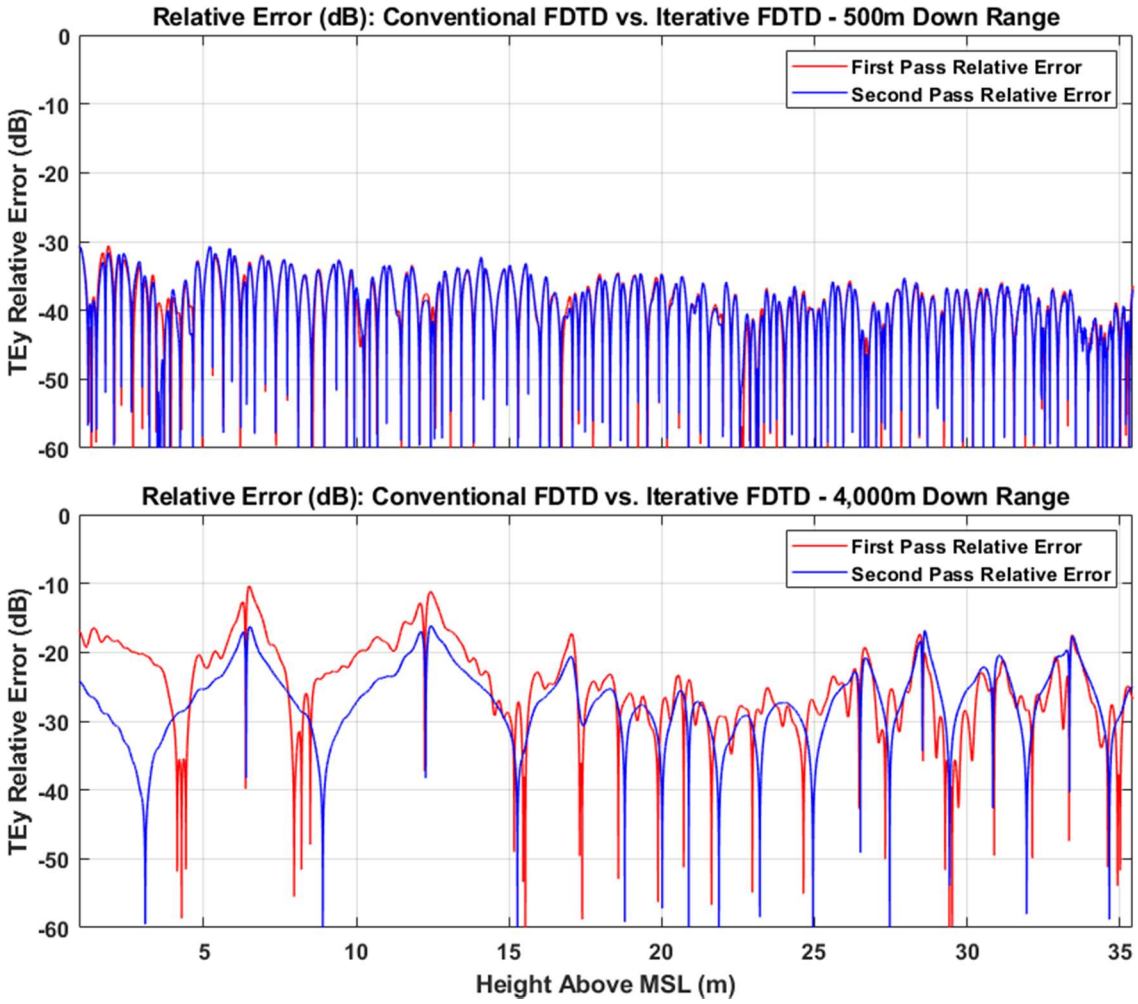

Figure 9.6: Relative error in Ey for the line cut comparisons 500m and 4,000m down range.

The second case, which involves the exaggerated linear duct above a flat sea, requires far more iterations than the elevated duct scenario. This is due to the multiple bounces off the flat sea which the source pulse undergoes, clearly visible in Figure 9.7. This represents, for all practical purposes, a worst-case scenario for the iterative method. The combined results of four iterations are shown compared to conventional FDTD, with the contributions of the first several iterations shown in Figure 9.8.

Figure 9.8 intuitively demonstrates how the iterative method accumulates a more correct answer with each iteration. It's easy to see how each "bounce" is re-injected back

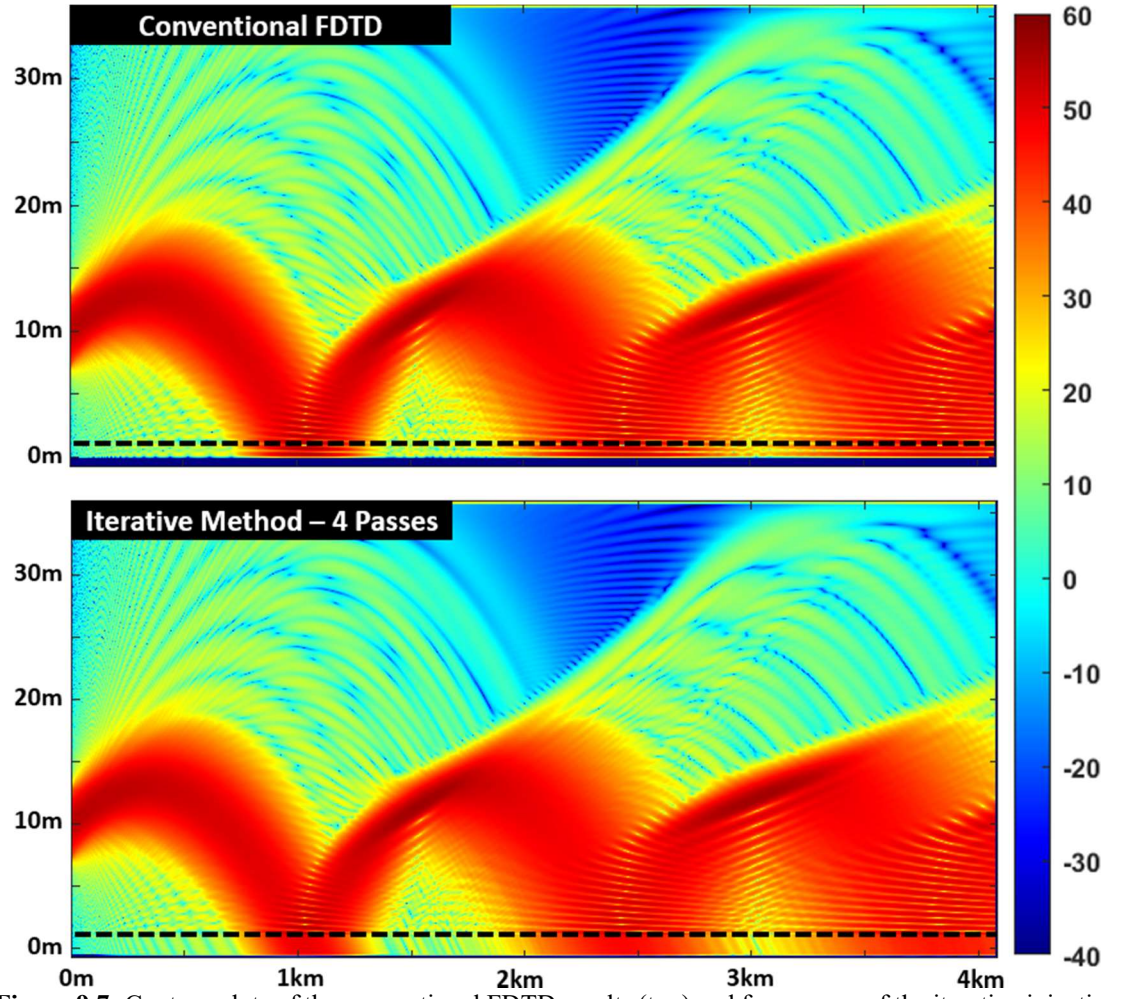

Figure 9.7: Contour plots of the conventional FDTD results (top) and four passes of the iterative injection method (bottom) for the second 2D test case containing a flat sea and exaggerated linear duct profile.

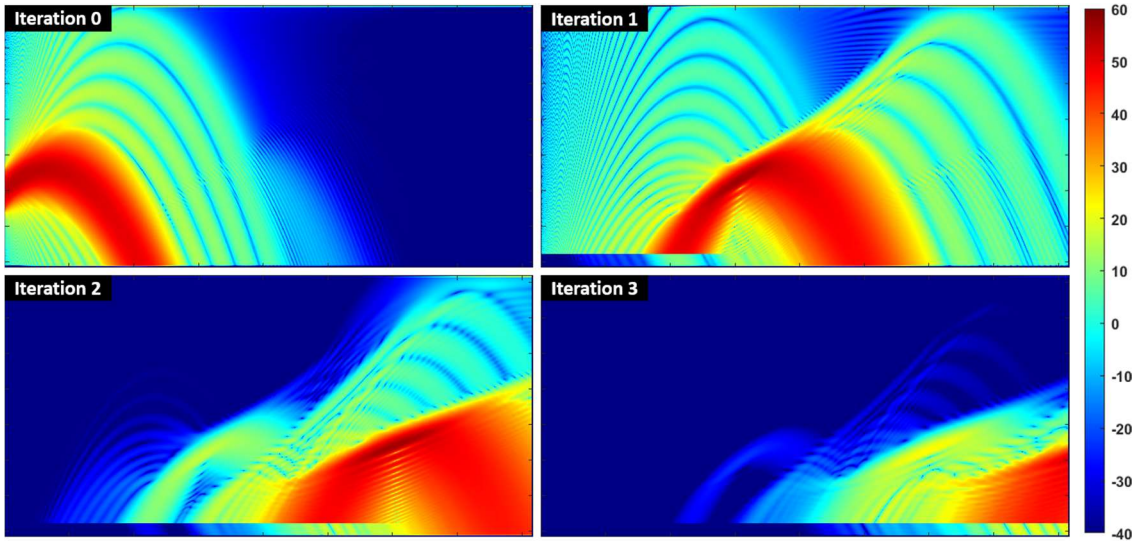

Figure 9.8: Contour plots of the duct space during the first few iterations of the second test case.

into the duct space, which leads to subsequent bounces. The sum total of each iteration constitutes the whole solution. Note that in Figure 9.4-Figure 9.9 a dashed line has been added to the iteration result to remind the reader that during iteration of the upper duct space, the region below the dashed line is empty and terminated in an absorbing boundary since the reflecting surface is contained in the "sea space".

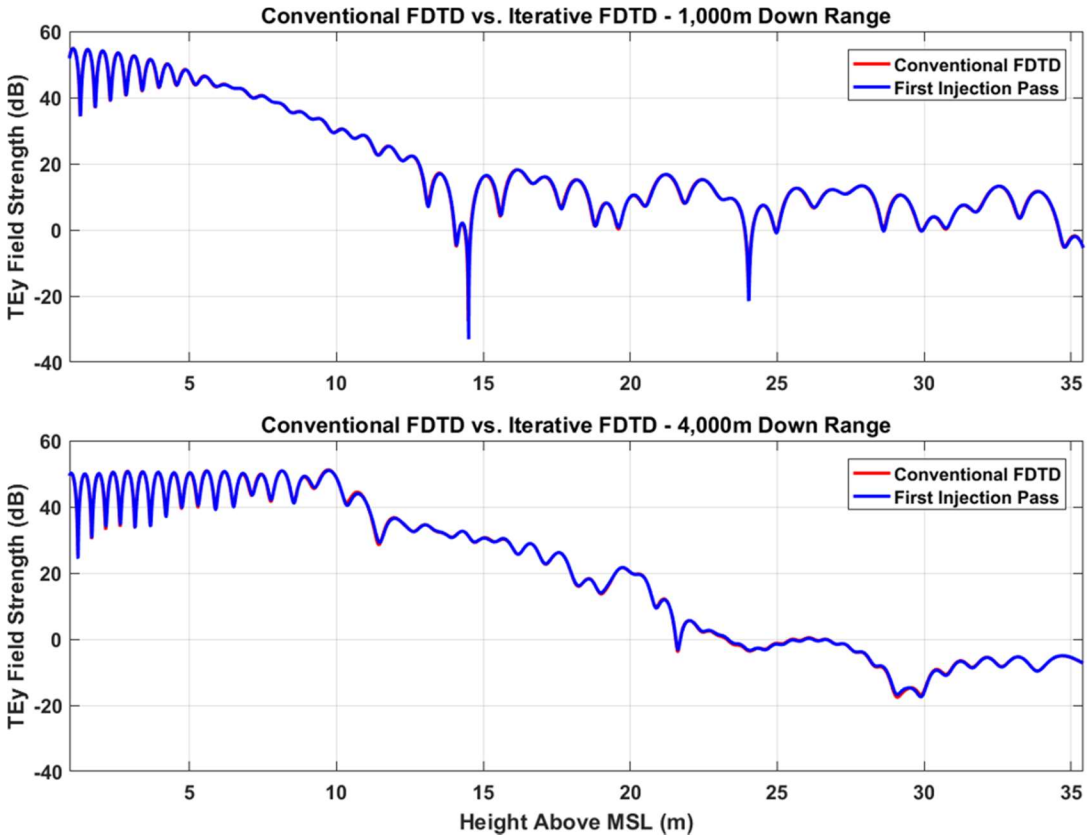

Figure 9.9: Line plot comparisons of E_y at 1km and 4km ranges for the exaggerated duct case.

It's evident, based on both test cases, the iterative method works exceptionally well when both the duct space and sea space are two-dimensional. It stands to reason that if both spaces were three-dimensional, the method would work equally as well. A method has yet to be completely devised for coupling a 3D sea space with a 2D ducting space, and is the subject of current investigation. The iterative technique is a likely candidate for incorporating ducting into the 3D sea corridor at a reasonable computational burden.

9.2 Hybridizing 2D and 3D: Bridging 2D FDTD into 3D FDTD

Though 3D corridors as long as 1km have been run, it is impractical to expect to run the entire problem (tens of kilometers) in 3D. However, based on our 2D results, we know the details of the field structure on a target plane 30m tall is almost entirely due to the scattering from the last kilometer of sea in front of the target plane. Therefore, the fields propagated from the source to a range of X kilometers can be considered as the incident fields which illuminate the last kilometer of sea in front of a target plane located at (X + 1) kilometers. And it is these fields which should be injected into a 1km 3D domain to develop the local 3D phenomena which impacts the target of concern. This means that the first X kilometers of sea can be simulated in 2D with the most efficient computational tool available (FDTD, VTRPE, etc.), and the last kilometer or so of the total space is simulated in 3D FDTD. If an inhomogeneous atmosphere is present, then additional hybridizing of the last kilometer with a 2D space as discussed in Section 9.1 could also be done.

