ABSTRACT

WOODRUFF, JOHNATHAN LUCAS. Subgrid Corrections in Storm Driven Coastal Flooding. (Under the direction of Casey Dietrich).

Coastal flooding models based on the numerical solution of the 2D shallow water equations are used widely to predict the timing and magnitude of inundation during storms, both in real-time forecasting and long-term design. Constraints on computing time, especially in forecasting, can limit the models' spatial resolution and hence their accuracy. However, it is desirable to have fast flooding predictions that also include the best-available representation of flow pathways and barriers at the scales of critical infrastructure. This need can be addressed via subgrid corrections, which use information at smaller scales to 'correct' the flow variables (water levels and current velocities) averaged over the model scale.

In this dissertation, subgrid corrections have been added to the ADvanced CIRCulation (ADCIRC) model, a widely used, continuous-Galerkin finite-element based, shallow water flow model. This includes the full derivation of averaged governing equations, closure approximations, and subgrid implementation into the source code. Testing of this new model was first performed on 3 domains: an idealized winding channel, a tidally influenced bay in Massachusetts, and a regional storm surge model covering Calcasieu Lake in Southwestern Louisiana with forcing from Rita (2005). By pre-computing the averaged variables from high-resolution bathy/topo data sets, the model can represent hydraulic connectivity at smaller scales. This allows for a coarsening of the model and thus faster predictions of flooding, while also improving accuracy. The implementation permits changing a logic-based wetting and drying algorithm to a more desirable logic-less algorithm, and requires averaging correction factors on both an elemental and vertex basis. This new framework further increases efficiency of the model, and is general enough to be used in other Galerkin-based, finite-element, hydrodynamic models. It is shown that the flooding model with subgrid corrections can match the accuracy of the conventional model, while offering a 10 to 50 times increase in speed.

Next, higher level corrections to bottom friction and advection were incorporated into the subgrid model, and the framework was expanded and tested at the ocean-scale. It was hypothesized that by adding higher-level corrections to the model and applying them to ocean-scale domains, accurate predictions of storm surge at the smallest coastal scales can be obtained. To accomplish this, higher-level corrections were derived and implemented

into the governing equations and extensive elevation and landcover data sets were curated to cover the South Atlantic Bight region of the U.S. Atlantic Coast. From there, the subgrid model was tested on an ocean-scale domain with tidal and meteorological forcing from Matthew (2016). The improvements in water level prediction accuracy due to subgrid corrections are evaluated at 218 observation locations throughout 1500 km of coast along the South Atlantic Bight. The accuracy of the subgrid model with relatively coarse spatial resolution ($E_{\text{RMS}} = 0.41$ m) is better than that of a conventional model with relatively fine spatial resolution ($E_{\text{RMS}} = 0.67$ m). By running on the coarsened subgrid model, we improved the accuracy over efficiency curve for the model, and as a result the computational expense of the simulation was decreased by a factor of 13.

Finally, subgrid corrections were systematically tested on a series of five ocean-scale meshes with minimum nearshore resolutions ranging from around 60 m on the highest resolution mesh to 1000 m on the coarsest mesh. This study aimed to find the mesh resolution that offered the best trade-off between accuracy and efficiency. The limitations of the subgrid model were explored and guidelines for future users were established. In all, it was found that the primary limitation to the subgrid model came from the aliasing of important flow-blocking features such as barrier islands in the coarsest resolution meshes. However, in areas without these features subgrid corrections can offer tremendous advantages while running on very coarse meshes.

The work completed in this dissertation moves the science of subgrid corrections forward by integrating the corrections into a widely used ocean-circulation and storm surge model. This work offers improvements to both hurricane storm surge forecasting and long term design by allowing for reduced run-times and increased accuracy on coarsened numerical meshes.

Subgrid Corrections in Storm Driven Coastal Flooding

by
Johnathan Lucas Woodruff

A dissertation submitted to the Graduate Faculty of
North Carolina State University
in partial fulfillment of the
requirements for the Degree of
Doctor of Philosophy

Civil Engineering

Raleigh, North Carolina
2023

APPROVED BY:

Andrew Kennedy	Rick Luettich
Helena Mitasova	Elizabeth Sciaudone

Casey Dietrich
Chair of Advisory Committee

DEDICATION

To my loving wife Ashley and my wonderful son Jack.
For their unconditional love and support.

BIOGRAPHY

Johnathan was born in Tallahassee, FL in 1992. Having been born and raised in Florida, he developed a love for the coastline and a passion for understanding and protecting it. During his undergraduate studies at the University of Florida, he took a few classes in coastal/water resources engineering and decided to pursue it further with a master's degree at Georgia Tech. There he specialized in coastal and water resources engineering and found his passion. At Georgia Tech, he took a particular interest in Coastal Hazards work which led him to the CCHT here at NC State working with Dr. Casey Dietrich. Outside academics, Johnathan enjoys running, biking, golf, tennis, surfing, skimboarding, skiing, diving, hiking and pretty much anything else outside. Apart from sports, Johnathan enjoys craft beer, playing guitar, hanging out with friends, and walking in the woods with his wife Ashley and son Jack.

ACKNOWLEDGEMENTS

I would like to acknowledge many people in my life that helped me complete this dissertation. First and foremost my wife Ashley and son Jack. Ashley, you have been there for me every step of the way, through thick and thin and you never cease to amaze me with your love, kindness, encouragement, and support. Jack, you offered me welcomed breaks in the last few months of my PhD in the form of refreshing outdoor walks and big smiles during playtime. I will be forever grateful for you two.

I would also like to acknowledge my parents Joy and Woody for always supporting and encouraging me in all my professional, academic, and life endeavours. My siblings Emily and Paden, you were my role models growing up, and always offered me sound advice (although you were somewhat blunt at times), and you have always been there for me when I needed you. I would also like to thank my in-laws, cousins, aunts, uncles, and grandparents who created a support system that never failed me and helped me pass any obstacle in my way.

This work would not have been possible without my advisor Dr. Casey Dietrich. Casey always supported and encouraged me from the start of my graduate journey. He mentored me and exemplified what it is to be a graduate research advisor, and he helped me become a better researcher, scientist, and engineer.

I am thankful for my committee for their support and expertise moving forward from my proposal and qualifying exams. I am also thankful for their patience and guidance while I was completing my dissertation and preparing for my defense.

Lastly, I would like to thank all of my friends in and outside the department and elsewhere around the world. You all have made my educational journey really amazing. Your friendships and encouragement got me across the finish line after all of these years. My friends in Raleigh especially have made the last 5 years really enjoyable, and have given Raleigh a very special place in my heart.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER

1

INTRODUCTION

1.1 Overview

The chapters of this dissertation focus on the improvement of coastal flooding predictions via subgrid correction factors in the widely used ADvanced CIRCulation (ADCIRC) model. Subgrid corrections use information at smaller spatial scales to correct flow variables (water levels and current velocities) at larger scales. These corrections have been used for several decades in hydraulic applications on small domains and can allow for accurate water level predictions on significantly coarsened computational grids. However, there are gaps in the research as to how effective subgrid corrections are for floods driven by hurricane winds and storm surge. In addition, there has been no implementation of subgrid corrections on an ocean-scale domain. Here, we describe the background of subgrid corrections and storm surge modeling and motivate the research in this dissertation.

1.2 Background

1.2.1 Hurricanes and Storm Surge Modeling

Storm surge is defined as the storm-induced rise in water level above the normal astronomical tide. In the Western Hemisphere, along the North Atlantic, Caribbean, and Gulf of Mexico coastlines, storm surge can be caused by a number of severe weather events. The most extreme and damaging storm surge is often caused by tropical cyclones. Tropical cyclones typically develop in the open ocean and are steered by Earth's prevailing winds. Along the U.S. Atlantic and Gulf coasts, these prevailing winds can steer hurricanes towards the U.S. mainland, which causes immense damage to both built and natural infrastructure and loss of life and property. Ian (2022) caused an estimated $112.9 billion in damages after making landfall in southwest Florida (National Centers for Environmental Information 2023) with more than 4.5 m of storm surge and wind speeds higher than 140 kt (Bucci et al. 2023). This devastating hurricane decimated the coastal communities of southwest Florida and caused the death of 66 people, of which 41 people were killed by storm surge, making it by far the deadliest hazard during the storm (Bucci et al. 2023). However, the impacts of Ian and many other hurricanes could have been worse had it had not been for hurricane storm surge forecasts, which alert emergency managers in coastal communities of incoming flood threats. These forecasts enable decision makers in these vulnerable areas to prepare for the incoming storm by planning evacuation zones and deploying appropriate emergency services personnel. Because the costs of evacuations (Whitehead 2003) and the risks associated with not evacuating vulnerable neighborhoods is extremely high (Borns 2022), obtaining an accurate and fast storm surge forecast is a top priority. These forecasts are produced by governments (National Hurricane Center 2021) and research institutions (Coastal Emergency Risks Assessment 2019) to help advise coastal communities during a storm.

As a tropical cyclone approaches the U.S. coast, the National Hurricane Center (NHC) issues advisories every 6 hours (National Hurricane Center 2023). These advisories are produced by a combination of large-scale meteorological models, observations, and NHC forecaster expertise, and contain up-to-date information about a cyclone's wind speed, location, forward speed, and pressure as well as forecast track and intensification information. An atmospheric model is then used to produce surface pressures and wind speeds for use in storm surge models as input into hydrodynamic equations that produce water level predictions along the coast (Figure 1.1). To be effective, the water level predictions must be completed quickly to help with decision making. Thus, it is imperative that these models run as efficiently as possible and produce accurate, realistic predictions that can then be passed on to communities that will potentially be affected by the storm. Depending on the hydrodynamic model in use, compute resources can range from a few computational cores on a laptop to several thousand cores on a large High Performance Computing (HPC) machine.

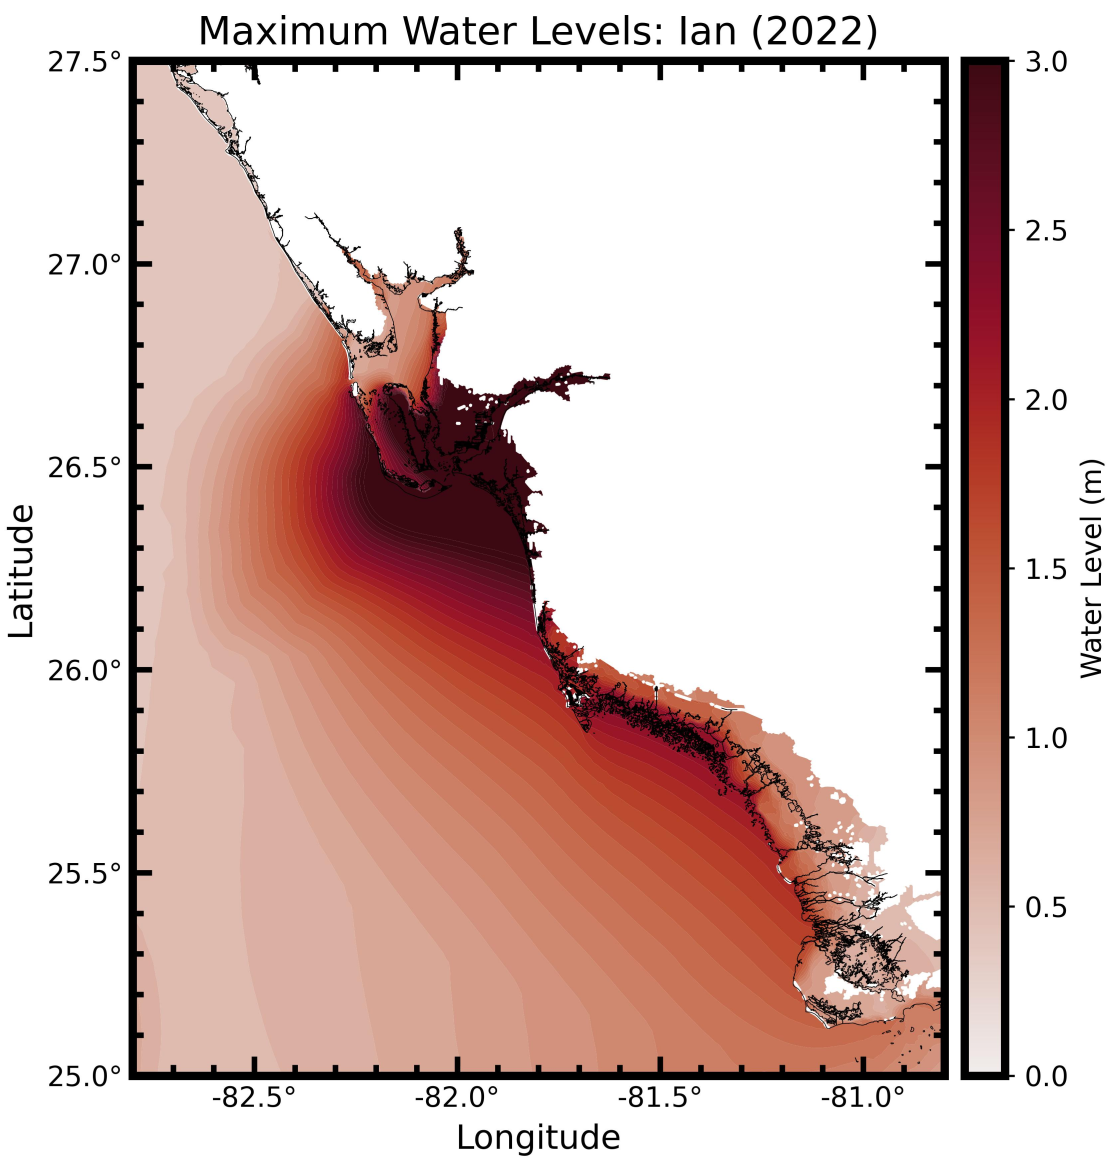

Figure 1.1: Maximum water levels prediction during Hurricane Ian (2022) produced by a storm surge prediction model for use in decision making prior the storms arrival.

In all storm surge models, there is a trade-off between accuracy and efficiency. While some models can run quickly and have minimal computational demand, they may be less accurate on a per-simulation basis (Kerr et al. 2013) and thus have to rely on hundreds of simulations to account for uncertainty in water level predictions and storm track (Taylor and Glahn 2008). These models often run on coarse, relatively small numerical grids that exclude a considerable amount of the geographic complexities that exist along the coastline. Other, more computational-intensive models take a different approach and run only a few perturbations in forecast. These models run on ocean-scale numerical grids with millions of computational cells that effectively describe important bathymetric and topographic features (Thomas et al. 2019). Many of these computational cells (50 to 90 percent) lie along inland flood plains so that the model can simulate inland penetration of surge during a storm (Roberts et al. 2021). Thus, the results produced from each simulation are often more accurate for a particular track and intensity forecast than the computationally inexpensive models (Fleming et al. 2008). Computationally intensive storm surge models require thousands of compute cores to complete a few simulations in the desired time (Dietrich et al. 2012). Ideally, these surge models could be made to run more efficiently while maintaining the accuracy of a high resolution simulation.

In recent years, there have been advancements to both of the modeling approaches mentioned above. Studies using the Probabilistic Surge (P-Surge) model which combines statistical analysis with an ensemble of Sea, Lake, and Overland Surges from Hurricanes (SLOSH) simulations have investigated improvements to the probabilistic approach (Kyprioti et al. 2021) and have sought to improve the underlying physics of the SLOSH model by incorporating high-resolution ground surface data (Begmohammadi et al. 2022). Other studies have progressed the more computationally expensive models such as the ADvanced CIRCulation (ADCIRC) hydrodynamic model by improving and streamlining large-scale mesh generation (Roberts et al. 2019b; Bilskie et al. 2020) and developing more stable numerical routines so that the model can be run with bigger time steps thereby decreasing model run times (Pringle et al. 2021). In addition, adaptive mesh switching has been used to improve run times by changing from coarse to fine meshes where the fine mesh is used to better resolve the area of the coast that will be impacted by the storm (Thomas et al. 2021). The advancement and increased use of machine learning have also led to studies that have replaced traditional, physics based modeling techniques with machine learning models that run in a fraction of the time (Lee et al. 2021; Lockwood et al. 2022).

1.2.2 Subgrid Corrections

Subgrid corrections have been used for decades to improve the accuracy and efficiency of hydrodynamic models by using high-resolution datasets of digital elevation models (DEMs) and landcover (Defina 2000; Casulli and Stelling 2011; Volp et al. 2013; Kennedy et al. 2019). These datasets inform the model of small-scale bathymetric and bottom roughness variation that would typically be aliased, because resolving features at the resolution of modern DEMs (≈ 1 m) (Danielson et al. 2018) would be too computationally expensive. These improvements allow for coarsened numerical grids to represent hydraulic connectivity through the smallest flow pathways represented in the DEM. Figure 1.2 demonstrates the advantages of using a subgrid corrections over conventional methodologies. Here, the subgrid model is able to represent flow through narrow inland canals and small tidal creeks that the conventional model cannot resolve.

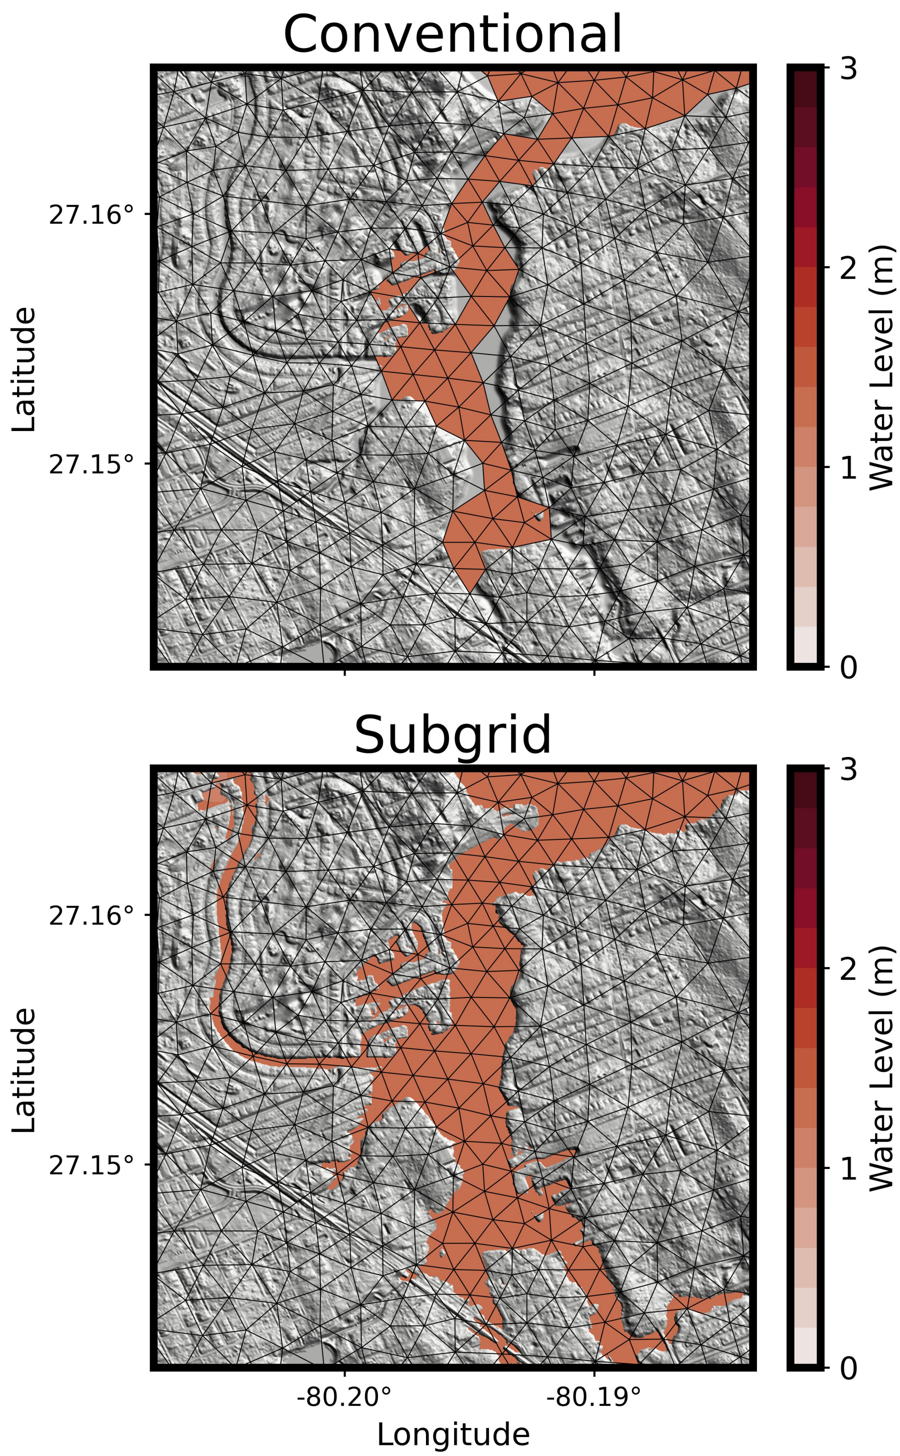

Figure 1.2: Maximum water levels predicted during Matthew near St. Lucie Inlet, FL produced on a relatively coarse ocean-scale numerical mesh overlayed on a high resolution DEM of the area. Note the ability of the subgrid model to push water into channels that are narrower than the resolution of the mesh.

Subgrid corrections are typically incorporated into hydrodynamic models by averaging the governing shallow water equations for mass and momentum (Defina 2000). These averaged equations are comprised of averaged flow variables that represent the integration and area-average of water surface elevation, velocity, and bottom stress present across an entire computational cell. In many subgrid studies, these averaged variables are pre-calculated and stored in lookup tables for a range of water levels, which are referenced during a simulation for a particular water level at a given timestep (Wu et al. 2016; Kennedy et al. 2019; Woodruff et al. 2021). The use of lookup tables enhances computational performance because no additional computations have to be performed while running the model.

In the past, subgrid studies have generally focused on relatively small (less than a few hundred square kilometers) regional domains. Forcing for subgrid models has typically been idealized or realistic tidal signals aimed to both develop and improve the fundamentals of subgrid correction factors in numerical models (Defina 2000; Casulli 2009; Casulli and Stelling 2011; Wu et al. 2016; Kennedy et al. 2019). A few of these studies have applied storm forcing from both wind and surge to their subgrid models (Sehili et al. 2014; Wang et al. 2014), and recently there has been an increase in efforts to incorporate subgrid corrections into widely used hurricane storm surge and ocean circulation models like ADCIRC, SFINCS, or SLOSH (Leijnse et al. 2021; Woodruff et al. 2021; Begmohammadi et al. 2022; Woodruff et al. 2023).

Subgrid corrections allow for significant (1 to 2 orders of magnitude) decreases in computational expense by running on coarsened grids while maintaining accurate flood predictions. It is critical that these methods be incorporated into widely used flood prediction models to not only improve accuracy for design studies, but also to reduce the time it takes to deliver predictions to coastal and inland communities during a storm event.

1.3 Motivation

The primary motivation driving the research in this PhD dissertation is the extension of subgrid corrections into the ocean circulation and storm surge model ADCIRC. Until recently, subgrid corrections had not been introduced into a widely used storm surge model, and had not been implemented on ocean-scale domains, which are required for accurate simulation of hurricane storm surge. This dissertation will investigate the necessary steps involved in incorporating subgrid corrections into a continuous Galerkin finite-element model, expand the subgrid areas to cover thousands of miles of coastline, and analyze the effects of decreasing resolution on ocean-scale flood predictions. The successful use of subgrid corrections in ADCIRC will enable storm surge model results to be delivered to emergency managers and decision makers faster leading to better preparedness prior to storm landfall. Subgrid corrections in ADCIRC will also offer cost savings when using simulations for design purposes like creating flood hazard maps and designing flood protection systems.

1.4 Dissertation Road-map

In this dissertation, we will explore ways to improve tropical cyclone storm surge predictions via subgrid correction factors with computationally intensive storm surge models. These corrections allow for accurate predictions on significantly coarsened numerical meshes thereby reducing computational time considerably. The theory behind these corrections, as well as their implementation on a wide range of both regional and ocean-scale domains, are explained and analyzed.

In Chapter 2, we will develop each step of implementing subgrid corrections into ADCIRC. This includes the procedure for averaging the governing shallow water equations, discretization of subgrid areas in a finite element model, calculation of subgrid correction factors, lookup table design, and testing on both idealized and realistic regional domains. The computational expense of adding subgrid corrections to ADCIRC is addressed, and the overall efficiency gains by running on coarsened meshes are analyzed.

In Chapter 3, we expand subgrid corrections to an ocean-scale storm surge model of the South Atlantic Bight using tidal and hurricane wind forcing from Matthew (2016). This chapter also includes the incorporation of higher level subgrid corrections to bottom friction and advection, and researches solutions to incorrect hydraulic connectivity in the subgrid model. High water marks and hydrographs from gauge data collected during the storm are used to validate the subgrid model, and results are compared to a high-resolution simulation.

Chapter 4 investigates how the systematic decrease in mesh resolution affects subgrid results when compared to traditional simulations. For this chapter, five ocean-scale, unstructured, triangular meshes were designed and tested with forcing from Matthew (2016) and Florence (2018). The pros and cons of increases in computational efficiency when compared to decreases in prediction accuracy are weighed, and a set of best practices is created to guide future subgrid ADCIRC users on maximum and minimum resolutions depending on the application.

Finally in Chapter 5, we will summarize the results from this dissertation and review the scientific contributions. The work in the dissertation has contributed to furthering the application of subgrid corrections in storm surge models, and developed technologies that can be widely used in the storm surge modeling community and increase the accuracy of flood predictions while decreasing the computational cost of running computationally expensive storm surge models.

CHAPTER

2

SUBGRID CORRECTIONS IN FINITE-ELEMENT MODELLING OF STORM-DRIVEN COASTAL FLOODING

2.1 Preface

In this chapter, we introduce subgrid corrections into the ADvanced CIRCulation (ADCIRC) hydrodynamic model. This includes the full derivation of the averaged governing equations, closure approximations, and subgrid implementation into the source code. Testing of subgrid ADCIRC was performed on 3 domains: an idealized winding channel, a tidally influenced bay in Massachusetts, and a regional storm surge model covering Calcasieu Lake in southwestern Louisiana with forcing from Rita (2005). This chapter has been published in *Ocean Modelling* (Woodruff et al. 2021).

2.2 Introduction

Storm surge, defined as the storm-induced rise in water above the normal astronomical tide, is the principal cause of loss of lives and damages to natural and built infrastructure during coastal storms. Storm surge can cause extensive flooding in regions with relatively flat coastal topography, such as the flooding of southeast Texas during Ike (2008), which pushed floodwaters up to 65 km inland (Hope et al. 2013). As storms become more intense due to climate change (Emanuel 2020), their associated flooding and impacts will be exacerbated. In the United States, about 7.1 million single-family and 250,000 multi-family residences are at risk of damage from storm surge, and the combined reconstruction costs, assuming complete destruction, of these structures has been estimated at nearly $1.8 trillion (CoreLogic 2020). There is a need to predict coastal flooding, both in real-time to aid in emergency management (Cheung et al. 2003), and between storms to aid in long-term planning and mitigation efforts (Helderop and Grubesic 2019).

Predictive numerical models must represent the evolution of storm surge over a wide range of spatial scales, from its generation in shallow shelfs, bays, and estuaries, to its conveyance into inland regions via narrow natural and man-made channels, to its interactions with hydraulic controls like dunes, levees, and raised roadways. The ADvanced CIRCulation (ADCIRC) modeling system (Luettich et al. 1992; Westerink et al. 2008) is widely used in coastal flooding predictions due partly to its use of unstructured, finite-element meshes, which can vary resolution from kilometers in the open ocean, to tens of meters in small-scale channels and inland regions. ADCIRC has been well-validated for predictions of storm surge along the U.S. Gulf and Atlantic coasts (Dietrich et al. 2011; Hope et al. 2013; Deb and Ferreira 2016; Cialone et al. 2017), often by using meshes with millions of elements to describe the coastal region of interest. However, this fine resolution (typically as small as 100 to 200 m) can lead to long simulation times. Although ADCIRC is highly scalable in high-performance computing environments (Tanaka et al. 2011; Dietrich et al. 2012), a typical ADCIRC storm surge simulation can require multiple hours of wall-clock time on hundreds (or thousands) of CPUs. Because of this, when ADCIRC is used for real-time forecasting (Fleming et al. 2008; Blanton et al. 2012; Dresback et al. 2013), it is limited typically to simulations of the consensus forecast and a few perturbations for each advisory. In contrast, other less computationally expensive models may consider an ensemble of storm scenarios to account for uncertainties in storm track, forward speed, and intensity. This method of ensemble forecasting is advantageous in that it gives researchers and emergency managers a broader view of potential storm impacts, thereby increasing their preparedness.

At the same time, because ADCIRC and other coastal models are used for predictions on regional (single or multiple state coastlines) domains, it has been computationally expensive for them to represent variability in topography and land cover at the highest available resolution. There has been significant improvement to both the quality and availability of topo/bathy data to describe the coastal zone. Databases, such as NOAA Digital Coast (National Oceanic and Atmospheric Administration 2020) and the USGS Coastal National Elevation Database (CoNED) (U.S. Geological Survey 2020a), offer high quality digital elevation models (DEMs) stretching across large swaths of coastline with resolutions typically ranging from 1 to 10 m. These geospatial data resolutions are much smaller than the mesh resolution used by flooding models. For model input, these data can be upscaled to identify the critical flow pathways and barriers that can be represented at the mesh scale (Bilskie et al. 2015), and with the model output, these data can be used to downscale the flooding predictions for decision support (Rucker et al. 2021). However, it has been cost-prohibitive to perform the model computations at the highest resolution of the geospatial data, thus limiting the accuracy of flooding predictions through the smallest channels and over the smallest roughness features.

Thus there is a need for faster flooding simulations that also represent flow pathways and barriers at the highest-resolution of geospatial data sets. This need can be addressed via *subgrid corrections*, which use information at smaller scales to 'correct' the flow variables (water levels, current velocities) averaged over the mesh scale.

Originally implemented to account for irregularities in model domains (Defina 2000), subgrid corrections have grown increasingly popular due to their abilities to improve accuracy, by better representing flows below the model scale, and/or efficiency, by enabling a similar prediction on a coarsened mesh. The governing shallow water equations are averaged to account for topography and bathymetry smaller than the model scale (Defina 2000; Casulli 2009; King 2001). These averaged equations contain variables that represent the integrated subgrid topography averaged over the computational cell area. Recent studies have shown a decrease in run time by 1 to 2 orders of magnitude when compared to simulations run on fine meshes, with the ability to decrease further if model time step were also increased (Sehili et al. 2014; Wu et al. 2016). Subgrid corrections have been demonstrated for synthetic domains to show proof of concept, and for relatively small, realistic domains like a tidally influenced marsh (Roig 1994; Bates and Hervouet 1999; Defina 2000; King 2001; Wu et al. 2016; Kennedy et al. 2019). Many of these studies forced their models with either a sinusoidal tidal curve, or with tidal data collected near the domain. Although some studies have forced a single flood wave (Viero 2019) and relatively minor storm surge events (Sehili et al. 2014), none have considered forcing due to hurricane winds, and thus there are remaining questions about the viability of subgrid corrections for storm-driven flooding.

We explore the use of subgrid corrections for predictions of coastal flooding in realistic domains using ADCIRC. *It is hypothesized that, even with a so-called 'Level 0' closure that corrects flow behavior only at the wet/dry front, the subgrid corrections will allow ADCIRC to better-represent the smallest flow pathways while using coarser resolution, thus improving both accuracy and efficiency.* We describe the implementation of subgrid correction factors into ADCIRC's governing equations. The performance of the model, with and without subgrid corrections, is evaluated on three test domains: an idealized winding channel domain, a small tidally influenced bay in Massachusetts, and a larger domain in southwestern Louisiana to provide a realistic storm surge scenario. It is shown that subgrid corrections can drastically improve storm surge predictions on coarse meshes. When tested on significantly coarsened meshes, subgrid ADCIRC can match the results of fine counterparts run with traditional methodology, while offering a 10 to 50 times increase in speed.