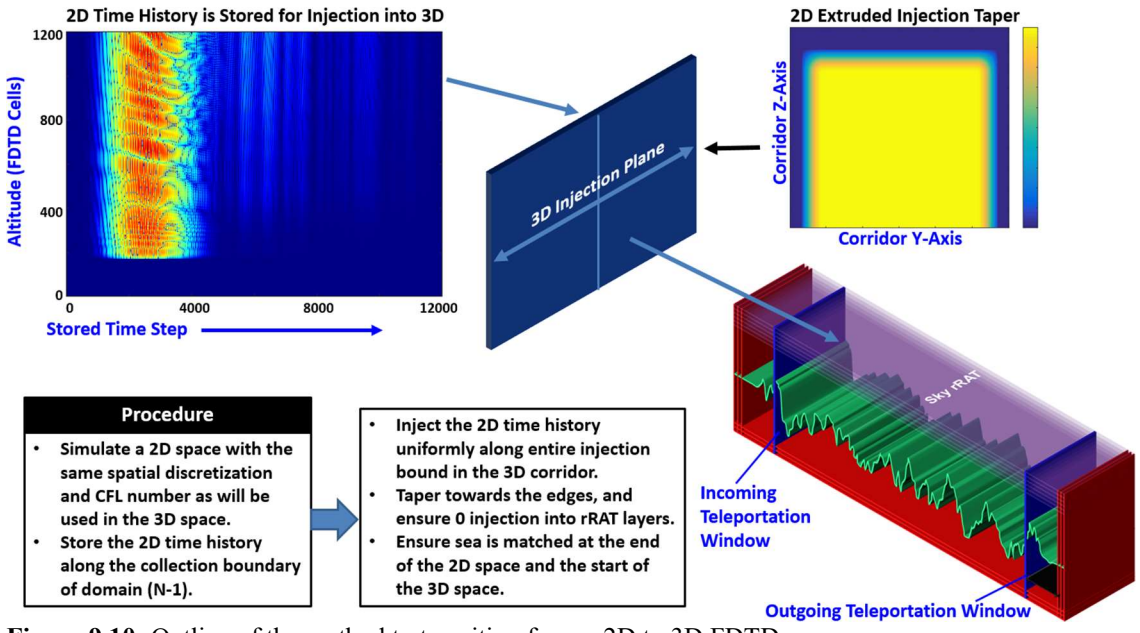

Figure 9.10: Outline of the method to transition from a 2D to 3D FDTD space.

In 2D FDTD to 3D FDTD coupling (outlined in Figure 9.10), first the 2D space is simulated at the same spatial step and CFL number as the 3D space, using the center contour of the realistic 3D sea as the sea surface. The time-history of the outgoing fields is collected from domain N-1 and stored. The stored 2D time-history is then injected along the starting plane of the 3D corridor using equivalent surface currents; the injected currents are extruded along the injection plane and tapered along the outside edges of the corridor to prevent excessive artificial scattering as well as injection directly into the rRAT absorbing boundaries. The seas of the 2D and 3D spaces are matched up to be continuous along the center contour at the injection plane.

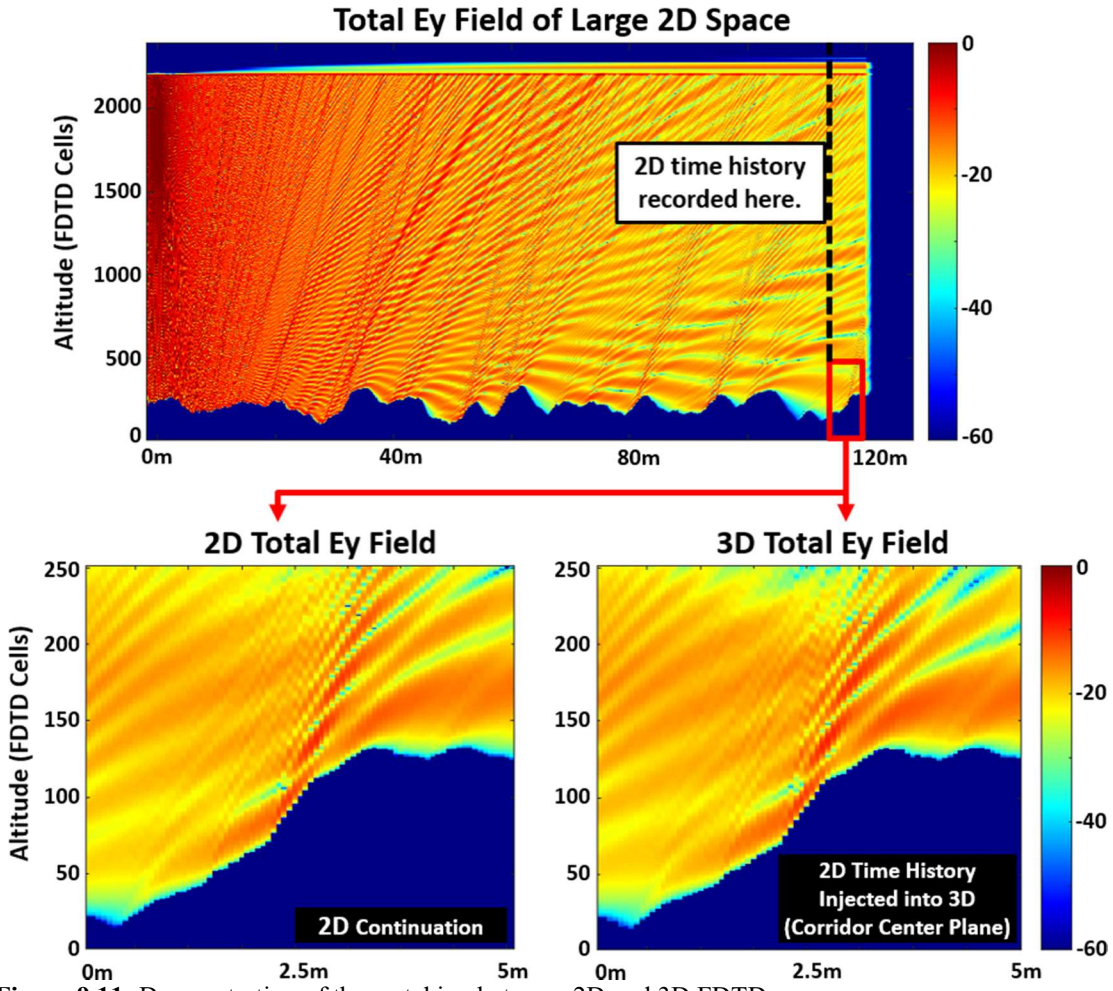

Figure 9.11: Demonstration of the matching between 2D and 3D FDTD spaces.

To test this idea a small domain has been run in which the last 5m section (150λ) can be selected as continuation of the 2D domain or replaced by a 3D domain. In the 5m overlap region we can compare the results. Figure 9.11 shows the matching of the 2D and 3D spaces. The total fields are approximately the same in both cases since 150 wavelengths is not sufficient for the total fields to incorporate the 3D nature of the sea into the 2D injected fields, but they differ to the extent that the scatterer and propagation physics are different in 3D and 2D. This approach will continue to be explored.

9.3 Conclusion

The analytic source injection VEP method is an absolute requirement to allow for the solution of sufficiently large sea surfaces in 3D FDTD since it so drastically reduces the computational burden of the space. Yet, even with such drastic reductions, the computational spaces of our corridors require well over 100GB of RAM, and about a week of runtime for a 5m-wide, 1km-long sea space. Brute-force inclusion of the 30m or more of altitude required to accommodate an inhomogeneous atmosphere unavoidably requires supercomputing. To overcome this, an approximate iterative technique is demonstrated which would allow hybridizing the present 3D FDTD computational code with 2D FDTD (or other codes) in order to include a realistic inhomogeneous atmosphere above the sea in the 3D computation. Additionally, the local 3D physics introduced by the sea may be able to be incorporated by simply simulating only the last kilometer of a very large space in 3D FDTD. A procedure to splice a long-range 2D computational domain (in principle tens of kilometers long) with the 3D FDTD model has been proposed and tested. Combination of this 2D-to-3D transition with (or potentially without) the iterative coupling is the subject

of current investigation and appears promising in its ability to incorporate the relevant 3D physics into a long-range sea-atmosphere space.

CHAPTER 10

CONCLUSION

10.1 Conclusions and Future Work

Both physical domain decomposition and parallel processing are mandatory when attempting to solve very large FDTD problems (hundreds of thousands of wavelengths long). The two absolute requirements for obtaining high fidelity results are: (a) seamless connection between the sub-domains, without excessive overlap regions, and (b) high performance absorbing boundary conditions for terminating the sub-domains. In the extreme case of grazing propagation along a very long surface, the second requirement becomes absolutely critical.

It has been shown that the re-Radiating Absorbing Termination (rRAT), a new FDTD domain termination based on the teleportation algorithm, can eliminate artificial reflections from the FDTD boundary extremely close to true grazing incidence. The accurate solution of 4km-long, 2.25m-tall 2D domains at X-band demonstrate the performance of the rRATs out to 0.03 degrees from grazing (the incident angles encountered in the analytic field injection VEP scattered field approach.) The rRAT works because (i) it is reflectionless in FDTD (since it uses exact discrete Schelkunoff currents), (ii) its combination of forward and backward teleportation planes makes it apparently impervious to numerical instability, therefore (iii) any number of rRAT layers can be stacked over each other to attain the desired attenuation, and (iv) said layers can be individually tuned to maximize performance at any selected angle of incidence by

adjustment of the time delay of the teleported fields. Finally, (v) it is computationally efficient since it is an object one dimension smaller than the space in which it operates.

Usage of dispersion-corrected tracking of the propagating (and spreading) time domain pulse, and domain decomposition via the teleportation algorithm in a Multi-GPU implementation of FDTD, allows a 122,500 wavelengths long, 1,100 wavelengths tall, 2D space (43.3 billion FDTD cells), segmented into 276 sub-domains containing a realistic sea surface and inhomogeneous atmosphere to be successfully compared to the parabolic equation method VTRPE. Successful comparisons to fullwave MoM and BIE have also been shown. Furthermore, our largest 2D space successfully simulated to date has consisted of a sea-atmosphere domain 450,000 wavelengths long, 1,240 wavelengths tall (178.3 billion FDTD cells), segmented into 1,013 sub-domains, which appears to be the largest 2D FDTD space in published literature.

The applicability of these techniques in 3D FDTD has been demonstrated by solving the same 110m sea surface problem in 2D and as a 3D extruded sea corridor. Maximum efficiency for problems involving a homogeneous atmosphere is attained with analytic source field injection directly into the material scatterers. Successful implementation of this feature by use of a dispersion- and anisotropy-corrected analytic source field has been demonstrated by cross-validating the conventional total field FDTD solution against the analytic field injection VEP scattered field approach. Application of the aforementioned techniques has allowed for successful simulation of enormous 3D spaces, the largest consisting of a 1,050m (31,500λ) long, 5m (150λ) wide, 2m (60λ) tall corridor of realistic 3D sea, totaling 283.5 million cubic wavelengths and 315.3 billion FDTD cells, which appears to be the largest 3D FDTD space in published literature.

Furthermore, a 225m-long, 5m-wide corridor of realistic 3D sea as well as extruded 3D sea excited by a horizontal electric dipole has been simulated in FDTD and yielded important scattering characteristics not able to be captured in 2D simulation spaces. A plane wave spectrum analysis of the scattered fields collected on the Schelkunoff surface just above the sea produced distinguishing characteristics between the two test cases of a realistic sea and an extruded sea (a 2D analogue). This demonstrated (i) the realistic sea scatters energy over a much wider range of angles, (ii) the transverse texture of the realistic sea geometry leads to scattered fields which contain a relevantly large TM component radiating toward the horizon in the principal propagation plane (x-z). These TM components are physically non-existent in the 2D scenario, and (iii) the realistic sea also generates strong flashes and nulls in different directions which are unobserved in the extruded or 2D cases.

he fields collected along the Schelkunoff surface also show strong localized spatial variations between the extruded and realistic 3D seas when observed off the center axis, outlining the effects of the local sea surface. To build upon the local field variations, the windowed PWS results show the multitudinous fluctuations in angular spread of the scattered fields, indicating the importance of the 3D scattering characteristics of the true sea. There is an important distinction between the conclusions that can be drawn from spatial field variation near the sea and the plane wave spectrum analysis of the radiation into the far field. It can be argued that a quasi-3D (2.5D) modeling technique would capture the field variations of the transverse offsets. However, as is demonstrated by the line plots of the PWS comparisons in Chapter 7, the energy which gets scattered back into a given 2D plane by the 3D texture of the sea would not be present in 2.5D, there is no mechanism

to communicate transversely directed scatter to adjacent planes. Furthermore, the absence of such a mechanism prevents the existence of cross-polarized fields in 2.5D as well. Energy which would never reach a target in a 2.5D modeling philosophy does in fact reach it in a realistic 3D model.

The PWS of the Schelkunoff surface in Chapter 7, which calculates far-field radiation, provides only part of the total picture. The fields which impinge upon a sea-based target are those propagating within the Fresnel zone of the sea. From the stored frequency-domain currents on the Schelkunoff surface along the whole 3D corridor of sea, the 3D scattered fields in the Fresnel zone can be reconstructed accurately in post processing via spatial integration techniques. The reconstruction from the Schelkunoff surface current sources is vulnerable to two primary sources of diffraction originating from (i) the sharp surface current discontinuity caused by the Schelkunoff surface's finite size, and (ii) from the fact that sub-domains don't send information back to previous sub-domains, which leads to accumulation of backward-travelling currents on the leading edge of some sub-domains containing large sea features. Via the removal of backward-travelling currents, and artificial extensions of the currents on the Schelkunoff surface, many of the discontinuities can be removed or improved.

We verify the accuracy of the spatial integration by direct comparison of the FDTD scattered fields with the reconstructed scattered fields on a plane inside the FDTD corridor. The relative error in the co-polarized field is very low, and peaks around -20dB. For the cross-polarized field, which is about 20dB weaker than the co-polarized field, the typical error is about -10dB, with the exception of a few faint regions of fields in which large errors can occur. The unimportance of large errors in regions of weak signal is demonstrated by

calculation of the plane wave spectrum to produce the angle of arrival data on the same comparison plane, which verifies the AoA is virtually identical between the FDTD and reconstructed fields. Small errors in the reconstructed fields do not compromise the integrity of the AoA calculations.

Analysis of the AoA on large target planes 150 wavelengths in width and in upwards of 900 wavelengths tall undoubtedly demonstrates 3D phenomenology unable to be captured in 2D or 2.5D spaces. As with the PWS of the Schelkunoff surface, the AoA of the target plane fields at 13m, 210m, and 450m down range reveal that significant amounts of information are scattered transversely at all ranges. The particularly striking results shown in Figure 8.27 and Figure 8.28 demonstrate how the transverse variations of the local sea in the last 40m leading up to the target plane can drastically redirect the scatter which reaches the target. Additionally, in all cases containing a realistic 3D sea, an appreciable amount of cross-polarized field is generated by the 3D texture of the sea. The impact of the 3D phenomenology demonstrated to be absent in 2D domains has yet to be fully quantified, and will be the subject of future investigation, but the groundwork for the analysis has now been laid.

Analytic source injection VEP techniques are an absolute requirement to allow for sufficiently large sea corridors in 3D FDTD. The ultimate goal is to incorporate inhomogeneous atmospheres into 3D simulation spaces, but the computational costs of the brute-force approach would require supercomputing. Since atmospheric ducts do not vary transversely an appreciable amount over the width of the sea corridor, and only vertical variation is of significance in the corridor problem, a 2D space could in principle be used to model the inhomogeneous atmosphere. The multiple reflections between the sea and

atmosphere could be coupled back into the 3D space via an iterative injection method, which has been demonstrated as highly accurate in 2D. Additionally, local 3D physics may be able to be incorporated into long sea spaces by hybridization of the present 3D FDTD code with other 2D methods. The first tens of kilometers, or more, of the sea space would be simulated with a 2D method to produce the fields to illuminate the last kilometer or so of the space, which would be simulated in 3D FDTD. It has been demonstrated that the 2D time history of the fields which exit the 2D space can be injected into the start of a new 3D space and will continue their propagation. Additional tests of the efficacy of this coupling method is the subject of further investigation. Coupling of the 2D space into the 3D corridor can be applied alongside the iterative injection method as well. Successful combination of these techniques would lead to massive reductions in the computational burden of the 3D simulation spaces and appears promising in its ability to incorporate the relevant 3D physics into a long-range sea-atmosphere space.

REFERENCES

[1] J. Wright, "Detection of ocean waves by microwave radar; The modulation of short gravity-capillary waves," *Boundary-Layer Meteorology,* vol. 13, no. 1, pp. 87-105, 1978.

[2] J. C. N. Borge, K. Hessner, P. Jarabo-Amores and D. de la Mata-Moya, "Signal-to-noise ratio analysis to estimate ocean wave heights from X-band marine radar image time series," *IET Radar Sonar Navigation,* vol. 2, no. 1, pp. 35-41, 2008.

[3] B. Lund, H. C. Graber and R. Romeiser, "Wind retrieval from shipborne nautical X-band radar data," *IEEE Transactions Geoscience Remote Sensing,* vol. 50, no. 10, pp. 3800-3811, Jan 2008.

[4] A. Ishimaru, "Wave Propagation and Scattering in Random Media and Rough Surfaces," *Proceedings of the IEEE,* vol. 79, no. 10, pp. 1359-1366, Oct 1991.

[5] F. J. Ryan, "User's Guide for the VTRPE (Variable Terrain Radio Parabolic Equation) Computer Model," Naval Ocean Systems Center, San Diego, 1991.

[6] D. Holiday, L. L. DeRaad and G. J. St-Cyr, "Forward Backward: A new method for computing low grazing angle scattering," *IEEE Transactions on Antennas & Propagation,* vol. 44, no. 5, May 1996.

[7] D. A. Kapp and G. S. Brown, "A new numerical method for rough surface scattering calculations," *IEEE Transactions on Antennas & Propagation,* vol. 55, no. 5, May 1996.

[8] C. Bourlier, H. Li and V. Fabro, "Radar propagation modeling using the boundary integral equations in a maritime environment with a duct," in *2014 International Radar Conference*, 2014.

[9] J. R. Smith and R. J. Burkholder, "Channeling Phenomenon in Electromagnetic Forward Scattering at Low Grazing," *IEEE Transactions on Geoscience and Remote Sensing,* vol. 42, no. 8, pp. 1731-1738, Aug 2004.

[10] K. B. Thiem, "A 3D Parabolic Equation (PE) based technique for predicting propagation path loss in an urban area," Naval Postgraduate School, Monterey, CA, 2001.

[11] R. S. Awadallah, J. Z. Gehman, J. R. Kuttler and M. H. Newkirk, "Modeling Radar Propagation in Three-Dimensional Environments," *Johns Hopkins APL Technical Digest,* vol. 25, no. 2, pp. 101-111, 2004.

[12] Z. H. Lai, J. F. Kiang and R. Mittra, "A Domain Decomposition Finite Difference Time Domain (FDTD) Method for Scattering Problem from Very Large Rough Surfaces," *IEEE Transactions on Antennas & Propagation,* vol. 63, no. 10, pp. 4468-4476, Oct 2015.

[13] P. C. Clemmow, The Plane Wave Spectrum Representation of Electromagnetic Fields, Pergamon Press, 1966.

[14] D. E. Barrick, "razing Behavior of Scatter and Propagation Above Any Rough Surface," *IEEE Transactions on Antennas & Propagation,* vol. 46, no. 1, pp. 73-83, Jan 1998.

[15] A. Karimian, P. G. C. Yardim, W. S. Hodgkiss and A. E. Barrios, "Multiple Grazing Angle Sea Clutter Modeling," *IEEE Transactions on Antennas & Propagation,* vol. 60, pp. 4408-4417, Sept 2012.

[16] M. E. Watts and R. E. Diaz, "Perfect Plane-Wave Injection into a finite FDTD domain through Teleportation of Fields," *Electromagnetics,* vol. 23, no. 2, pp. 187-201, Jan 2003.

[17] S. Kim and J. Choi, "Optimal design of PML absorbing boundary conditions for improving wide-angle reflection performance," *Electronics Letters,* vol. 40, no. 2, pp. 104-105, Jan 2004.

[18] M. F. Hadi, "Wide-Angle Absorbing Boundary Conditions for Low and High-Order FDTD Algorithms," *Applied Computational Electromagnetics Society Journal,* vol. 24, no. 1, pp. 9-15, Feb 2009.

[19] R. E. Diaz and I. Scherbatko, "A new multistack radiation boundary condition for FDTD based on self-teleportation of fields," *Journal of Computational Physics,* vol. 203, no. 1, pp. 176-190, Feb 2005.

[20] T. M. Elfouhaily and J. T. Johnson, "A new model for rough surface scattering," *IEEE Transactions on Geoscience and Remote Sensing,* vol. 45, no. 7, pp. 2300-2308, Jul 2007.

[21] A. Taflove, K. R. Umashankar and T. G. Jurgens, "Validation of FDTD modeling of the radar cross-section of three-dimensional structures spanning up to 9 wavelengths," *IEEE Transactions on Antennas & Propagation,* vol. 33, pp. 662-666, Jun 1985.

[22] G. A. Kriegsmann, A. Taflove and K. R. Umashankar, "A new formulation of electromagnetic wave scattering using an on-surface radiation boundary condition approach," *IEEE Transsactions on Antennas & Propagation,* vol. 35, pp. 153-161, Feb 1987.

[23] K. R. Umashankar, S. Nimmagadda and A. Taflove, "Numerical analysis of electromagnetic scattering by electrically large objects using spatial decomposition technique," *IEEE Transactions on Antennas & Propagation,* vol. 40, pp. 866-877, Aug 1992.

[24] S. T. Nguyen, B. J. Zmk and X. Zhang, "Distributed Computation of Electromagnetic Scattering Problems Using Finite-Difference Time-Domain Decompositions," in *Proceedings of the 3rd IEEE International Symposium on HPDC*, 1994.

[25] D. Yang, J. Xiong, C. Liao and L. Jen, "A Parallel FDTD Algorithm Based on Domain Decomposition Method Using the MPI Library," in *PDCAT 2003*, 2003.

[26] C. Jia, L. Guo and K. Li, "Graphics processor unit accelerated finite-difference time domain method for electromagnetic scattering from one-dimensional large scale rough soil surface at low grazing incidence," *Journal of Applied Remote Sensing,* vol. 8, no. 1, 2014.

[27] H. E. Abd-El-Raouf, R. Mittra and J. Ma, "A new domain decomposition finite-difference time domain for solving large electromagnetic problems," *Microwave and Optical Technology Letters,* vol. 48, no. 12, p. 2399–2405, Dec 2006.

[28] Z.-H. Lai, J. F. Kiang and R. Mittra, "Two-dimensional domain-decomposition FDTD method to simulate wave scattering by rough surfaces," in *2013 7th European Conference on Antennas and Propagation*, 2013.

[29] E. I. Thorsos, "Acoustic scattering from a “Pierson–Moskowitz” sea surface," *Journal of the Acoustical Society of America,* vol. 88, no. 1, pp. 335-349, July 1990.

[30] H. V. Hitney, J. H. Richter, R. A. Pappert, K. D. Anderson and G. B. Baumgartner, "Tropospheric radio propagation assessment," *Proceedings of the IEEE,* vol. 73, no. 2, pp. 265-283, Feb 1985.

[31] A. Valcarce, H. Song and J. Zhang, "Characterization of the Numerical Group Velocity in Yee’s FDTD Grid," *IEEE Transactions on Antennas & Propagation,* vol. 58, no. 12, pp. 3974-3982, Dec. 2010.

[32] R. E. Diaz and I. Scherbatko, "A simple stackable re-radiating boundary condition (rRBC) for FDTD," *IEEE Antennas and Propagation Magazine,* vol. 46, no. 1, pp. 124-130, Feb 2004.

[33] I. Scherbatko, G. Churyumov and I. M. M. Ustyantsev, "A comparison between the FDTD convolutional PML and reradiating boundary conditions for simple 3D test problem," in *12th International Conference on Mathematical Methods in Electromagnetic Theory*, 2008.

[34] J.-P. Berenger, "On the Huygens absorbing boundary conditions for electromagnetics," *Journal of Computational Physics,* vol. 226, no. 1, pp. 354-378, Sept. 2007.

[35] C. M. Rappaport, "Interpreting and Improving the PML Absorbing Boundary Condition Using Anisotropic Lossy Mapping of Space," *IEEE Transactions on Magnetics,* vol. 32, no. 3, pp. 968-974, May 1996.

[36] O. M. Ramahi and J. B. Schneider, "Comparative Study of the PML and C-COM Mesh-Truncation Techniques," *IEEE Microwave and Guided Wave Letters,* vol. 8, no. 2, pp. 55-57, Feb 1998.

[37] N. V. Kantartzis, T. V. Yioultsis, T. I. Kosmanis and T. D. Tsiboukis, "Nondiagonally Anisotropic PML: A Generalized Unsplit Wide-Angle Absorber for the Treatment of the Near-Grazing Effect in FDTD Meshes," *IEEE Transactions on Magnetics,* vol. 36, no. 4, pp. 907-911, July 2000.

[38] M. F. Hadi, "Near-Field PML Optimization for Low and High Order FDTD Algorithms Using Closed-Form Predictive Equations," *IEEE Transactions on Antennas & Propagation,* vol. 59, no. 8, pp. 2933-2942, Aug 2011.

[39] W. Sun, H. Pan and G. Videen, "General finite-difference time-domain solution of an arbitrary electromagnetic source interaction with an arbitrary dielectric surface," *Applied Optics,* vol. 48, no. 31, pp. 6015-6025, Nov. 2009.

[40] V. H. Rumsey, "Reaction Concept in Electromagnetic Theory," *Physical Review,* vol. 94, no. 6, pp. 1483-1491, June 1954.

[41] R. J. Burkholder and J. Johnson, *Provided MoM Simulations,* 2014.

[42] F. J. Ryan, *Provided VTRPE Simulations,* 2014.

[43] C. Bourlier and N. Pinel, "Spatial Green function of a constant medium overlying a duct with linear-square refractive index profile," *IEEE Transactions on Antennas & Propagation,* vol. 61, no. 6, pp. 3172-3181, June 2013.

[44] U. Oguz and L. Gurel, "Interpolation techniques to improve the accuracy of the plane wave excitations in the finite difference time domain method," *Radio Science,* vol. 32, no. 6, p. 2189–2199, 1997.

[45] G. Singh, E. L. Tan and Z. N. Chen, "Analytic Fields with Higher-Order Compensations for 3-D FDTD TF/SF Formulation with Application to Beam Excitations," *IEEE Transactions on Antennas & Propagation,* vol. 59, no. 7, pp. 2588-2598, July 2011.

[46] S.-K. Jeng, "An Analytical Expression for 3-D Dyadic FDTD-Compatible Green's Function in Infinite Free Space via z-Transform and Partial Difference Operators," *IEEE Transactions on Antennas & Propagation,* vol. 59, no. 4, pp. 1347-1355, April 2011.

[47] J. B. Schneider and Z. Chen, "Incorporating the G-TFSF Concept into the Analytic Field Propagation TFSF Method," *IEEE Transactions on Antennas & Propagation,* vol. 59, no. 9, pp. 3296-3304, Sept. 2011.

[48] T. P. Stefansk, "Applications of the discrete green's function in the finite-difference time-domain method," *Progress In Electromagnetics Research,* vol. 139, pp. 479-498, May 2013.

[49] J. Frechot, "Realistic simulation of ocean surface using wave spectra," in *Proceedings of the First International Conference on Computer Graphics Theory and Applications*, 2006.

APPENDIX A

IMPLEMENTATION: 2D SCATTERED FIELD RECONSTRUCTION UNDER VOLUME EQUIVALENCE PRINCIPLE

A.1 Introduction

The volume equivalence principle (VEP) technique (Section 2.4) , which allows for source field injection directly into the materials of the scatterers in the FDTD space, in conjunction with dispersion-corrected analytic source fields (Sections 3.4 and 4.1) allows for an enormous reduction in computational burden for both 2D and 3D FDTD spaces for the case of a homogeneous atmosphere. The sky boundary can be brought down to just above the sea surface (Figure 2.8), and all the dead space in the sky can be eliminated. The relative reduction in size for the 2D problem is pictured in Figure A.1. The scattered fields which would otherwise propagate through the homogeneous atmosphere, which is now excluded from the FDTD space, can be reconstructed from the induced currents gathered on the Schelkunoff surface, which is a plane that sits just above the peak wave height of the sea surface. A spatial integration technique which utilizes a dispersion-corrected Green's function, is presented in Section 8.1. This appendix covers the implementation-specific details for the scattered field reconstruction in 2D. Though the computational burden is comparatively enormous in 3D, the logical principles presented transfer directly.

The fundamental breakdown of the field reconstruction's computational problem is shown in Figure A.2. The tangential scattered fields (H_x and E_y for TEy, E_x and H_y for TMy) are collected along the Schelkunoff surface of the FDTD space, and a discrete Fourier transform applied during run time to convert the fields crossing the plane to complex frequency domain fields at any frequency of interest. The scattered field samples are taken from a single row of FDTD cells (a plane in 3D) at a given height within the FDTD space, which spans the full length of the total simulation domain. Thus, the resulting

output fields in 2D are in the form of a single row of discretely sampled cells of size equal to the FDTD grid spacing.

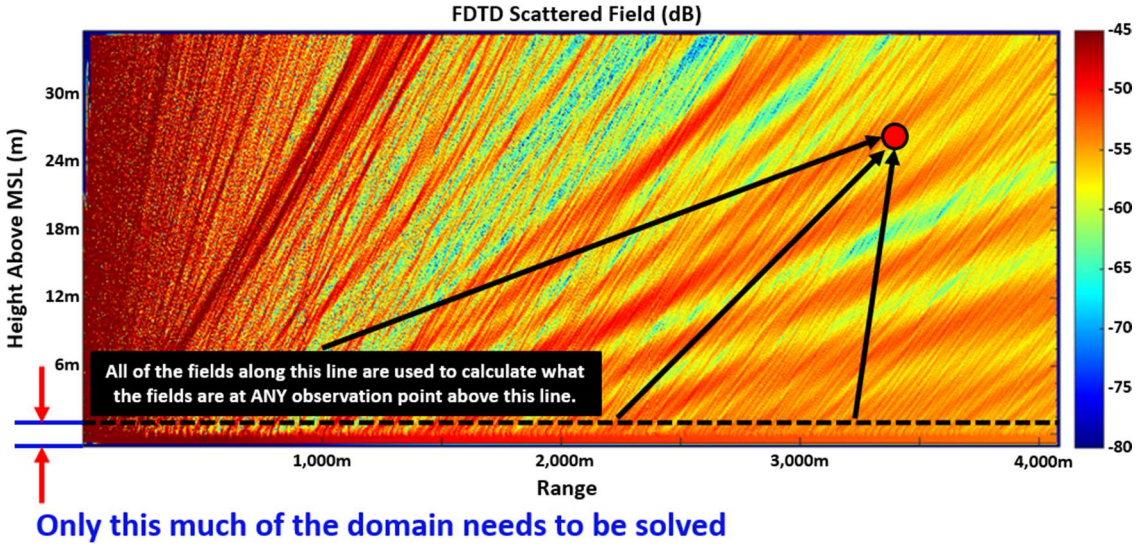

Figure A.1: The basic problem setup for the field reconstruction. The source fields are collected along the Schelkunoff surface and radiated through space to the observation point to reconstruct the fields.