2.3 Methods

2.3.1 ADvanced CIRCulation (ADCIRC)

ADCIRC uses the continuous-Galerkin, finite-element method with linear C^0 triangular elements to numerically solve the 2D Shallow Water Equations (SWE). This set of equations consists of the depth-averaged continuity and momentum equations, which are solved for water surface elevations ζ and depth-averaged velocities U and V for coastal circulation (Luettich and Westerink 2004). ADCIRC solves the Generalized Wave Continuity Equation (GWCE), a reformulation of the primitive continuity equation into a generalized second-order wave equation, to avoid spurious oscillations associated with the primitive form of the equation (Kinnmark 1986). This study uses the so-called 'conservative' form of the momentum equations, in which the dependent variables are the fluxes UH and VH (where H is the total water depth), to ease the implementation.

The subgrid corrections will have their greatest effect in partially wet regions, and thus their implementation will require a revision to ADCIRC's wetting and drying algorithm. Traditional ADCIRC uses a complicated but robust system of logic to determine whether mesh vertices are wet or dry (Luettich and Westerink 1995b). It analyzes not only the values of total water depth but also water surface gradients and current velocities to update a wet/dry status of finite-element vertices during the simulation. These checks occur in the middle of each time-marching step, i.e. after the GWCE is solved for updated water surface elevations but before the momentum equations are solved for updated current velocities. A vertex becomes wet if a sufficient water surface gradient is large enough to allow a wetting velocity to its location, and it remains wet if its total water depth is sufficiently large. An element is considered wet only if its three vertices are wet; otherwise it is dry. Thus there cannot be any partially wet vertices or elements, in contrast to other algorithms (see Medeiros and Hagen (2013) for a review of various wetting/drying algorithms). This can lead to inaccuracies in the wet/dry front, especially if it is not resolved sufficiently at the mesh scale. However, Dick et al. (2013) showed in 1D that ADCIRC's wetting and drying algorithm is amenable to a partially wet scheme.

ADCIRC converts wind velocity to wind stress using the drag formulation from Garratt (1977). Wind stress is then applied to vertices in the momentum solver when solving for flow velocity. In this work, this formulation was revised to reduce the wind stress magnitudes in regions with shallow water depths, to mitigate the possibly unstable situation when high winds are blowing over a thin film of water. The wind stress is multiplied by a wind limiter (C_τ) in the form of a hyperbolic tangent function (Equation 2.1):

$$C_\tau = \tanh\left(\frac{\rho g H}{C_{ws}|\tau_s|}\right), \tag{2.1}$$

in which τ_s is the unaltered wind stress, ρ is the density of seawater, g is the acceleration due to gravity, H is the total water depth (which can be grid-averaged as defined below), and C_{ws} is a dimensionless constant ($C_{ws} = 2.5\text{e}6$ in this study). This limiter asymptotes to unity for low wind speeds and large water depths, but decreases to zero as water level decreases and wind speed increases.

2.3.2 Averaged Variables

We follow the methodology from Kennedy et al. (2019), which formalizes various aspects of earlier subgrid corrections in the context of SWE with unresolved bed profile at the model scale (Defina 2000; Casulli 2009; Volp et al. 2013). Flow variables, including the water surface elevation ζ above mean sea level, the total water depth $H = \zeta + h$ (in which h is the bathymetric depth), and the depth-averaged horizontal velocity components U and V, are averaged to the mesh-scale. It is noted that previous studies have used related but distinct approaches; the flow variable is first integrated over the subgrid cells in the area of interest, and then it is either area-averaged (Defina 2000) or left as a volume quantity (Casulli 2009). In this study, we perform an area-averaging.

Kennedy et al. (2019) describe a 'Level 0' closure, in which the mesh-scale areas are allowed to be partially wet. This requires the *a priori* computation of mesh-scale wet areas A_W, which are related to the mesh-scale total areas A_G via the wet-area fraction ϕ:

$$\phi = \frac{A_W}{A_G}. \tag{2.2}$$

Wet area fractions are pre-computed from a given high resolution topographical dataset typically available as a Digital Elevation Map (DEM). For a possible water surface elevation ζ, wet DEM cells are identified as being within the averaging area and having a positive total water depth. The number of wet cells divided by the total number of cells within the area is taken to be a wet area fraction ϕ. This process is repeated for the full range of possible water surface elevations, thus providing a look-up table to connect wet area fractions ϕ to water surface elevations ζ at every element and vice versa. With the wet area fraction ϕ, we can convert between wet-averaged and grid-averaged quantities. For any flow variable Q, the conversion is:

$$\langle Q \rangle_G = \phi \langle Q \rangle_W, \tag{2.3}$$

in which the angle brackets $\langle \cdot \rangle$ indicate an averaging to the wet (W) or total (G) area:

$$\langle Q \rangle_G = \frac{1}{A_G} \int_{A_W} Q \; \mathrm{d}A \quad \text{and} \quad \langle Q \rangle_W = \frac{1}{A_W} \int_{A_W} Q \; \mathrm{d}A. \tag{2.4}$$

There is a challenge to represent the averaged flow variables for an unstructured triangular mesh within a continuous-Galerkin, finite-element framework, due to its vertex-based placement of unknowns (ζ, U, V). This challenge is overcome via the use of representative areas for both elements and vertices (Figure 2.1). Elements are sub-divided into three sub-areas, with each sub-area corresponding to the area nearest a vertex. The elemental sub-areas surrounding a vertex are then combined to form a vertex area.

Averaged total water depth $\langle H \rangle$, averaged Manning's $\langle n \rangle$, and wet area fraction ϕ are pre-computed from a high-resolution DEM and land cover data for a range of possible water surface elevations (with an increment of 0.05 m in this study). The values are stored in lookup tables, and then referenced at every time step during the simulation.

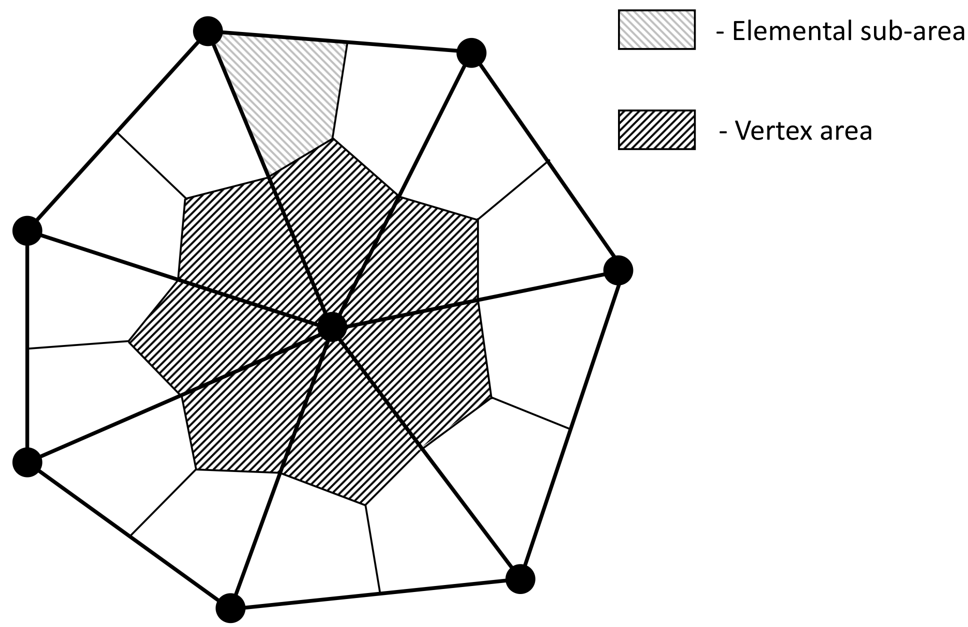

Figure 2.1: Schematic of: elemental sub-areas, which are created by dividing each element into three equal pieces; and vertex areas, which are created by combining the elemental sub-areas surrounding each vertex.

2.3.3 Averaged Governing Equations

In this work, we consider the governing equations arising from applying the formal averaging technique (Whitacker 1999) to the standard 2D SWE written in the conservative form (see detailed derivation in A.1). These equations involve averaged flow variables, namely the surface water level $\langle\zeta\rangle_W$, grid-averaged x- and y-directed fluxes $\langle UH\rangle_G$ and $\langle VH\rangle_G$; more precisely, they consist of the averaged horizontal x- and y-momentum equations in the conservative form:

$$
\begin{aligned}
&\frac{\partial\langle UH\rangle_G}{\partial t}+g\langle H\rangle_G\frac{\partial\langle\zeta\rangle_W}{\partial x}=-\frac{\partial\langle U\rangle\langle UH\rangle_G}{\partial x}-\frac{\partial\langle V\rangle\langle UH\rangle_G}{\partial y}+\\
&f\langle VH\rangle_G-g\langle H\rangle_G\frac{\partial P_A}{\partial x}+\phi\left\langle\frac{\tau_{sx}}{\rho_0}\right\rangle_W-\frac{C_f|\langle\mathbf{U}\rangle|\langle UH\rangle_G}{\langle H\rangle_W}\\
&+\frac{\partial}{\partial x}\left(\widetilde{E}_h\frac{\partial\langle UH\rangle_G}{\partial x}\right)+\frac{\partial}{\partial y}\left(\widetilde{E}_h\frac{\partial\langle UH\rangle_G}{\partial y}\right),
\end{aligned}
\tag{2.5}
$$

$$
\begin{aligned}
&\frac{\partial \langle VH\rangle_G}{\partial t} + g\langle H\rangle_G \frac{\partial \langle \zeta\rangle_W}{\partial y} = -\frac{\partial \langle U\rangle\langle VH\rangle_G}{\partial x} - \frac{\partial \langle V\rangle\langle VH\rangle_G}{\partial y} - \\
&f\langle UH\rangle_G - g\langle H\rangle_G \frac{\partial P_A}{\partial y} + \phi\left\langle \frac{\tau_{sy}}{\rho_0}\right\rangle_W - \frac{C_f|\langle \mathbf{U}\rangle|\langle VH\rangle_G}{\langle H\rangle_W} \\
&+\frac{\partial}{\partial x}\left(\widetilde{E}_h \frac{\partial \langle VH\rangle_G}{\partial x}\right) + \frac{\partial}{\partial y}\left(\widetilde{E}_h \frac{\partial \langle VH\rangle_G}{\partial y}\right),
\end{aligned} \tag{2.6}
$$

and the averaged continuity equation recast into the GWCE form:

$$
\begin{aligned}
\phi \frac{\partial^2 \langle \zeta\rangle_W}{\partial t^2} &+ \frac{\partial \phi}{\partial t}\frac{\partial \langle \zeta\rangle_W}{\partial t} + \tau_0 \phi \frac{\partial \langle \zeta\rangle_W}{\partial t} \\
&- \frac{\partial}{\partial x}\left(g\langle H\rangle_G \frac{\partial \langle \zeta\rangle_W}{\partial x}\right) - \frac{\partial}{\partial y}\left(g\langle H\rangle_G \frac{\partial \langle \zeta\rangle_W}{\partial y}\right) \\
&+ \frac{\partial \langle \tilde{J}_x\rangle_G}{\partial x} + \frac{\partial \langle \tilde{J}_y\rangle_G}{\partial y} - \langle UH\rangle_G \frac{\partial \tau_0}{\partial x} - \langle VH\rangle_G \frac{\partial \tau_0}{\partial y} = 0,
\end{aligned} \tag{2.7}
$$

where:

$$
\langle \tilde{J}_x\rangle_G = \text{RHS of (2.5)} + \tau_0 \langle UH\rangle_G \quad \text{and} \quad \langle \tilde{J}_y\rangle_G = \text{RHS of (2.6)} + \tau_0 \langle VH\rangle_G,
$$

in which f is the Coriolis parameter, g is the acceleration due to gravity, τ_{sx} and τ_{sy} are surface stresses, ρ_0 is a reference density, C_f is the bottom friction coefficient, E_h is the lateral stress coefficient, and τ_0 is a positive (spatially varying) parameter weighting the primitive continuity equation. In the above equations, the grid-averaged total water depth $\langle H\rangle_G$ (and $\langle H\rangle_W = \phi\langle H\rangle_G$) is assumed known for a given value of $\langle \zeta\rangle_W$. For the depth-averaged velocity, instead of using the formal definition of the averaged quantity (as in Equation 2.4), the averaged $\langle \mathbf{U}\rangle = (\langle U\rangle, \langle V\rangle)$ corresponds to the so-called volume-averaged velocity, more specifically $\langle \mathbf{U}\rangle = \langle \mathbf{U}H\rangle_G / \langle H\rangle_G$. This definition reduces to a point-wise definition of velocity in the limit of the averaging area approaching zero $A_G \to 0$; see A.1 for more detailed discussion.

Note that Equations 2.5, 2.6, and 2.7 are structurally similar to the form of the shallow water equations considered in ADCIRC except for the additional parameter ϕ and term $\partial\phi/\partial t$, the latter representing the time rate of change of the wet area fraction. The spatial and temporal discretization of this term is described in A.1.4. It is noted that these equations are nonlinear, both before and after the averaging; however, we avoid solving this nonlinear system through the time discretization scheme, which converts the equations into a linear algebraic system. The addition of the time derivative term in ϕ was an extra linearization step. As demonstrated later, it is important to note that C_f must be determined carefully, because a straightforward mesh-scale average formula does not necessarily ensure satisfactory results. Indeed, this aspect is the focus of ongoing research (Sehili et al. 2014; Viero 2019; Volp et al. 2013).

The GWCE is solved implicitly via the use of a global mass matrix, while the momentum equations are solved semi-implicitly. In this study, the ADCIRC solvers were kept the same, but averaged variables were substituted for their non-averaged counterparts. Both element- and vertex-based quantities are used in these solutions. On each time marching step, the GWCE (Equation 2.7) uses elementally-averaged quantities to find a vertex-averaged water surface elevation $\langle\zeta\rangle_W$. This quantity is then used to look up the corresponding vertex-averaged total water depth $\langle H\rangle_G$ and wet area fraction ϕ, which are used along with elementally-averaged quantities to solve Equations 2.5 and 2.6 for the vertex-averaged water velocities. Because we are solving averaged equations, the solutions for $\langle\zeta\rangle_W$, $\langle U\rangle$, and $\langle V\rangle$ are appropriately averaged. Therefore, no further manipulation to the solutions is required.

A primary contribution of this work is the use of a logic-free wet/dry algorithm. The new algorithm determines the wet/dry state by enforcing a minimum wet area fraction of the element:

$$\phi > \phi_{\min}. \tag{2.8}$$

This minimum fraction $\phi_{\min}$ is set by the user and can be adjusted depending on the application, e.g. a minimum wet area fraction $\phi_{\min} = 0.05$ would require that only 5% of an element must be submerged for it to be active and included in calculations. This new algorithm improves the code in several ways: replaces the existing algorithm and its extensive logic statements, gives a more accurate representation of the wet/dry front, smooths the transition between wet and dry elements and vertices, and allows ADCIRC to resolve subgrid hydraulic features.

2.3.4 Test Cases

Three test cases are used to evaluate the effectiveness of ADCIRC with subgrid corrections. The first test case is a plane sloping beach with a small winding channel of width 250 m in the middle of the domain. The domain is described by a synthetic 10-m DEM, which is then used to develop meshes with varying resolution to either fully or inadequately resolve the channel (Figure 2.2).

The second test case is a tidal simulation for Buttermilk Bay, Massachusetts. This domain is chosen because it has several well-defined, small-scale channels, which must be represented in numerical models for accurate predictions of flows into back bays (Kennedy et al. 2019). Coarse and fine meshes are generated for this domain, with bathymetry interpolated from a 3-m DEM (Figure 2.4). The topo/bathy data are obtained from NOAA Digital Coast (National Oceanic and Atmospheric Administration 2020).

The third test case is chosen as a realistic scenario for storm surge predictions. Using a 3-m DEM from USGS CoNED (U.S. Geological Survey 2020a), two ADCIRC meshes are created for Calcasieu Lake and the connected Bayou Contraband in southwestern Louisiana. Its location along the Gulf of Mexico, low-lying topography, and shallow, flat bathymetry make it highly vulnerable to storm surge. There are also numerous well-defined, small-scale channels in this region including Calcasieu Pass, Bayou Contraband, and intra-coastal waterways. With traditional ADCIRC, this domain requires a fine mesh (with resolution down to 50 m) to represent the hydraulic connectivity. There also exist water elevation data both at the coast and far up the bayou, which will serve to validate the results of the subgrid model.

2.3.5 Error Metrics

The accuracy and efficiency of the model will be evaluated in each test case. To evaluate accuracy with and without the sub-grid corrections on coarse meshes, we select three error metrics that are focused on the conveyance of tides and flood waters through channels below the model scale. First, for tides, we compute the duration (in hours) that channel locations are wet during one tidal cycle. We compare to predictions from a fine-mesh simulation, and thus an optimal result is a perfect match between durations on the coarse and fine meshes. Second, for flood waters, we consider the predicted peak water levels at channel locations. We compare to either the results from a fine-mesh simulation or to gauge observations, and an optimal result is a zero difference between peaks. Third, for both tides and flood waters, we consider the predicted maximum water levels along channel thalweg transects, i.e. the line connecting the deepest parts of the channel, again to examine the conveyance. We compare to results from a fine-mesh simulation by computing a root-mean-square error (E_{RMS}) using all points along the transect, and thus an optimal result is an $E_{\mathrm{RMS}} = 0$.

$$E_{\mathrm{RMS}} = \sqrt{\frac{\sum_{i=1}^{N}(x_i - \hat{x}_i)^2}{N}}, \tag{2.9}$$

in which N is the number of points along the transect, and x and $\hat{x}$ are the predicted maximum water levels from simulations on coarse and fine meshes, respectively. With these three error metrics, we assess the accuracy of predictions of flow through small-scale channels to inland locations.

Model efficiency was measured by wall-clock timings. Simulations were run on Intel Xeon E5-2650 v2 processors, which have 8 dual-thread cores per processor, 20MB of cache, and a frequency of 2.60 GHz. The processors are connected via an IB6131 Infiniband switch in the High-Performance Computing Services at North Carolina State University, but all simulations were run in serial to remove the inter-core communication times from the comparisons. For the timing comparisons, each simulation was run in triplicate, and the average wall-clock time was reported.

2.4 Results

2.4.1 Winding Channel

The first test has a 12-km by 12-km plane sloping beach with a 250-m winding channel (Figure 2.2). A synthetic DEM was created with a resolution of 10 m and with minimum and maximum elevations of −5 m and 2 m, respectively. The channel thalweg is always 1 m below the surrounding ground surface, and it was included to test the ability of the subgrid ADCIRC to represent flows below the mesh scale.

Two meshes are developed (Figure 2.2): a coarse mesh with average element side length of 1000 m, and a fine mesh designed to fully resolve the winding channel with a minimum resolution of 50 m and maximum of 500 m. The coarse mesh has 192 vertices and 334 elements, while the fine mesh has 12,475 vertices and 24,852 elements. Thus, the number of degrees of freedom of the coarse mesh is approximately 65 times less than that of the fine mesh. The bathymetry for both meshes is set using Inverse Distance Weighted (IDW) interpolation from the DEM (Burrough and McDonnell 1998).

We consider simulations of three run configurations: (1) fine traditional, (2) coarse traditional, and (3) coarse subgrid. Each simulation is forced by a 5-day diurnal tidal signal with amplitude of 1 m, with a 2-day ramp to prevent abrupt introduction of elevation

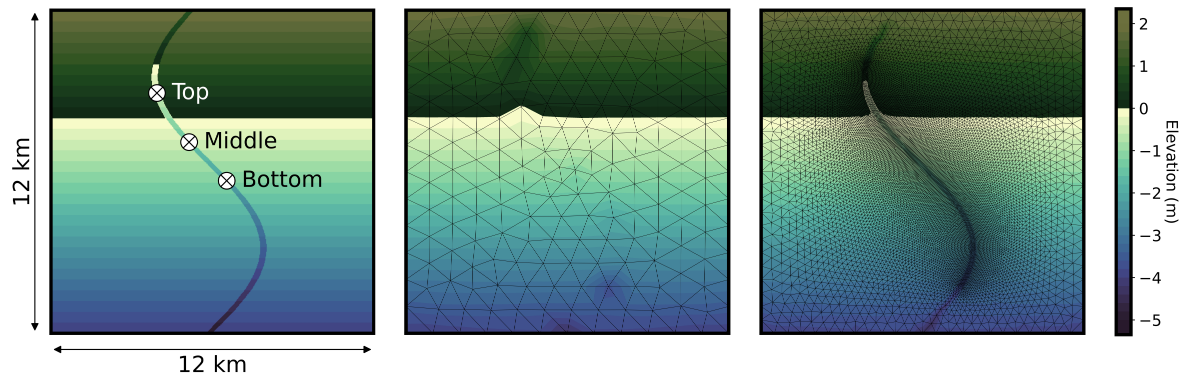

Figure 2.2: For the winding channel test case, (left) DEM and locations where water surface elevations were recorded, (center) coarse-resolution mesh, and (right) fine-resolution mesh. Contours indicate the ground surface elevations (m relative to mean sea level).

forcing. Bottom friction is computed with a constant Manning's coefficient of $n = 0.012$, and horizontal eddy viscosity is set to a constant value of $\tilde{E}_h = 20$ m^2/s. For traditional simulations, the wet and dry states are controlled by requiring a minimum wetting velocity of 0.1 m/s and a minimum water depth of 0.1 m, respectively. For the subgrid simulation, the minimum wet area fraction $\phi_{\min} = 0.05$.

Predicted water levels were recorded at stations along the channel thalweg and near the Top, Middle, and Bottom of the tidal range (Figure 2.2). The hydrographs show the ability of the subgrid corrections to represent the tidal behavior in this small channel (Figure 2.3). At the station near the top of the tidal range, the ground surface is −0.5 m relative to mean sea level. Because the domain is small enough to prevent a significant lag between the boundary forcing and the water levels within the domain, this station should be wetted for the 16 hr surrounding each peak tide. However, considering the fourth tidal peak (when the model forcing is at full strength), this wet duration is varied among the simulations (Table 2.1). The fine traditional simulation can represent about 12.5 hr, wetting when the water level rises to −0.04 m and drying when the water level falls to −0.10 m. The inability of the fine traditional simulation to represent the full 16 hr of the tidal peak at this location is likely due to: inaccuracies introduced when upscaling the synthetic ground surface to its 50-m resolution; and the binary nature of traditional ADCIRC's wet/dry algorithm, which can limit the predictions of the wetting front. The coarse traditional simulation can represent less of the high tide, or about 11.25 hr, wetting when the water level rises to 0.1 m and drying when the water level falls to 0.096 m. In contrast, the coarse subgrid simulation is

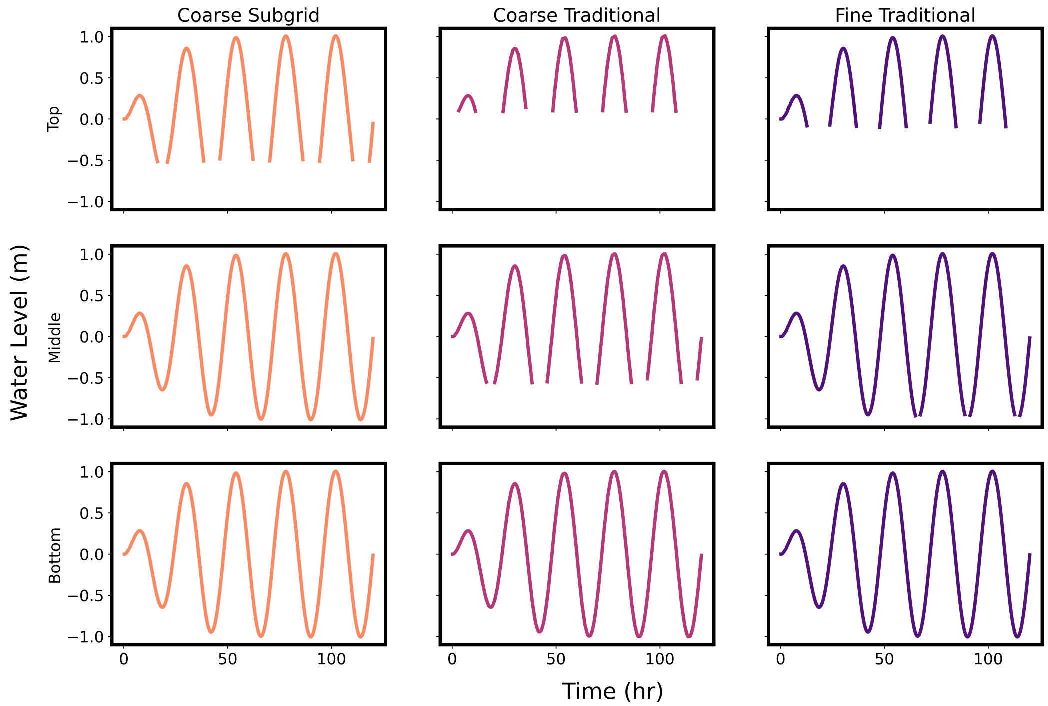

Figure 2.3: For the winding channel test case, time series of water levels (m relative to mean sea level). Order of plots from top to bottom matches the station locations in Figure 2.2.

able to represent the full 16 hr of high tide, wetting when the water level rises to −0.51 m and drying when the water level falls to −0.49 m.

The middle station is located where the ground surface is −1.45 m relative to mean sea level. This station should stay wet throughout the duration of the tidal cycle. However, both the fine traditional and the coarse traditional simulations become dry at the middle station. The fine traditional simulation represents 21.5 hr of the signal, becoming wet with the flood tide at a water surface elevation of −0.95 m and drying with the receding tide when the water level falls past the same elevation of −0.95 m. The coarse traditional simulation represents only 16.5 hr of the tidal cycle. The middle station becomes wet at a water level of −0.57 m and dries when the water level falls back to −0.56 m. The coarse subgrid simulation is able to represent the full tidal cycle at the middle station and does not dry at any time.

The bottom station is located where the ground surface elevation is −2.165 m relative to mean sea level. This station lies well beneath the lowest part of the tidal signal, and should

Table 2.1: For the winding channel test case, accuracy results for: wet duration (hr) during the fourth tidal period for each station, peak-to-peak difference (m) between coarse simulation and fine simulation, and E_{RMS} (m) of maximum water level along main channel thalweg between coarse simulations and fine simulations.

Simulation	Wet Duration (hr)			Peak-to-Peak Difference (m)			E_{RMS} (m)
	Top	Middle	Bottom	Top	Middle	Bottom	
Coarse Subgrid	16	24	24	8.0e-4	3.0e-4	8.9e-4	7.4e-5
Coarse Traditional	11.25	16.5	24	7.0e-4	2.6e-3	3.0e-3	1.2e-3
Fine Traditional	12.5	21.5	24	–	–	–	–
Theoretical	16	24	24	–	–	–	–

Table 2.2: For all test cases, wall-clock times (sec) for ADCIRC simulations on a serial processor, and ratios of wall-clock times. The average time of three simulations was reported.

	Winding Channel	Buttermilk Bay	Calcasieu Lake
Wall-Clock Times (sec)			
Coarse Subgrid	107	508	5,248
Coarse Traditional	62	277	3,728
Fine Traditional	5,787	4,176	167,514
Ratios of Wall-Clock Times			
Coarse Subgrid / Coarse Traditional	1.73	1.83	1.41
Fine Traditional / Coarse Subgrid	54.1	8.22	31.9

never dry. All three simulations were able to represent the full tidal range at the bottom station.

For the peak-to-peak differences and thalweg E_{RMS} relative to the fine mesh (Table 2.1), the values were about one order of magnitude smaller with the subgrid corrections, e.g. the channel thalweg $E_{\text{RMS}} = 7.4\text{e-}5$ for the coarse subgrid, but $E_{\text{RMS}} = 1.2\text{e-}3$ for the coarse traditional. However all of these peak-to-peak differences and thalweg E_{RMS} were very small for both simulations.

The subgrid corrections add computational time when compared to traditional ADCIRC simulations on the same mesh (Table 2.2). The increase in run time is attributed to reading the lookup tables, referencing to the tables at every time step of the simulation, and interpolating between table increments. For the coarse winding channel test case, subgrid ADCIRC ran 73% more slowly than its traditional counterpart. The efficiency of the subgrid implementation can likely be increased with better coding practices and smaller lookup table file sizes. However, the subgrid ADCIRC allowed flooding in the winding channel for more of the tidal cycle than a traditional simulation on a mesh with 65 times finer resolution, and it produced results 54 times faster. Thus the decrease in efficiency at the same mesh resolution is more than overcome by the increase in accuracy at coarser mesh resolutions for the subgrid corrections.

2.4.2 Buttermilk Bay

Buttermilk Bay is a small bay near the community of Bourne, Massachusetts (Figure 2.4). It is connected via the Cape Cod Canal to Cape Cod and Buzzards Bay to the north and south, respectively. A channel with a width of 250 m connects into a main bay with surface area of 1.54 km^2. From the main bay, a smaller channel with a width of 50 m connects into a smaller inner bay with a surface area of 0.42 km^2. Thus it is a good test to represent the propagation of tidal flows through channels below the model scale.

A high-resolution, 3-m DEM from NOAA Digital Coast is used to represent the bathymetry and topography, and two unstructured meshes are developed from this DEM (Figure 2.4). In the coarse mesh, the elements are 'paved' over the region, with no attempt to align their locations or sizes with the ground contours. The average element side length for the coarse mesh is about 100 m. In the fine mesh, vertices are aligned with the 0 m elevation contour to ensure that channels and coastlines are properly defined. The fine mesh has a minimum element side length of 10 m and a maximum of 50 m. The coarse mesh has 830 vertices and $1,569$ elements, while the fine mesh has $4,795$ vertices and $9,412$ elements.

The model parameters for the Buttermilk Bay simulations are similar to the winding channel test case. A diurnal tidal signal of 1 m amplitude with a 2 -day ramping period is forced at the ocean boundary. Constant Manning's $n = 0.022$ is applied over the entire domain. Horizontal eddy viscosity is set to $\tilde{E}_h = 2.0$ m^2/s for the fine simulation and $\tilde{E}_h =$ 50 m^2/s for the coarse simulation. For the traditional ADCIRC, the wet/dry parameters of minimum water depth and minimum velocity are set to 0.1 m and 0.1 m/s, respectively. For subgrid ADCIRC, the minimum wet area fraction $\phi_{\min} = 0.05$.

Water level results are evaluated at three stations in Buttermilk Bay (Figure 2.4). These stations are selected to evaluate the ability of subgrid ADCIRC to predict flow through regions with hydraulic features that are smaller than the resolution of the coarsened mesh. The Main station, located in the fully wet area of the domain, serves as a baseline to show all models were forced properly. The Arm station is in a small, tidally influenced stream that is between 5 m and 10 m wide. The Back station lies in Little Buttermilk Bay and is separated from the main bay by a 50-m wide channel.

Table 2.3: For the Buttermilk Bay test case, accuracy results for: wet duration (hr) during the fourth tidal signal for each station, peak-to-peak difference (m) between coarse simulation and fine simulation, and E_{RMS} (m) of maximum water level along the arm and back bay thalweg between coarse simulations and fine simulations.

Simulation	Wet Duration (hr)			Peak-to-Peak Difference (m)			E_{RMS} (m)	
	Arm	Main	Back	Arm	Main	Back	Arm	Back
Coarse Subgrid	13.75	24	24	1.4e-5	5.2e-4	1.6e-3	6.5e-4	8.9e-4
Coarse Traditional	0	24	24	–	5.3e-4	1.0e0	6.8e-4	5.5e-1
Fine Traditional	24	24	24	–	–	–	–	–

At the Main station, the water level time series is matched in all three simulations in both amplitude and phase (Figure 2.5 and Table 2.3). However, only the coarse subgrid and fine traditional simulations can capture hydraulic connectivity to the stations located in or near small channels.