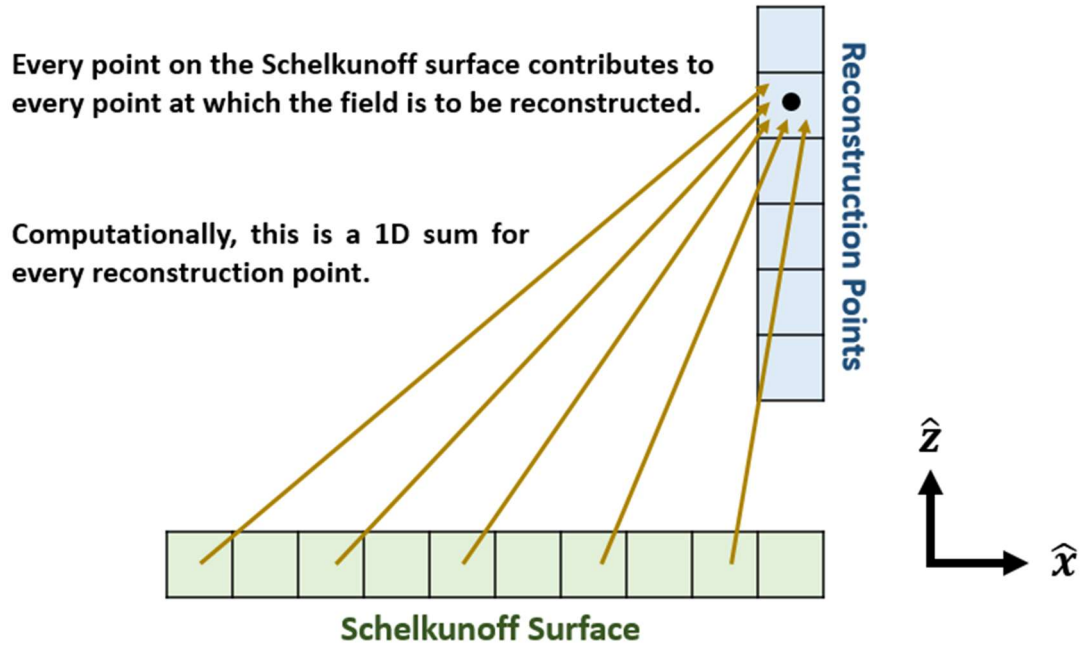

Figure A.2: The general computational problem involves summing the radiated field contributions from each point on the Schelkunoff surface for all points at which the field is to be reconstructed.

For ease of explanation only the TE case will be considered here; the TM case is nearly identical. Additionally, only the reconstruction of co-polarized field, E_y, is

considered since the reconstruction of the other components is nearly identical as well. Also E_y in the TE case (H_y for TM) provides all the information required to compute the angle of arrival. Note, in 2D, the only Schelkunoff surface fields required to reconstruct the scattered co-polarized fields anywhere in space are simply the co-polarized fields of the Schelkunoff surface. For example, as will be shown in this appendix, only E_y on the Schelkunoff surface is required to reconstruct E_y at any point above the surface. The 3D case is more complicated in the sense of bookkeeping because it requires all four tangential field components to reconstruct any one scattered field component, as well as special pre-processing techniques, all of which is discussed in Section 8.1.

The frequency domain E_y elements stored along the Schelkunoff surface produce magnetic surface currents in the x-direction (eq. 8.1). The magnetic current elements are radiated into space using the Green's function of a 2D infinite line source, a Hankel function (eq. A.4), which is multiplied by the appropriate element factor (eq. A.1). We apply the same anisotropic phase velocity correction as used in analytic source injection to apply the appropriate numerical wavenumber at any propagation angle (eqs. A.2 and A.3).

$$EF = \cos(\theta) = \frac{z}{r} \tag{A.1}$$

$$v_{p2D}(\theta) = \tilde{v}_{pmin2D} + [2\,cos(\theta)\,sin(\theta)]^2 \cdot \Delta v_{p2D} \tag{A.2}$$

$$\tilde{k}(\theta) = \frac{\omega}{v_{p2D}(\theta)} \tag{A.3}$$

$$E_y(x,z) = \frac{1}{2}\Delta s\; k_0\; H_1^1\big(\tilde{k}(\theta)r\big) \cdot EF \tag{A.4}$$

The fundamental ideas of the spatial integration-based reconstruction have now been established, the remaining discussion is centered on the implementation-specific

considerations for our OpenCL FDTD domain decomposition program, as well as parallelization of the reconstruction process with both OpenCL and OpenMP. The first consideration regarding the FDTD domain decomposition method is that information flows between domains in exclusively the forward direction (Figure A.3). Thus, fields scattered from the sea in domain D+1 are not re-radiated back into domain D, they will only forward into domain D+2.

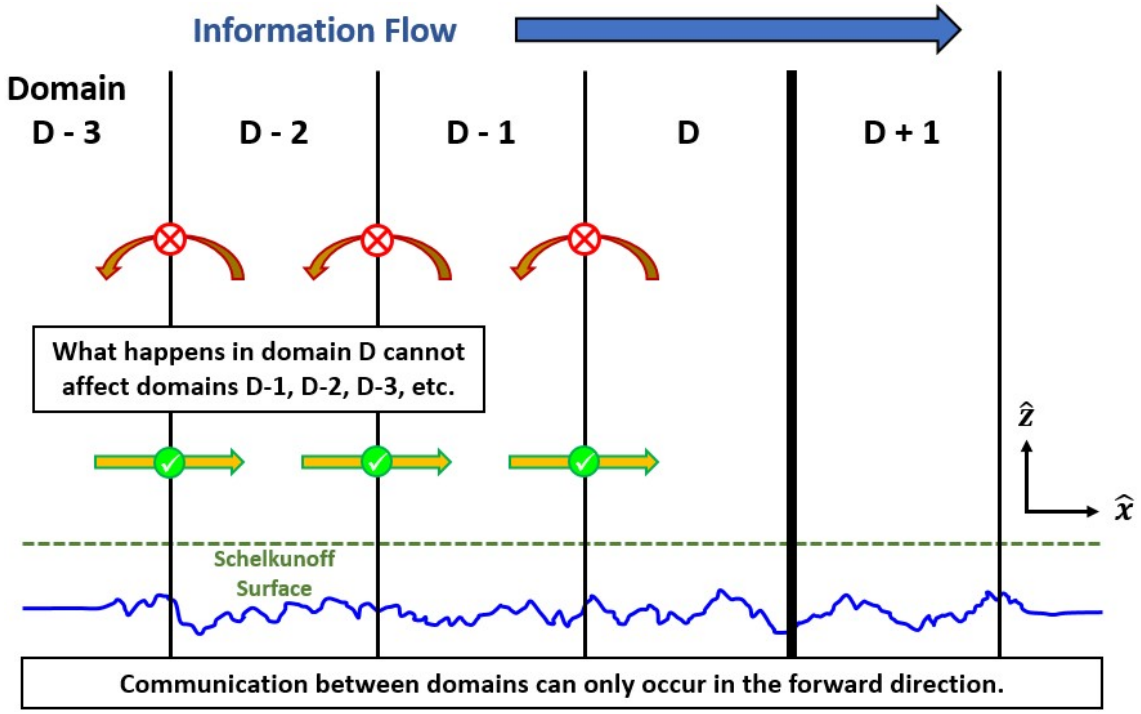

Figure A.3: Field information is transmitted between domains exclusively in the forward direction.

Since the FDTD domains can't communicate backwards during the FDTD simulation, in order to faithfully reconstruct the FDTD fields the parts of the Schelkunoff surface residing in domains further down range than the domain of the current reconstruction point cannot be included (Figure A.4). We now label the domain number in which the current reconstruction point resides in as domain D and the number of Schelkunoff surface sample points per domain as L. Then, the total number of Schelkunoff

surface points to include as part of the reconstruction (N_P) at the point in domain D is simply given by eq. A.5. In most cases the fields we would like to reconstruct consist of several points (M points) along a vertical plane, at a given range (Figure A.5). Often, it's desirable to reconstruct multiple target planes placed at various ranges in the 2D space at once, or in fact we would like to reconstruct the entire 2D space.

$$N_P = D \cdot L \qquad \text{A.5}$$

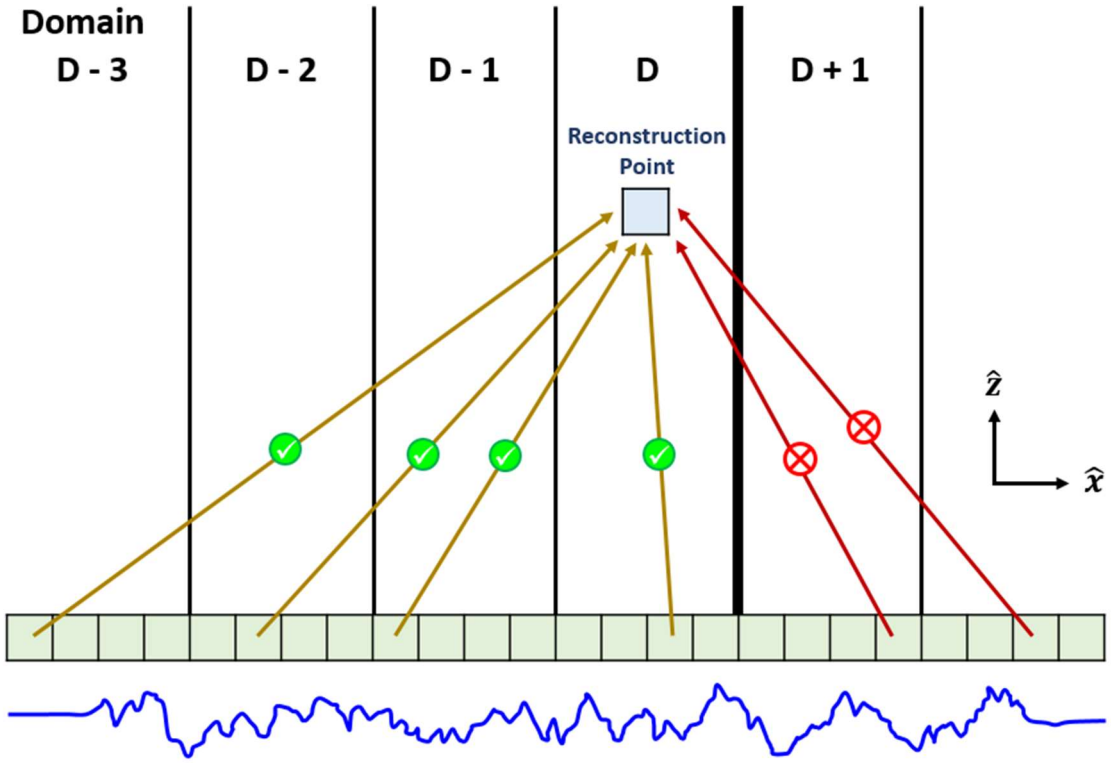

Figure A.4: Parts of the Schelkunoff surface in later domains can't contribute to the reconstruction.

The basic reconstruction computational problem at a single point is a one-dimensional space of size N_P. The contributions from the Schelkunoff surface are summed only once per reconstruction point. For reconstruction along a vertical plane, the problem size increases by one dimension. If the plane consists of M reconstruction points, the total

problem size becomes $N_P \cdot M$. If we want to reconstruct the fields on P different target planes, the problem increases in dimension yet again to a size of $N_P \cdot M \cdot P$ (Figure A.6).

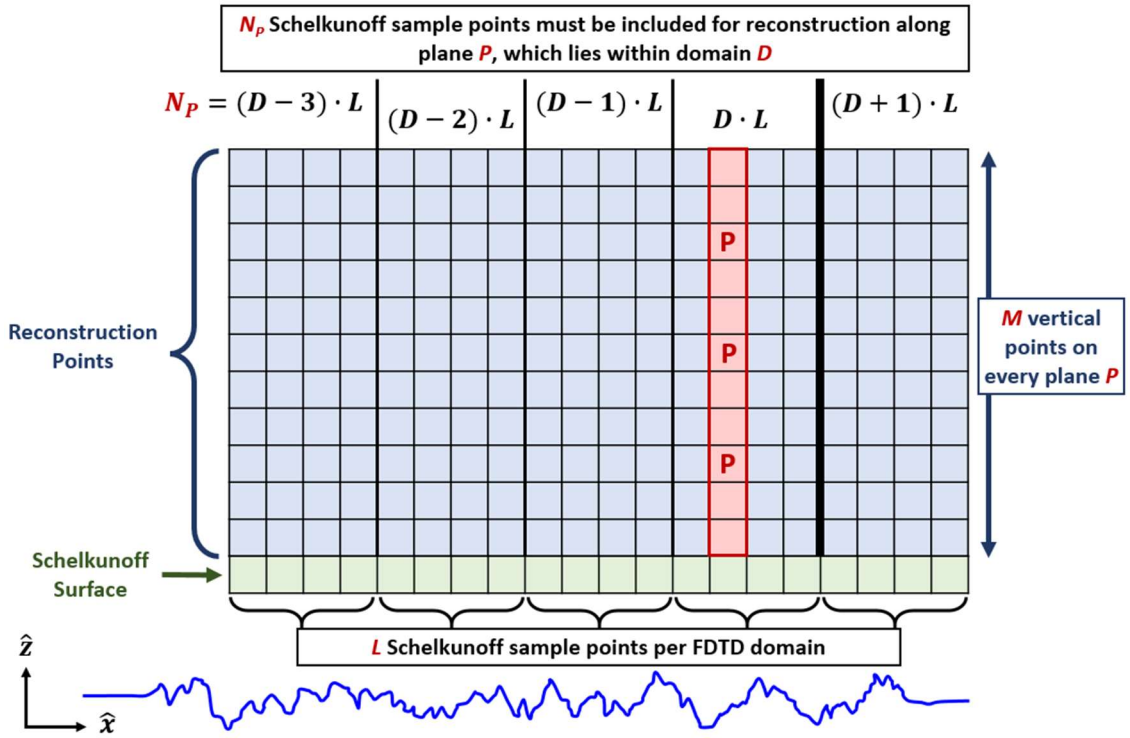

Figure A.5: The number of Schelkunoff surface points to include for reconstruction depends on the reconstruction plane's range.

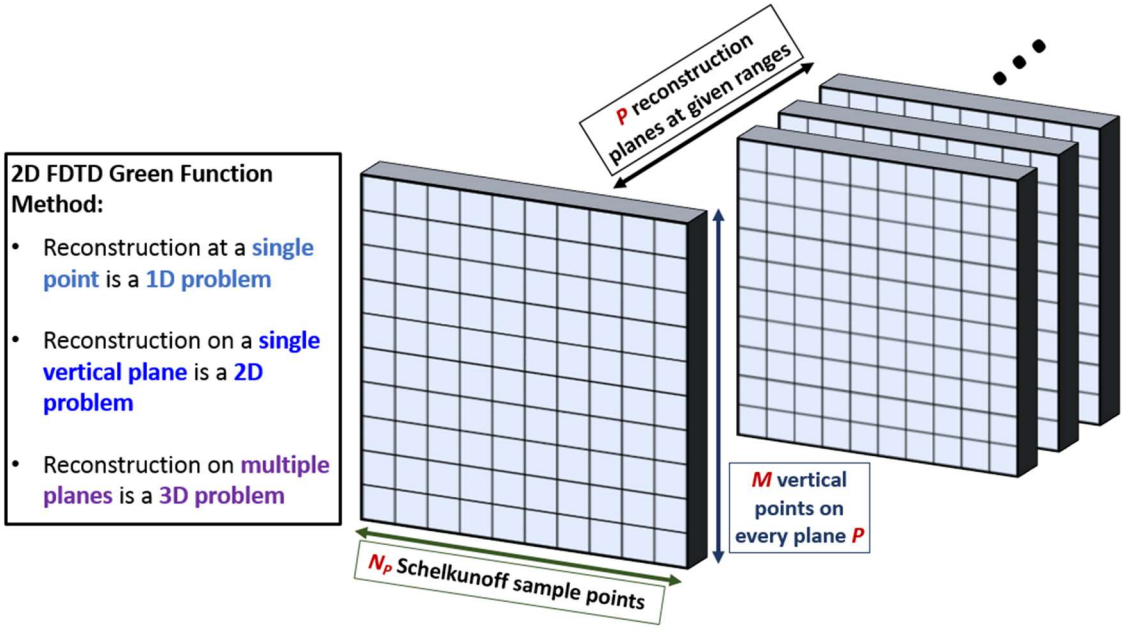

Figure A.6: Computational problem dimensionality.

The serial implementation of the 2D FDTD spatial integration reconstruction is shown as a pseudocode representation in Figure A.7. It consists of three for-loops, each accounting for a dimension of the problem space. The outermost loop iterates over the target planes, the second-tier loop iterates over the surface points of every plane, and the inner most loop iterates over every cell in the Schelkunoff surface and sums up their radiated contributions at every point on each plane.

Serial Reconstruction Loop

```
for(p = 0; p < P; p++)                    //Iterate over all planes
{
    for(m = 0; m < M; m++)                //Iterate over all cells in each plane
    {
        *Calculate height of current cell*

        for(n = 0; n < Np; n++)           //Iterate over all Schelkunoff cells
        {
            *Load Schelkunoff surface fields*

            *Calculate distance between Schelkunoff cell
             and reconstruction point, correct phase velocity*

            *Calculate contribution from current Schelkunoff
             cell at current reconstruction point (Bessel functions)*

            *Cumulative sum of all Schelkunoff contributions*
        }

        *Store cumulative sum at current reconstruction point m*
    }
}
```

Figure A.7: Pseudocode for the serial reconstruction loop.

Understanding the way in which the problem changes and grows is important to create the most general algorithm for reconstructions of any dimensionality and is of paramount importance during parallelization of the algorithm to ensure efficient utilization of resources. Additionally, the computational burden of the reconstruction will increase as the target plane range increases because it will require a larger number of Schelkunoff surface points, which is due to the relation in equation A.5. This has implications for load

balancing when utilizing multiple compute devices for reconstruction, which will be addressed later.

A.2 Parallelization with OpenMP

We now consider how to CPU-parallelize the code for the serial implementation shown in Figure A.7. Common reconstruction scenarios generally involve a small number of target planes, each containing thousands or tens of thousands of reconstruction points. The number of Schelkunoff surface points included for a given target plane regularly spans into the realm of hundreds of thousands, or in 3D, hundreds of millions of points. There are three places in which parallelization can be implemented – the p-loop, m-loop, or the n-loop – and each has different implications.

The p-loop generally contains a smaller number of iterations than any of the other loops, usually fewer than 10 and often just 1. Most modern CPUs are capable of 8-way simultaneous multithreading (SMT), so if the reconstruction contains just a single plane, parallelization of the p-loop does nothing and remains as a serial execution. If the CPU is capable of 8 threads and there are 10 reconstruction planes, all 8 threads will be occupied initially for the first 8 planes, but 6 of the 8 threads will be idle during the calculation of the last 2 planes. Clearly this isn't ideal for the most common scenarios, particularly since CPUs are relatively slow and reconstruction along many planes would not be feasible to begin with.

The n-loop contains far more points than the p-loop, and as such is easily parallelizable without the need to consider idle threads. As an example, consider a CPU with 8-way SMT and let's say the n-loop contains 800,000 iterations, which when divided

up distributes 100,000 iterations to each thread. For this particular problem, 100,000 iterations in the n-loop is near trivial for the CPU thread and will complete in a fraction of a second. Now consider that for every iteration of the m-loop OpenMP will have to re-branch the main thread at the start of the n-loop and re-join it at the end. The re-branching/joining time is quick, on the order of milliseconds, but may be on the order of up to 10% of the total time required to complete the n-loop iterations across all threads. This is a non-trivial amount of overhead. Additionally, parallelizing in the manner leads readily to memory collisions; where multiple threads write to the same memory location at once. This leads to unpredictable results and a requirement to make the program thread safe, which adds additional overhead and program complexity.

OpenMP Parallel Reconstruction Structure

```
for(p = 0; p < P; p++)        //Iterate over all planes
{
    #OpenMP Parallel Block (#threads)   //Create thread pool based CPU capacity
    {
        t_num = get_thread_num();                       //Get local thread ID
        block_size = integer_round_up(M / #threads);    //Divide m space into #threads groups

        m_start = t_num * block_size;       //Local m start value
        m_end = (t_num + 1) * block_size;   //Local m end value

        if(m_end > M) -> then(m_end = M)    //Check if m_end is in bounds

        for(m = m_start; m < m_end; m++)                //Iterate over all cells in each plane
        {
            *Calculate height of current cell*

            for(n = 0; n < Np; n++)       //Iterate over all Schelkunoff cells
            {
                *Load Schelkunoff surface fields*

                *Calculate distance between Schelkunoff cell
                 and reconstruction point, correct phase velocity*

                *Calculate contribution from current Schelkunoff
                 cell at current reconstruction point (Bessel functions)*

                *Cumulative sum of all Schelkunoff contributions*
            }

            *Store cumulative sum at current reconstruction point m*
        }
    }
}
```

Figure A.8: Pseudocode for the OpenMP parallel CPU-based reconstruction.

The winning scheme is to parallelize the m-loop. Hundreds or thousands of iterations of the m-loop per CPU thread is sufficient to load the CPU enough to avoid the overhead of thread creation. Additionally, it's easier to prevent memory conflicts since each thread is assigned a specific, non-overlapping range of cells along the current target plane, P. The parallelization shown in Figure A.8 is performed explicitly. The total number of points, M, in plane P are divided up into blocks. The size of each block is equal the total number of points divided by the number of threads to be used, rounded up to the nearest integer. Each thread is then assigned a specific range of m-values to iterate over.

A.3 Parallelization with OpenCL

OpenCL is similar, but significantly different to OpenMP. OpenCL enables device-agnostic functionality, which includes not only support for CPUs, but GPUs and other compute devices as well. The OpenCL pseudo code of Figure A.9 is similar to the OpenMP code in Figure A.8 in that the thread IDs are in both instances generated from a thread pool – a thread group containing a pre-defined number of threads and their associated ID numbers, each of which are spawned and completed only once. The primary difference in the algorithms is both the p- and m-loops are eliminated completely and replaced by thread pools.

A set of thread IDs encompassing all integer values between 0 an M are generated in one dimension, and similarly between 0 and P in the second dimension. As a rule of thumb, higher dimensions in the OpenCL thread pool should be kept smaller than lower dimensions. The total number of thread IDs created for the code in Figure A.9 is $M \cdot P$. The thread IDs can be considered as ordered pairs (eq. A.6), which account for all

reconstruction points m along plane p, for all P target planes of the full reconstruction space. For reasons beyond the scope of this writing (regarding OpenCL work items and work groups), the thread IDs generated will almost always extend beyond the maximum dimensions M and P, so the indices must be explicitly checked to avoid extending beyond memory bounds.

Thread ID ordered pairs: (m, p), $0 \leq m < M$, $0 \leq p < P$ A.6

OpenCL Parallel Reconstruction Structure

```
int m = get_global_id(0);      //Get a thread ID in thread pool 0
int p = get_global_id(1);      //Get a thread ID in thread pool 1

int global_p = p + current_GPU_start_p_value;   //Where the current p on the current GPU
                                                //fits into the global P range

if(m < M && p < local_P && global_p < P)        //Check indices are in bounds
{
    *Calculate height of current cell*

    for(n = 0; n < Np; n++)                     //Iterate over all Schelkunoff cells
    {
        *Load Schelkunoff surface fields*

        *Calculate distance between Schelkunoff cell
         and reconstruction point, correct phase velocity*

        *Calculate contribution from current Schelkunoff
         cell at current reconstruction point (Bessel functions)*

        *Cumulative sum of all Schelkunoff contributions*
    }

    *Store cumulative sum at current reconstruction point m*
}
```

Figure A.9: Pseudocode for the OpenCL parallel reconstruction.

A common question arises based on the pseudo code of Figure A.9. Why not replace the n-loop by adding a third thread pool to the OpenCL kernel? The issue lies with synchronicity, particularly with the cumulative sum of the Schelkunoff surface contributions. Many threads end up both reading from and writing to a single memory location simultaneously, much in the same way as in the n-loop parallelization of the OpenMP code. This causes unpredictable results. Synchronization points can be added to

the kernel, but they introduce tremendous overhead for this problem space, in addition to program complexity. The P and M dimensions are typically large enough to sufficiently parallelize the reconstruction without the serial nature of the n-loop impeding the process. If P and M are small enough to the point at which the serial n-loop creates overhead, then the total problem is very likely so trivial as to finish within fractions of a second anyway.

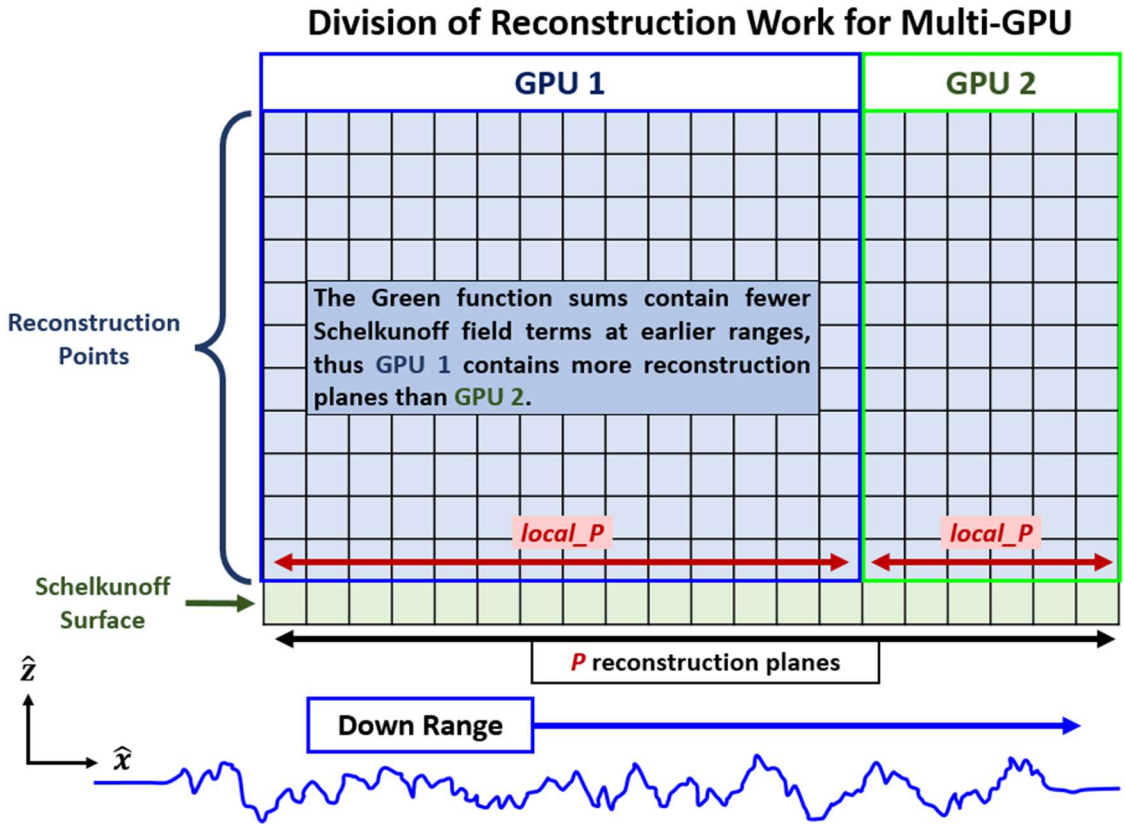

Figure A.10: Possible division of work between two GPUs.

OpenCL is not limited to one compute device at a time, multiple devices can be simultaneously queued. Enqueing multiple devices in OpenCL is not as straightforward as utilizing multiple CPUs in OpenMP, the division of work for each device must be explicitly defined. For the problem of the spatial integration reconstruction using GPUs, this division of work is most readily performed by dividing up the P total number of target planes into smaller sets of target planes to send to each GPU. This leads to a simpler program

implementation than by dividing up the M-dimension and is also more readily conducive to fighting the watchdog timer problem associated with Windows, which will be described later. Recall, the reconstruction workload on a given target plane increases with range since the number of contributing Schelkunoff surface points increases, thus the division of the P total planes in the reconstruction space across two GPUs may appear uneven, like Figure A.10 shows.

The total number of reconstruction planes, P, is divided up into sets (*local_p* planes) according to the total number Schelkunoff surface points are required for each set of planes. The workload is directly proportional to the number of source points. The value of *local_p* is different for each GPU. For instance, a reconstruction space divided into 200 planes and calculated on two GPUs may have *local_p = 150* for GPU1, and local_p = 50 for GPU2. The pseudo code for determining the value of local_p is demonstrated in Figure A.11.

```
cumulative_schelk_cells = 0;      //Stores the total number of Schelkunoff points used
local_p[num_GPUs];  //Stores number of planes to include for each GPU

//Sum up the total number of Schelkunoff surface points used for every reconstruction
//plane
for(p = 0; p < P; p++)
{      cumulative_schelk_cells = cumulative_schelk_cells + Np(p);  }

//Calculate equal work divisions for each GPU
global_count = 0;

for(GPU_num = 0; GPU_num < num_GPUs; GPU_num++)
{
       sum_points = 0;
       count = 0;

       //Sum Schelkunoff points to determine local_p for each GPU
       while(sum_points < cumulative_schelk_cells/num_GPUs && global_count < P)
       {
              sum_points = sum_points + Np(global_count);
              global_count = global_count + 1;
              count = count + 1;
       }

       //Store the number of planes to be calculated by each GPU
       local_p[GPU_num] = count;
}
```

Figure A.11: Pseudocode to determine the value of the parameter local_p, used to load balance GPUs.

A general concern exists for any application of GPU computing in that many operating systems implement a feature which regularly polls devices to check if they are responsive, and if they are not responsive their drivers are restarted, this is the roll of the watchdog timer. In Windows, devices are typically given around 10 seconds to respond to the watchdog timer polls. Large workloads for the GPUs may take minutes, hours, or longer to complete, this is where the watchdog timer becomes relevant, particularly for GPUs which are connected as primary output to displays. To avoid a system crash caused by the watchdog timer, the workloads of each individual GPU should be divided further by queuing up no more than a pre-determined number of reconstruction planes at any single time. The size of the divisions will be dependent upon the size of the problem, and the specifications of the current computing system.

A.4 Accelerating the Hankel Function Calculations

The Hankel functions which serve as the Green function for the field reconstruction are simply composed of a Bessel and Neumann function (eqs. A.7 and A.8).

$$\mathrm{H}_{\alpha}^{1}(x) = J_{\alpha}(x) + iY_{\alpha}(x) \tag{A.7}$$

$$\mathrm{H}_{\alpha}^{2}(x) = J_{\alpha}(x) - iY_{\alpha}(x) \tag{A.8}$$

The integral definitions of the Bessel and Neumann functions are shown in eqs. A.9 and A.10. Many programming languages incorporate built-in Bessel functions, c++ and MATLAB included. While the built-in functions are relatively fast and accurate, they are very simply approximated under certain conditions.

$$J_\alpha(x) = \frac{1}{\pi}\int_0^\pi \cos(\alpha\tau - x\sin(\tau))\,d\tau \tag{A.9}$$

$$Y_\alpha(x) = -\frac{1}{\pi}\int_0^\pi \sin(\alpha\tau - x\sin(\tau))\,d\tau - \frac{1}{\pi}\int_0^\infty [e^{\alpha t} + (-1)^\alpha e^{-\alpha t}]e^{x\sinh(t)}dt \tag{A.10}$$

Due to the size of the FDTD discretization and the center-frequency's wavenumber for most of our simulations, the Bessel function arguments are almost never smaller than $x = k_0 r > 0.35$, and most points exist in the range x >> 1. For large arguments, the Bessel functions of both kinds yield simple approximations (eqs. A.11 and A.12). These approximations become more accurate the larger the argument becomes.

$$J_\alpha(x) \cong \sqrt{\frac{2}{\pi x}}\cos\left(x - \frac{\alpha\pi}{2} - \frac{\pi}{4}\right) \qquad 1 \ll x \tag{A.11}$$

$$Y_\alpha(x) \cong \sqrt{\frac{2}{\pi x}}\sin\left(x - \frac{\alpha\pi}{2} - \frac{\pi}{4}\right) \qquad 1 \ll x \tag{A.12}$$

Since extreme accuracy beyond about 3 digits is not required, we can pre-determine the argument value at which the approximations of eqs. A.11 and A.12 become accurate to within 3 digits (around x > 500-600) for first order Bessel functions. Then during the n-loop in Figure A.7, the argument value can be checked before the Bessel calculations to determine whether the more accurate functions should be used instead of the faster approximations. For the typical case of our X-band spaces with k0 = 188.6, this means all propagation radii longer than roughly 3m can use the faster approximations with negligible

loss in accuracy. For a medium or large problem, this leads to more than a doubling of performance.