At the Arm station, again considering the fourth tidal peak (when forcing is at its full strength), there is variability in the predictions. The coarse traditional simulation was unable to represent any water at the Arm station throughout the duration of the tidal signal. The fine traditional simulation is able to represent connectivity to the Arm for about 9 hr during the crest of the fourth tidal peak. It loses hydraulic connectivity from the Arm to the Main Bay at hour 82.5 and maintains a steady water surface elevation of 0.39 m for 15 hr until the return of the flood tide at hour 97. The surface elevation is maintained because, after connectivity is lost, water becomes trapped and cannot drain to the Main Bay, therefore this station remains wet throughout the simulation (Table 2.3). The coarse subgrid simulation maintains connectivity for 13.75 hr during the fourth tidal peak. Its Arm was fully dried at hour 85 and water surface elevation of −0.2 m for 10.5 hr until hour 95.5, when it floods again with the incoming high tide. Thus, at the Arm, the coarse subgrid simulation shows improved connectivity when compared to the coarse traditional and fine traditional simulations (Table 2.3).

At the Back station, the coarse traditional simulation indicates that there is water but no tidal flow, and thus the peak-to-peak difference is 1 m and the thalweg $E_{\mathrm{RMS}} = 5.5\text{e-}1$ m. The coarse subgrid and fine traditional simulations are able to represent flow through the small channel that connects from the main bay. For the coarse subgrid simulation, the errors are reduced by three orders of magnitude; the peak-to-peak difference is 1.6e-3 m and the $E_{\mathrm{RMS}} = 8.9\text{e-}4$ m.

The subgrid corrections increase the computational time when compared to traditional ADCIRC simulations on the same mesh (Table 2.2). For the coarse mesh, subgrid ADCIRC ran 83% more slowly than its traditional counterpart. However, it produced results that showed greater connectivity through small channels than a traditional simulation run on a mesh with 6 times the resolution, and its results were produced more than 8 times faster. The coarse mesh can likely be coarsened further, but was constrained by the width of the lateral boundary where the tidal forcing was applied. If this constraint was not present, further efficiency gains between the coarse subgrid and fine traditional simulations could be achieved.

2.4.3 Calcasieu Lake

The storm used in this test case was Rita (2005), which made landfall near the Texas/Louisiana border as a Category 3 hurricane on the Saffir-Simpson scale (Knabb et al. 2005). Lake Calcasieu and its neighboring communities were highly impacted by this storm due to their position in the northeast quadrant of the wind field and their low lying, flat topography. Maximum water levels reached 4.7 m along the coast with flood waters extending as far as 80 km inland (Dietrich et al. 2010; Berenbrock et al. 2008).

Similar to Buttermilk Bay, a coarse-resolution mesh is paved over the domain with no consideration of bathymetric details. The average element side length for the coarse mesh is 2000 m. A fine mesh is created with a minimum element side length of 50 m and a maximum of 2000 m. Vertices in the fine mesh are aligned along the 0-m elevation contour to ensure that channels and coastlines were properly defined. The fine mesh has a similar resolution and development as in larger studies of storm surge in the same region (Hope et al. 2013). The coarse resolution mesh has 1,236 vertices and 2,370 elements, while the fine mesh has 40,816 vertices and 81,321 elements (Figure 2.6).

The model parameters for the Calcasieu Lake meshes are interpolated from an ocean-scale, fine mesh available for this region. These model parameters include wind reduction factors derived from land-use/land-cover data, horizontal eddy viscosities in classes of $\tilde{E}_h = 2, 20, 50$ m^2/s, and values for the primitive weighting in the GWCE in classes of $\tau_0 = 0.005, 0.02, 0.03$. Manning's n coefficients for the meshes were derived from a 2006 Coastal Change Analysis Program (C-CAP) regional land cover dataset downloaded from the NOAA Digital Coast (National Oceanic and Atmospheric Administration 2021a). Values were interpolated onto the mesh vertices using a harmonic average of the Manning's n values contained in the surrounding vertex-elements. For the subgrid simulation, wet averaged

Manning's n values were computed prior to the simulation and looked up based on water surface elevations (Equation 2.10):

$$\frac{g\langle n\rangle_W^2 |\langle \mathbf{U}\rangle \langle UH\rangle_G}{\langle H\rangle_W^{4/3}}, \tag{2.10}$$

in which $\langle n\rangle_W$ is the wet averaged Manning's n. This was done to prevent overestimation of bottom friction in the subgrid model. Traditional simulations use a minimum water depth and a minimum velocity for wetting of 0.1 m and 0.1 m/s, respectively, while the subgrid model uses a threshold $\phi_{\min} = 0.05$.

The model is forced along its ocean boundary with water surface data taken from an ocean-scale ADCIRC simulation of Rita and winds produced by a Generalized Asymmetric Holland Model (GAHM) of the same storm (Gao et al. 2017). At every vertex, GAHM computes wind velocities and surface atmospheric pressures; the wind velocities are then scaled based on surface roughness and canopy cover present in the area. Parametric models such as GAHM can generate a reasonable representation of a hurricane wind field provided that proper wind parameters are used (Lin and Chavas 2012), and in this case, GAHM will provide a realistic forcing with which to evaluate the subgrid ADCIRC. The simulation is run for a total of 23 days, with water surface elevations recorded from locations in the mesh corresponding to USGS gauges deployed prior to the storm (U.S. Geological Survey 2020b), as well as locations spaced every 2000 m along the main channel thalweg from the Gulf of Mexico to Lake Charles, LA.

Predicted water levels are compared with hydrographs at the USGS gauges (Figure 2.7). Water levels at gauge stations LA12, LC7, LC8a, LC9, and LC12 were similar between simulations with differences less than 15 cm (Table 2.4). These gauges are located near the open coast, so when the 5 m storm surge propagated in, connectivity and subgrid corrections played less of a role in altering the overall water level. However, this is not the case for gauges LC2a, LC5, and LC6a, which are located further inland. At these locations, the coarse subgrid outperforms the coarse traditional simulation by more than 20 cm. Again, this is expected because, as the surge propagates further inland, the influence of subgrid features and flow connectivity have greater effects on the flow.

The most notable difference between the coarse subgrid and coarse traditional simulation is at the LC2a gauge located north of Calcasieu Lake (Figure 2.7). This gauge is farthest from the open coast and is connected via the narrow Bayou Contraband, and it recorded a maximum water level of 2.55 m during the storm. At this location, the coarse traditional simulation goes dry at 1100 UTC 24 September at a water level of 0 m and then rapidly wets

	Coarse Subgrid	Coarse Traditional	Fine Traditional
LA12	0.065	0.028	0.060
LC2a	0.423	1.328	0.152
LC5	0.281	0.538	0.435
LC6a	0.898	1.095	0.940
LC7	0.006	0.048	0.002
LC8a	0.312	0.327	0.412
LC9	0.180	0.192	0.155
LC12	0.202	0.182	0.206

Table 2.4: For the Calcasieu Lake test case, peak water level differences (m) for all simulations compared to the recorded gauge water levels during Rita (2005).

at 1400 UTC 24 September during the peak of the storm surge. The maximum water level of the coarse traditional simulation remains more than 1 m below the maximum surge predicted by the fine traditional simulation at this gauge, and is hydraulically disconnected from Calcasieu Lake.

The fine traditional and coarse subgrid simulations predicted a peak surge of 2.45 m and 2.18 m, respectively. Thus the coarse subgrid results are too low by about 0.27 m at this location when compared to the fine traditional results, likely due to high winds pushing water out of Calcasieu Lake (causing an excessive draw down), and a minimum wet threshold of $\phi_{\text{min}} = 0.05$, which may not fully capture the subgrid processes in Bayou Contraband.

To further evaluate the three simulations, maximum water levels were taken along the main channel thalweg from the Gulf of Mexico to Lake Charles, LA (Figure 2.8). From the north end of Lake Calcasieu to Lake Charles, the maximum water levels from the coarse subgrid simulation are 0.25 m below that from the fine traditional simulation, while the coarse traditional simulation underpredicts water levels by more than 1 m compared to the fine simulation. For the E_{RMS} along the main channel thalweg, the coarse subgrid E_{RMS} = 0.220 m, while the coarse traditional E_{RMS} = 0.564 m. This further demonstrates the superiority of the subgrid simulation at conveying flows through narrow channels.

For these simulations, the subgrid corrections add about 40% to the run-time when compared to the coarse traditional simulation (Table 2.2). However, the coarse subgrid was about 32 times faster than the fine traditional and was able to connect flow from the Gulf of Mexico, through Lake Calcasieu, and up the Contraband Bayou.

2.5 Discussion

In these test cases, the subgrid ADCIRC consistently out-performs its traditional counterpart in terms of hydraulic connectivity and maximum water level accuracy, and it allows for efficiency gains by using coarser meshes to represent coastal regions. These advancements have implications for the prediction of storm surge and coastal flooding, both in real-time forecasting and for long-term planning.

The subgrid corrections can be used for predictions with realistic storm forcing in realistic coastal domains. This is an extension of recent subgrid modeling studies, which have used water levels applied at the open boundary from idealized sinusoidal tidal curves or water level data from field measurements (Defina 2000; Casulli 2009; Kennedy et al. 2019; Wu et al. 2016). Sehili et al. (2014) used atmospheric forcing from a storm event in the North Sea; however this storm event was not on the scale or power of a tropical cyclone.

In our third test case, subgrid ADCIRC was forced with hurricane-strength winds and storm surge from Rita (2005). The model was able to represent the storm's effects on flow at the coast, more specifically, the flooding of the low-lying topography of southwest Louisiana, and the flow through channels smaller than the model scale. The largest discrepancy between coarse traditional and subgrid simulations was at the LC2a gauge where the model resolution was about seven times larger than the 300-m-wide Bayou Contraband.

Subgrid ADCIRC also allows for a coarsening of the meshes used to describe the coastal region. For the winding channel test case, nearly identical maximum water levels were predicted in the channel by the coarse subgrid and the fine traditional simulations, with improved connectivity in the coarse subgrid simulation (Figure 2.3), despite the coarse subgrid simulation having 65 times fewer degrees of freedom and a minimum resolution that was 20 times coarser. The simulation of Buttermilk Bay also showed virtually no difference in maximum water levels between the coarse subgrid and fine traditional simulations, despite the coarse subgrid simulation having almost 6 times fewer degrees of freedom and a minimum resolution that was 10 times coarser. Again, the subgrid showed better hydraulic connectivity, especially in locations in small-scale channels, than the fine traditional. For the Calcasieu Lake test case, the coarse subgrid and fine traditional showed good comparison to gauge observations from Hurricane Rita (2005), despite the coarse subgrid simulation having 33 times fewer degrees of freedom and a minimum resolution that was 40 times coarser. The subgrid simulation was able to represent flows to the inland LC2a gauge, because it allowed flow through the Bayou Contraband below the model scale. These results are similar to those by Kennedy et al. (2019), Sehili et al. (2014), and Wu et al.

(2016) who found that the subgrid corrections allowed for a coarsening of meshes by at least 1 order of magnitude.

These advancements have implications for real-time forecasting and long-term engineering and design. When ADCIRC is run traditionally with fine-resolution meshes, each simulation can require thousands of compute cores and hours of wall-clock time (Hope et al. 2013). During a storm event, this requirement can limit its use in a probabilistic forecasting framework, which can account for slight variations in storm track, intensity, and timing (Fleming et al. 2008), and which is used by other forecast models like the Sea, Land, and Overland Surges from Hurricanes (SLOSH) model (National Hurricane Center 2020). Subgrid ADCIRC may enable probabilistic forecasting. Between storms, agencies like the United States Army Corps of Engineers (USACE) and the Federal Emergency Management Administration (FEMA) use ADCIRC to better prepare coastal cities and communities from future flooding events (U.S. Army Corps of Engineers 2015; Federal Emergency Management Agency 2019), typically by simulating hundreds of synthetic storm surge scenarios to produce flood hazard maps for state and municipalities. Subgrid ADCIRC could drastically reduce these studies' computational and monetary cost.

These results do indicate paths for future work, specifically in the drawdown and underprediction of water levels at the LC2a gauge, and the consistent underprediction of water levels by the subgrid simulation along the main channel thalweg in the Calcasieu Lake test case (Figure 2.7 and 2.8). In those tests, the coarse subgrid consistently under-predicted water levels when compared to the fine traditional. This may be attributed to an over-estimation of friction by the subgrid model, and increases in Manning's n values from interpolation to the coarsened mesh. Volp et al. (2013) presented a scheme to correct this over-prediction and take advantage of high resolution roughness data. Implementation of a friction correction should lend itself well to the current subgrid framework present in the code. The drawdown that occurred at this location as the storm made landfall is present in both the fine traditional and coarse subgrid simulations; the water levels are decreased to -1.0 m, or about 1.3 m below the gauge data. These differences can largely be attributed to the gauge installation. The LC2a gauge was a barotropic pressure sensor mounted sub-aerially at 0.303 m NAVD88. Therefore, the sensor was not able to measure a drawdown below 0.303 m. The flattening of the gauge data from 0000 UTC to 0500 UTC 24 September indicates that the water level dropped below the gauge mount elevation, thus there is no way of verifying prediction accuracy during this time period. Other factors that could have affected model accuracy include poor representation of vertical features like roadways and levees that lie along the lake's edge and act as hydraulic barriers to keep water in the lake during the storm. These hydraulic features can be better represented with *cell clones*, which prevent flow between non-hydraulically connected features. Previous implementations of cell clones have used numerical schemes in which the velocities are located along the cell edge, which allows for connectivity and/or blocking of flows within the cell (Begmohammadi et al. 2021; Casulli 2019). This capability will be challenging to implement in ADCIRC, because the model defines the flow variables at the vertices of each element, and thus it is not straight-forward to identify connectivity for each clone. However, the capability would better represent the blocking of flow due to subgrid obstacles.

2.6 Conclusions

In this study, subgrid corrections were implemented in the widely-used ADCIRC model for storm surge and coastal flooding. These corrections were tested on a variety of domains and showed promising results both for idealized and realistic tides and storm surge. Subgrid ADCIRC is able to capture hydraulic connectivity and water level calculations on coarsened meshes in which small hydraulic features are not resolved at the mesh scale. This improvement is attributed to subgrid ADCIRC's ability to represent small hydraulic features contained within partially wet elements. Without the use of sufficiently small element sizes, traditional ADCIRC cannot resolve these features. The inclusion of partially wet elements to solve for water levels and velocities was achieved by redesigning the wetting and drying routine within the code to solely rely on the wet area fraction (ϕ) when determining the wet/dry state of an element or vertex.

The main contributions and findings of this study are:

1. *Extension of subgrid corrections using the widely used ADCIRC storm surge model with hurricane strength forcing.* The addition of subgrid corrections to ADCIRC's governing equations allowed for use of partially wet elements and vertices. This permits modified storm forcing at the wet/dry boundary by way of the wet area fraction. Testing on the realistic Calcasieu Lake domain using forcing from Rita (2005) demonstrated that these modifications give good overall matches to gauge hydrographs when run on coarsened meshes.

2. *Subgrid corrections in ADCIRC allow for increases in accuracy and hydraulic connectivity when running on significantly coarsened meshes.* In a forecasting scenario, this would give emergency managers and decision makers a more-accurate prediction

of when flood waters will arrive and recede. This will allow them to use the best information possible when deciding evacuation times and coordinating search and rescue missions.

3. *For a given grid, introducing subgrid corrections to ADCIRC increases computational cost to the code; however, these costs are small when compared to the efficiency gained by running on coarsened meshes.* In our current implementation, the coarse subgrid storm surge simulation on Calcasieu Lake is approximately 40% slower than its coarse traditional counterpart. Nevertheless, it ran 32 times faster than the fine simulation and produced comparable results, reducing the simulation run time from 42.2 hours to 1.3 hours.

With these additions, subgrid ADCIRC has the potential to predict coastal flooding at a fraction of the computational cost. Further investigation is needed as to whether this efficiency can be further increased with adjustments to the model time step. Future work will include tests of subgrid ADCIRC on ocean-scale domains, the use of ensemble frameworks to forecast storm surge, and the use of additional correction such as friction to further improve model results.

2.7 Acknowledgments

This study was supported by National Science Foundation grants ICER 1664037 and 1664040. The support of the National Science Foundation and the US Army Corps of Engineers is greatly appreciated by all contributors to this project.

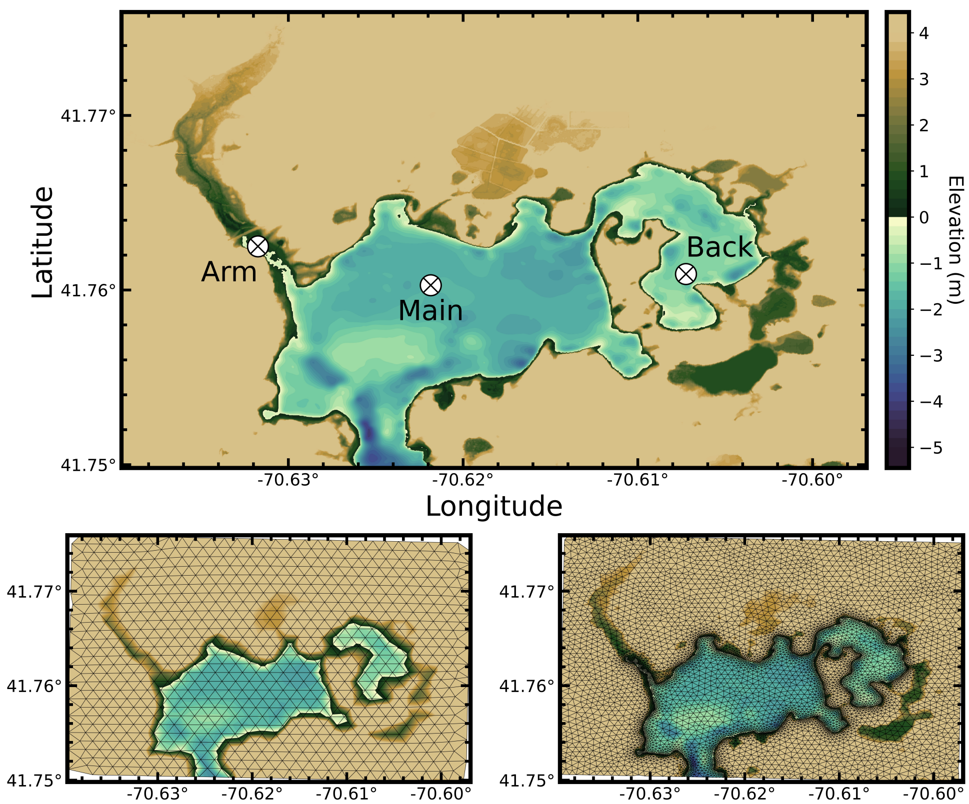

Figure 2.4: For the Buttermilk Bay test case, (top) DEM in WGS84 coordinates and locations where water surface elevations were recorded, (bottom left) coarse-resolution mesh, and (bottom right) fine-resolution mesh. Contours indicate the ground surface elevations (m relative to NAVD88).

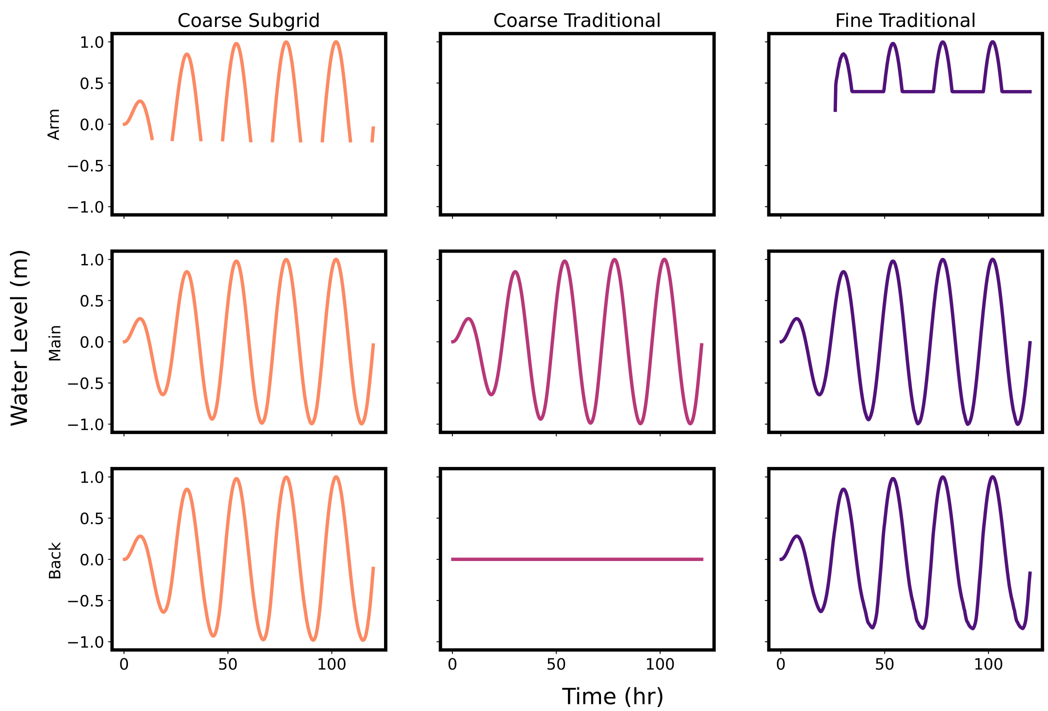

Figure 2.5: For the Buttermilk Bay test case, time series of water levels (m relative to mean sea level). Station locations are indicated in Figure 2.4. For the coarse traditional simulation, the Arm station is always dry, while the Back station is wet but disconnected from the tidal forcing.

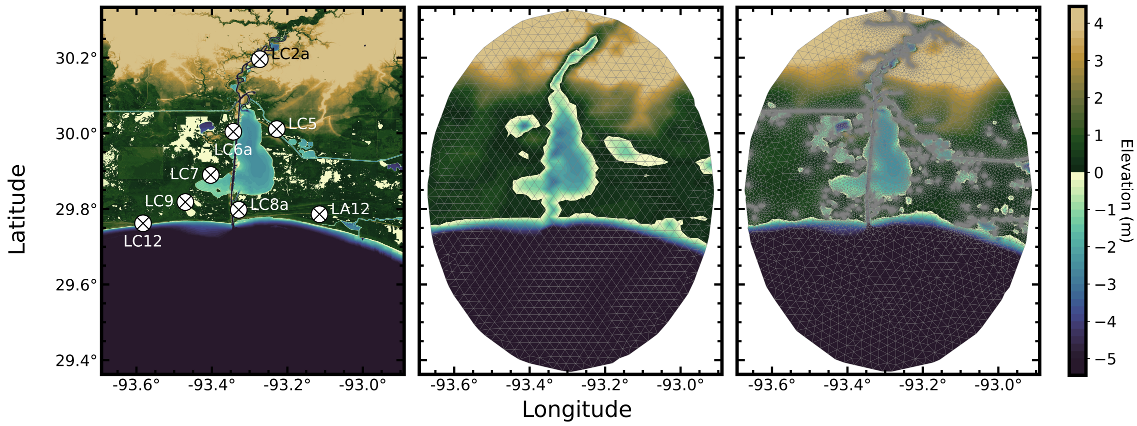

Figure 2.6: For the Calcasieu Lake test case, (left) 3-m DEM from USGS CoNED with gauge numbers and locations indicated by the crossed circles, (center) coarse-resolution mesh; and (right) fine-resolution mesh. Contours indicate the ground surface elevations (m relative to NAVD88).

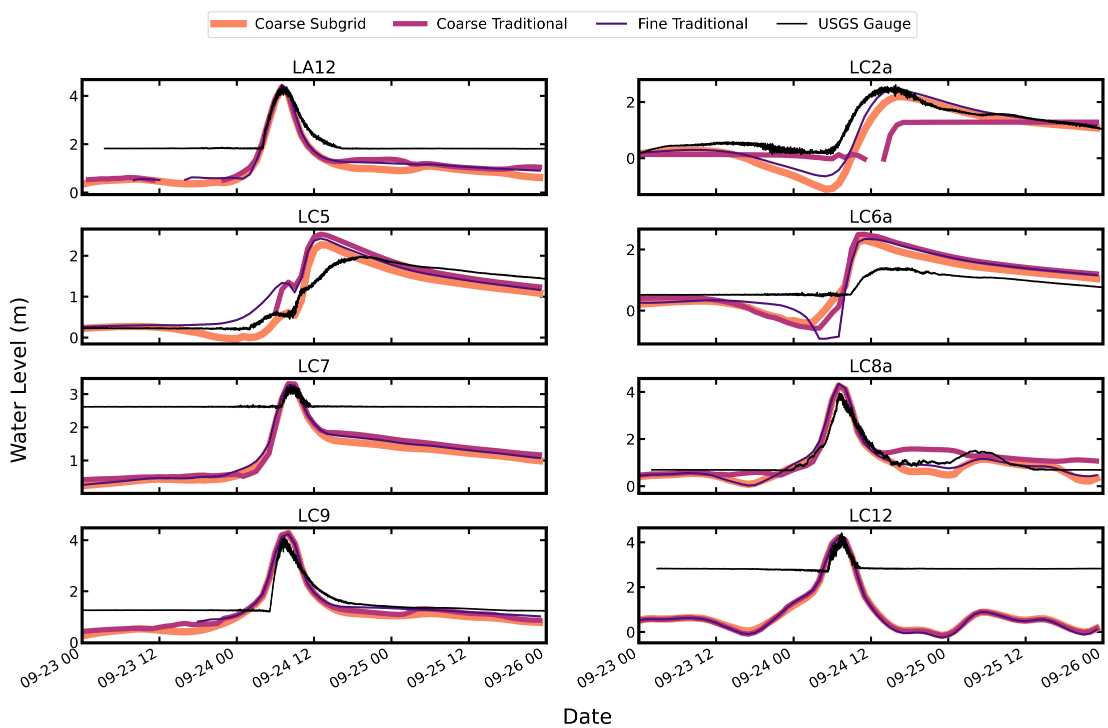

Figure 2.7: For the Calcasieu Lake test case, time series of water levels (m relative to mean sea level) at USGS gauges with locations shown in Figure 2.6.

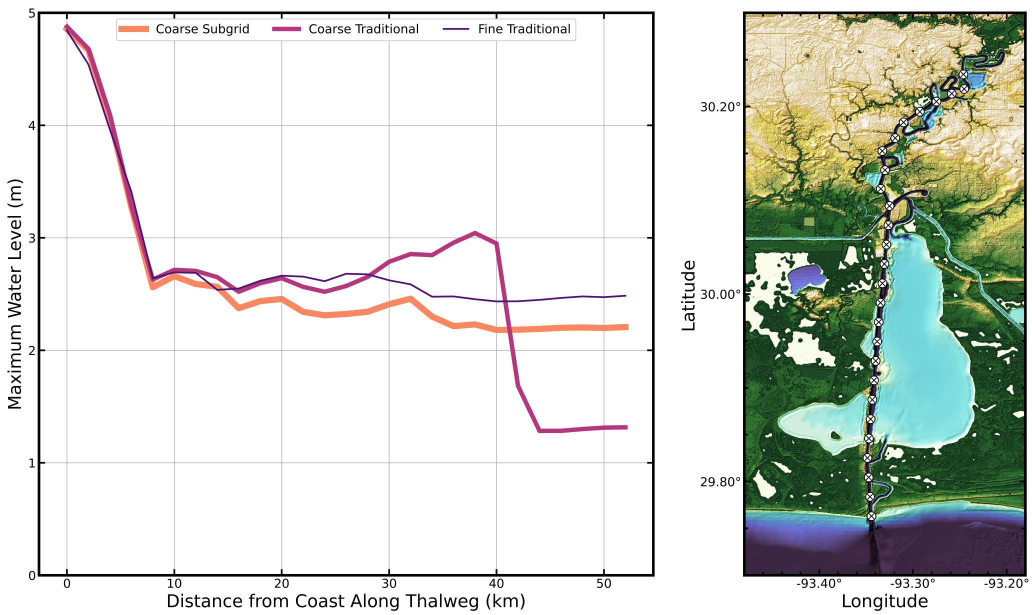

Figure 2.8: For the Calcasieu Lake test case, (left) maximum water levels (m relative to mean sea level) along the main channel thalweg, and (right) location of thalweg along the Calcasieu shipping channel and into Bayou Contraband.

CHAPTER

3

STORM SURGE PREDICTIONS FROM OCEAN- TO SUBGRID-SCALES

3.1 Preface

In this chapter, we expand the work done in Chapter 2 by incorporating higher level corrections into ADCIRC and implementing subgrid corrections on an ocean-scale numerical mesh. Herein, we derive the higher level corrections to bottom friction and advection, and test them on a synthetic winding channel. From there, we test subgrid ADCIRC on an ocean-scale domain with tidal and meteorological forcing from Matthew (2016). Results are compared to high-resolution simulations using traditional ADCIRC and the implications of these additions are discussed. This research has been published in *Natural Hazards* (Woodruff et al. 2023).

3.2 Introduction

Tropical cyclones and other coastal storms can cause storm surge, which is the rise in water levels above the normal astronomical tides (Harris 1963). Storm surge can damage infrastructure along the coast and far inland (Lin et al. 2010a; Tomiczek et al. 2014; Needham et al. 2015). The extent to which storm surge can propagate overland depends on topographic and bathymetric controls, including natural and built channels and barriers, as well as varying friction associated with land cover (Stark et al. 2015; Herdman et al. 2018). Storm surge and flooding can be predicted with computational models generally based on the numerical solution of 2D shallow water equations (Westerink et al. 2008; Lin et al. 2010b; Leijnse et al. 2021). The computational model can generate detailed maps of inundation levels and extents, which are used to support decision making for coastal communities (Xian et al. 2015; Ramirez et al. 2016; Rucker et al. 2021). Therefore, during a storm event, it is essential that model predictions are the best possible representation of flooding in coastal regions.

Storm surge models are applied typically along large stretches of coastline, to predict water levels, currents, and flooding extents in the full region affected by the storm. These models represent the coastal environment with numerical grids (or meshes) with varying spatial resolution – often finer near the coast, where the bathymetry/topography varies significantly and predictive accuracy is critical. The Sea, Land, and Overland Surges from Hurricanes (SLOSH) model (Jelesnianski et al. 1992) resolves length scales on the order of hundreds of meters and includes limited topographical complexity (Zhang et al. 2008). The ADvanced CIRCulation (ADCIRC) model (Luettich et al. 1992; Westerink et al. 2008) uses finite-element meshes with small length scales on the order of tens of meters to represent hydraulic and topographic information with high fidelity (Bunya et al. 2010; Hope et al. 2013). Typically, the inclusion of more geospatial information allows for more accurate flooding predictions, because the model can represent flows at very small scales (Bilskie and Hagen 2013; Kerr et al. 2013). However, such inclusion can lead to a model with a large number of grid cells that is too computationally expensive to obtain results in a timely manner. In addition, due to technologies like Light Detection and Ranging (LiDAR), there will be a gap between any model and the best-available representation of the coastal environment (e.g. Digital Elevation Models (DEMs) with spatial resolutions less than 1 m) (Danielson et al. 2018). Traditionally, trade-offs between spatial accuracy and computational efficiency must be considered when developing a model.

Subgrid corrections can bridge the gap between model and finer scales. The governing shallow water equations are averaged (Defina 2000) to introduce closure terms, which can include higher-resolution data to correct flow variables (flow accelerations, water levels, and current velocities) at the grid scale. Bates and Hervouet (1999) used subgrid corrections to improve predictions of the moving wet/dry boundary. Defina (2000) applied corrections to the convective accelerations to incorporate small-scale changes in flow due to ground irregularities. Volp et al. (2013) corrected bottom friction to account for variable bathymetry and roughness by assuming uniform flow direction and constant friction slope. These and other studies aim to attain highly accurate results on grids with length scales that are several orders of magnitude larger than those of the topographic datasets (Casulli and Stelling 2011).

Subgrid corrections have only recently been used for storm surge predictions. Sehili et al. (2014) incorporated winds and storm surge into a regional Unstructured Tidal, Residual, Intertidal Mudflat ($UnTRIM^2$) subgrid model of the North Sea. Daily meteorological predictions were used to force an operational subgrid model of the Elbe Estuary to predict water levels, velocities, and salinity transport. They achieved accurate predictions with the subgrid model while decreasing computational expense by a factor of 20. Wang et al. (2014) used the ocean-scale Semi-implicit Eulerian Lagrangian Finite Element (SELFE) hurricane storm surge model to provide water levels to $UnTRIM^2$ to predict water levels in New York City during Hurricane Sandy in 2012. High-resolution elevation data were used to predict street-scale water levels comparable to observations taken during the storm. Woodruff et al. (2021) added subgrid corrections into ADCIRC with real hurricane winds and storm surge forcing. Using a relatively small domain focused near the landfall location of Hurricane Rita in 2005, which caused extensive flooding in southwest Louisiana, the authors obtained accurate results while running ADCIRC with subgrid corrections on a coarse computational mesh that decreased run time by a factor of 32 when compared to a high resolution counterpart that had nearly 40 times more grid cells.

There are remaining challenges to the implementation of subgrid corrections for storm surge simulations on large domains. One challenge is to account for small-scale variations in bottom roughness and advection, which can significantly affect predicted water depths during a storm (Rego and Li 2010). Bottom friction is the primary contributor to storm surge attenuation as it flows overland (Resio and Westerink 2008), and small uncertainties in bottom friction can lead to large errors in predicted surge elevations (Akbar et al. 2017). The averaging of topographic and bathymetric features can lead to over-estimations of bottom friction (Defina 2000; Volp et al. 2013; Kennedy et al. 2019), which may inhibit propagation. Advection, due to storm surge interaction with the astronomical tides and shelf geometry, can affect predicted water levels by as much as 1 m in some locations (Thomas et al. 2019). However, when averaging to larger grid-scale areas, subgrid models may not represent small-scale variations in nonlinear advection in complex coastal environments (Defina 2000; Kennedy et al. 2019).

Another challenge is that most previous subgrid studies have focused on demonstrating the performance of the approach on small regional domains (Roig 1994; Bates and Hervouet 1999; Defina 2000; Wu et al. 2016; Kennedy et al. 2019) with areas less than 500 km^2. However, for storm surge applications, small domains can lead to significant under-predictions and undue boundary influences (Blain et al. 1994; Pringle et al. 2018). Larger domains are necessary to represent the storm's effects in open water, its interactions with the complex coastal environment, and the development of surge forerunners and shelf edge waves (Westerink et al. 1994). Because of their relatively high number of grid cells, large, ocean-scale hydrodynamic models require high-performance computing (HPC) systems and parallelized coding practices to reduce computing times (Tanaka et al. 2011; Roberts et al. 2021).

The ability of subgrid models to be parallelized and scaled to large domains, the availability of high-resolution data, and data processing limitations have been contributing factors to why previous research studies have not applied subgrid models to larger domains. In this study, we investigate the extension of subgrid models for storm surge on ocean-scale domains. *It is hypothesized that accurate predictions of storm surge at the smallest coastal scales can be obtained if: (1) higher-level subgrid corrections to bottom friction and advection are implemented into a widely used storm surge model, (2) the extensive datasets needed to describe subgrid information can be efficiently processed, and (3) storm surge predictions are corrected for flows in complex coastal environments.* We extend subgrid corrections in ADCIRC and comprehensively evaluate their performance compared to conventional methodologies. First, we introduce higher-level corrections to advection and bottom friction and demonstrate their benefits for controlled flow on a synthetic domain. Then, we extend to a domain of the Western North Atlantic Ocean, to simulate the storm surge generated by Matthew in 2016. This ocean-scale model will use hundreds of DEMs and landcover datasets to represent coastal regions from south Florida to the North Carolina Outer Banks. Finally, by comparing with observations from the storm, we demonstrate improvements in the subgrid model's ability to predict water levels over a range of spatial scales along the coast.

3.3 Methods

3.3.1 Extension of Subgrid Corrections in ADCIRC

Closures for Bottom Friction and Advection

ADCIRC is a coastal circulation model with applications in predictions of tides (Luettich et al. 1992; Westerink et al. 1992; Blain et al. 1998), density-driven circulation (Dresback et al. 2010; Blain et al. 2012; Cyriac et al. 2020), and storm surge (Westerink et al. 2008; Bunya et al. 2010; Dietrich et al. 2010; Weaver and Luettich 2010; Sebastian et al. 2014). ADCIRC uses the continuous-Galerkin, finite-element method to solve shallow water equations that consist of the depth-integrated mass equation reformulated into the Generalized Wave Continuity Equation (GWCE) and conservative momentum equations to predict water levels and current velocities at vertices in an unstructured mesh. Using volume-averaging techniques from Whitacker (1999), these equations were averaged to obtain the subgrid system for locally averaged flow variables. Such a system (see the full detailed derivation in Woodruff et al. 2021) consists of the averaged momentum equations:

$$\begin{aligned}\frac{\partial \langle UH\rangle_G}{\partial t} + g C_\zeta \langle H\rangle_G \frac{\partial \langle \zeta\rangle_W}{\partial x} = &-\frac{\partial C_{UU}\langle U\rangle\langle UH\rangle_G}{\partial x} - \frac{\partial C_{VU}\langle V\rangle\langle UH\rangle_G}{\partial y} \\ &- f\langle VH\rangle_G - g\langle H\rangle_G \frac{\partial P_A}{\partial x} + \phi\left\langle \frac{\tau_{sx}}{\rho_0}\right\rangle_W - \frac{C_{M,f}\langle \mathbf{U}\rangle\langle UH\rangle_G}{\langle H\rangle_W} \\ &+ \frac{\partial}{\partial x} E_h \frac{\partial \langle UH\rangle_G}{\partial x} + \frac{\partial}{\partial y} E_h \frac{\partial \langle UH\rangle_G}{\partial y},\end{aligned} \tag{3.1}$$

and:

$$\begin{aligned}\frac{\partial \langle VH\rangle_G}{\partial t} + g C_\zeta \langle H\rangle_G \frac{\partial \langle \zeta\rangle_W}{\partial y} = &-\frac{\partial C_{UV}\langle U\rangle\langle VH\rangle_G}{\partial x} - \frac{\partial C_{VV}\langle V\rangle\langle VH\rangle_G}{\partial y} \\ &- f\langle UH\rangle_G - g\langle H\rangle_G \frac{\partial P_A}{\partial y} + \phi\left\langle \frac{\tau_{sy}}{\rho_0}\right\rangle_W - \frac{C_{M,f}\langle \mathbf{U}\rangle\langle VH\rangle_G}{\langle H\rangle_W} \\ &+ \frac{\partial}{\partial x} E_h \frac{\partial \langle VH\rangle_G}{\partial x} + \frac{\partial}{\partial y} E_h \frac{\partial \langle VH\rangle_G}{\partial y},\end{aligned} \tag{3.2}$$

and the averaged GWCE:

$$
\begin{aligned}
\phi \frac{\partial^2 \langle\zeta\rangle_W}{\partial t^2} + \frac{\partial \phi}{\partial t}\frac{\partial \langle\zeta\rangle_W}{\partial t} + \tau_0 \phi \frac{\partial \langle\zeta\rangle_W}{\partial t} & \\
- \frac{\partial}{\partial x}\left(g\langle H\rangle_G \frac{\partial \langle\zeta\rangle_W}{\partial x}\right) - \frac{\partial}{\partial y}\left(g\langle H\rangle_G \frac{\partial \langle\zeta\rangle_W}{\partial y}\right) & \\
+ \frac{\partial \langle \tilde{J}_x\rangle_G}{\partial x} + \frac{\partial \langle \tilde{J}_y\rangle_G}{\partial y} - \langle UH\rangle_G \frac{\partial \tau_0}{\partial x} - \langle VH\rangle_G \frac{\partial \tau_0}{\partial y} = 0, &
\end{aligned}
\tag{3.3}
$$

where:

$$
\langle \tilde{J}_x\rangle_G = \text{RHS of (3.1)} + \tau_0 \langle UH\rangle_G \quad \text{and} \quad \langle \tilde{J}_y\rangle_G = \text{RHS of (3.2)} + \tau_0 \langle VH\rangle_G .
$$

In the above equations, $\langle s\rangle_G$ and $\langle s\rangle_W$ denotes the grid-average and wet-average, respectively, of the flow variables s ($s = H, UH, VH, \zeta$), and $\phi(\langle\zeta\rangle_W)$ is the wet area fraction, U and V are the water velocities in the x and y directions, $H = \zeta + h$ is the water depth, ζ is the water surface elevation, h is the bathymetric depth, f is the Coriolis parameter, g is the acceleration due to gravity, P_A is the atmospheric pressure, τ_{sx} and τ_{sy} are surface stresses, E_h is the lateral stress coefficient, τ_0 is a positive (spatially varying) parameter weighting the primitive continuity equation. The unknown solution variables to be computed are the surface elevation $\langle\zeta\rangle_W$ and the averaged x- and y- directed flux per unit width $\langle UH\rangle_G$, $\langle VH\rangle_G$. The grid-averaged water depth $\langle H\rangle_G$, wet-averaged water depth $\langle H\rangle_W$, and wet fraction ϕ are found using look-up tables (as elaborated below) for a given surface elevation $\langle\zeta\rangle_W$. The grid-scale velocity $\langle \boldsymbol{U}\rangle$ corresponds to the ratio $\langle \boldsymbol{U} H\rangle_G / \langle H\rangle_G$.

Several closure coefficients are introduced by the averaging and are indicated in red font in Equations 3.1, 3.2, and 3.3. These closure coefficients include: the convective acceleration coefficients C_{UU}, C_{VU}, C_{UV}, and C_{VV}; the friction coefficient $C_{M,f}$; and the water surface gradient coefficient C_ζ. To only apply corrections to the wetting and drying (denoted a 'Level 0' correction by Kennedy et al. 2019), these closure terms would be represented as the following:

$$
C_{UU}, C_{VU}, C_{UV}, C_{VV} = 1, \qquad C_{M,f} = \left\langle \frac{g n^2}{H^{1/3}} \right\rangle_W , \qquad C_\zeta = 1, \tag{3.4}
$$

where n is a Manning's roughness coefficient, assigned typically from land-use/landcover data sets.

In this work, we retain the water surface gradient correction as $C_\zeta = 1$, but we extend corrections for the advection and bottom friction coefficients (denoted as 'Level 1' corrections by Kennedy et al. 2019). For the bottom friction, Volp et al. (2013) created a weighted friction coefficient by applying the conveyance method, which was generalized for a two-dimensional setting by assuming uniform flow direction and friction slope at the subgrid level across a coarsened computational cell. This flow assumption results in the correction factor adjusting the conventional friction coefficient as follows (Kennedy et al. 2019):

$$C_{M,f} = \langle H \rangle_W R_v^2 \quad \text{where:} \quad R_v = \frac{\langle H \rangle_W}{\left\langle H^{3/2} C_f^{-1/2} \right\rangle_W}, \tag{3.5}$$

where the dimensionless friction coefficient C_f is calculated using Manning's equation:

$$C_f = \frac{g n^2}{H^{1/3}}. \tag{3.6}$$

In addition, the above mentioned canonical flow assumption at the subgrid level leads to the following advection correction coefficients (Defina 2000; Kennedy et al. 2019):

$$C_{UU} = C_{VU} = C_{UV} = C_{VV} = \frac{1}{\langle H \rangle_W} \left\langle \frac{H^2}{C_f} \right\rangle_W R_v^2. \tag{3.7}$$

Note that the equations (3.5) and (3.7) depend only on subgrid water depth and bottom roughness. Therefore, these coefficients can be pre-computed for a range of surface elevation values and stored as look-up tables. With the look-up table, evaluating these coefficients can be done efficiently and is independent of the subgrid resolution, as the procedure reduces to retrieving relevant entries in the look-up tables and requires only $O(1)$ operations.

Wetting and Drying

The conventional ADCIRC requires extensive logic to determine whether an element or vertex is considered 'wet' (included in the computations) or 'dry' (ignored) (Dietrich et al. 2004; Medeiros and Hagen 2013). This algorithm requires elements and vertices to be either fully wet or fully dry (i.e. to be wet, a triangular element must have nonzero water depths at all three of its vertices). Logic decisions are based on local water depths, water surface gradients, and bottom friction. ADCIRC's wet/dry algorithm has been applied for accurate predictions of flooding in many overland regions (e.g. Westerink et al. 2008; Hope et al. 2013; Thomas et al. 2019). However, this approach can lead to inaccuracies in the location of the wet/dry interface (Rucker et al. 2021), which can be represented only at the grid scale, and to numerical instabilities, partly due to the addition/subtraction of entire elements to the local conservation of mass and momentum.

In subgrid ADCIRC, an alternative wetting and drying algorithm was devised. The algorithm is based solely on the local wet area fraction (Woodruff et al. 2021); more specifically, mesh quantities (elements and vertices) are considered wet if they satisfy the one condition:

$$\phi > \phi_{\min} \tag{3.8}$$

where ϕ_{min} is a minimum wet area fraction set by the user. Typical values for $\phi_{\min}$ are in the range of 0.01 to 0.1. This setting allows for partially wet elements and vertices, thus smoothing the transition between the active and inactive parts of the domain.

In previous subgrid studies with a similar wetting criterion (Defina 2000; Casulli and Stelling 2011; Wang et al. 2014; Kennedy et al. 2019; Woodruff et al. 2021), one challenge has been the inability of the subgrid model to block flow between hydraulically disconnected regions. This is a result of assuming that flow variables are constant over their respective regions (Casulli 2019). Therefore, where a coarsened computational element spans a raised feature (e.g. dune, levee, or barrier island) between two water bodies, the coarsened subgrid model would not be able to "see" that these two water bodies are not hydraulically connected. One solution is to incorporate cell clones that predetermine connectivity between computational cells (Casulli 2019; Begmohammadi et al. 2021); however, this solution cannot be implemented readily in ADCIRC because of its vertex-based numerical scheme. Cell clones rely on pre-computing flow across cell edges, instead of along them as is done in ADCIRC. Thus, the addition of cell clones to ADCIRC is left for future work.

Instead, to prevent flows across small barriers, we allow the wetting criterion $\phi_{\min}$ to vary spatially. In most of the domain, it retains its typically small values, thus allowing computations in areas that are barely wet. However, in areas that include a dune crest, levee, or barrier island (Figure 3.1), the criterion can be set to higher values, thus preventing flows until the water level in these areas reaches an elevation sufficient for over-topping. These areas are identified manually before the simulation based on high points in the topography, and the higher $\phi_{\min}$ values are specified for select vertices and elements. In practice, this can be done relatively quickly by using a polygon shapefile that outlines the areas where variable ϕ_{min} is desired. From there, a simple Python script can be used to identify and store the elements and vertices contained in the polygon to be looked up later when ADCIRC is running. Herein, we use a value of $\phi_{\min} = 0.05$ in most of the domain, and a higher value of $\phi_{\min} = 0.99$ in select areas to represent flow barriers.

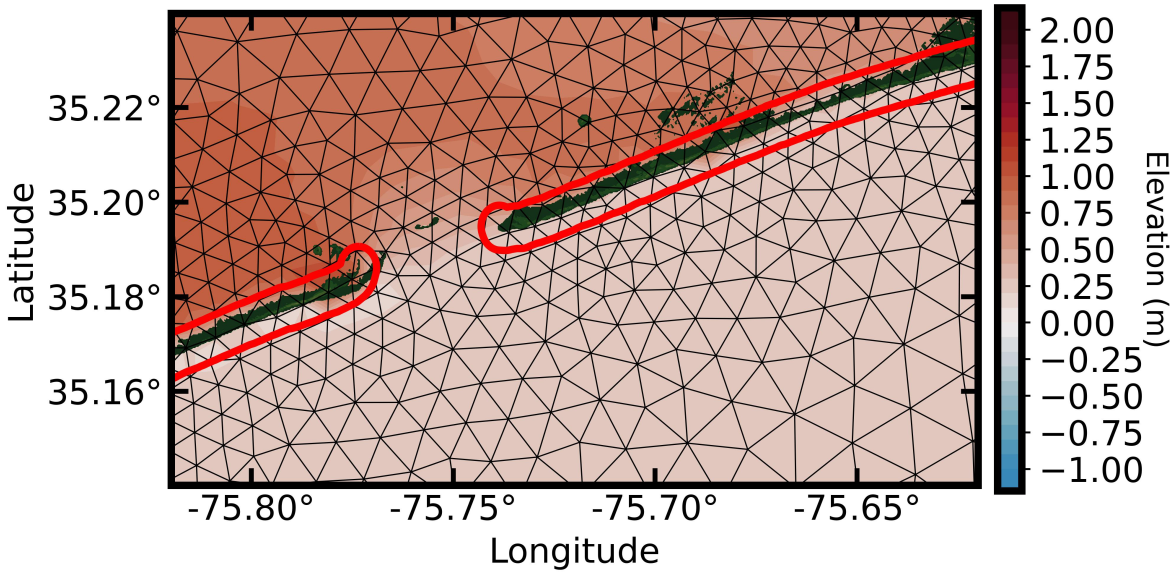

Figure 3.1: Water level contours (m relative to NAVD88) along the North Carolina Outer Banks during a simulation of Hurricane Matthew in 2016 on the SABv2 mesh. The barrier islands, where the wetting criterion $\phi_{\min} = 0.99$ has a higher value, are outlined in red.

Precomputing Corrections at the Grid Scale

ADCIRC uses the continuous-Galerkin, finite-element method to solve for water levels and current velocities at the vertices of triangular finite elements within an unstructured mesh. Accordingly, subgrid corrections are included by averaging quantities over areas corresponding to the elements and vertices (Woodruff et al. 2021). Element-averaged quantities include: $\langle H \rangle_G$, C_{UU}, C_{VU}, C_{UV}, C_{VV}, and ϕ, whereas vertex-averaged quantities include: $\langle H \rangle_G$, $\langle H \rangle_W$, $C_{M,f}$, and ϕ. These averaged quantities are computed from elevation and landcover raster datasets at a much higher resolution than the model grid. Note that $\langle H \rangle_G$ and ϕ must be represented as both vertex- and element-averaged quantities, due to how ADCIRC uses these variables to solve the governing shallow water equations.

For the element-averaged quantities, each element is split into three sub-elements (Figure 3.2), the raster cells within each sub-element area are located, and then the averaged quantities are computed for a given water surface elevation. For vertex-averaged quantities, the values from the surrounding element sub-areas are integrated and area-averaged to the vertex. Additionally, depending on the closure approximation used for each term in the governing-equation, either a grid-averaged $\langle \cdot \rangle_G$ or a wet-averaged representation $\langle \cdot \rangle_W$ was computed. For grid-averaged quantities, the entire averaging area is taken into account, whereas for wet-averaged quantities, only the subgrid areas that are underwater for a particular water surface elevation are included.

For computational efficiency, we use look-up tables in subgrid models while performing calculations. These tables are created by pre-computing subgrid quantities prior to the start of the simulation using information from bathymetric and landcover datasets. Upon initiation of the model, these arrays are read into memory to be accessed by the code for the duration of a simulation. In Woodruff et al. (2021), for element-based quantities, the size of the look-up table of each quantity was:

$$N_\zeta \times N_{\mathrm{VE}} \times N_{\mathrm{E}}, \tag{3.9}$$

where N_ζ is the number of water surface elevations used in the look-up table (e.g. 401 possible elevations between −20 m and +20 m), $N_{\mathrm{VE}} = 3$ is the number of vertices in a triangular element, and N_{E} is the number of elements in a computational mesh. For vertex-based quantities, the size of the look-up table is:

$$N_\zeta \times N_V, \tag{3.10}$$

where N_V is the number of vertices in a mesh.

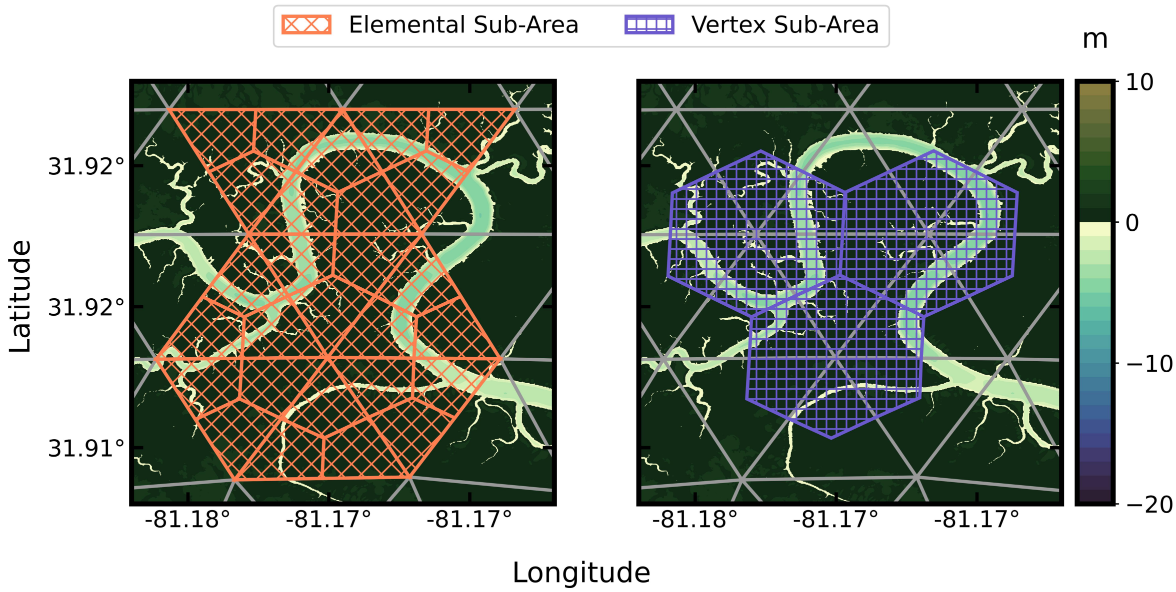

Figure 3.2: Examples of element- and vertex-averaged areas: (left) averaged quantities from sub-elements are combined for each triangular element, and (right) averaged quantities from sub-elements are combined for each vertex for a tidal creek near Savannah, GA.

Although look-up tables are used commonly to support subgrid corrections (Sehili et al. 2014; Wu et al. 2016; Kennedy et al. 2019; Woodruff et al. 2021; Begmohammadi et al. 2021), their expansion to ocean-scale domains created challenges for both their creation and their use at runtime. For the precomputations to create the look-up table in this study, the in-memory data storage of ocean-scale elevation and landcover data was managed carefully. The creation of the subgrid look-up tables was done entirely in Python (`https://github.com/ccht-ncsu/subgridADCIRCUtility`). Although datasets were split into manageable sizes, the task is challenging with CPUs because relatively small ($0.25° \times 0.25°$) DEM tiles with 1/9 arc-second resolution contain almost 66 million pixels. To speed up the calculations, within the coverage of the datasets, subgrid calculations were performed using a Graphical Processing Unit (GPU) accelerated Python library called CuPy (Ryosuke et al. 2017), by looping through each of the elevation and landcover datasets. Outside the coverage of the datasets (e.g. in open water), the look-up tables were constructed to mimic the conventional ADCIRC, by allowing elements and vertices to be either fully wet or fully dry. These look-up tables were then saved in a NetCDF-formatted file to be read by subgrid ADCIRC.

For the new mesh described below, the NetCDF-formatted file would have been larger than 10 GB, which has implications for file input and memory usage at runtime. For meshes with higher resolution, the look-up table sizes would become untenable. To reduce their sizes, the look-up tables were changed to consider a set of possible ϕ values, representing a state of dryness ($\phi = 0$) to fully wet ($\phi = 1$) at an evenly spaced increment. Averaged variables were then calculated by using the water surface elevations that would correspond to these wet area fractions. This reduced the size of the look-up tables to:

$$N_\phi \times N_{VE} \times N_E \tag{3.11}$$

where N_ϕ is the number of possible ϕ values, which is determined by the user. This is contrasted with the number of possible surface elevations, N_ζ, in the earlier version of the look-up tables in Equation 3.9. Whereas N_ζ was set typically to larger values (e.g. 401 possible water surface elevations), N_ϕ can be set to smaller values and still represent the variability in wet area fraction at each location. Herein, we use $N_\phi = 11$, which decreased the look-up table size. With this new scheme, a new look-up table of N_ϕ water surface elevations (ζ) corresponding to each ϕ increment was derived so that the ϕ for a particular element could be found depending on the water surface elevation in the element, the ζ look-up table, and the evenly spaced ϕ increments.

3.3.2 Storm Simulations with Subgrid Corrections

South Atlantic Bight

Subgrid ADCIRC is extended for storm surge predictions along the South Atlantic Bight (SAB) on the southeast U.S. coast. This region stretches from West Palm Beach, Florida (FL), to Cape Hatteras, North Carolina (NC), and includes more than 1000 km of coastline with a maximum shelf width of 200 km (Atkinson and Menzel 1985). The SAB's location and wide continental shelf make it particularly vulnerable to storm surge caused by tropical cyclones. Many studies have sought to understand the complex behavior of tides and circulation in this region (e.g. Redfield 1958; Blumberg and Mellor 1983; Chen et al. 1999). Tidal prediction along the SAB is particularly challenging due to the amplification that occurs as the tide propagates from the shelf break toward the coast, and due to the dramatic dissipation in energy as it interacts with the complex estuarian and riverine geometry present in the region (Blanton et al. 2004; Bacopoulos and Hagen 2017).

Due to the large size of the SAB, a substantial amount of elevation and landcover data are available to describe its coastline. A total of 830 elevation and landcover datasets (415 of each) were identified for use in this study (Figure 3.3). Elevation datasets were collected from the National Oceanic and Atmospheric Administration (NOAA) through the NOAA Digital Coast platform, and from The National Map (TNM) from the United States Geological Survey (USGS). The NOAA datasets are comprised of 1/9 arc-second and 1/3 arc-second digital elevation model (DEM) tiles of nearshore bathymetry and topography from the Continuously Updated Digital Elevation Model (CUDEM) produced by the NOAA National Centers for Environmental Information (CIRES 2014). These datasets were merged using QGIS with 1/3 arc-second DEM tiles from TNM for inland regions. Land-use and landcover are represented by Coastal Change Analysis Program (C-CAP) regional 1 arc-second resolution datasets. It should be noted that, although the model domain will extend beyond the SAB to also represent the western North Atlantic Ocean, Caribbean Sea, and Gulf of Mexico, subgrid corrections were applied only to the SAB region where flooding is expected and where detailed water level analysis is desired. This reduced the overall amount of data, but even so, these elevation and landcover datasets amounted to more than 197 GB of compressed raster-formatted data.

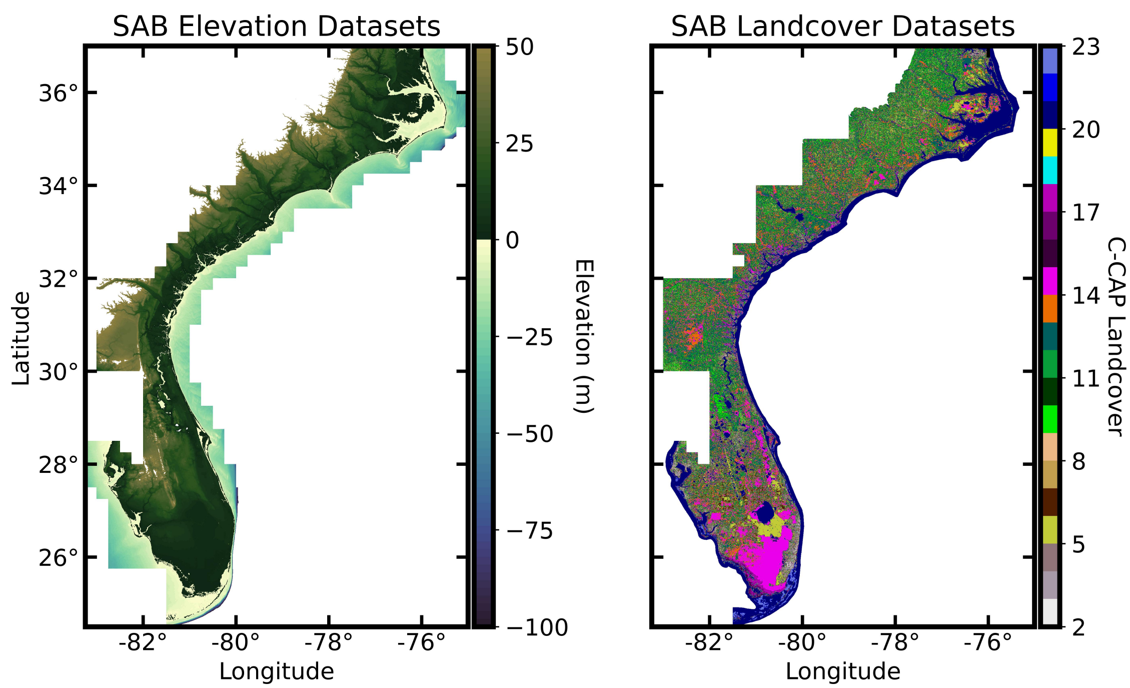

Figure 3.3: Merged rasters containing the 415 elevation and 415 landcover datasets for the SAB.

Mesh Development

Several meshes have been developed for ADCIRC-related studies in the SAB. Blanton and Luettich (2008) created a high-fidelity mesh to resolve important topographic and bathymetric features in North Carolina, and extended the mesh inland to the 15 m contour. A mesh with high resolution along the South Carolina coast was developed to resolve features with the size on the order of 100 m (URS Corporation 2009). Bender (2013, 2014, 2015) developed meshes to cover the region from South Florida to Georgia. These meshes were used to develop storm surge and flooding risk maps for their areas of coverage. Thomas et al. (2022) merged these regional meshes to describe the entire SAB with a mesh with about 5.5 million vertices and an average resolution in coastal areas of about 100 m. Apart from high-resolution mesh development for the use in floodplain mapping studies, coarser meshes like the Hurricane Storm Surge Operational Forecast System (HSSOFS) mesh are used during active storm events to forecast water levels along a coast. The HSSOFS mesh consists of about 1.8 million vertices, has an average resolution in coastal areas of about 500 m, with floodplain coverage from Southern Texas along the Gulf of Mexico to the North Carolina Outer Banks (Riverside Technology and AECOM 2015).

Recently, a mesh was developed as part of the South Atlantic Coastal Study (SACS; U.S. Army Corps of Engineers 2021), with the goal to understand vulnerability and flooding risks along the entire coastline. The SACS mesh has a minimum element edge length of about 20 m, and thus it resolves most of the hydraulically significant channels along the SAB. This mesh was validated for seven historical storms and then used in a study involving ensembles of thousands of synthetic storms (Owensby et al. 2020). However, the SACS mesh is expensive, with 12,288,247 elements and 6,179,416 vertices. Although this mesh is the state-of-the-art for storm surge predictions in the SAB, there is an opportunity for subgrid corrections to offer comparable accuracy on a coarser (more-efficient) mesh.

In this study, a new mesh was developed to test subgrid corrections for storm surge predictions in the SAB. This mesh provides coverage at ocean scales but with relatively coarse resolution along the SAB. It consists of $772,268$ elements and $392,358$ vertices, with a maximum element edge length of 50 km in open water and a minimum of 500 m in the nearshore. These maximum and minimum resolutions were chosen so the mesh would resolve large-scale channels and bathymetric features, but would alias many subgrid-scale features like small tidal channels, raised roadways, or intercoastal waterways at the grid level. Elements and vertices were aligned with a high-resolution coastline along the SAB (Contreras et al. 2021) and the US medium shoreline (National Oceanic and Atmospheric

Administration 2021c) in other regions. The mesh was bounded at the same inland locations as the SACS mesh, which has an inland boundary that aligns with either the 10 m or 20 m topographical contour (Owensby et al. 2020). Thalweg data were used to align elements and vertices with important hydraulic features like major rivers, inlets, and inland waterways. Steep bathymetric gradients in the offshore were resolved to ensure proper tidal propagation (Roberts et al. 2019b). The mesh was designed using Oceanmesh2D (Roberts et al. 2019a). It should be noted that this new SABv2 mesh has 15 times fewer computational grid cells than the existing SACS mesh. The highest disparities in resolution between the two meshes occur at inland locations along the complex coastline of the SAB (Figure 3.4).