A.5 Performing the Bessel and Neumann Integrals

For the c++ implementation of the reconstruction, the Bessel functions are built into the programming language, simplifying the program dramatically. OpenCL on the other hand does not have built-in Bessel functions so the integral definitions of eqs. A.9 and A.10 must be programmed. We can readily pre-determine the argument value at which to switch over to the large argument approximations, now we need to determine how to efficiently program the integral and determine what parameter limits will meet the desired accuracy requirements in the small argument limit. For field reconstruction purposes, we're only interested in the first order Bessel functions (i.e. $\alpha = 1$). More complex calculation methods exist for functions of any order, but they are difficult to program and offer little speedup. The accuracy of the integrals is obviously dictated by the discretization of the integral. There is negligible difference between utilizing a trapezoidal method over a standard rectangular method for either integral. The integral parameters for 2- or 3-digit accuracy are shown in the tables below.

TABLE A-1: PARAMETERS TO OBTAIN A GIVEN $J_1(X)$ FUNCTION ACCURACY

Accuracy	Min Arg. Value (x)	Max Arg. Value (x)	τ Discretization
2 digits	0	50	π/32
3 digits	0	500	π/288

TABLE A-2: PARAMETERS TO OBTAIN A GIVEN $Y_1(X)$ FUNCTION ACCURACY

Accuracy	Min Arg. (x)	Max Arg. (x)	τ Discretization	t Discretization	t Upper Limit
2 digits	0.001	50	π/48	0.2	10
3 digits	0.001	600	π/416	0.025	10

A noteworthy difference between the calculations of $J_1(x)$ and $Y_1(x)$ is firstly the value at which the large argument approximation takes over for 3-digit precision is 500 for $J_1(x)$ and 600 for $Y_1(x)$. Additionally, the first order Neumann function has a lower bound on the argument accuracy limit (Table A-2). This lower bound shrinks as the upper limit on t is increased and grows when the upper limit on t is decreased. The lower argument bound for the Neumann function (x = 0.001) is unlikely to ever be encountered in most practical field reconstruction scenarios but can be adjusted as necessary by increasing the upper limit on t. Comparisons of the built-in Bessel functions to the integral calculations and large argument approximations are shown in the figures below, as well as the errors for each method.

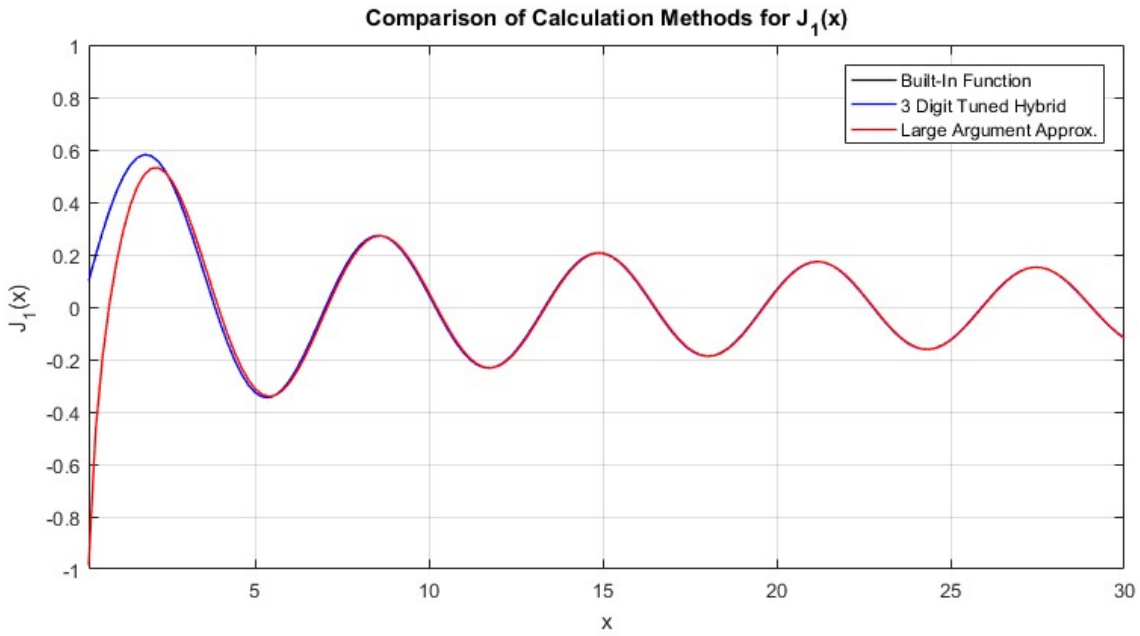

Figure A.12: Comparison of J1(x) calculation methods.

Figure A.13 and Figure A.**15** show the relative error of the large argument approximation and the numerical integral definition of the $J_1(x)$ and $Y_1(x)$ functions, respectively, against the "exact" built-in c++ functions. The red curves show the large angle

approximation, green is the hybrid method tuned to two digits of accuracy, blue is the hybrid at 3 digits of accuracy, and black is the numerical integral solution. The integral definition of $Y_1(x)$ is not near the double precision noise floor as $J_1(x)$ is, this is likely a result of the second complicated integral in equation A.10, but the error is sufficiently low.

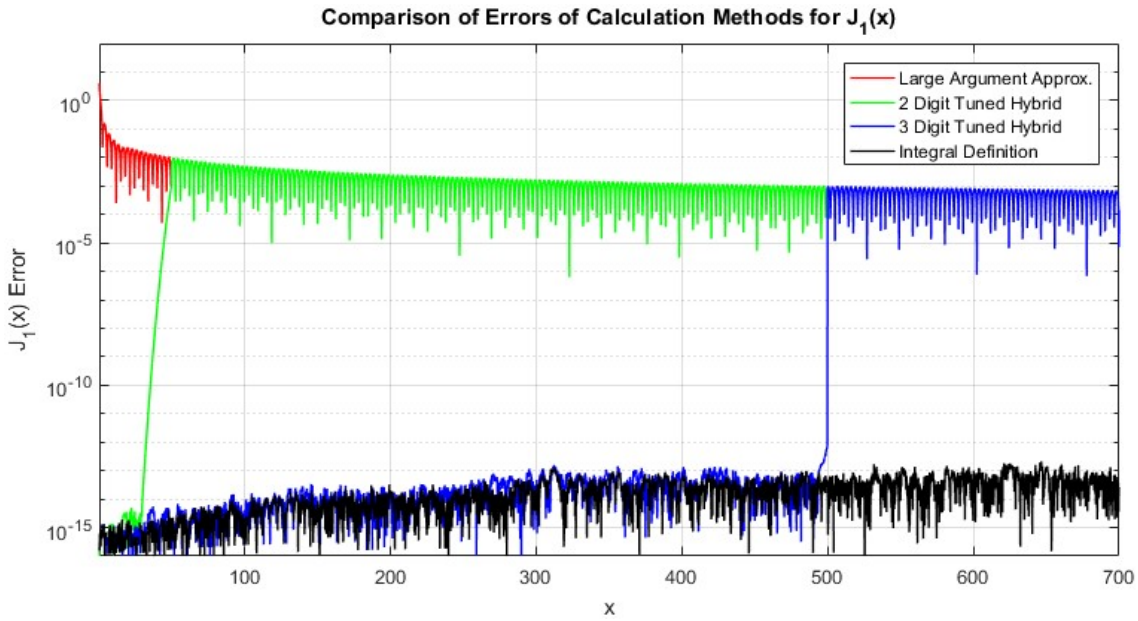

Figure A.13: Error of J1(x) calculation methods.

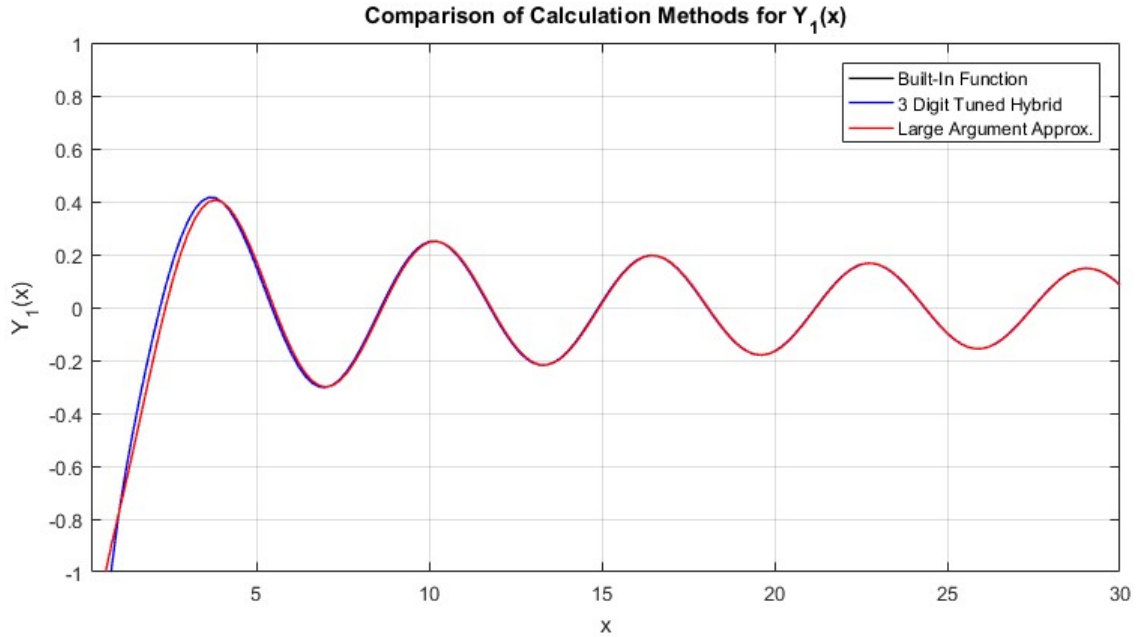

Figure A.14: Comparison of Y1(x) calculation methods.

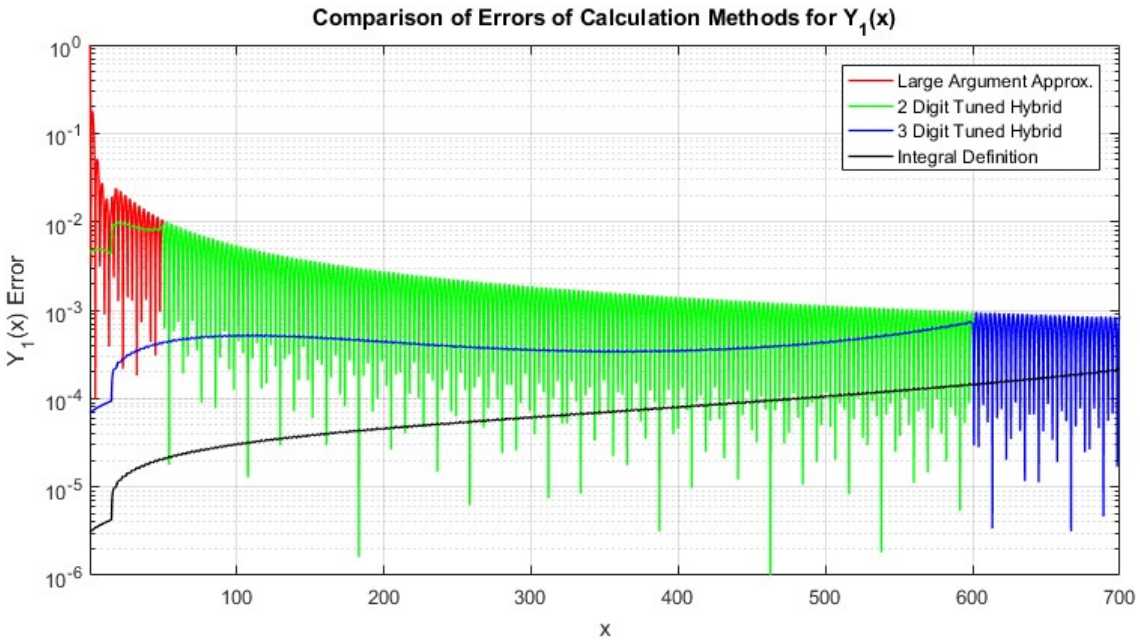

Figure A.15: Error of Y1(x) calculation methods.

A.6 Benchmark Results – CPU vs. GPU

As a comparison of high-end workstation CPUs (2x Xeon E5-2690, 32 threads) against a pair of middle-high-end consumer GPUs (2x AMD Radeon R9 290X), two scenarios are presented:

Scenario 1: Reconstruction of 9GHz scatter from a 1km sea state 3 sea

- P = 5, planes placed at 0m, 250m, 500m, 750m, and 1,000m
- M = 20,000, or 37.5m high at a discretization of 1.875mm

Scenario 2: Reconstruction of 9GHz scatter from a 1km sea state 3 sea

- P = 1022, planes placed 1m apart between source and end of simulation space
- M = 20,000, or 37.5m high at a discretization of 1.875mm

Table A-3 below shows a comparison between the speeds of using the exact built-in c++ Bessel functions in comparison to the hybrid method by performing scenario 1 on the CPU only (no built-in Bessel or Neumann functions exist in OpenCL.)

TABLE A-3: COMPARISON OF FIELD RECONSTRUCTION RUN TIMES FOR EXACT AND HYBRID BESSEL CALCULATIONS

Calculation	Run Time (seconds)	Speed Increase (Over Exact)
Scenario 1 – Exact	214.4	1.0x
Scenario 1 – Hybrid	78.7	2.72x

Table A-4 and Table A-5 show scenarios 1 and 2, respectively, each using the hybrid method, in order to compare CPU performance to single- and multi-GPU performance.

TABLE A-4: COMPARISON OF RUN RIMES FOR SCENARIO 1 ON CPUS AND GPUS

Calculation	Run Time (seconds)	Speed Increase (Over CPU)
Scenario 1 – CPU	78.7	1.0x
Scenario 1 – Single GPU	3.6	21.9x
Scenario 1 – Dual GPU	2.31	34.1x

TABLE A-5: COMPARISON OF RUN RIMES FOR SCENARIO 2 ON CPUS AND GPUS

Calculation	Run Time (seconds)	Speed Increase (Over CPU)
Scenario 2 – CPU	16,349	1.0x
Scenario 2 – Single GPU	406.2	40.2x
Scenario 2 – Dual GPU	235.8	69.3x

As is particularly evidenced by Table A-5, the GPUs are staggeringly quick in comparison to the CPUs, which can vastly outweigh the slight increase in complexity of the GPU reconstruction code. Provided the GPUs are sufficiently loaded with work, the number crunching capabilities far outweigh those of the CPU. This is further illustrated in APPENDIX B.

A.7 Conclusions on Implementation

In Section 3.4 we demonstrated the capability to exactly recreate the conventional FDTD scattered fields within a homogeneous space via spatial integration with a dispersion-corrected Green's function. In this appendix, the details of the method were expounded upon for the reconstruction of the co-polarized field, E_y, for the TEy case. The H_y fields for the TMy case can analogously be reconstructed as well. Under our 2D FDTD teleportation-based domain decomposition paradigm, the field information required from the Schelkunoff surface must include all fields up to and within the sub-domain which contains the target plane, including the fields in the sub-domain which may fall beyond the range of the target plane.