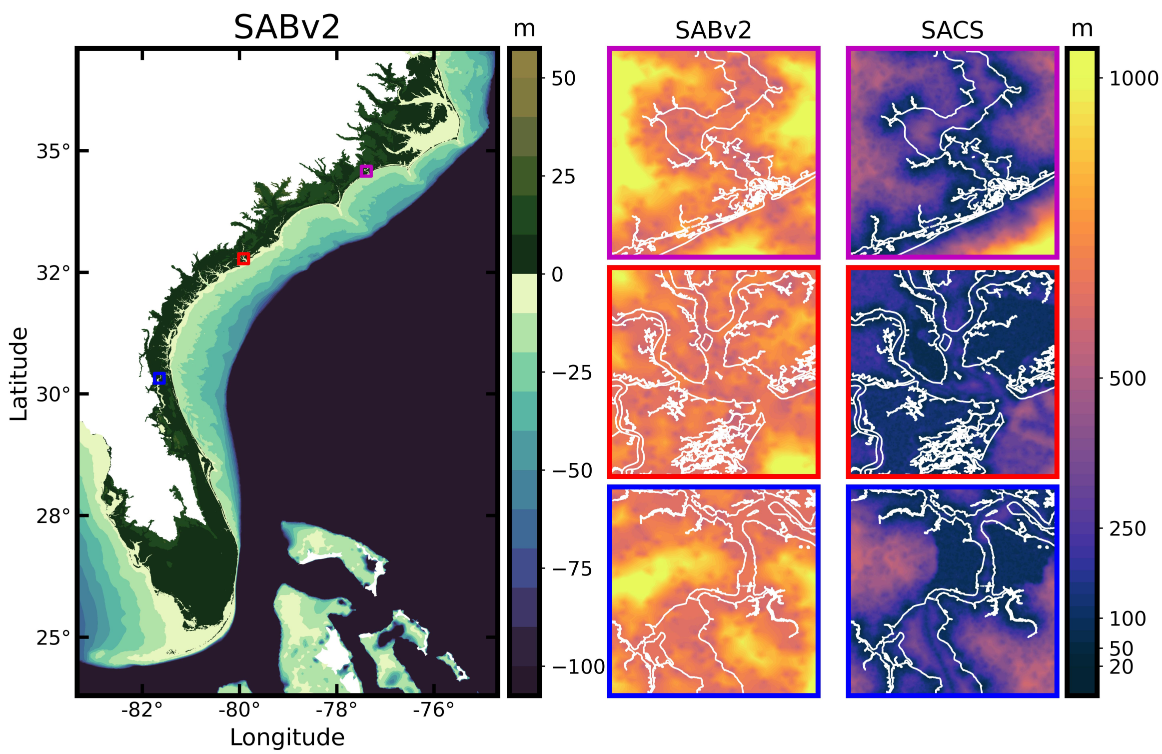

Figure 3.4: SAB portions of SABv2 and SACS meshes: (left) SABv2 mesh bathymetry along the SAB, with colored boxes (magenta, red, blue) to indicate locations for (right) comparison between mesh resolutions (m) for SABv2 and SACS, with the coastline shown as a white line. Note that both the SABv2 and SACS meshes extend beyond what is shown in this figure.

Hurricane Matthew in 2016

Hurricane Matthew evolved in the western North Atlantic Ocean over a 15-day period during 2016. The storm started as a strong tropical wave below 10° N latitude on September 23 off the western coast of Africa, and strengthened during the next few days until it became a tropical storm on September 28 just north of Barbados in the Lesser Antilles of the West Indies (Stewart 2017). Matthew rapidly intensified between September 30 and October 1 to a category-5 hurricane on the Saffir-Simpson scale with a peak intensity of 145 knots. During the eye wall replacement cycle, Matthew downgraded to a category 4 storm before making landfall near Les Anglais, Haiti, on October 4. Matthew weakened as it passed through the Bahamas, making landfall near West End on Grand Bahama Island on October 7 with category-3 status. Continuing northward, the storm ran shore-parallel between 30 and 50 nautical miles from Florida to North Carolina with category 2 and 1 intensity before starting its extra-tropical transition on October 9, moving eastward of Cape Hatteras, North Carolina, and dissipating in the North Atlantic (Stewart 2017).

The storm's effects on coastal water levels were observed by a total of 232 temporary gauges and long-term stations throughout the SAB. NOAA operates 22 stations, typically at the open coast. The USGS deployed 204 temporary gauges both at the coast and along inland waterways, and the Flood Inundation Mapping and Alert Network (FIMAN) operated 6 permanent stream gauges during the event (Figure 3.5). These observations describe how the surge varied with the storm. The maximum observed storm surge in the United States during Matthew was 2.35 m above normal tide at Fort Pulaski, Georgia (GA). Similarly high water levels occurred along the SAB from Fernandina Beach, Florida (FL), Charleston, South Carolina (SC), and Hatteras, North Carolina (NC), with high water levels measuring 2.12 m, 1.89 m, and 1.85 m at these respective locations. Inundation extended inland in this region, with many locations experiencing 0.6 to 1.5 m of surge above ground level. The Racy Point gauge along the St. Johns River, FL, recorded a maximum storm tide of 1.4 m above mean higher high water (MHHW) generated by the combined effect of storm surge and freshwater input from rainfall. These flood levels varied significantly in NC with the highest levels recorded on the sound side of the Outer Banks (Stewart 2017).

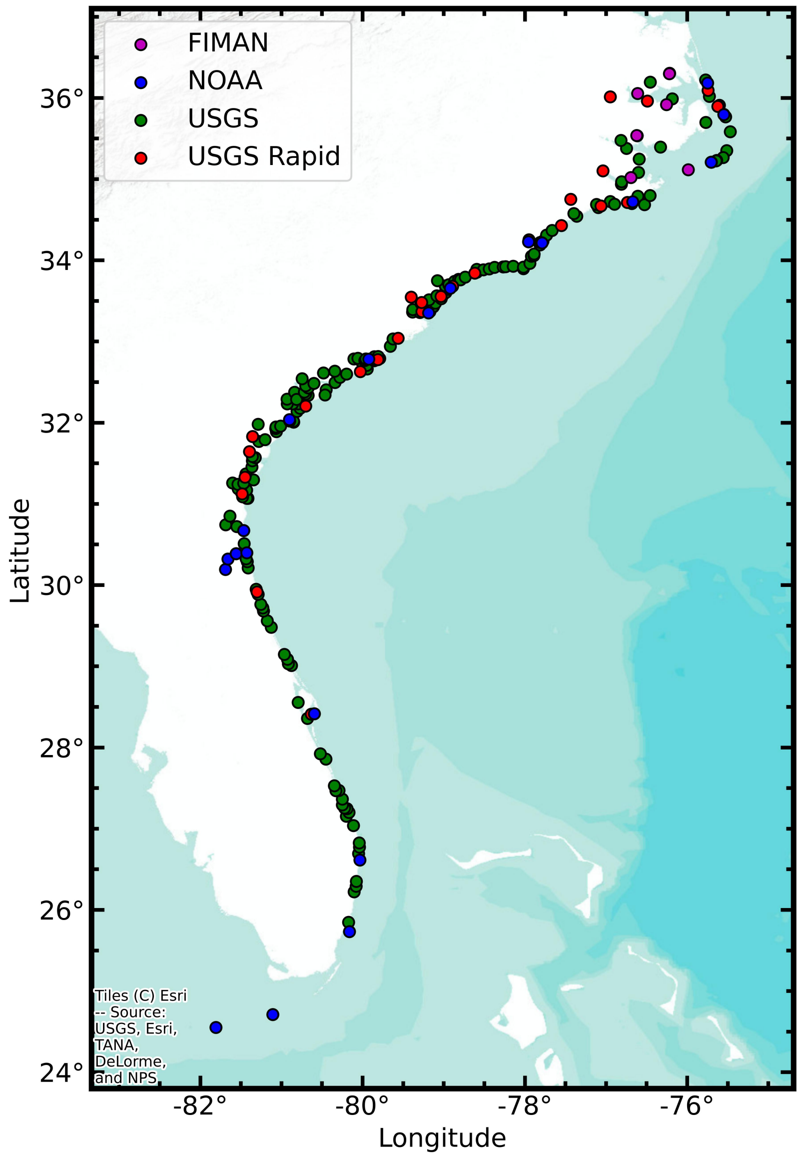

Figure 3.5: Gauge locations for Matthew (2016).

Simulations

Wind fields and surface pressures of Matthew (2016) from Oceanweather Inc. (OWI) were used in this study to force the simulations. The OWI fields are produced from weather station, buoy, aircraft, ship, and satellite stations and are considered highly accurate for use in hurricane storm surge hindcasts (Oceanweather Inc. 2018). For Matthew, forcing data are described on a lower-resolution, basin-scale grid covering 5° N to 47° N and 99° W to

55° W with a resolution of 0.25°, and a higher-resolution, regional inset grid from 15° N to 40° N and 82° W to 68° W with resolution 0.05°. The OWI data cover a time period between 0000 UTC 01 October 2016 to 0000 UTC 11 October 2016 with a 15 minute time interval (Thomas et al. 2019).

ADCIRC simulations were performed using conventional ADCIRC on both the high-resolution SACS and coarse SABv2 meshes, and subgrid ADCIRC on only the SABv2 mesh. Each simulation covered 25 days starting at 0000 UTC 16 September 2016 and ending at 0000 UTC 11 October 2016. All simulations used a 1 s timestep and started with a tidal ramp period of 5 days.

Time-dependent tidal elevation prescribed along the open ocean boundary is computed from the harmonic constituents of the TPXO tidal model (Egbert and Erofeeva 2002). For the SAB region of the SABv2 mesh, 2016 Coastal Change Analysis Program (C-CAP) regional landcover datasets were downloaded from NOAA Digital Coast (National Oceanic and Atmospheric Administration 2021b) were converted to Manning's *n* values (Owensby et al. 2020). Areas outside the SAB region had a constant Manning's *n* value of 0.02. These values were then interpolated to mesh vertices for the SABv2 coarse mesh and used for bottom friction and advection corrections in the subgrid simulation on the SABv2 mesh. Other parameters used in simulations on the SABv2 mesh included: spatially constant horizontal eddy viscosity of 50 m^2/s, variable primitive weighting coefficient (τ_0) between 0.005 and 0.03, sea surface height above geoid of 0.284 m to represent the pre-storm rise in water levels due to long-term atmospheric and oceanographic effects (Gill and Niiler 1973; Ferry and Reverdin 2000; Shin and Newman 2021), surface directional effective roughness length, and the surface canopy coefficient. The sea surface height above geoid value was determined from the mean monthly water level taken at water level gauges during the months when historical storms Hugo, Andrew, Fran, Frances, Matthew, Irma, and Florence occurred. The surface directional effective roughness and surface canopy coefficient in the SABv2 mesh were also derived from C-CAP landcover datasets, and were interpolated onto the mesh using the ADCIRCModules toolkit (Cobell 2020). The SABv2 simulations use an implicit formulation to compute the complete gravity wave term, which is available in ADCIRC v55 (Pringle et al. 2021).

Simulations on the SACS mesh used seven nodal attributes including: sea surface height above the geoid, primitive weighting coefficient, Manning's n, internal tide friction, surface directional effective roughness length, advection state, and surface canopy coefficient (Owensby et al. 2020). The same sea surface height above geoid used for the SABv2 simulations was used for the SACS simulation. NOAA's 2010 30 m C-CAP landcover dataset was used for Manning's n values and surface directional effective roughness length. Primitive weighting in the continuity equation was set to 0.03 for areas immediately surrounding the coast and 0.005 elsewhere. In areas where wooded canopy was present and likely to prevent the momentum transfer of wind to the water surface, the surface stresses were disabled via surface canopy. Internal tide friction was used in deep areas with steep bathymetric gradients (shelf breaks) to partially account for dissipation from the conversion of the barotropic tides to baroclinic tides not directly considered for in barotropic ADCIRC (Owensby et al. 2020). The SACS simulations were performed with the lumped explicit formulation (Tanaka et al. 2011).

3.3.3 Error Metrics

Water level data from the 218 gauge locations were used to evaluate the accuracy of the ADCIRC simulations (Figure 3.5). The accuracy of the simulations will be evaluated relative to observations of peak water levels and hydrographs. Peak-to-peak analysis of high water levels were compared to observations including root-mean-square-error (E_{RMS}):

$$E_{\mathrm{RMS}} = \sqrt{\frac{\sum_{i=1}^{N}\left(\zeta_i - \hat{\zeta}_i\right)^2}{N}}, \tag{3.12}$$

coefficient of determination (R^2):

$$R^2 = 1 - \frac{\sum_{i=1}^{N}\left(\zeta_i - \hat{\zeta}_i\right)^2}{\sum_{i=1}^{N}\left(\zeta_i - \bar{\zeta}\right)^2}, \tag{3.13}$$

and mean normalized bias (B_{MN}):

$$B_{\mathrm{MN}} = \frac{\frac{1}{N}\sum_{i=1}^{N}\left(\hat{\zeta}_i - \zeta_i\right)}{\frac{1}{N}\sum_{i=1}^{N}|\zeta_i|} \tag{3.14}$$

in which N is the number of observation stations, ζ and $\hat{\zeta}$ are the water surface elevations from the observations and predictions, respectively, and $\bar{\zeta}$ is the mean water surface elevation from the observations. Another metric was best-fit slope (m in $\hat{\zeta} = m\zeta$) from a linear regression fit between peak water levels from predictions and observations. Ideal values for E_{RMS} and B_{MN} are zero, indicating a perfect match in peak water level prediction. Ideal values for R^2 and m are unity, indicating an ideal 1-to-1 match between predictions and observations. In addition to performing error metrics on high water measurements at the gauge locations, E_{RMS} and B_{MN} statistics were calculated for the entire water level hydrographs (i.e. time series) at select locations along the coast for each simulation. For the time-series analysis, in Equations 3.12-3.14, N denotes the number of surface elevations in the time series, and ζ_i and $\hat{\zeta}_i$ are the model and observed surface elevation at Δt_s with Δt_s being the sampling time.

Computational efficiency was quantified by comparing the minimum run times of the subgrid and conventional ADCIRC simulations on the SABv2 and SACS meshes. The ocean-scale test case was run in parallel on 256 cores contained on 4 AMD Epyc "Milan" processors, each processor has 64 cores, and each node has 2 processors with 256 GB of memory and a clock speed of 2.45 GHz. The processors are connected via an Infiniband switch in the Anvil high-performance computing cluster at Purdue University. Each simulation was run in triplicate, and the minimum wall-clock time was used for timing comparisons.

3.4 Results

3.4.1 Level 1 Corrections in Synthetic Winding Channel

The Level 1 corrections in ADCIRC will first be evaluated via simulations of controlled flow in a synthetic compound channel. This synthetic domain (Figure 3.6, left) has dimensions of 120 m by 800 m, and it consists of a planar floodplain and a deeper, trapezoidal winding channel. The planar floodplain has a slope of 0.0001 m/m, and its ground surface elevations vary linearly from −1 m at the top of the domain to −1.08 m at the bottom of the domain. The channel is always 1 m deeper than its surrounding floodplain, with sloping banks that are each 5 m wide and 1 m deep and a bed that is 5 m wide. This test was designed to not include wetting and drying. Water levels are applied as boundary conditions at the top and bottom of the domain, and then discharges (through the channel and down the beach) will be compared with and without the Level 1 corrections.

Two finite-element meshes were created to represent this problem (Figure 3.6, middle and right). The coarse mesh has 224 vertices and 370 elements, the high-resolution mesh has 2052 vertices and 3734 elements, and thus the high-resolution mesh has approximately 9 times more grid cells. The coarse mesh has an average element side length of 24 m, while the high-resolution mesh has element side lengths ranging from 20 m on the floodplain to 5 m in the channel. The high-resolution mesh aligns elements and vertices with the channel banks and bed to fully resolve flow. The coarse mesh was created so that interpolation of elevation data to the vertices would alias the channel. The ground surface elevations for both meshes were interpolated using linear interpolation. Bottom friction is calculated using a constant Manning's coefficient of $n = 0.02$, and the horizontal eddy viscosity is set to $E_h = 2\ \mathrm{m}^2/\mathrm{s}$.

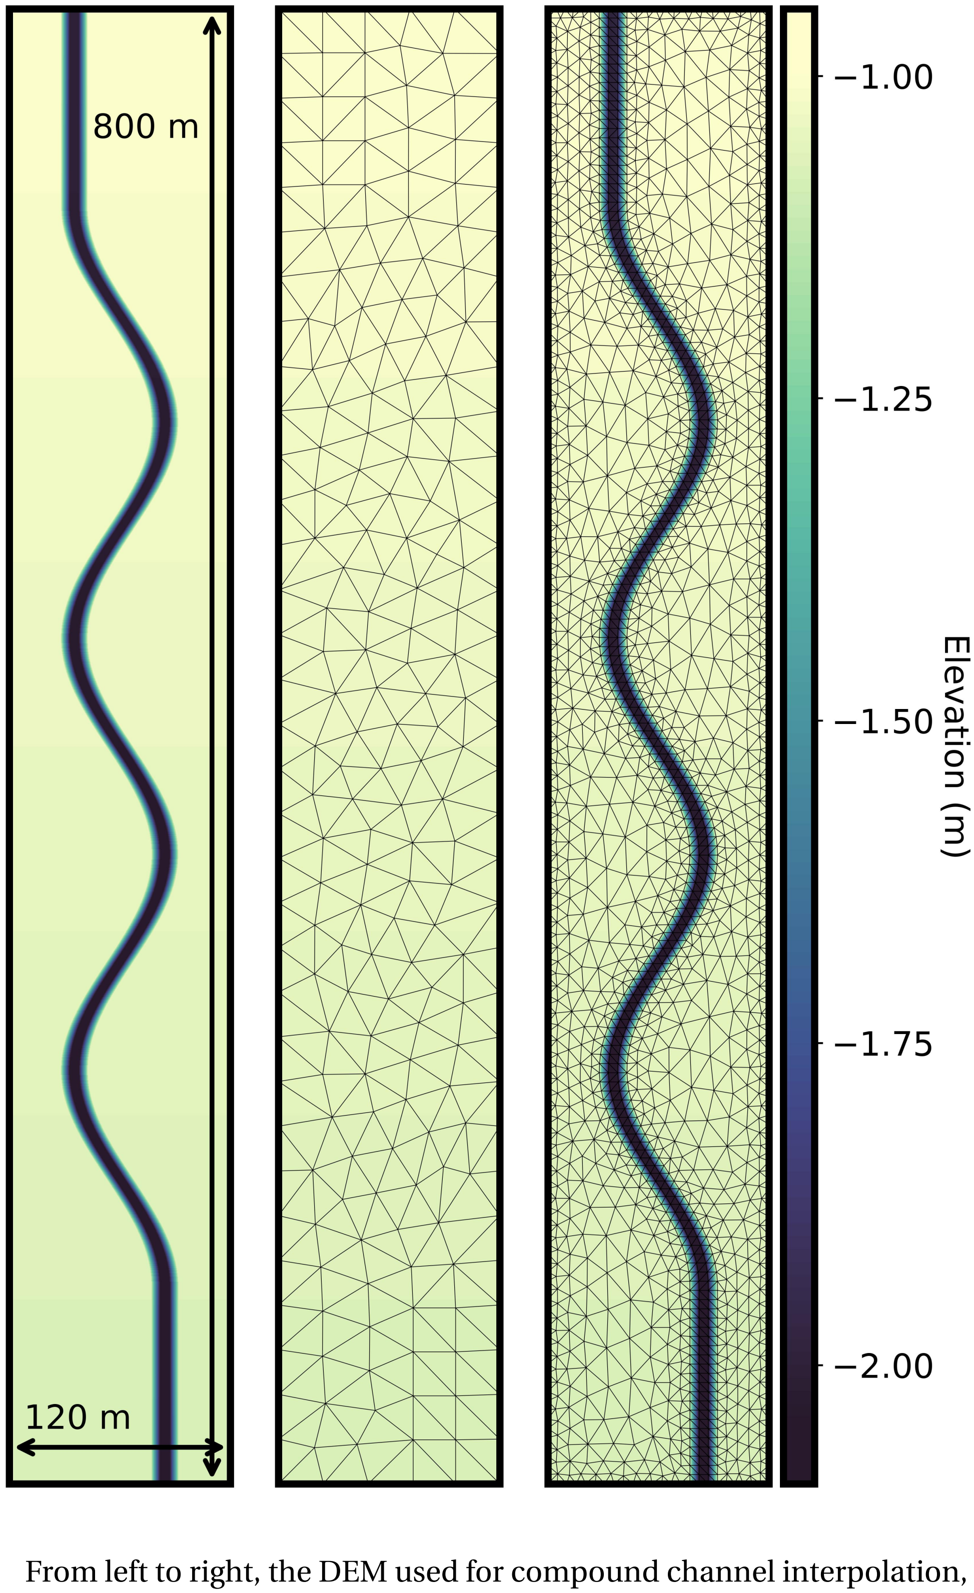

Figure 3.6: From left to right, the DEM used for compound channel interpolation, coarse mesh, high-resolution mesh.

The water surface elevations at the top and bottom boundaries are fixed to create a constant water surface gradient. The water surface gradient is the same as the bottom slope, but the flow depth above the floodplain is varied from 0.15 m to 1.5 m. The simulations are run for 10 days to achieve steady conditions. For each simulation, the total discharge is computed by integrating the discharge across the bottom boundary. The Level 1 corrections are evaluated via deviations from the 'truth' discharge from the high-resolution simulations:

$$Q_{\text{dev}} = \frac{Q_{\text{coarse}} - Q_{\text{truth}}}{Q_{\text{truth}}} \times 100\% \tag{3.15}$$

in which Q_{dev} is the discharge deviation (expressed as a percent difference), and Q is a discharge (units of m^3/s). Smaller deviations in discharge indicate a better representation of flow along the compound channel.

For this test case, five simulations are conducted: *High-Resolution Conventional*, in which ADCIRC is applied without any corrections on the high-resolution mesh, to provide a 'truth' solution for comparison; *Coarse Conventional*, in which ADCIRC is applied without any corrections on the coarse mesh; *Coarse Level 0*, in which subgrid ADCIRC is applied with Level 0 corrections on the coarse mesh; *Coarse Level 1 Only Advection*, where only the Level 1 correction to advection is added; and *Coarse Level 1*, in which subgrid ADCIRC is applied with Level 1 corrections on the coarse mesh. For small depths, where friction forces dominate, it is expected that the Level 0 corrections will overestimate bottom friction and thus limit artificially the total discharge. The Level 1 corrections to bottom friction on the other-hand should be a better representation of flow processes. However, as water depths increase, the improvements offered by corrections to bottom friction should diminish as frictional effects reduce, and will give way to the corrections to advection.

All simulations under-predict discharge across the bottom boundary of the domain, when compared to a reference simulation on a high-resolution mesh (Figure 3.7). The under-predictions are worse at small water depths; at a depth of 0.15 m, *Coarse Conventional* has $Q_{\text{dev}} = -25.6\%$, the *Coarse Level 0* has $Q_{\text{dev}} = -17.7\%$, and the *Coarse Level 1 Only Advection* has $Q_{\text{dev}} = -17.6\%$. These results indicate that the bottom friction is limiting the flows predicted on the coarse mesh. This effect is reduced for *Coarse Level 1*, which has smaller discharge deviations for all water depths, notably $Q_{\text{dev}} = -14.9\%$ at a water depth of 0.15 m. Therefore, at small water depths, Level 1 corrections offer improvement to simulated flow, and this improvement can mainly be attributed to enhancements in bottom friction representation.

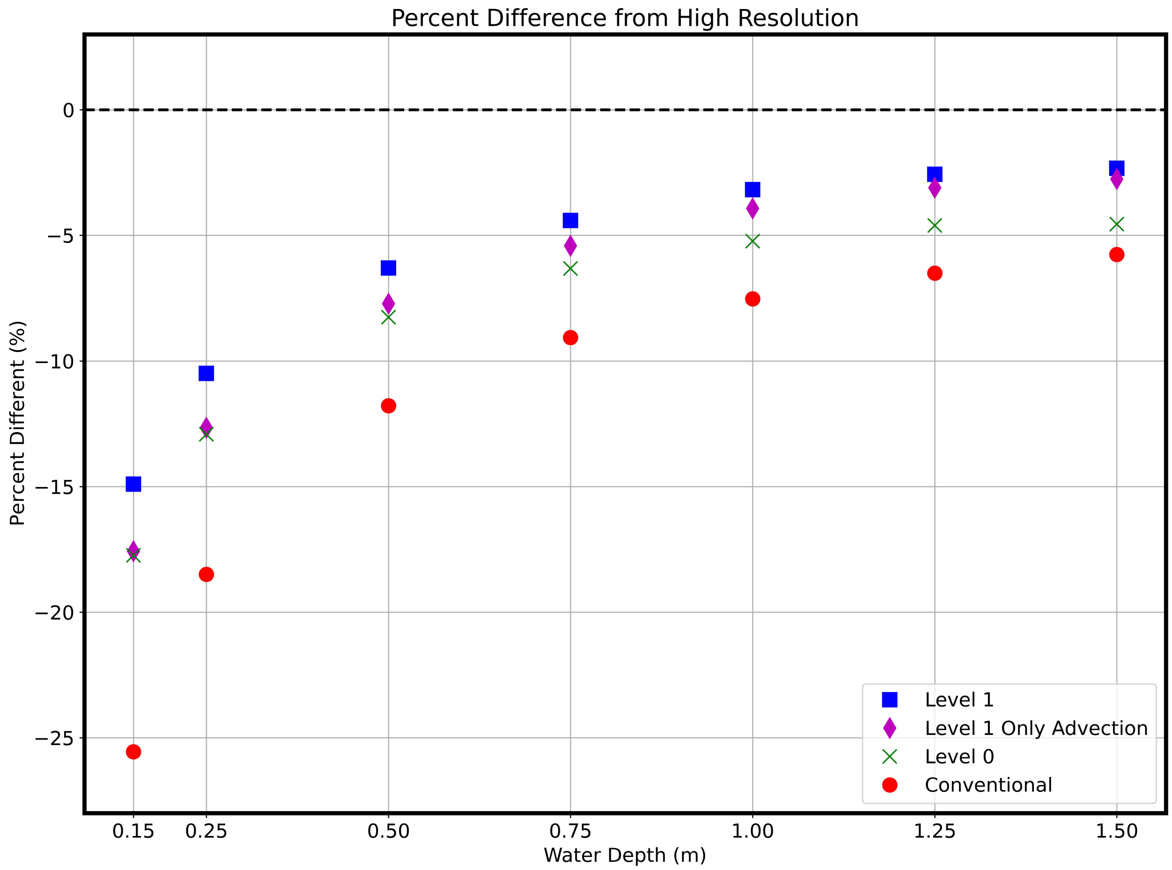

Figure 3.7: Discharge deviation of the *Coarse Conventional* (red circle), *Coarse Level 0* (green X), *Coarse Level 1 Only Advection* (magenta diamond) , and *Coarse Level 1* (blue square) from the high resolution simulation (dashed line).

As water depths increase, the contribution of the advection correction begins to develop. For instance, the percent difference in velocity magnitude at a water depth of 0.5 m above the floodplain shows the relative contribution of each Level 1 correction to the flow when compared to the Level 0 simulation (Figure 3.8). Here, a positive percent difference (colored red) indicates a higher velocity magnitude in the Level 1 simulations. From Figure 3.8, it can be seen that although the bottom friction correction dominates along the channel, the advection correction influences flow where there are abrupt changes in the channel direction.

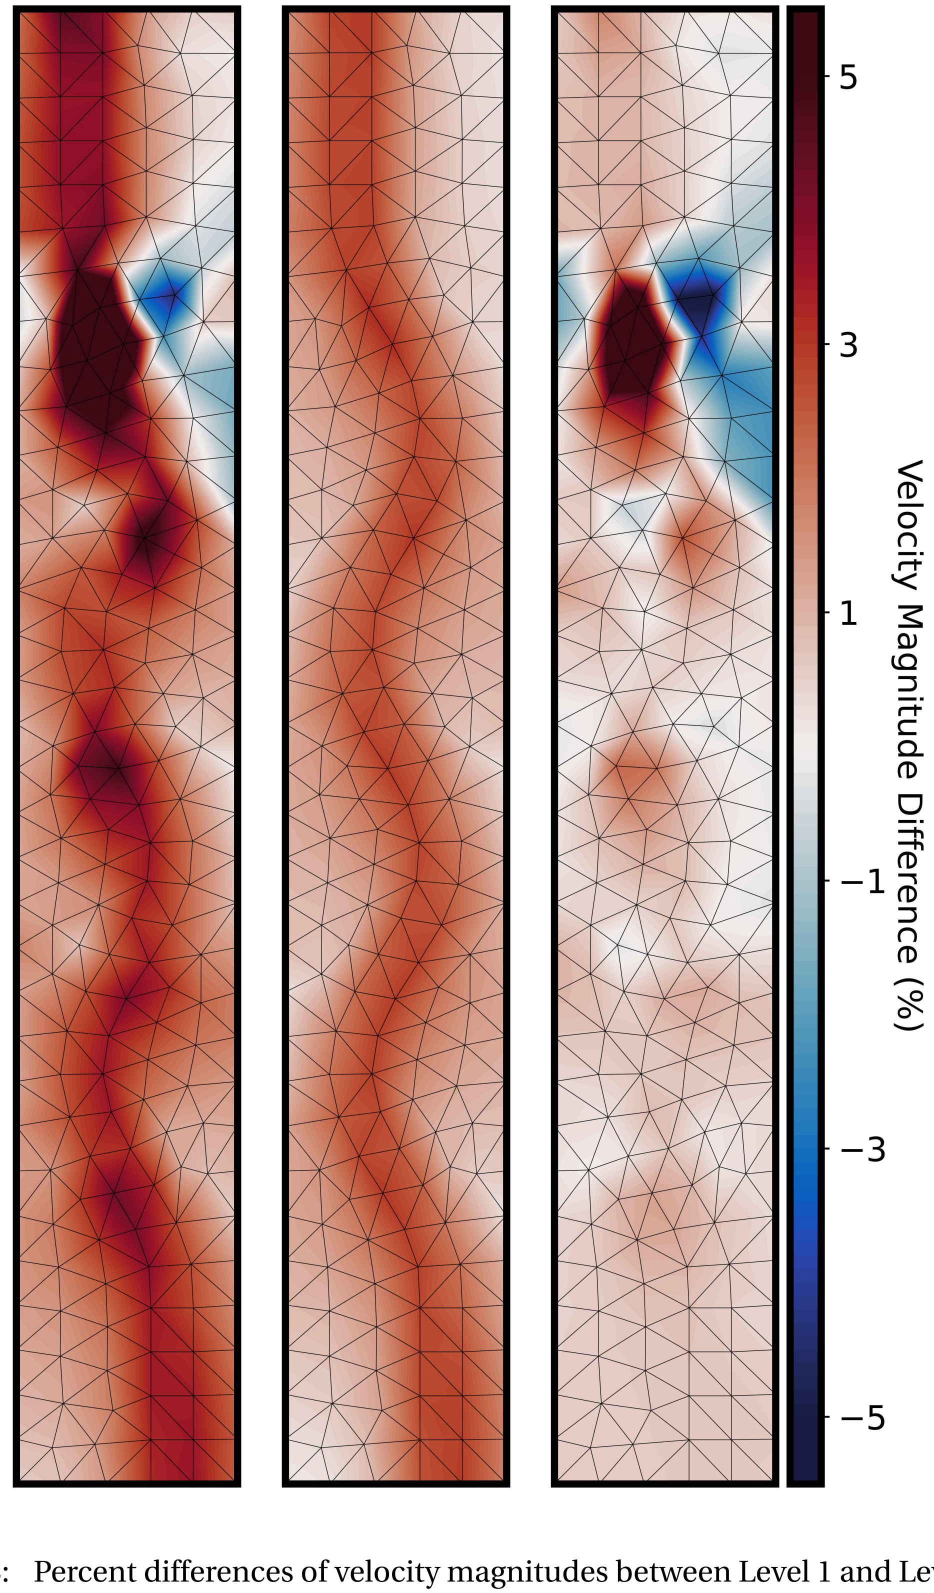

Figure 3.8: Percent differences of velocity magnitudes between Level 1 and Level 0 (left), Level 1 and Level 1 Only Advection (center), and Level 1 Only Advection and Level 0 (right) at a water level of 0.5 m above the floodplain.

As water depths increase further, the improvements to discharge predictions offered by Level 1 corrections begin to be dominated by the corrections to advection. At a water depth of 1.5 m, the *Coarse Level 1 Only Advection* and the *Coarse Level 1* simulations only differ by 0.43% with discharge deviations of −2.76% and −2.33% respectively. These results indicate that Level 1 corrections to advection can be non-negligible and can still offer improvements to simulated flow.

3.4.2 Storm Surge Predictions in the South Atlantic Bight

Ocean-scale simulations are performed on coarse SABv2 and high-resolution SACS meshes of the Western North Atlantic with an emphasis on the SAB. Three simulations are performed: *SACS Conventional*, *SABv2 Conventional*, and *SABv2 Subgrid* with Level 1 corrections. Water levels predicted with tidal and atmospheric forcing from Matthew in 2016 are compared against observations of water levels for a 25-day period surrounding the storm.

Matthew affected water levels along the SAB from FL through NC (Figure 3.9). Beginning October 7, the storm moved north from the Bahamas and started influencing water levels along the south FL coast. In this area, high water levels of 1.75 m are predicted along barrier islands and inland waterways and estuaries. High water levels of 3 m in the small canals that line the eastern FL coast are predicted in *SACS Conventional* and *SABv2 Subgrid*. Matthew continued north, steered closer to the coast throughout October 7, and began affecting the Georgia/South Carolina coast on October 8. The GA and SC coasts have riverine delta systems with streams, channels, and tributaries that experienced elevated water levels of about 3 m as the storm passed. Matthew tracked parallel to the coastline and caused flooding in NC on the morning of October 9 before moving offshore of the NC Outer Banks later that afternoon. Much like FL, the NC coast is characterized by barrier island and lagoon systems. As Matthew approached, storm surge was driven through tidal inlets and affected locations several kilometers from the open coast, with water levels of over 1.5 m in the Neuse River estuary. For a more detailed description of the storm's synoptic history, refer to Thomas et al. (2019).

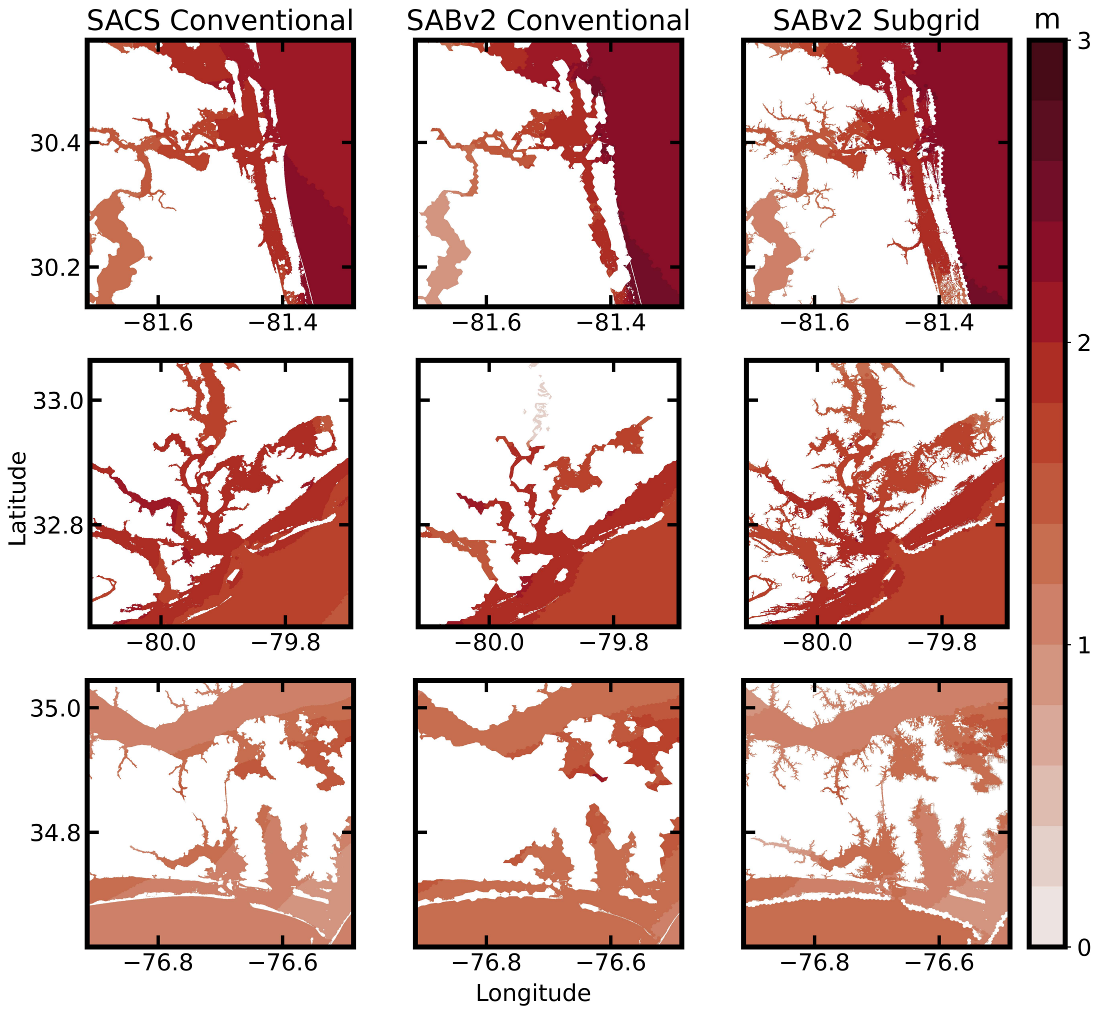

Figure 3.9: Maximum water levels in the *SACS Conventional* (left), *SABv2 Conventional* (middle), and *SABv2 Subgrid* (right) simulations along the SAB as Matthew moved up the coast. From the top row to the bottom row, the locations pictured are in the regions surrounding Jacksonville, FL, Charleston, SC, and Carteret County, NC.

Matthew's effects on coastal water levels were observed at 232 stations and gauges throughout the SAB. Of these 232 stations, 218 were used for water level analysis. To evaluate the performance of subgrid ADCIRC, we focus on 6 stations with representative results (Figure 3.10). Most of these 6 stations were located along inland waterways and far from the open coast, to highlight the subgrid ADCIRC's ability to represent surge propagation to inland areas on a coarsened computational mesh.

Storm surge gauge FLMAR03742 (last row in Figure 3.10) was located along a 40-m-wide canal near the Saint Lucie Inlet, FL. A peak water level of 0.84 m at 0441 GMT 07 October 2016 was observed at this gauge during the storm. In this area, the SACS mesh has 103-m resolution, and the SABv2 mesh has 775-m resolution. (Mesh resolutions at station and gauge locations were determined from element edge lengths connected to the nearest mesh vertex.) *SABv2 Conventional* was unable to resolve water levels at this station due to insufficient resolution in the surrounding area. *SABv2 Subgrid* predicted a peak water level of 1.26 m at 0400 GMT 7 October 2016, which is an over-prediction of 0.42 m. *SACS Conventional* briefly became wet at this location and recorded a peak water level of 1.17 m at 0400 GMT 7 October 2016, over-predicting the observations by 0.33 m. Although the *SACS Conventional* was able to simulate high water levels for a couple of hours surrounding the peak of the surge event, *SABv2 Subgrid* captures flow in this canal and thus had much better hydraulic connectivity in the area. Therefore, this location is an example of how subgrid corrections can improve connectivity without costly mesh refinements. In addition, *SABv2 Subgrid* had a comparable prediction error to the *SACS Conventional* at this location (Table 3.1).

NOAA station 8720357 (fifth row in Figure 3.10) is located on the St. Johns River near the I-295 Buckman Bridge in FL. At this station, the maximum observed water level during Matthew was 1.04 m at 0000 GMT 8 October 2016. The St. Johns River in this location is nearly 5 km wide and is resolved by the SABv2 and SACS meshes with local resolutions of 674 m and 106 m, respectively. *SACS Conventional* predicted a maximum elevation of 1.37 m at 2300 GMT 7 October 2016, and over-predicted water levels at this location by 0.33 m. *SABv2 Conventional* experienced a rapid flooding event during the storm that increased water levels to 0.83 m. This water then became trapped and could not drain back to the ocean due to insufficient mesh resolution at a 350 m wide channel constriction approximately 15 km downstream on the St. Johns River toward the coast. *SABv2 Subgrid* predicted a peak water level of 1.05 m at 0000 GMT 8 October 2016, over-predicting the peak observation by 0.01 m. Therefore, although local resolution in the SABv2 mesh was sufficient to predict water levels along a wider section of the St. John's River, subgrid corrections were necessary to resolve tidal propagation and storm surge from the open coast to the station. This was also indicated by reductions in both E_{RMS} and B_{MN} at this station in the *SABv2 Subgrid* simulation.

Gauge SCGEO14322 (fourth row in Figure 3.10) was placed along a 175 m-wide section of the Sampit River in Georgetown County, SC. The station was mounted at 0.88 m NAVD88 and thus could not observe the full tidal range leading up to the storm. This station recorded a peak water surface elevation of 1.82 m at 1557 GMT 8 October 2016. Mesh resolutions in this area vary from 666 m in the SABv2 mesh to 100 m in the SACS mesh. At this station, both *SACS Conventional* and *SABv2 Subgrid* were able to resolve storm-induced flows in this channel. However, due to insufficient local resolution in its mesh, *SABv2 Conventional* was unable to resolve flow. *SACS Conventional* predicted a maximum water level of 1.80 m at 1500 GMT 8 October 2016, within 0.02 m of the observed peak. *SABv2 Subgrid* predicted a peak of 1.74 m at 1600 GMT 8 October 2016, or within 0.08 m of the observed peak. Although both the *SACS Conventional* and *SABv2 Conventional* predicted water levels very close to the observed peak, the subgrid simulation produced better overall results with around a third of the E_{RMS} and smaller value B_{MN}. This gauge location was chosen because it offered a direct comparison of results produced by the high-resolution SACS mesh and the coarse SABv2 mesh with subgrid corrections. Without the added corrections, the SABv2 mesh was unable to resolve flow.

Gauge SCHOR14326 (third row in Figure 3.10) was located on a very small tidal creek near the border between NC and SC. SCHOR14326 recorded a maximum water surface elevation of 2.25 m at 1702 GMT 8 October 2016. This gauge was mounted at an elevation of 1.83 m NAVD88, and thus could not measure the tidal fluctuations leading up to the storm. The tidal creek has a width of less than 2 m. The resolution surrounding this location is roughly 663 m and 255 m in the SABv2 and SACS meshes, respectively. At this gauge, only *SABv2 Subgrid* can represent flow, neither *SACS Conventional* nor *SABv2 Conventional* have the necessary resolution to resolve this small-scale channel. *SABv2 Subgrid* predicted a peak water level of 2.31 m at 1700 GMT 8 October 2016, which was within 0.06 m of the observed peak. The small channel near the gauge would be difficult to resolve in an ocean-scale model, because the resolution required would be expensive with respect to simulation wall-clock time.

Gauge NCONS13068 (second row in Figure 3.10) was located near Jacksonville, NC, along the New River adjacent to Camp LeJeune Marine Base. Station NCONS13068 recorded a maximum water surface elevation during the storm of 0.92 m at 1854 GMT 8 October 2016. The SABv2 and SACS meshes have resolution in this area of 625 m and 56 m respectively. The New River is connected to the Atlantic Ocean through a channel that is 250 m wide, and the channel surrounding the gauge location is 100 m wide. Thus, storm surge and tidal propagation to the gauge cannot be simulated with the SABv2 mesh due to insufficient resolution. This is evident in the hydrograph, where *SABv2 Conventional* does not show any water level results for the location. The coarse subgrid and SACS simulations are able to resolve flow though the New River estuary. *SABv2 Subgrid* predicted a maximum water level at the station of 1.12 m at 2200 GMT 8 October 2016 and *SACS Conventional* had a result of 1.27 m at 2100 GMT 8 October 2016.

NOAA station 8658163 (top row in Figure 3.10) is located at the end of the Johnnie Mercers Fishing Pier in Wrightsville Beach, NC. At this location, the maximum observed water level reached 1.28 m at 1700 GMT 8 October 2016. This sensor is located along the open coast and was exposed to unobstructed tides and surge during Matthew. At this location, the SABv2 mesh has resolution of 760 m, whereas the SACS mesh has resolution of 130 m. At the open coast, this resolution should be sufficient to fully resolve flow. All of the simulations predicted water levels at this location, with *SABv2 Subgrid* predicting 1.54 m at 1700 GMT 8 October 2016, *SACS Conventional* predicting 1.53 m at 1600 GMT 8 October 2016, and *SABv2 Conventional* predicting 1.56 m at 1700 GMT 8 October 2016.

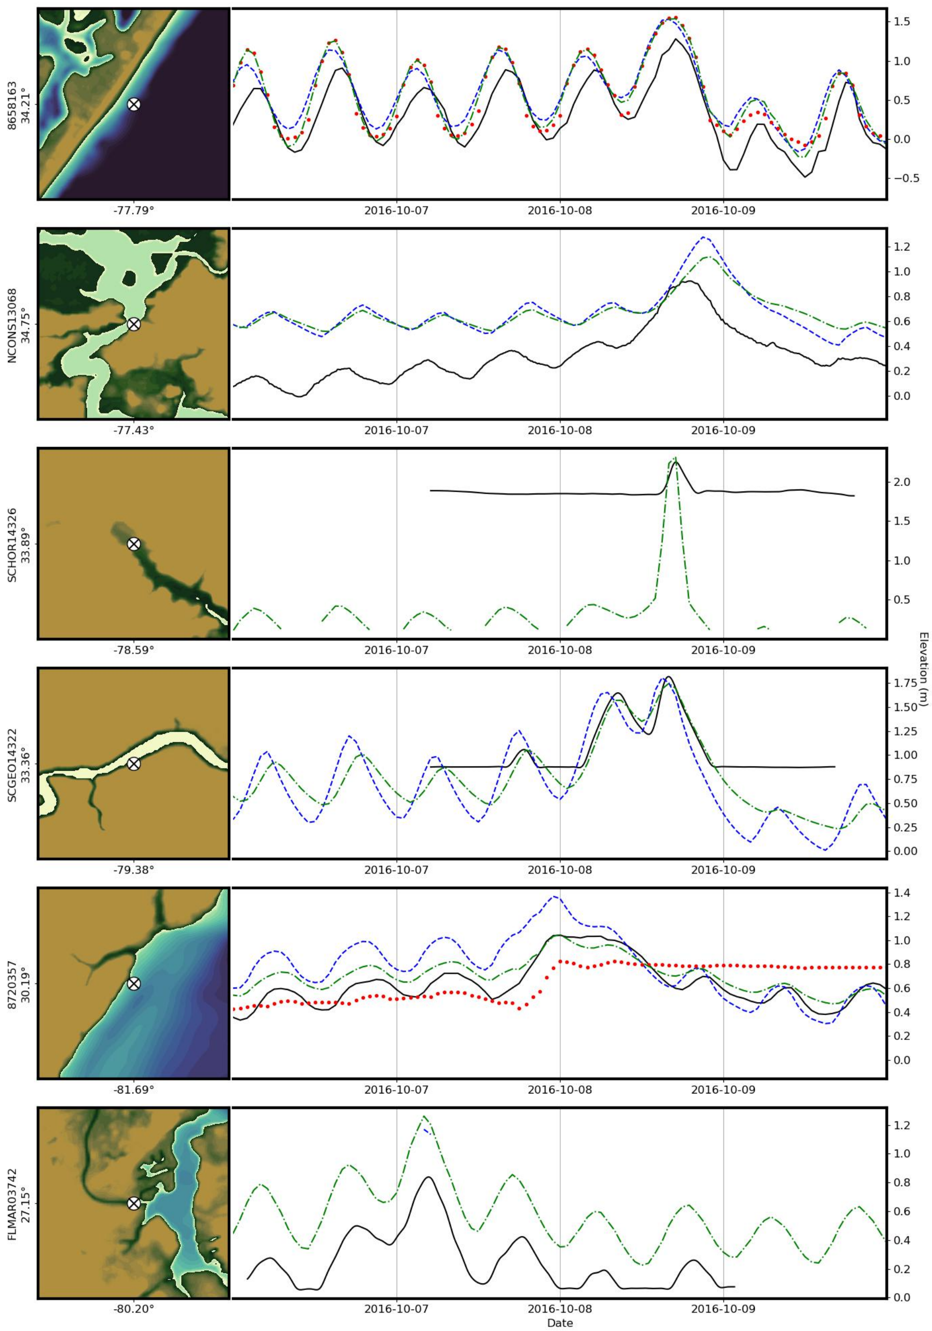

Figure 3.10: Station location (left) and hydrograph comparisons (right) between observation (black solid), coarse subgrid (green dash dot), coarse conventional (red dot), SACS conventional simulations (blue dash) relative to NAVD88 datum. These stations are in order from North to South, starting on the top row with station 8658163 in Wrightsville Beach, NC and ending on the last row with station FLMAR03742 in southeast Florida.

Table 3.1: Error statistics comparing coarse and fine simulation hydrographs to observations during Matthew (2016).

Station	SACS Conventional		SABv2 Conventional		SABv2 Subgrid	
	E_{RMS}(m)	Bias	E_{RMS}(m)	Bias	E_{RMS}(m)	Bias
8658163	0.33	0.5821	0.33	0.5745	0.33	0.5253
NCONS13068	0.35	0.9257	–	–	0.35	0.9402
SCHOR14326	–	–	–	–	0.88	-0.2559
SCGEO14322	0.20	0.0083	–	–	0.08	0.0016
8720357	0.24	0.4369	0.21	0.2679	0.21	0.3872
FLMAR03742l	0.33	0.4002	–	–	0.40	1.2167

Maximum water levels from simulations were compared to observations across the SAB using 1:1 plots (Figure 3.11), E_{RMS}, R^2, B_{MN}, number of dry stations, and best-fit slope (Table 3.2) to evaluate the performance of each simulation at predicting peak water levels. When considering all 218 high water levels, *SABv2 Subgrid* out-performed both *SABv2 Conventional* and *SACS Conventional* across almost every error metric. Hydraulic connectivity was improved significantly in *SABv2 Subgrid*, in which only two observation locations were not wetted during the storm simulation. There was a reduction in variance (R^2) off the 1:1 line in *SABv2 Subgrid*, indicating that the accuracy of the model results were improved. *SABv2 Conventional* and *SACS Conventional* produced negative R^2 values, indicating that their predictions had a higher variance from the 1:1 than the mean of the observations. This high variance can be attributed to the large number of dry observation stations during their simulations: 14 dry stations for *SACS Conventional*, and 25 dry stations for *SABv2 Conventional*. The E_{RMS} of the peak water levels from *SABv2 Subgrid* was 47% and 39% lower than that of *SABv2 Conventional* and *SACS Conventional*, showing that there was an increase in accuracy offered by the subgrid model. Finally, the best-fit slope of *SABv2 Subgrid* shows a near-perfect fit to the observational data, meaning that the simulation and observations are highly correlated in the subgrid model.

Separately, peak water level statistics of the stations that became or remained wet in all simulations were analyzed to examine how the models predicted water levels when dry stations are not factored into calculations. Discounting the dry stations reduced peak water level prediction error (E_{RMS}) of all simulations, most notably in *SABv2 Conventional* and *SACS Conventional*, which showed a 44% and 39% improvement, respectively. *SABv2 Subgrid* showed improvements to E_{RMS} and R^2, but the simulation best-fit slope moved

Table 3.2: Statistics from peak-to-peak analysis of subgrid and non-subgrid simulations when compared to observational data taken during Matthew (2016)

Simulation	E_{RMS} (m)	R^2	Best Fit Slope (m/m)	B_{MN}	Dry Stations
SACS Conventional All Stations	0.67	-0.13	0.99	0.0316	14
SABv2 Conventional All Stations	0.77	-0.52	0.93	-0.0126	25
SABv2 Subgrid All Stations	0.41	0.57	1.02	0.0535	2
SACS Conventional Wet Stations	0.41	0.56	1.1	0.1287	–
SABv2 Conventional Wet Stations	0.43	0.53	1.10	0.1442	–
SABv2 Subgrid Wet Stations	0.35	0.68	1.05	0.0826	–

slightly off 1.02. The only-wet analysis also significantly increased the B_{MN} of the *SABv2 Conventional* and *SACS Conventional*, which indicates that the model consistently over-predicts water levels.

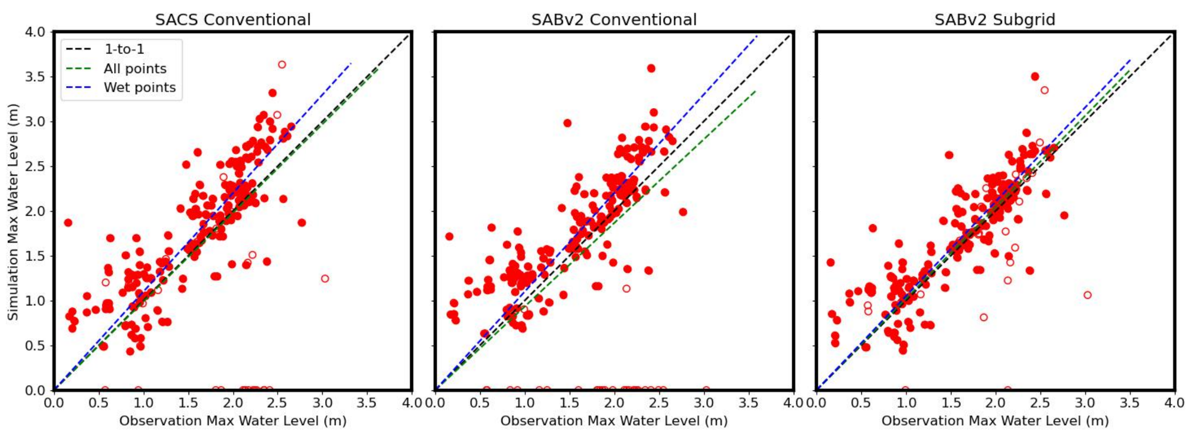

Figure 3.11: Peak water level comparison between observations and simulations for *SACS Conventional*, *SABv2 Conventional*, and *SABv2 Subgrid* in relation to NAVD88 datum. The green line in the plots represents the linear regression best fit for all of the stations, the blue line represents the linear regression best fit for only the wet stations in the simulation. The solid red dots represent observation stations that were wet in all simulations, and the empty dots represent observation comparison for the particular simulation.

Computing times of the three models used in this study show that the subgrid additions added computational expense to the model (Table 3.3). When compared to conventional

ADCIRC run on the same computational mesh, the subgrid model ran about 13% slower. However, the subgrid model ran over 13 times faster than the high-fidelity model and achieved comparable results. Thus, there was a significant efficiency gain by running subgrid ADCIRC on a coarsened mesh.

Table 3.3: Wall-clock times (sec) for ADCIRC simulations on 256 CPUs, and ratios of wall-clock times. The average time of three simulations was reported.

Wall-Clock Time (sec)	
SACS Conventional	82415
SABv2 Conventional	5433
SABv2 Subgrid	6083
Wall-Clock Time Ratio	
SABv2 Subgrid / SABv2 Conventional	1.13
SACS Conventional / SABv2 Subgrid	13.4

3.5 Discussion

The addition of Level 1 corrections and the expansion of subgrid ADCIRC to ocean-scale domains has resulted in considerable improvements to accuracy and efficiency of storm surge predictions when running on coarsened meshes. These improvements were demonstrated on a synthetic compound channel test case, and then implemented in a realistic storm surge simulation of Matthew in 2016 on an ocean-scale domain with emphasis on the South Atlantic Bight. In this section, we discuss the implications of the additional accuracy provided by the subgrid corrections, and the remaining challenges for future work.

Ocean-scale subgrid corrections improved hydraulic connectivity throughout the SAB region when compared to the conventional model on the given coarse mesh. As an example, for the water level predictions on the SABv2 mesh near the town of New Bern, NC (Figure 3.12), the subgrid corrections capture hydraulic connectivity in the waterways surrounding the town. *SABv2 Conventional* does not predict water in the Trent River (which flows into the Neuse River and is approximately 300 m wide in this area), whereas *SABv2 Subgrid* fully inundates this waterway and smaller connected channels through the domain. The maximum water levels are predicted to be about 1.0 m at the confluence of the Trent and Neuse Rivers. This additional accuracy is important; New Bern fared relatively well during Matthew, but there was flooding due to storm surge in its downtown near the confluence.

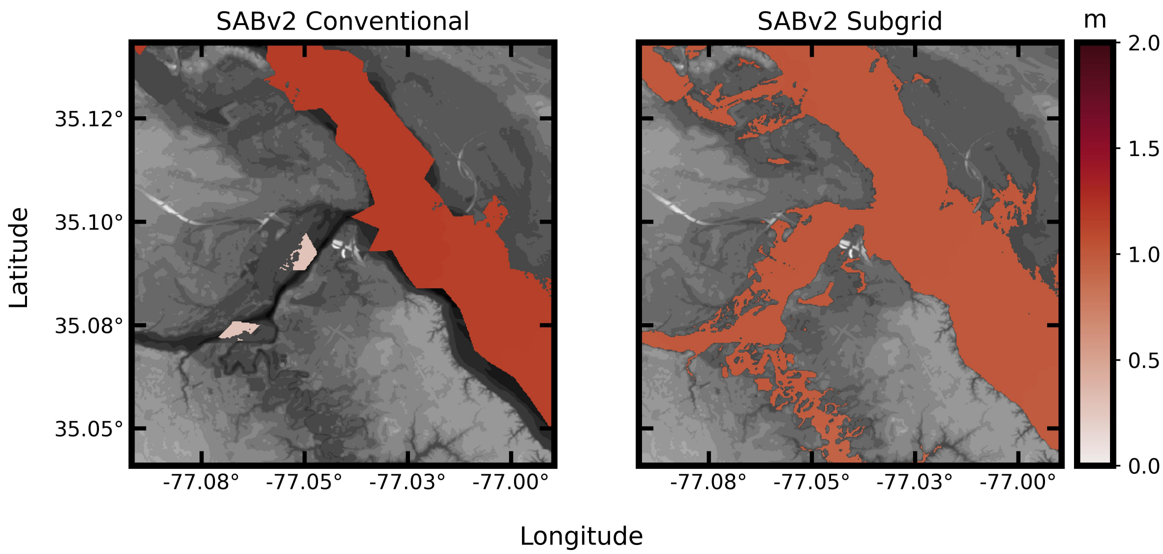

Figure 3.12: Maximum water levels (m NAVD88) during Matthew for the area surrounding New Bern, NC, along the Neuse River and adjacent waterways, as predicted by *SABv2 Conventional* (left) and *SABv2 Subgrid* (right).

Level 1 corrections improve subgrid ADCIRC by allowing the model to account for subgrid changes to bottom roughness and bathymetry to account for the effect of subgrid bathymetry and bottom roughness, which can affect bottom friction and advection terms. These additions can affect maximum storm surge height and inland penetration (Resio and Westerink 2008; Rego and Li 2010; Akbar et al. 2017; Thomas et al. 2019). The synthetic compound channel test case demonstrated that Level 1 corrections improved discharge through the channel when compared to conventional ADCIRC run on the same computational mesh, by better representing bottom friction in the model. This enhancement is important when modeling storm surge because bottom friction is often one of the main influences on inland surge propagation. The Level 1 correction to advection also influences flow by accounting for flow contractions and expansions. In the case of the New River Inlet, NC (Figure 3.13), Level 1 corrections to advection increased flow velocities though the narrow inlet. This is what we would expect. As flow enters the narrow channel, it accelerates though the contraction. However, when the bathymetry of the channel is discretized onto a coarse computational mesh, the contraction is made wider and shallower and therefore will not accelerate the flow as much. By accounting for these subgrid changes in bathymetry with the Level 1 corrections to advection, we can improve the prediction of flow acceleration and velocity. In addition, the Level 1 correction to bottom friction allows flow to pass through the channel faster by reducing the bottom friction coefficient in the deeper channel. These improvements to the representation of local advection and bottom friction are a contributor to the improvements in predictive accuracy offered by the subgrid ADCIRC.

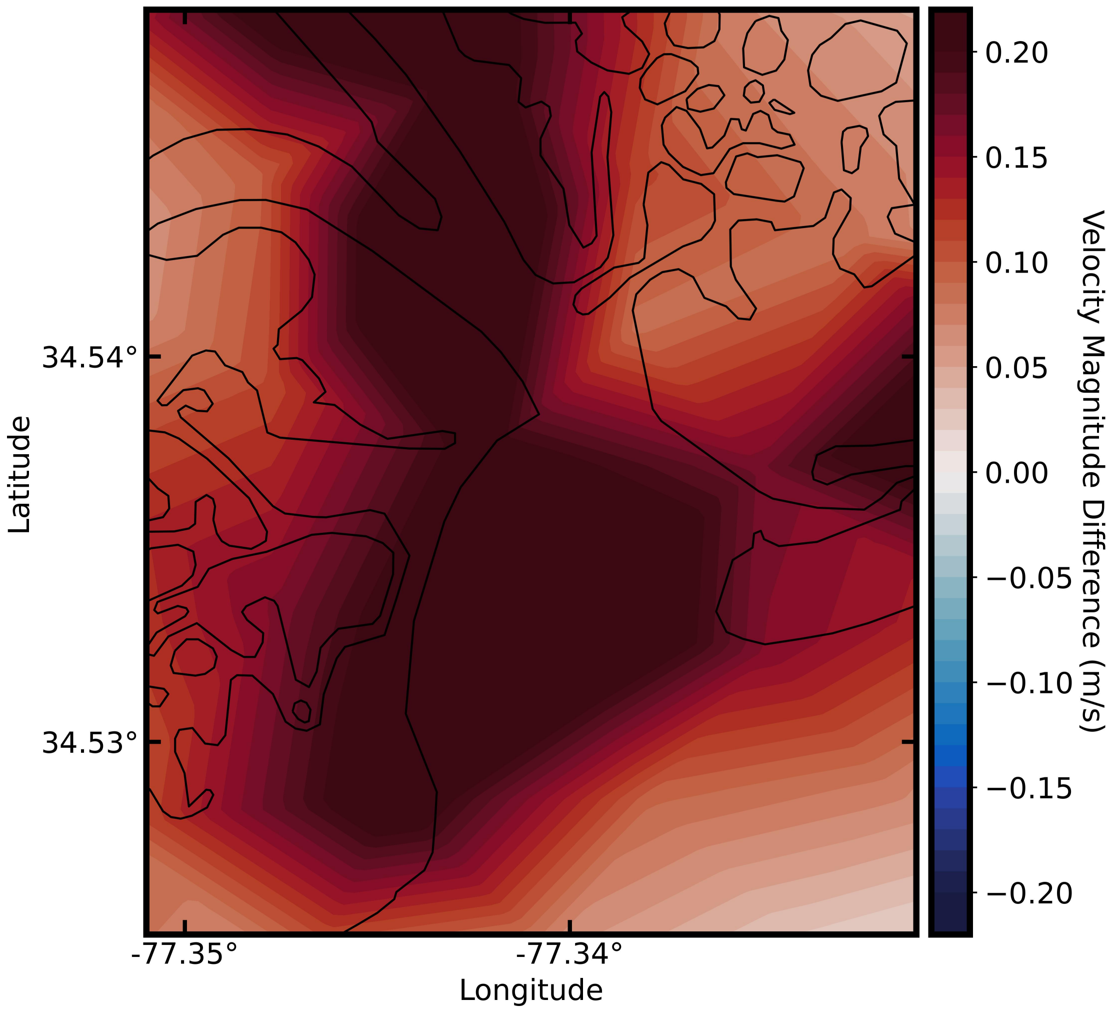

Figure 3.13: Difference in velocity magnitude between *SABv2 Subgrid* simulations run with Level 1 and Level 0 corrections at the New River Inlet, NC during the height of the Matthew (2016).