Next, the general considerations for parallelization were outlined with pseudocode examples for OpenMP-based CPU and OpenCL-based multi-GPU reconstruction. Since OpenCL has no built-in Bessel or Neumann functions, we blend together the integral and large argument definitions, and switch between them for appropriate argument values. We can tune the combination of hybrid and large argument definitions to attain a pre-determined accuracy, and we subsequently develop a recipe for the required integral definition parameter values and bounds on argument size. The hybrid approach speeds up calculation times by almost 3-fold for CPU-based reconstruction. Lastly, we demonstrated the astronomical speed advantage of the GPU-accelerated reconstruction compared to the CPU reconstruction, which far outweighs the minor complexity increase of the coding.

APPENDIX B

PROGRAM RUN TIME AND COMPUTER SYSTEM COMMENTS

B.1 Memory and Performance Gains from Domain Decomposition

Our domain decomposition method, presented in Section 2.2, reduces both memory and computational burdens of the FDTD space, and the size of the subdivision itself introduces further memory and performance considerations. In general, the smaller the sub-domain size, the lower the memory requirements, and the faster the overall run time of the FDTD space becomes. Reduction of the size of a sub-domain obviously reduces the field grid size, and thus lowers the memory required as well as the total number of calculations performed during each time step. Less obviously, it can reduce the total number of time steps required to fully simulate the sub-domain, as well as reduce the size of the time-history storage for the fields teleported between sub-domains. As an example, we will take a 20,000-cell long, 4,000-cell tall 2D domain and compare the memory requirements and run times for no subdivisions all the way up to 8 subdivisions.

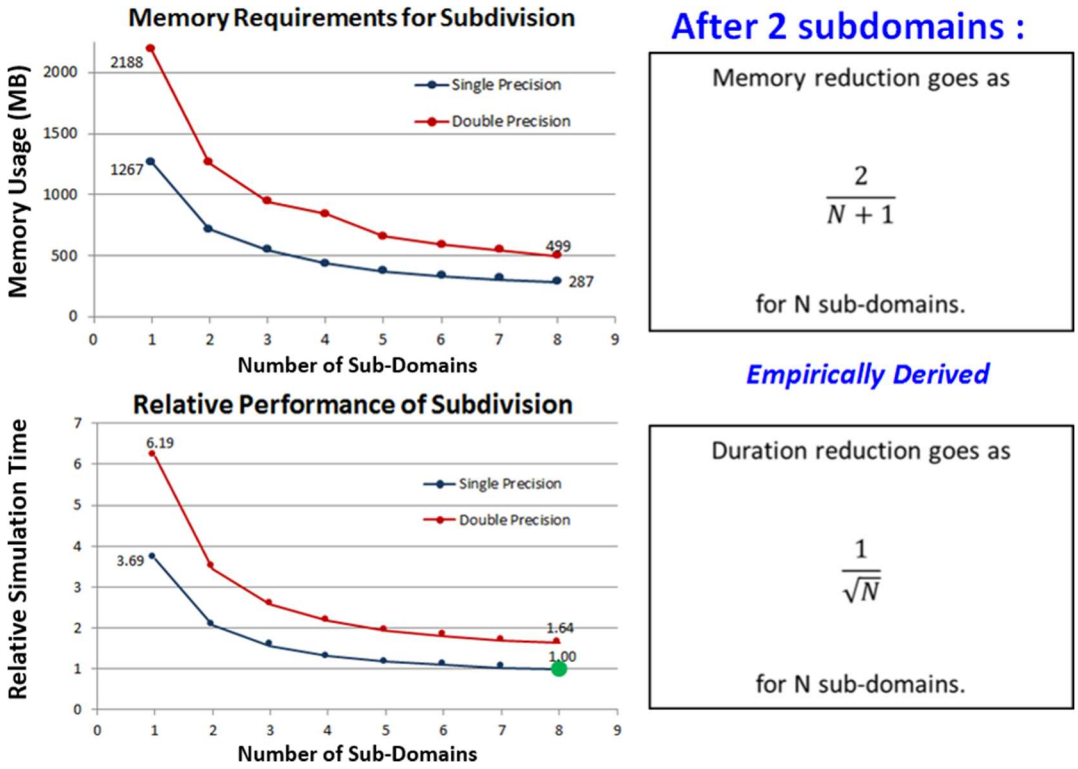

Figure B.1: Memory and performance comparisons for 8 different subdivision sizes of a large space.

Figure B.1 shows the performance and memory results for numerous subdivision sizes. The memory reduction is proportional to the inverse of the number of subdivisions of the total space, and the performance increases as the inverse square root of the number of subdivisions. Note, this test was performed on CPUs, which typically have much better relative double precision to single precision performance (usually double precision operates at 1/2 single precision performance) than typical consumer GPUs. Most professional GPU offerings from AMD and Nvidia have 1/2 or 1/3 double precision performance, much like CPUs. However, most AMD consumer-grade cards have 1/8 double performance, while Nvidia consumer cards run at 1/24 double performance.

Another important aspect of the implementation of the domain decomposition method lies in how the size of the teleportation field arrays increase for very large 2D and 3D sub-domains. In the case of the 2D FDTD and VTRPE comparisons of Section 3.2, each sub-domain was 8,000 cells long, 19,600 cells tall, and required about 38,000 time steps per sub-domain. The time history is collected in this case starting at roughly time step 11,000, so the dimensions for each component of the time history storage end up at 27,000 by 19,600 cells, which is more than three times the size of each field component of the FDTD grid and consumes about 65% of the total memory of the simulation. GPU memory is highly limited compared to CPU-side memory, thus we opt to store the time history in CPU memory. This is done either by direct PCI Express (PCI-e) bus transfers, or via mapping or pinning of CPU-side memory by the GPU. Since the time history storage arrays are accessed relatively infrequently there is negligible performance impact introduced.

B.2 Run Time and Computer System Comments for Several Problem Spaces

Several different computer systems were utilized for different simulation cases over the years, and the system specifications for the most commonly used computers are outlined in Table B-1. To provide a partial set of benchmarks a rough set of performance expectations for certain hardware configurations, several FDTD scenarios will be presented, and the run-times will be displayed for the systems utilized on those problem spaces.

TABLE B-1: SPECIFICATIONS OF COMMONLY UTILIZED COMPUTER SYSTEMS

	System 1	**System 2**	**System 3**	**System 4**
CPU	Intel Core i7 4770k	Intel Core i7 4770k	2x Intel Xeon E5-2690	2X Intel Xeon E5-2687W
System RAM	32GB DDR3-2133	32GB DDR3-2133	128GB DDR3-1600	256GB DDR3-1600
GPU	2x AMD Radeon R9 390X 8GB	2x AMD Radeon RX Vega 64 8GB	2x AMD Radeon R9 290X 4GB	3x Nvidia Titan X Pascal 12GB

The simulation test spaces and run times are shown in Table B-2 through Table B-11. Not all simulations were run on all systems.

TABLE B-2: FDTD COMPUTATIONAL TEST SPACE 1

Simulation Type	Conventional 2D FDTD
Total Domain Dimensions	4.1km long, 37m tall, 276 sub-domains
Total Time Steps	3.1 million
Sub-Domain Dimensions	8,002 cells long, 19,602 cells tall
Sub-Domain Time Steps	41,000
Simulation Run Time	System 1: 29 hours System 2: 22 hours System 3: 34 hours System 4: 21 hours

TABLE B-3: FDTD COMPUTATIONAL TEST SPACE 2

Simulation Type	Analytic Injection VEP 2D FDTD
Total Domain Dimensions	4.1km long, 3m tall, 276 sub-domains
Total Time Steps	3.1 million
Sub-Domain Dimensions	8,002 cells long, 1,602 cells tall
Sub-Domain Time Steps	18,000
Simulation Run Time	System 1: 2.9 hours System 2: 2.2 hours System 3: 3.5 hours System 4: 3.0 hours

TABLE B-4: FDTD COMPUTATIONAL TEST SPACE 3

Simulation Type	Conventional 2D FDTD
Total Domain Dimensions	15km long, 41m tall, 1,013 sub-domains
Total Time Steps	11.4 million
Sub-Domain Dimensions	8,000 cells long, 22,002 cells tall
Sub-Domain Time Steps	42,500
Simulation Run Time	System 4: 92 hours

TABLE B-5: COMPUTATIONAL TEST SPACE 4

Simulation Type	Conventional 2D FDTD
Total Domain Dimensions	1km long, 202m tall, 34 sub-domains
Total Time Steps	750,000
Sub-Domain Dimensions	16,002 cells long, 108,003 cells tall
Sub-Domain Time Steps	182,000
Simulation Run Time	System 4: 110 hours

TABLE B-6: COMPUTATIONAL TEST SPACE 5

Simulation Type	Conventional 2D FDTD
Total Domain Dimensions	4.1km long, 55m tall, 276 sub-domains
Total Time Steps	3.1 million
Sub-Domain Dimensions	8,000 cells long, 29,504 cells tall
Sub-Domain Time Steps	53,700
Simulation Run Time	System 3: 84 hours

TABLE B-7: COMPUTATIONAL TEST SPACE 6

Simulation Type	Analytic Injection VEP 3D FDTD
Total Domain Dimensions	1,050m long, 5.1m wide, 1.95m tall, 352 sub-domains
Total Time Steps	565,000
Sub-Domain Dimensions	1,002 cells long, 1,536 cells wide, 584 cells tall
Sub-Domain Time Steps	11,300
Simulation Run Time	System 4: 205 hours

TABLE B-8: COMPUTATIONAL TEST SPACE 7

Simulation Type	Analytic Injection VEP 3D FDTD
Total Domain Dimensions	176m long, 3.4m wide, 9.3m tall, 137 sub-domains
Total Time Steps	91,500
Sub-Domain Dimensions	320 cells long, 1,024 cells wide, 2,802 cells tall
Sub-Domain Time Steps	5,900
Simulation Run Time	System 4: 71 hours

TABLE B-9: COMPUTATIONAL TEST SPACE 8

Simulation Type	Analytic Injection VEP 3D FDTD
Total Domain Dimensions	114m long, 1.5m wide, 2.25m tall, 68 sub-domains
Total Time Steps	105,500
Sub-Domain Dimensions	1,002 cells long, 800 cells wide, 1,200 cells tall
Sub-Domain Time Steps	9,000
Simulation Run Time	System 4: 36 hours

TABLE B-10: COMPUTATIONAL TEST SPACE 9

Simulation Type	2D Total Space Reconstruction
Reconstruction Space	Total domain Ey field reconstruction with 1,023 target planes 37.5m in height on spaced 1m apart beginning at the source.
Total Target Plane Cells	20.46 million
Total Schelkunoff Cells	561,000
Reconstruction Run Time	System 1: 8 minutes System 2: 4 minutes System 3: 10 minutes

TABLE B-11: COMPUTATIONAL TEST SPACE 10

Simulation Type	3D Target Plane Reconstruction
Reconstruction Space	Ey and Ez reconstructed on a single 30m tall, 4.7m wide target plane placed 1,003m down range.
Total Target Plane Cells	12.69 million
Total Schelkunoff Cells	437.5 million
Reconstruction Run Time	System 1: 68 hours System 2: 49 hours System 3: 83 hours

It's worth mentioning a brief note on multi-socket CPU configurations and how those types of systems impact GPU performance. Multi-socket CPU systems have a set number of PCI-e lanes per CPU, and the lanes from each CPU are assigned to different

physical PCI-e slots. It is advantageous to ensure that the command queue for a given GPU is executed on the same CPU whose physical bus the GPU is connected to. GPU performance may be impacted if the CPU is unable to feed commands faster than the GPU completes them, and accounting for the physical PCI-e connection in this regard can improve performance through reduction of latency.

Furthermore, Multi-socket CPU systems typically have independent memory buses connected to each socket and generally utilize two modes for accessing memory: Uniform Memory Access (UMA) and Non-Uniform Memory Access (NUMA). Basically, UMA pools the memory connected to both sockets, so it appears as one large pool with a single average latency. NUMA on the other hand treats the memory connected to each CPU as its own pool. When a NUMA-enabled CPU accesses its local memory, it experiences lower latency than under UMA, but if it accesses memory connected to another socket, a latency penalty is incurred, and the result are latencies higher than UMA. These topics are advanced and well beyond the scope of this text. The memory access pattern of the system becomes relevant for multi-GPU configurations when we share boundary field information between GPUs twice during every time step. The boundary field sharing constitutes a very small, blocking data transfer which all GPUs use as a synchronization point. Therefore, obtaining low latency for the transfer translates to notable real-world performance gains (up to 20% in some 2D cases). Each GPU transfer involves intermediary storage in CPU memory. As such, it's worth considering the physical layout of the computer system when transferring data between GPUs. Professional-grade GPUs can alleviate this latency by providing DMI pathways directly between GPUs, skipping the intermediary CPU, however they are many times more expensive than their consumer-grade counterparts.

APPENDIX C

ADDITIONAL 3D SCATTERED FIELD RECONSTRUCTION PLOTS

C.1 Additional 3D Field Reconstruction Comparison Plots

This section contains additional reconstructed field comparison plots at ranges of 250m and 1,000m for the 1,050m 3D corridor simulation of Section 8.1.

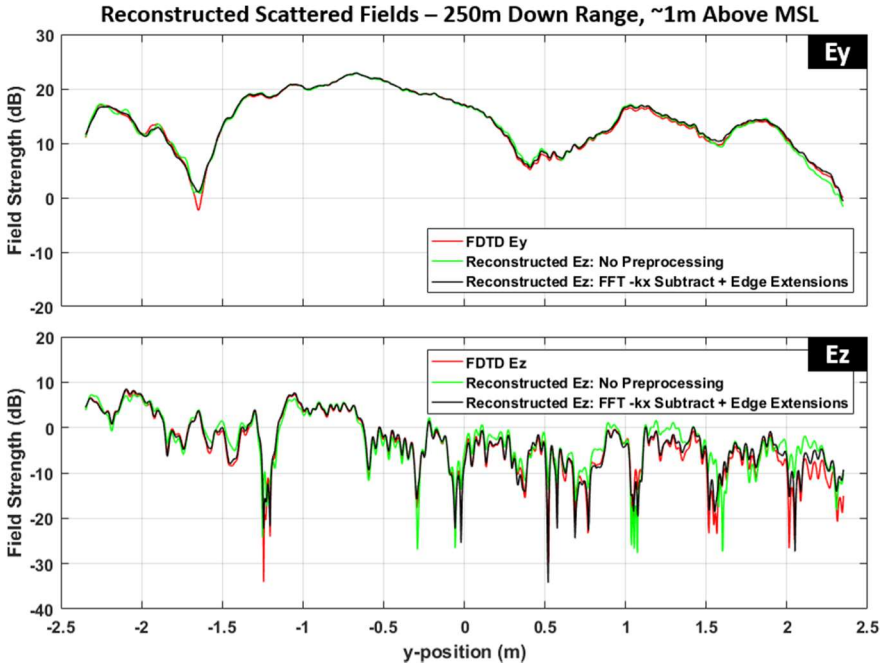

Figure C.1: Reconstruction results compared to FDTD at 250m down range.

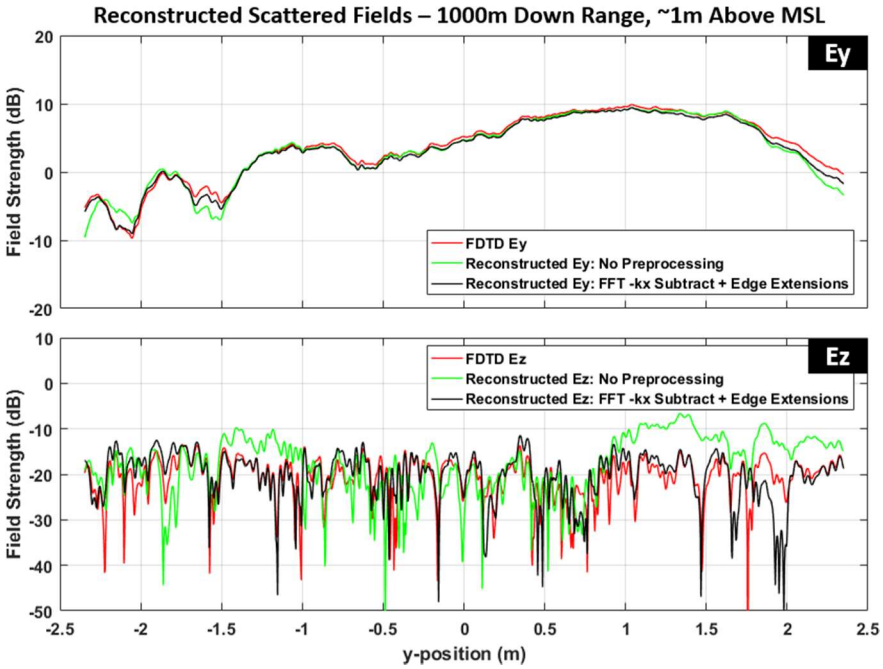

Figure C.2: Reconstruction results compared to FDTD at 1,000m down range.