A spatially variable wetting criterion $\phi_{\min}$ was introduced into subgrid ADCIRC as a way to limit flow in regions that were hydraulically disconnected, but had elements that spanned a flow-blocking feature like a narrow barrier island. At a regional scale, this addition worked well to prevent flows from the ocean passing across barrier islands and into estuaries and other inland water ways. For example, Cape Canaveral, FL (Figure 3.14, left), is characterized by a narrow barrier island (with widths as small as 100 m) and a wide lagoon. If the mesh elements are too coarse to resolve the island at the model scale, it would not act as a barrier to flows. However, with the spatially variable ϕ_{min}, the larger storm tides are prevented from passing from the ocean into the lagoon, at least until the island is submerged. However, this criterion is not perfect, as it can limit artificially the flows at small inlets and channels near the flow-blocking feature. In some parts of the domain, like the area surrounding Port Canaveral, FL (Figure 3.14, right), the spatially variable ϕ_{min} also inadvertently limited flow in the small channel into the port. This resulted in a lack of water level predictions from *SABv2 Subgrid* at a NOAA station 8721604 located in the port. Thus, there is need for continued improvement on the wetting and drying algorithm in subgrid ADCIRC to improve hydraulic connectivity.

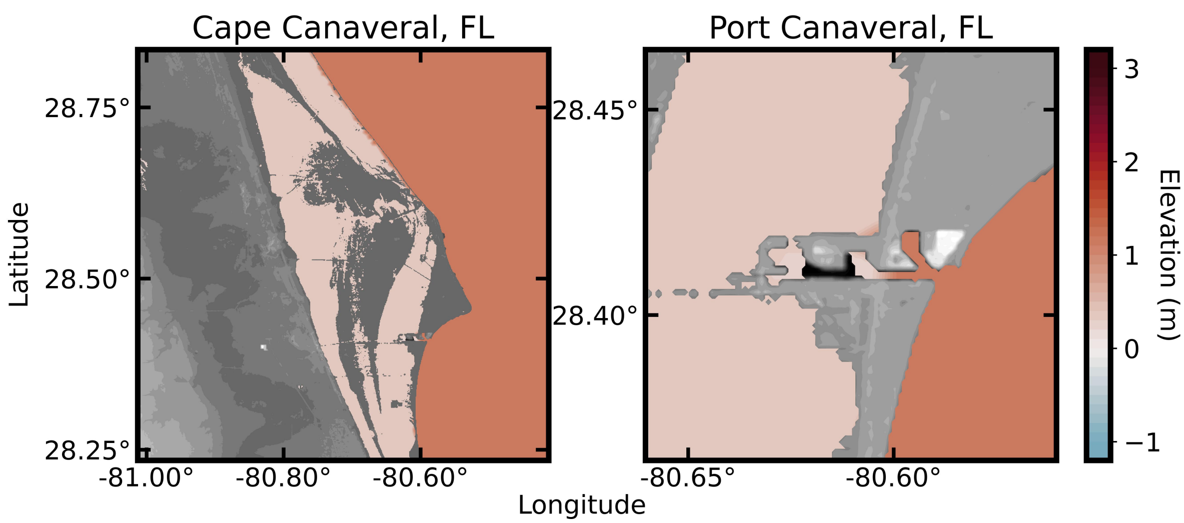

Figure 3.14: Examples of flow blocking due to spatially variable ϕ_{min} at 2 locations along the SAB: Cape Canaveral, FL (left) and Port Canaveral, FL (right).

3.6 Conclusion

In this study, higher order subgrid corrections to advection and bottom friction were implemented in an ocean-scale storm surge model for the South Atlantic Bight. It was found that subgrid corrections allowed for accurate predictions of water levels across this domain during a simulation of Matthew in 2016 on the coarsened SABv2 computational mesh by resolving subgrid flow processes using high-resolution bathymetric and topographic data.

The main contributions and findings of this study are:

1. *The subgrid approach performs better with bottom friction and advection corrections from higher-resolution datasets.* Corrections for these processes were added to the governing equations and look-up tables. These corrections improved discharge calculations in the synthetic winding channel test case by 11% when compared to the conventional model by better representing bottom friction in the model. Advection corrections in the ocean-scale storm surge model increased flow velocity magnitudes through inlets and winding channels, allowing for better predictions of flows to inland locations.

2. *Subgrid corrections can be extended to ocean-scale domains, but only with careful design of the pre-processor to compute the corrections, the reduced representation of corrections for a range of wet area fractions, and efficient handling of corrections in memory during the simulation.* The expansion of subgrid corrections to an ocean-scale domain required extensive elevation and landcover datasets with resolution down to 1/9 arc second to cover the entire SAB. This large amount of data required the use of HPC and advanced memory management to pre-compute look-up tables and process these tables in subgrid ADCIRC.

3. *Ocean-scale subgrid corrections improved the accuracy of a hindcast storm surge simulation of Matthew in 2016 while running on a coarsened computational mesh.* Water level predictions were validated with 218 permanent and temporary station and gauge locations from south Florida to the North Carolina Outer Banks. Peak water levels and hydrographs were analyzed and showed that subgrid corrections on the SABv2 mesh produced results with 39% less error than the SACS mesh and ran over 13 times faster.

The extension of subgrid ADCIRC to ocean-scale domains has the potential to improve accuracy and reduce computational cost for forecast and design studies of hurricane storm surge. These improvements are critical when evaluating flood risk. Coastal city planners and emergency managers need to understand which areas have the highest likelihood of flooding, often with resolution to the level of critical infrastructure. This information is necessary when designing flood control structures, creating and managing evacuation routes, and making decisions during the event. However, the necessary resolution is not feasible for a conventional model, especially when trying to predict over a large region during an active storm event. Thus, subgrid corrections are a viable option for providing this information at a fraction of the computational cost of a high-resolution conventional model.

Future work may include the addition of cell clones to properly resolve flow blocking features and hydraulically disconnected elements, and changing the wetting and drying threshold to rely on grid-averaged water depth $\langle H \rangle_G$ instead of the wet area fraction ϕ. Additionally, subgrid corrections in ADCIRC could play an important role in incorporating compound flooding from rainfall into the model, because it accounts for the total wet area within a computational cell.

3.7 Statements & Declarations

3.7.1 Funding

This study was supported by the National Science Foundation grants ICER 1664037 and 1664040. This work used Anvil at Purdue University through allocation MCA08X013 from the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) program, which is supported by National Science Foundation grants 2138259, 2138286, 2138307, 2137603, and 2138296. The support of the National Science Foundation is greatly appreciated by all contributors to this project.

3.7.2 Competing Interests

The authors have no financial or non-financial competing interests to disclose.

3.7.3 Author Contributions

All authors contributed to the study conception and design. Material preparation, data collection and analysis were performed by Johnathan Woodruff, Casey Dietrich and Damrongsak Wirasaet. The first draft of the manuscript was written by Johnathan Woodruff and all authors collaborated on versions of the manuscript. All authors read and approved the final manuscript.

3.7.4 Data Availability

After this manuscript is accepted for publication, the final data and source codes will be archived in a public repository at DesignSafe-CI.

CHAPTER

4

RESOLUTION SENSITIVITIES FOR SUBGRID MODELING OF COASTAL FLOODING

4.1 Preface

In this chapter, we explore how mesh resolution affects subgrid model predictions of storm surge and coastal flooding from Matthew (2016) and Florence (2018). Five ocean-scale meshes with varying resolution were created with an emphasis on the South Atlantic Bight. Subgrid simulation results were then compared to results produced using conventional methodologies without subgrid corrections. The work done in this chapter will be submitted to a scientific journal like Coastal Engineering.

4.2 Introduction

Coastal regions, including low-lying cities near mean sea level (MSL), are vulnerable to flooding caused by coastal storms. The most destructive flooding is caused by tropical cyclones and their associated storm surge, which is the storm-induced rise in water levels above the normal astronomical tides. These elevated water levels can be quite large, exceeding 10 m along the Mississippi coastline during Katrina (2005) (Fritz et al. 2007), and can cause widespread destruction as they propagate into bays and estuaries. To mitigate damages, prevent unnecessary loss of life, and better prepare coastal communities for incoming storm threats, public and private agencies use numerical models to predict storm surge and coastal flooding, both during storms for real-time decision support (Jelesnianski et al. 1992; Fleming et al. 2008) and between storms for infrastructure design and planning (Blanton and Luettich 2008; URS Corporation 2009; Bender 2013, 2014, 2015; U.S. Army Corps of Engineers 2021). Local, state, and federal stakeholders use this information to plan evacuation routes, set insurance rates, and design flood protection structures. For these predictions to be useful, they must be accurate (matching real water levels within tens of centimeters) and efficient (completing within 1 to 2 hours).

Storm surge can vary significantly in coastal regions, even between areas that are geographically close, due to the slope of the shelf, bathymetry and geometry of coastal features, and track of the storm (Park et al. 2022). Numerical models use computational grids (or meshes) to represent coastal domains with millions of computational cells (or elements), which can vary in size to represent complexities in coastal geometry and ground surface elevations. One of the largest contributors to loss in accuracy in hydrodynamic numerical models is improperly representing bathymetric (i.e. rivers and smaller channels) and topographic (i.e. dunes, ridges, and other raised features) flow controls. If the grid resolution is too coarse, then it may disallow the inclusion of small-scale flow controls and thus have detrimental effects on prediction accuracy. However, the grid resolution can also be too high, especially if it increases the computational expense beyond what is required in real time, even if the grid produces highly accurate results. Cell placement and resolution are critical because a small change in minimum cell size can dramatically increase the number of computational cells needed to represent the domain (Hagen et al. 2001; Bilskie et al. 2020). Thus, there is a trade off between the accuracy and efficiency of the model based on the amount of topographical and bathymetric complexity included in the model grid (Kerr et al. 2013).

This trade-off can be overcome via subgrid models, which account for small-scale changes in ground surface and landcover on coarsened numerical grids. Whereas the model scale (i.e. grid resolution) may be on the order of tens or hundreds of meters, the subgrid scale (e.g. digital elevation model) may be 1 m or smaller, and thus it can represent features that would be aliased by the model using conventional methodology. Subgrid modeling has been developed for decades and has shown that drastic improvements in computation time and accuracy can be achieved by running hydrodynamic models with subgrid information. Bates and Hervouet (1999) and Defina (2000) were some of the first hydrodynamic studies to include subgrid corrections that allowed for partially wet computational cells in their regional-scale tidal models. Defina (2000) also included a subgrid correction to advection, which enhanced the model's ability to account for small-scale changes in accelerations due to rapidly varying bathymetry. Casulli (2009) and Casulli and Stelling (2011) incorporated subgrid corrections to a 2D and 3D finite volume model that improved the wetting and drying of the model and the mass balance. Volp et al. (2013) used a 2D, depth averaged, finite volume model with subgrid based additions to bottom roughness and high resolution bathymetry to properly account for small scale variations in these parameters in the model. Sehili et al. (2014) used principles from Casulli (2009) and Casulli and Stelling (2011) in the semi-implicit finite difference UnTRIM^2 3D model. This study used subgrid additions to correct wetting and drying as well as bottom friction overestimation in a storm surge model of the Elbe Estuary, Germany.

Since these original studies, there have been many advancements and expansions of subgrid corrections in the modeling community. Wang et al. (2014) used the Semi-implicit Eulerian Lagrangian Finite Element (SELFE) model as boundary forcing to a regional subgrid model of New York City during Sandy (2011). Wu et al. (2016) combined the reformulated 2D shallow water equations used in Defina (2000) with the improvements to friction from Volp et al. (2013) in a numerical model of a tidal marsh. This study used pre-calculated subgrid variables to further improve the efficiency of the model and stored them in lookup tables. Kennedy et al. (2019) formalized many of these previous subgrid studies and added a subgrid correction to the water surface gradient terms in the governing equations which allowed for a better representation of the timing of the flow in some situations. Casulli (2019) and Begmohammadi et al. (2021) added cell clones to their subgrid models which allows the model to account for flow blocking features within a coarsened computational cell. This technology resolves the problem of incorrect hydraulic connectivity often seen in subgrid models. Woodruff et al. (2021) implemented subgrid corrections on a regional-scale storm surge model of Calcasieu Lake, LA with forcing from Rita (2005). This study was the first to implement subgrid corrections on the widely used ADCIRC hurricane storm surge and ocean circulation model. Similarly, Begmohammadi et al. (2022) added subgrid corrections into the well established Sea, Lake, and Overland Surges from Hurricanes (SLOSH) model, which is used by the National Weather Service (NWS) in real-time forecasting. Deb et al. (2023) implemented the averaging techniques described in Defina (2000) into the FVCOM hydrodynamic model as well as a marsh porosity slot algorithm, which allowed the model to hydraulically connect areas of the marsh that are connected by microbathymetric features such as cuts and rills. Woodruff et al. (2023) was the first to expand subgrid corrections to the ocean-scale and performed simulations of Matthew (2016) using ADCIRC on a unstructured mesh of the Western North Atlantic. Thus, subgrid corrections can offer improvements to model efficiency and accuracy by helping coarsened models resolve small scale changes.

As subgrid models have been applied for increasingly complex simulations (e.g. storm surge on an ocean-scale grid, Woodruff et al. 2023), there is a remaining question: how coarse is too coarse for real applications? Several studies have used two numerical grids to analyze how subgrid corrections performed and improved hydraulic connectivity on coarsened numerical grids. Defina (2000) used a fine resolution grid with an average resolution of 30 m and a coarse resolution grid with an average resolution of 260 m. Wu et al. (2016) used a high resolution grid and a coarse resolution grid with resolutions of 2 m and 8 m respectively. Begmohammadi et al. (2021) used two grids with resolutions 4 m and 256 m. Woodruff et al. (2021) created two unstructured finite element meshes surrounding Calcasieu Lake, LA with minimum resolutions of 50 m and 2000 m. Begmohammadi et al. (2022) created two polar grids representing eastern North Carolina that had averaged resolutions of 1.84 km and 3.67 km. These studies found comparable results could be achieved on grids with 1 to 2 orders of magnitude fewer computational cells. Although using two levels of resolution demonstrates that subgrid corrections enhance hydraulic connectivity on coarsened grids, it does not fully examine the extent to which a grid can be coarsened and still maintain accurate predictions. Other subgrid studies have used multiple levels of coarseness to show how results degrade with increased cell spacing. Casulli and Stelling (2011) used four resolutions of 25 m, 50 m, 100 m, and 300 m. Volp et al. (2013) also used four resolutions ranging from 5 m in the finest grid to 50 m in the coarsest. Kennedy et al. (2019) used grids with resolutions of 8 m, 16 m, 32 m, 64 m, 128 m, 256 m, and 512 m to demonstrate how subgrid simulations degrade with coarser and coarser grids. Developing and testing various levels of coarseness was done fairly easily with a small number of land-use and ground surface datasets, since these studies were completed on regional scales (with domains smaller than a few hundred square kilometers).

For simulations of tropical cyclones and storm surge on ocean-scale domains, the model must represent flows over a broad range of spatial scales. Ocean-scale grids must define features as small as a 50-m-wide tidal channel to as large as a 200-km-wide continental shelf. The extent to which subgrid corrections can sufficiently resolve this range of scales has not been thoroughly examined and tested. Exploring the resolution required to adequately resolve these features to maintain accurate predictions would allow for further increases in computational efficiency, because guidelines could be established about relative coarseness of model resolution.

It is hypothesized that if a set of guidelines for large-scale subgrid model resolution can be established, then more efficient and accurate water level predictions can be achieved on an optimized computational grid. To test this hypothesis, we will use an ocean-scale hurricane storm surge model to run subgrid simulations on sequentially coarsened meshes to find the maximum computational cell spacing at which we can achieve accurate water level predictions when compared to high resolution simulations. We will explore the numerical mechanisms that limit the coarsening of the model by reviewing water levels, current velocities, and hydraulic connectivity within the mesh to find where and why coarsened simulations produce inaccurate predictions. This work will provide a path forward for designing and running ocean-scale subgrid storm surge simulations through analysis of benchmark times between various model resolutions and justification of any minor losses in accuracy with increases in computational efficiency through the use of coarsened meshes.

4.3 Methods

4.3.1 Subgrid ADCIRC

We use the ADvanced CIRCulation (ADCIRC) coastal circulation model with subgrid additions that were developed by Woodruff et al. (2021) and (Woodruff et al. 2023). ADCIRC has been used for the last 30 years to predict tide-, atmospheric-, and density-driven circulation using a continuous-Galerkin, finite-element framework (Luettich et al. 1992; Westerink et al. 2008).

Averaged Equations

ADCIRC uses modified versions of the shallow water equations to compute water levels and current velocities at vertices of triangular elements in an unstructured mesh. For the subgrid additions, these governing equations are averaged using techniques from Whitacker (1985), and approximations to the boundary integrals from the averaging are found using various closure approximations from Defina (2000), Volp et al. (2013), and Kennedy et al. (2019). After the averaging, the governing equations include closure terms, e.g. for the conservative momentum equation in the x-direction:

$$\begin{aligned}\frac{\partial \langle UH\rangle_G}{\partial t} + g\, C_\zeta \langle H\rangle_G \frac{\partial \langle \zeta\rangle_W}{\partial x} = &-\frac{\partial\, C_{UU}\langle U\rangle\langle UH\rangle_G}{\partial x} - \frac{\partial\, C_{VU}\langle V\rangle\langle UH\rangle_G}{\partial y} \\ &- f\langle VH\rangle_G - g\langle H\rangle_G \frac{\partial P_A}{\partial x} + \phi \left\langle \frac{\tau_{sx}}{\rho_0}\right\rangle_W - \frac{C_{M,f}\langle \mathbf{U}\rangle\langle UH\rangle_G}{\langle H\rangle_W} \\ &+ \frac{\partial}{\partial x} E_h \frac{\partial \langle UH\rangle_G}{\partial x} + \frac{\partial}{\partial y} E_h \frac{\partial \langle UH\rangle_G}{\partial y},\end{aligned} \tag{4.1}$$

where the averaged variables are in brackets $\langle\cdot\rangle$, and closure terms are indicated in red font. Refer to Woodruff et al. (2021) and (Woodruff et al. 2023) for a complete derivation and discretization of the averaged governing equations.

The averaging must be computed over areas at the model scale, which can be challenging for models like ADCIRC, in which the computed quantities are defined at the vertices of the triangular elements (instead of the element centers). Elemental areas are defined by splitting each element into three sub-elements, where bathymetric and land-cover data are integrated and averaged into the flow variables described in the governing equations. Vertex areas are defined by combining the sub-elements surrounding each vertex; quantities are area-averaged to the vertex (Figures 2.1 and 3.2). Refer to Woodruff et al. (2021) and Woodruff et al. (2023) for a detailed explanation of the averaging areas.

Stability in Wetting and Drying

Previous subgrid ADCIRC studies have used a minimum wet area fraction as the wetting and drying criteria for the model. If the wet area fraction for an element or vertex was above some minimum ($\phi > \phi_{min}$) set by the user, that element or vertex would be wet and included in the computations, otherwise it would be dry and excluded. However, in some regions, this method can allow extremely small total water depths (on the order of millimeters or smaller), which can lead to instabilities when water velocities are calculated as volume flux divided by the total water depth. This problem is fairly common in numerical modeling schemes for wetting/drying and can also lead to spurious water slopes and subsequent instabilities, even without subgrid additions (Luettich and Westerink 1995a; Bates and Hervouet 1999; Bradford and Sanders 2002; Dietrich et al. 2004; Medeiros and Hagen 2013).

With subgrid additions to ADCIRC, when the subgrid model was relatively coarse (as in Woodruff et al. 2021), any numerical instabilities were not significant enough to cause the model to stop running or affect the result. However, in this study with increases in mesh resolution, the instabilities caused by extremely small total water depths can lead to simulation failure. Therefore, a new wetting criteria was developed for this study and based on the grid-averaged water depth $\langle H \rangle_G$. For an element or vertex to be considered wet, it must have a $\langle H \rangle_G > \langle H \rangle_{G_{min}}$ set by the user. (For this study, $\langle H \rangle_{G_{min}} = 0.1$ m.) This new wetting criteria improves model stability and does not significantly change model results. This scheme also enables subgrid ADCIRC to have a wetting and drying scheme that quasi-depends on the volume of water in a subgrid area.

Efficiency in Lookup Tables

One key requirement for subgrid models is the use of lookup tables, which define relationships between wet area fractions and averaged quantities for use during the simulation, e.g. to find the specific grid-averaged water depth $\langle H \rangle_G$ corresponding to a wet area fraction $\phi = 0.42$ at a given vertex. These tables are pre-computed and then read into memory at the start of the simulation. For simulations with high spatial resolution and/or large domain coverage, the memory requirements must be managed carefully.

Previous ADCIRC subgrid studies have used lookup tables that cover a large range of potential water surface elevations, e.g. −20 m to 20 m. In those studies, for each water surface elevation, both element- and vertex-averaged variables were computed. For each element, three sub-element averaged quantities were computed for a range of possible water surface elevations. Similarly, for each vertex, the average of the surrounding sub-element quantities were found for each water surface elevation in the range. Thus, the elemental lookup tables had size:

$$N_\zeta \times N_{VE} \times N_E, \tag{4.2}$$

where N_ζ is the number of possible water surface elevations, N_{VE} is the number of sub-elements corresponding to each vertex in a triangular element, and N_E is the number of elements in the mesh. Similarly, the vertex lookup tables have size:

$$N_\zeta \times N_V, \tag{4.3}$$

where N_V is the number of vertices in the mesh. For meshes with a relatively small number of elements and vertices (fewer than a couple hundred thousand), this lookup table structure can be managed in memory during the simulation.

However, as meshes surpass several hundred thousand elements and vertices, these lookup tables get prohibitively large to read into memory (especially with a large range of surface elevations). Therefore, in this study, to reduce the size of the lookup tables, a set of possible ϕ values were used to represent the water depths at which an element or vertex is fully dry, partially wet, or fully wet. These ϕ values are spaced at an even increment from $\phi = 0$ (fully dry) to $\phi = 1$ (fully wet), which reduces the size of the lookup tables to:

$$N_\phi \times N_{VE} \times N_E \tag{4.4}$$

and

$$N_\phi \times N_V, \tag{4.5}$$

where N_ϕ is the number of possible ϕ values determined by the user. For this study, $N_\phi = 11$.

4.3.2 Meshes with Varying Resolution of U.S. Atlantic Coast

Datasets

This study will focus on the South Atlantic Bight (SAB), which stretches from the southern tip of Florida through the North Carolina Outer Banks. This coastline contains sandy beaches, riverine deltaic systems, and complex lagoons and tidal inlets (Blanton et al. 2004). In addition, the region's shallow, wide continental shelf is especially vulnerable to storm surge amplification. The SAB also experiences the effects of several tropical cyclones each year (Zarillo et al. 2012), and thus it is a relevant area of interest for this study.

Ground surface data sets were collected from the NOAA Digital Coast Database and the United States Geological Survey (USGS) National Map (Figure 3.3). For nearshore and coastal regions along the SAB, the Continuously Updated Digital Elevation Model (CUDEM) data sets were used. This ground surface data has bathymetric and topographic information with resolutions ranging from 1/9 arc-second (3.4 m) in the nearshore- and coastline-adjacent overland regions to 1/3 arc-second (10.3 m) in the offshore. The National Map (TNW) data sets were used for further inland regions and in areas where there were gaps in coverage of the CUDEM data sets. TNW data provide ground surface information at 10 m resolution for the entirety of the U.S. Atlantic coast. For areas farther offshore and regions outside the SAB, the Shuttle Radar Topography Mission (SRTM) global ground surface data were used. These open-source data provide continuous coverage of the entire globe at about 462 m resolution.

Landcover data sets were collected from the 2016 Coastal Change Analysis program (C-CAP) standardized remotely sensed landcover data set, which covers coastal intertidal areas, wetlands, and adjacent upland portions of the United States. This data set has 30 m resolution and extends about 5 km offshore. For areas far from the SAB, landcover was assumed as open water. These values were converted to Manning's n values (Table 4.1), which are used in the calculation of the bottom friction correction and the normal bottom friction calculation with Manning's formula (Equation 3.6) used in traditional ADCIRC.

Many of the ground surface data sets are patchy in their coverage of the coastline and inland areas. Therefore, hundreds of smaller data sets were combined to create a mosaic representation of nearshore and onshore coastal bathymetry and topography (Figure 3.3). When calculating the subgrid variables for storage in the lookup tables, the pre-processing code uses a hierarchy of data sets, so that the highest quality/resolution data sets are given priority over the lower resolution data sets.

For mesh development, a set of channel thalweg shapefiles were created from the mosaic of ground surface data sets. These thalwegs were derived using a geospatial flow routing technique that accounts for the gradients in DEMs to calculate flow path, flow path length, and flow accumulation (Cho 2020). The result of this process is a network of streams indicating the natural flow path from upland to coastal areas (Figure 4.1).

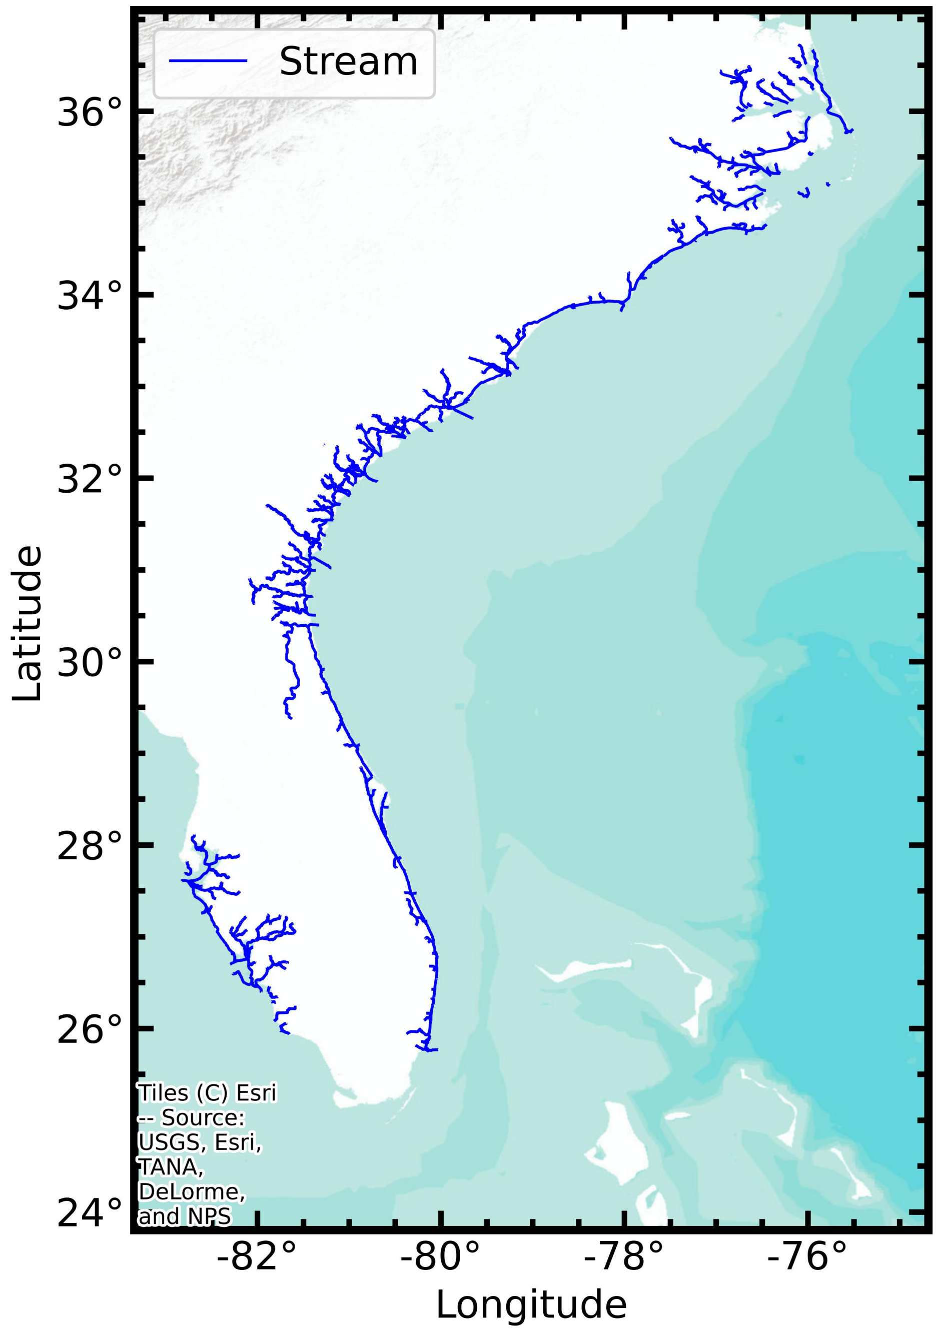

Figure 4.1: Stream network generated using flow routing technique to find channel thalwegs. Note some smaller streams and intercoastal water ways had to be manually drawn.

Table 4.1: C-CAP landcover to Manning's n conversion table.

Classification	C-CAP	Manning's n
Open Water	0	0.025
High Intensity Developed	2	0.120
Medium Intensity Developed	3	0.100
Low Intensity Developed	4	0.070
Developed Open Space	5	0.0350
Cultivated Land	6	0.01
Pasture/Hay	7	0.055
Grassland	8	0.035
Deciduous Forest	9	0.16
Evergreen Forest	10	0.18
Mixed Forest	11	0.17
Scrub/Shrub	12	0.08
Palustrine Forested Wetland	13	0.15
Palustrine Scrub/Shrub Wetland	14	0.075
Palustrine Emergent Wetland	15	0.07
Estuarine Forested Wetland	16	0.15
Estuarine Scrub/Shrub Wetland	17	0.07
Estuarine Emergent Wetland	18	0.05
Unconsolidated Shore	19	0.03
Bare Land	20	0.03
Open Water	21	0.025
Palustrine Aquatic Bed	22	0.035
Estuarine Aquatic Bed	23	0.03
Tundra	24	0.090
Perennial Ice/Snow	25	0.010

Mesh Development

The meshes used in this study were designed using OceanMesh2D (Roberts et al. 2019a). This software uses coastline geometry, mesh size gradients, and bathymetric data to design high quality, unstructured, triangular element meshes for coastal circulation models. Channel thalwegs were used to ensure proper representations of important hydraulic features along the coast such as rivers, shipping channels, intercoastal waterways, and steep bathymetric gradients offshore for accurate tidal propagation from the deep ocean onto the continental shelf and into bays, rivers, and estuaries (Roberts et al. 2019b).

To test how subgrid corrections perform at varying resolutions, five meshes were developed (Table 4.2) These meshes were designed so that each coarsening level roughly halved the number of vertices and elements in the mesh, which effectively reduces a given simulation time by a factor of 2 from one level to the next. For this study, the highest resolution mesh was designed as a 'forecast-grade' mesh, meaning that the computational expense of the mesh is similar to that of a mesh used in real time forecasting. This mesh has a minimum resolution of about 60 m, which is adequate to resolve most major inlets, rivers, intercoastal waterways, and flow blocking features such as barrier islands. The next level of resolution has a minimum element edge length of about 100 m, which was chosen so that most major inlets and waterways would still be resolved; although, some bathymetric features would be aliased such as narrow sections of intercoastal waterways and smaller tidal creeks. Following these two higher resolution meshes, each subsequent mesh resolves fewer and fewer bathymetric and topographic features. The minimum 200 m resolution mesh will only resolve major inlets and waterways, the minimum 400 m resolution mesh will alias many major rivers and inlets, and the minimum 1000 m resolution mesh will create gaps in flow blocking features such as barrier islands.