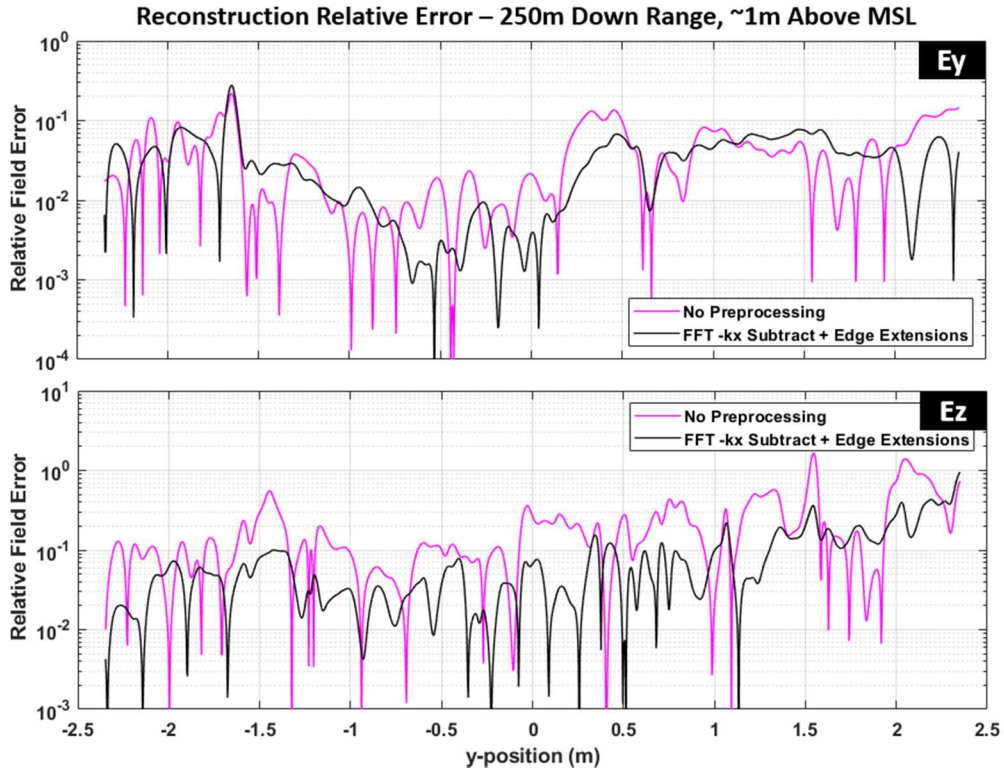

Figure C.3: Reconstruction relative field error to FDTD, compared at 250m down range.

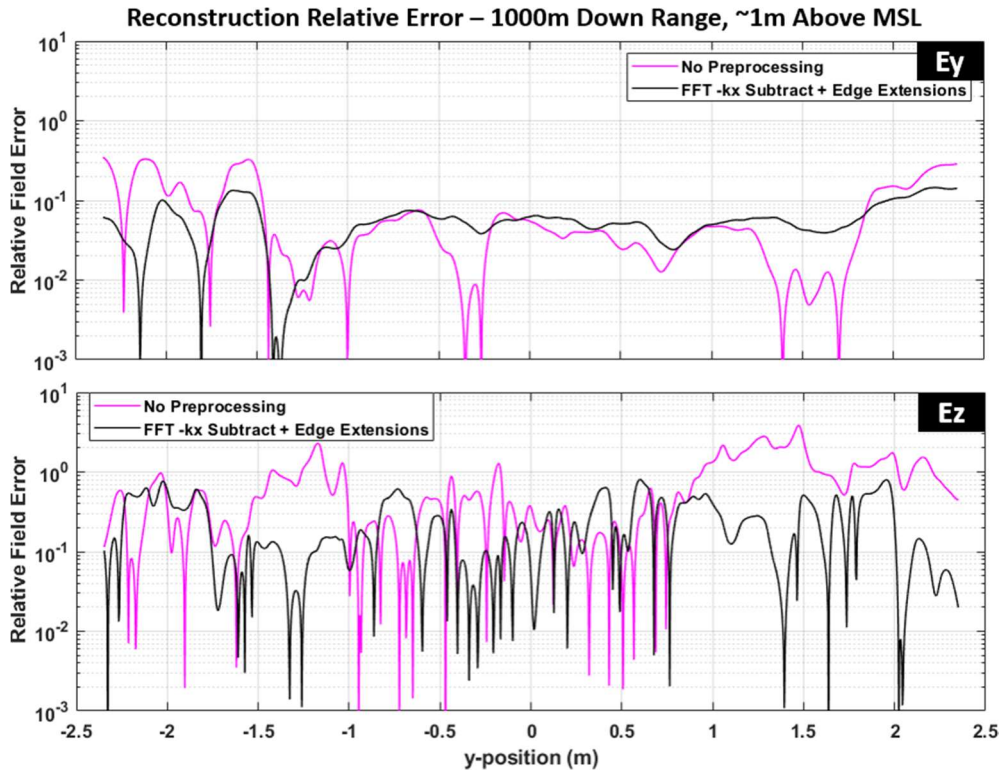

Figure C.4: Reconstruction relative field error to FDTD, compared at 1,000m down range.

C.2 Additional Error Plots for FDTD vs. Reconstruction Angle of Arrival

This section contains error plots for the AoA comparison in Section 8.2. The relative errors are those of the constant θ and ϕ line cuts of Figure 8.16-Figure 8.19.

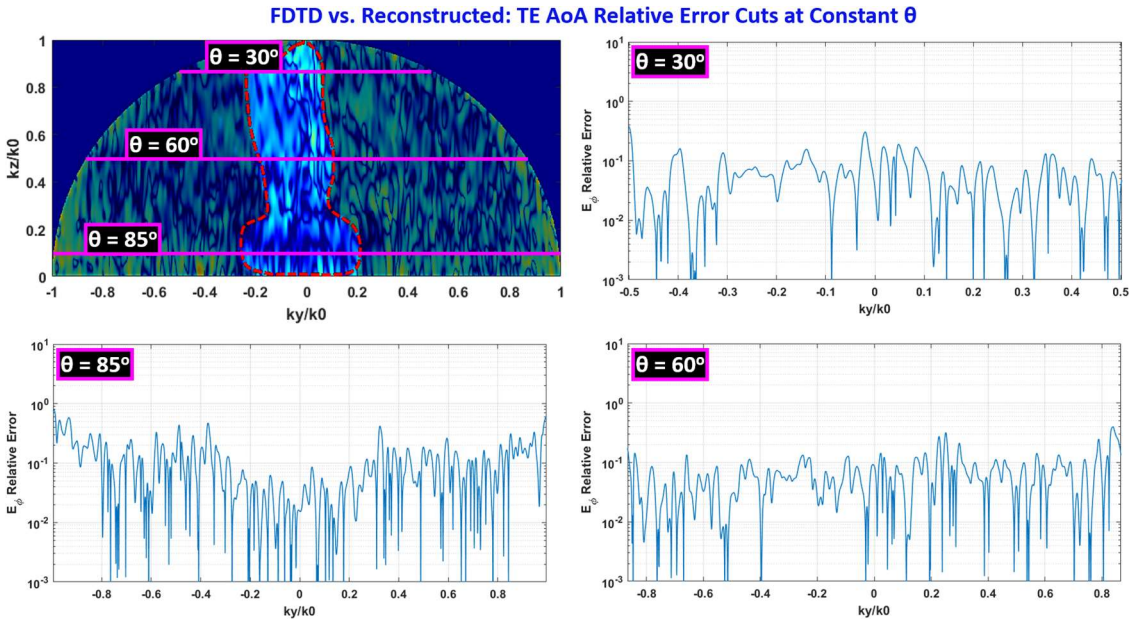

Figure C.5: AoA E_ϕ (TM) relative error cuts at constant values of θ, over all k_z in the visible region.

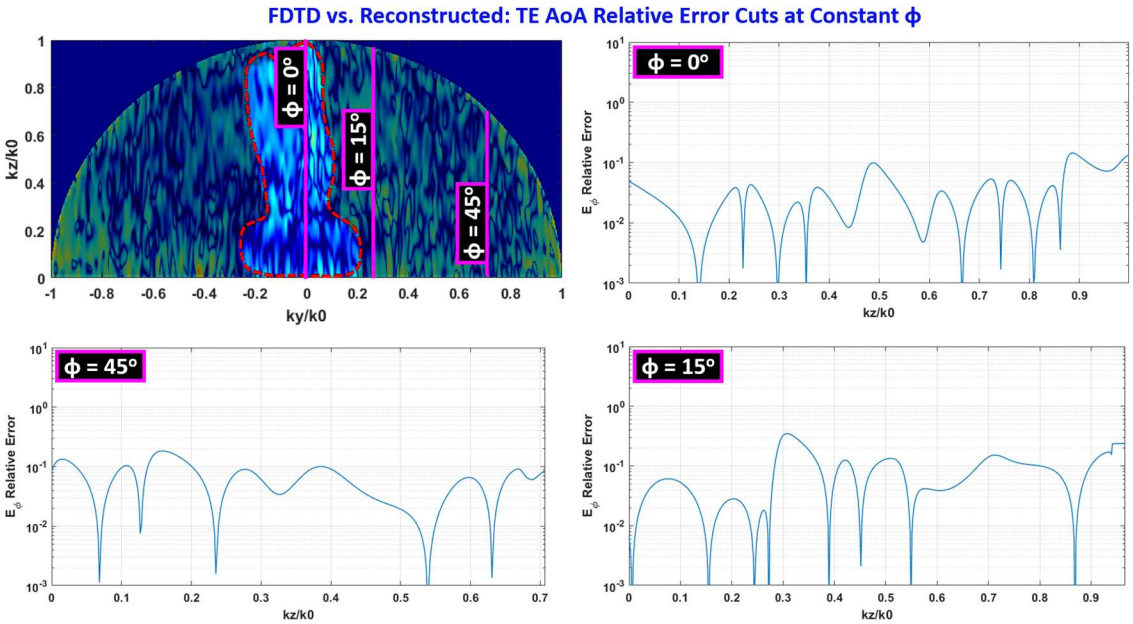

Figure C.6: AoA E_ϕ (TM) relative error cuts at constant values of ϕ, over all k_z in the visible region.

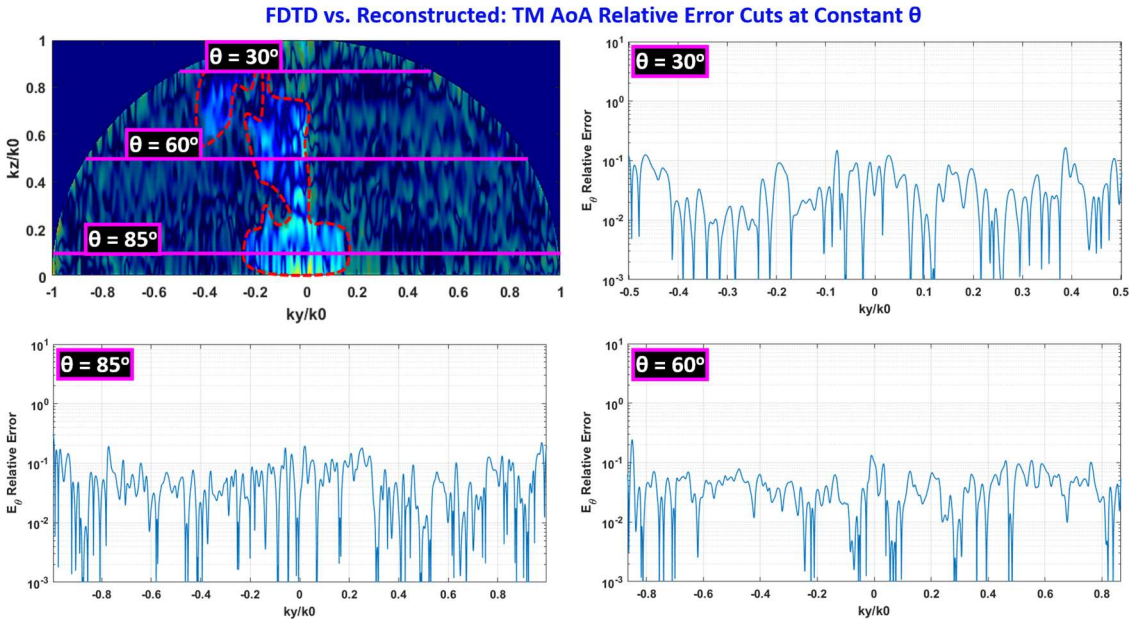

Figure C.7: AoA E_θ (TM) relative error cuts at constant values of θ, over all k_z in the visible region.

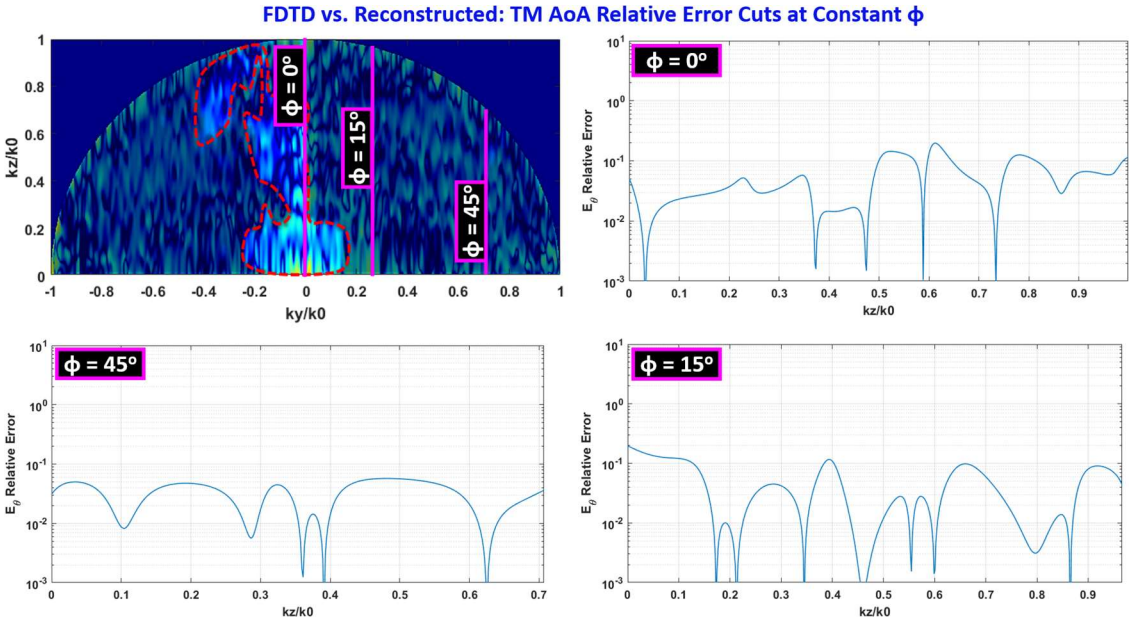

Figure C.8: AoA E_θ (TM) relative error cuts at constant values of ϕ, over all k_z in the visible region.