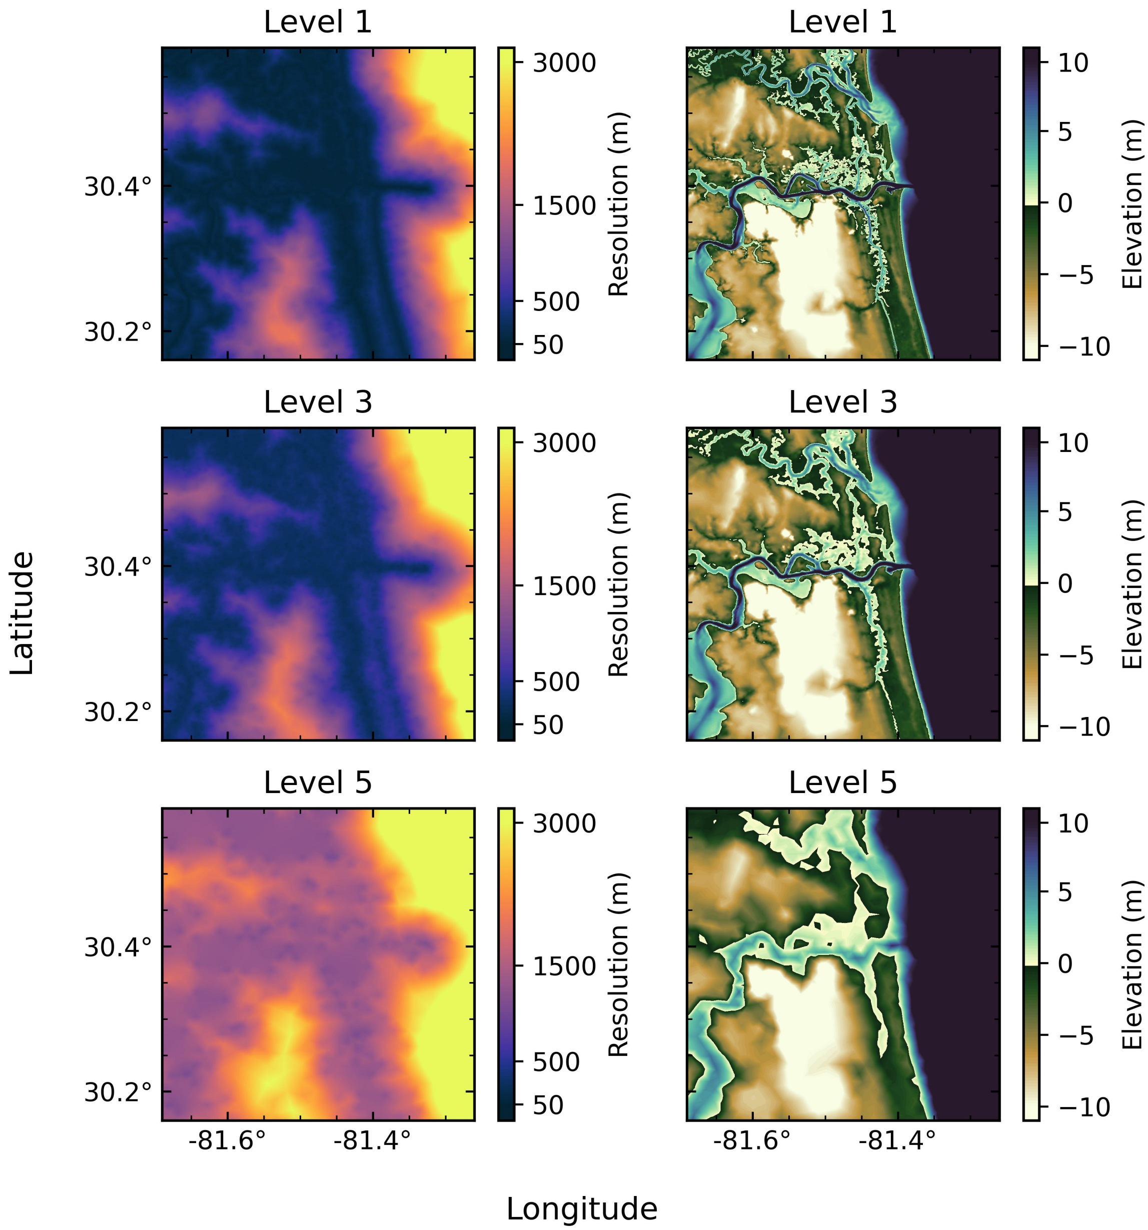

Figure 4.2: Mesh resolution and bathymetric interpolation differences between the Level 1, Level 3, and Level 5 meshes. The area shown encompasses Jacksonville, FL and the northern portion of the St. John's River from –81.7° to –81.25° W and 30.15° to 30.6° N.

Table 4.2: Mesh name, minimum resolution (min), maximum resolution nearshore (maxns), and maximum resolution (max) for each level of mesh along with the number of vertices and elements contained in the mesh.

Mesh name	min (m)	maxns (m)	max (m)	# of vertices	# number of elements
SABv3-60m	60	100	50,000	3,531,883	6,812,980
SABv3-100m	100	250	50,000	1,751,839	3,490,798
SABv3-200m	200	300	50,000	754,159	1,496,184
SABv3-400m	400	500	50,000	359,086	706,623
SABv3-1000m	1,000	2,000	50,000	178,398	346,137

4.3.3 Storm Simulations

Matthew (2016) and Florence (2018)

Matthew (2016) and Florence (2018) were chosen as atmospheric forcing in this study. Matthew (2016) was a strong, category-5, tropical cyclone that tracked shore parallel to the U.S. Atlantic coast from 7-9 October 2016. This storm was selected for this study because it affected an expansive section of coastline and thus allows for extensive testing of subgrid corrections across a wide range of coastal environments. This storm is described by observations at 218 locations for use in assessing model results. In contrast, Florence (2018) approached the coast as a category-4 hurricane on a shore normal trajectory. Because of this trajectory, Florence produced a more localized storm surge constrained mostly to southeast North Carolina. Florence is described by observations at 114 locations (Figure 4.3) and will allow us to see how subgrid corrections perform in a shore normal landfall scenario. Both the Florence and Matthew observations were collected at NOAA permanent tidal gauges, USGS permanent tidal gauges, and USGS rapid deployment pressure sensors in the path of the incoming storm.

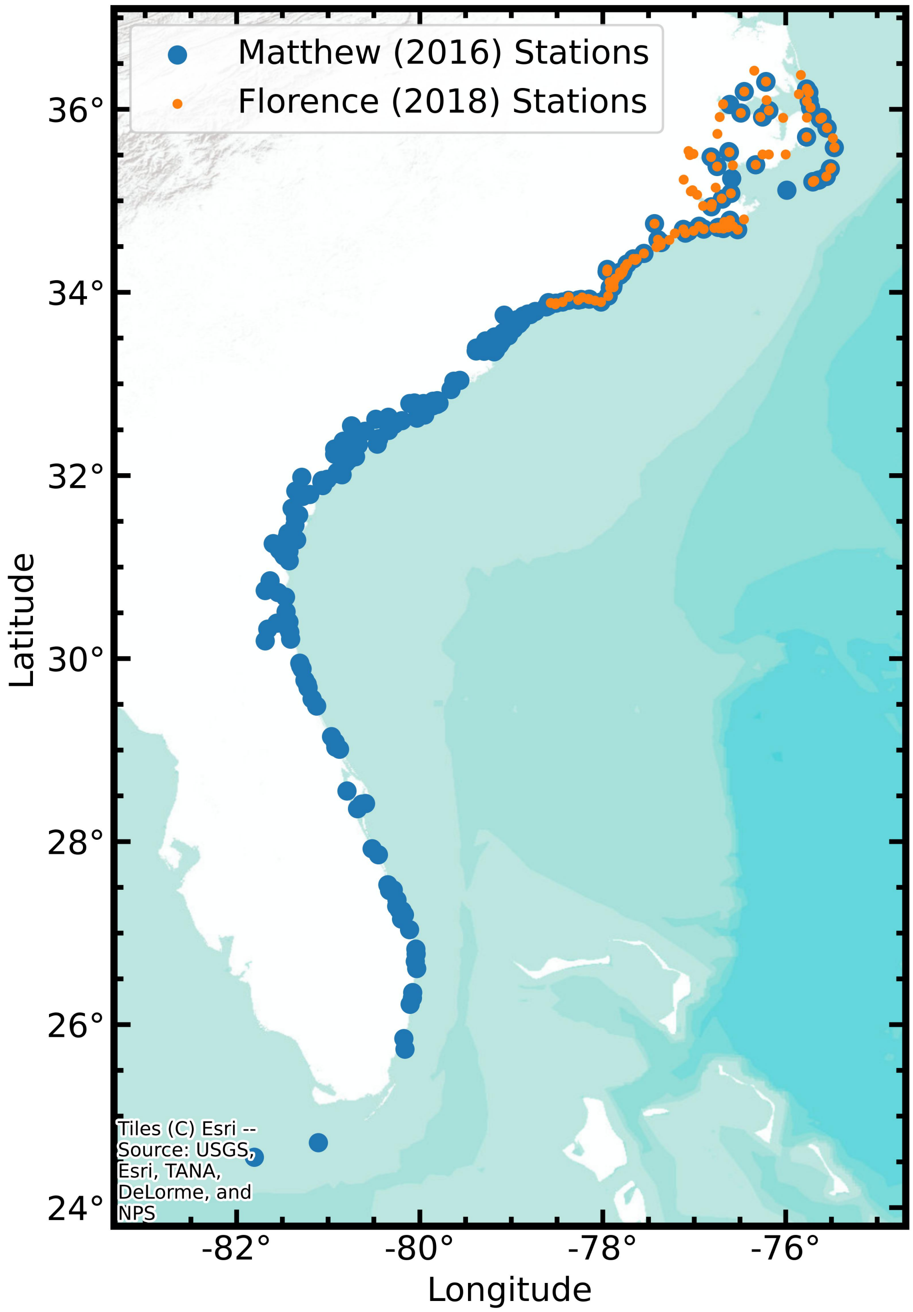

Figure 4.3: Hurricane Matthew (2016) and Florence (2018) observation locations.

Matthew (2016) caused increases in water level from southeast Florida through the North Carolina Outer Banks. The maximum surge observed by a pressure sensor was 2.56 m

NAVD88 at Fort Matanzas Beach. Elsewhere in northeast Florida, peak water levels ranged between 1.5 to 2.1 m NAVD88. In Georgia, Matthew's storm surge reached peaks between 0.9 to 1.5 m NAVD88 with a maximum of 1.54 m NAVD88 recorded at Fort Pulaski, GA. Similar to Georgia, much of the South Carolina coastline experienced 0.9 to 1.5 m NAVD88 of peak surge. Across North Carolina, surge heights varied significantly depending on location due to the wide ranging coastal geometry. Along the barrier islands from the southern border to Cape Hatteras, North Carolina water levels peaked between 0.9 to 1.5 m NAVD88. However, the highest surges in the state occurred in the sound side of the Outer Banks and measured 1.2 to 1.8 m NAVD88 (Stewart 2017).

In contrast to the widespread effects of Matthew, Florence had more localized surge constrained to coastal North Carolina. Maximum surge for this storm occurred along the Neuse River estuary, with observed peaks between 2.4 to 3.4 m NAVD88. Along the barrier islands from Cape Fear through Cape Lookout, maximum surge was generally between 1.5 to 2.4 m NAVD88 with a maximum of 2.74 m observed in Wrightsville Beach, North Carolina (Stewart and Berg 2019).

ADCIRC has been used to predict storm surge for Matthew and Florence both in real time forecasting and hindcast studies (Asher et al. 2019; Thomas et al. 2019; Rucker et al. 2021; Thomas et al. 2022). These studies produced highly accurate results when compared to observation data and aimed to replicate the real storm surge and flooding as best as possible. In this study, although observation data are used to calibrate the highest resolution mesh, the central focus will not necessarily be on mimicking the real ocean response as recorded at a still-limited number of observation locations, but rather on quantifying how the subgrid predictions change overall with incremental decreases in resolution.

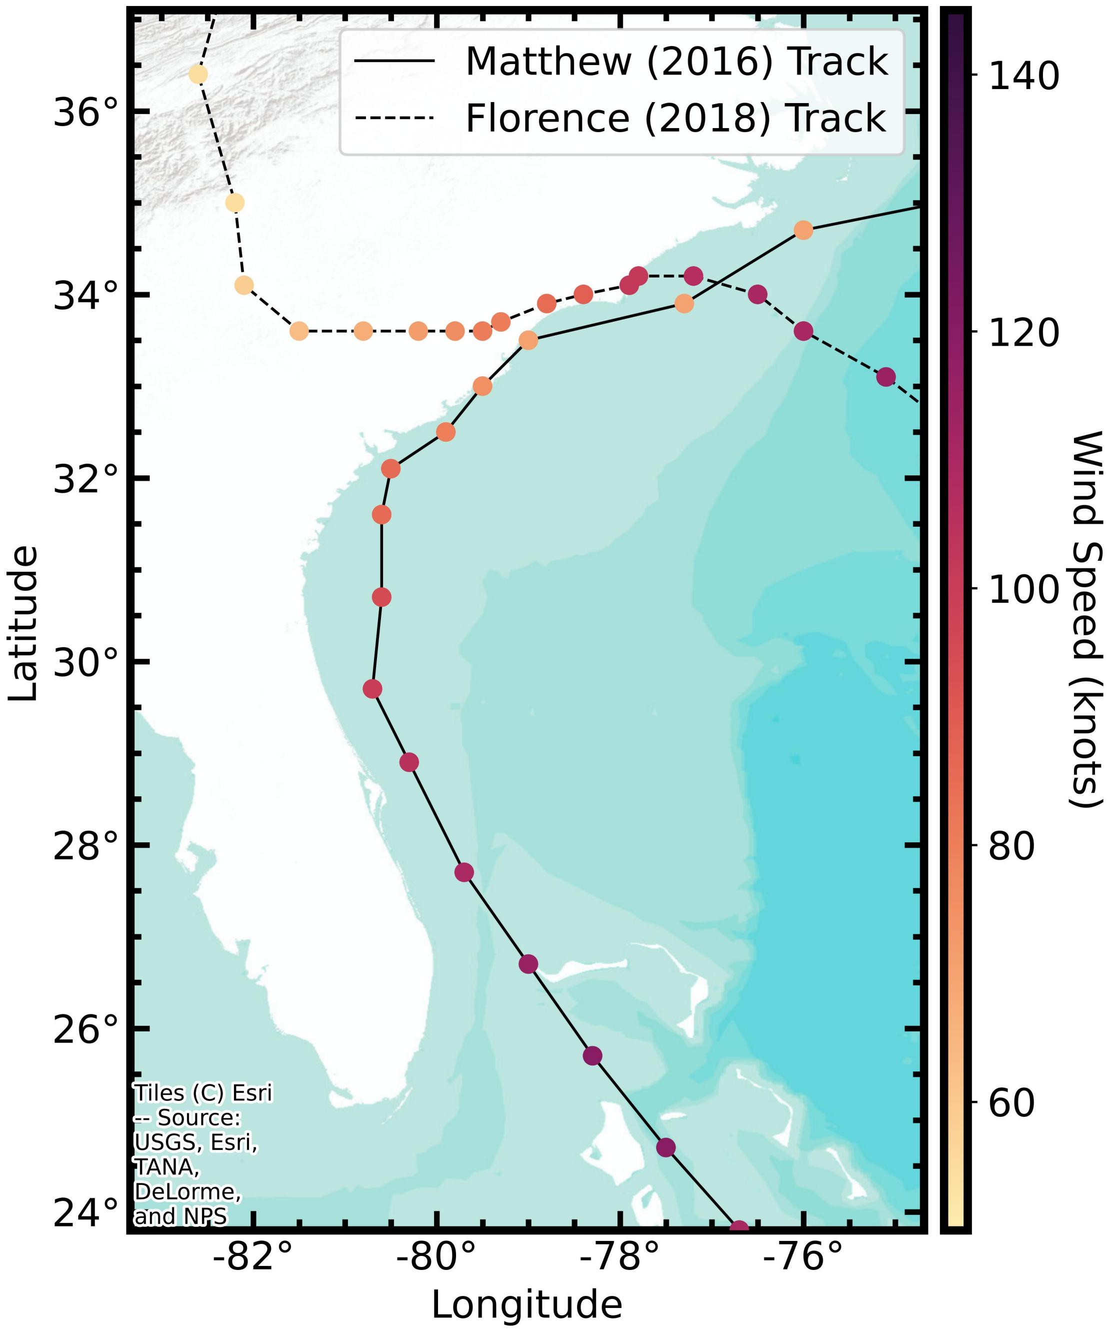

Figure 4.4: Hurricane Matthew (2016) and Florence (2018) track and intensity.

Model Parameters

The atmospheric forcing (surface pressure and wind fields) for Matthew (2016) and Florence (2018) were provided from Oceanweather Inc. (OWI) and Ratcliff (2022). OWI fields are created using data from weather station, buoy, aircraft, ship, and satellite stations (Oceanweather Inc. 2018). For both storms, the OWI wind fields consist of a coarse, basin scale grid with resolution 0.25° stretching from 5° N to 47° N and 99° W to 55° W. In addition to the coarser resolution wind field, each storm had a 0.05° high resolution wind field stretching from 15° N to 40° N and 99° W to 55° W for Matthew and 31° N to 37° N and 82° W to 74° W for Florence. Wind speeds are scaled at the ground level using coefficients determined from landcover data sets. In ADCIRC, these coefficients are the surface directional effective roughness length and the surface canopy coefficient.

For Matthew (2016), a tidal spin up was first applied for 15 days (0000 GMT 16 September 2016 to 0000 GMT 1 October 2016). The model was then run for an additional 10 days (0000 GMT 1 October 2016 to 0000 GMT 11 October 2016) with both atmospheric and tidal forcing. The simulations using Florence (2018) were run with an 11-day tidal spin up (0000 GMT 21 August 2018 to 0000 GMT 1 September 2018), followed by a 17-day (0000 GMT 1 September 2018 - 0000 GMT 18 September 2018) simulation with both atmospheric and tidal forcing. Model parameters including primitive weighting coefficient in the continuity equation (τ_0), horizontal eddy viscosity ($ESLM$), advection state, and Manning's roughness coefficient (n) were similar for all simulations but varied slightly depending on mesh resolution and whether or not subgrid corrections were used. Primitive weighting in the continuity equation was spatially varying and determined based on depth at a particular vertex and its distance from connected vertices. The Manning's roughness coefficient was spatially varying and derived from the C-CAP landcover data sets mentioned earlier, and the horizontal eddy viscosity is kept constant throughout each mesh. Sea surface height above geoid was set to 0.284 m for Matthew (2016) and 0.196 m for Florence (2018). These heights represent the background water levels in the Western North Atlantic around the time of each storm (Owensby et al. 2020).

Metrics

For each mesh and each storm, we will evaluate the ability of subgrid ADCIRC to predict the magnitude and timing of peak water levels, the upstream propagation of flood waters, and the overall flooded area. For the magnitude of peak water levels, a 1:1 analysis will be performed and root mean square error (E_{RMS})

$$E_{\mathrm{RMS}} = \sqrt{\frac{\sum_{i=1}^{N}\left(\zeta_i - \hat{\zeta}_i\right)^2}{N}}, \tag{4.6}$$

the coefficient of determination (R^2):

$$R^2 = 1 - \frac{\sum_{i=1}^{N}\left(\zeta_i - \hat{\zeta}_i\right)^2}{\sum_{i=1}^{N}\left(\zeta_i - \bar{\zeta}\right)^2}, \tag{4.7}$$

mean normalized bias (B_{MN}):

$$B_{\mathrm{MN}} = \frac{\frac{1}{N}\sum_{i=1}^{N}\left(\hat{\zeta}_i - \zeta_i\right)}{\frac{1}{N}\sum_{i=1}^{N}|\zeta_i|} \tag{4.8}$$

and best fit slope will be calculated. For this analysis, we want to investigate how the decrease in resolution affected maximum surge height at the observation locations. Prior to running the analysis on the coarsened mesh storm simulations, we first performed this analysis on the Level 1 subgrid simulation using observation results taken during the storm at stations along the coast to demonstrate that this simulation provides adequate model results to serve as a benchmark to which coarsened mesh storm simulations could be compared to. Thus for validation of the Level 1 simulation, $\hat{\zeta}_i$ is the maximum observed water level at a particular station, and ζ_i is the maximum water level of the Level 1 subgrid simulation for the same location. For all other 1:1 comparisons performed between meshes with varying resolutions, $\hat{\zeta}_i$ will be the maximum water level in the Level 1 subgrid simulation and ζ_i is the maximum water level from one of the coarsened mesh storm simulations.

The E_{RMS} will allow us to investigate the overall magnitude of elevation difference between the Level 1 subgrid and coarsened simulations. The R^2 value with respect to the 1:1 line will quantify how much the coarsened simulations varied from the Level 1 subgrid, thereby giving and indication of how well the coarsened models performed. The B_{MN} determines if the coarsened models under- or over-predicted, and the best fit slope is a linear regression indicating how well the coarsened results fit to the Level 1 subgrid results (a perfect fit would be a slope of 1.0).

When analyzing storm surge results, it is important to both match the peak water level of observations, as well as the time at which the peak water level occurred. To accomplish this, water level time series were taken at the beginning, middle, and end of continuous channel thalwegs at various locations along the coast affected by each storm. For this analysis, only the E_{RMS} and B_{NM} will be calculated. This will give us a representation of how model resolution affects the timing and magnitude of surge propagation from the open coast to inland locations. Additionally, to demonstrate how storm surge propagates up a channel, the maximum water level along several channel thalwegs will be plotted against distance along the thalweg. This will help show where various choke points are along a channel and how mesh resolution plays a roll in storm surge ability to propagate inland through narrowing sections of a waterway.

4.4 Results

The results for the simulations described in the methods section are contained in the following paragraphs. First, we will validate that our highest resolution subgrid simulation can serve as 'truth' to which we can compare coarsened simulation results. next, we analyze how mesh resolution affects maximum water level results when compared to our 'truth'. Then, we show how the propagation of storm surge changes with mesh resolution along waterways in the SAB by analyzing maximum water levels along the channel thalwegs. This is followed by water level time series analysis at stations along select channels. Finally, we calculate the total wet area predicted by the conventional and subgrid simulations to see how it changes with varying resolution.

4.4.1 SABv3-60m Simulations as 'Truth'

Although the main focus of this study is the relative performance of the SABv3 meshes as their resolution is degraded, it is necessary to first establish a baseline performance for the highest-resolutions simulations on the SABv3-60m mesh. Simulations of Matthew and Florence were first performed on the SABv3-60m mesh with the subgrid ADCIRC. These simulations will offer the best hydraulic connectivity because they can represent subgrid features that exist below the minimum resolution of the mesh (60 m). Maximum water level predictions were compared to observations at station locations along the SAB (Figure 4.5), and error statistics were calculated (Table 4.3).

For Matthew, E_{RMS} is 0.47 m for all of the stations and 0.34 m if only wet stations are included. These results are adequate for the purposes of this study. Thomas et al. (2019) predicted maximum water levels with an E_{RMS} of 0.28 m; however, the simulations in that study also account for waves and thus were a more accurate representation of conditions during the storm. The R^2, best fit slope, and B_{MN} values are also similar to the values found by Thomas et al. (2019). In addition, only 4 out of the 218 stations were dry in this simulation on the SABv3-60m mesh, demonstrating the subgrid ADCIRC is able to represent flooding to inland locations.

For Florence, the performance was better with an E_{RMS} of 0.25 m; this result is comparable to studies done on much higher resolution simulations (Thomas et al. 2022). Similarly, the results for B_{MN}, best fit slope, R^2, and the number of dry stations all show that the subgrid ADCIRC simulation does an adequate job of producing high water levels along the coast. Thus, for the purposes of this study, the subgrid ADCIRC simulations on the SABv3-60m mesh can act as "truth" for which we can compare the results of the coarsened meshes.

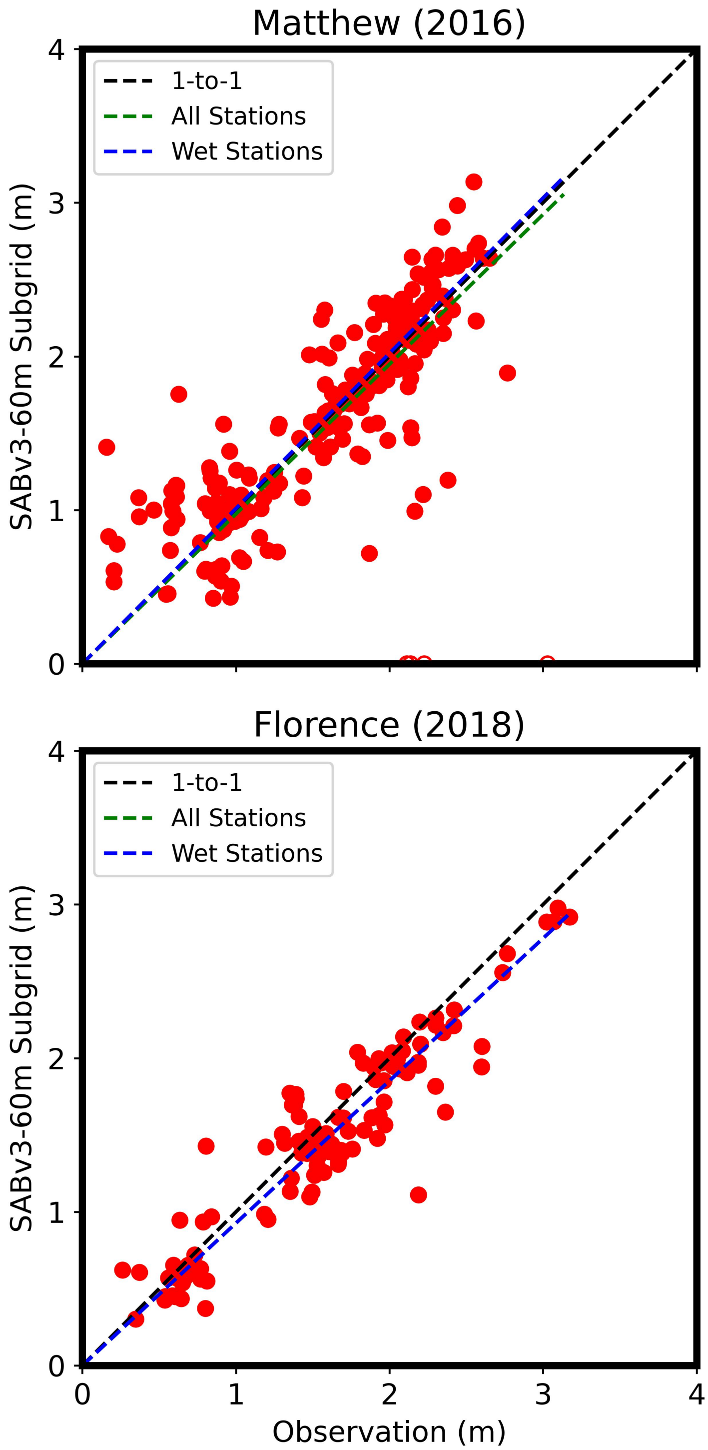

Figure 4.5: SABv3-60m subgrid simulation maximum water levels compared to observation taken during Matthew (2016) and Florence (2018).

Table 4.3: Statistics of the maximum water level comparison between observations taken during Matthew (2016) and Florence (2018) and subgrid ADCIRC simulations on the SABv3-60m mesh.

Simulation	E_{RMS} (m)	R^2	Best Fit Slope (m/m)	B_{MN}	Dry Stations
Matthew All Stations	0.47	0.44	0.97	0.0046	4
Matthew Wet Stations	0.34	0.70	1.01	0.0325	–
Florence All Stations	0.25	0.86	0.93	-0.0651	0
Florence Wet Stations	0.25	0.86	0.93	-0.0651	–

4.4.2 Sensitivity for Maximum Water Levels

To analyze the performance of the coarsened resolution meshes, storm simulations using forcing from Matthew (2016) and Florence (2018) were completed for each level of mesh coarseness, both with and without subgrid corrections. Predicted maximum water levels were interpolated at the observation locations, and then compared to the maximum water levels from the subgrid ADCIRC simulation on the SABv3-60m mesh, which acted as our “truth”. It is emphasized that, from this point forward, we do not compare to observed water levels. Instead, all errors are computed relative to the subgrid simulations with the finest SABv3-60m mesh.

For the Matthew simulations, the traditional simulation performed as expected: as mesh coarseness increased, the number of dry stations increased, and the E_{RMS}, R^2, B_{MN}, and best-fit slope became worse (Figure 4.6 and Table 4.4). In addition, moving from the subgrid simulation to the conventional simulation on the SABv3-60m mesh, there was a 0.49-m increase in E_{RMS} when considering all of the stations, and 17 dry stations were introduced. This worsening in model predictions was also reflected in the R^2 value which dropped significantly from the ideal value of 1.0 to 0.4. The best fit slope and the B_{MN} both decreased in the SABv3-60m to −0.0737 and 0.93 which would indicate that the model started under-predicting maximum water depth. However, the decreased performance of the SABv3-60m Conventional simulation is likely due to the increased number of dry stations in the simulation; since, for the most part, the statistics for the wet stations did not suffer significantly in the SABv3-60m Conventional simulation. This pattern was repeated as mesh resolution decreased. Although, when considering all of the stations for Matthew the

Conventional simulations suffered from poor accuracy and high numbers of dry stations, when only considering the wet stations, there was not a huge difference in the Conventional and Subgrid simulations at their respective levels of coarseness. For instance, in the coarsest simulations (SABv3-1000m), the Conventional and Subgrid model only differed by 0.05 m in E_{RMS}, 0.06 in R^2, 0.0127 in B_{MN}, and 0.01 in best fit slope. Therefore, the increase in accuracy of the subgrid model at the various levels of coarseness can mainly be attributed to the increase in hydraulic connectivity indicated by the lack of dry stations in the subgrid simulations. The increase in hydraulic connectivity does come at the cost of aliasing some flow blocking features, and is reflected in the B_{MN} and best fit slope metrics which show that the subgrid model tends to over-predict as mesh resolution is coarsened.

The results from the Florence simulations mirrored many of the patterns found in the Matthew simulations (Figure 4.7 and Table 4.5). Removing subgrid corrections from the simulation on the SABv3-60m mesh introduced an E_{RMS} of 0.38 m, 10 dry stations, and significantly decreased the R^2 value relative to the SABv3-60m Subgrid simulation. As mesh resolution was coarsened, the error statistics for the traditional simulations continued to worsen, and the number of dry stations increased to 14 in the traditional simulation on the coarsest SABv3-1000m mesh. In contrast, when looking at the subgrid simulations, the E_{RMS} peaked at 0.48 m compared to 0.79 m in the traditional simulation, and the R^2 remained positive indicating that even the coarsest subgrid simulation maintained relatively small differences to the target data set. Similar to the Mathew data, the statistics for the Florence simulations suffered most when considering all of the station. When only considering the wet stations, the subgrid simulations did provide a better replication of the SABv3-60m subgrid results than the conventional simulations, but the differences were not as large. This again can be attributed to the fact that the majority of the subgrid simulations had no dry stations, with the exception of the simulation on the SABv3-100m mesh, which did not calculate water levels at 2 stations located on Pea Island and Hatteras Island along the North Carolina Outer Banks. Similar to the Matthew simulations, the Florence simulations with coarsened meshes tended to over-predict maximum water levels in both the subgrid and conventional simulations. Again, this can be attributed to the aliasing of flow blocking features in the coarsened meshes, especially along the barrier islands of North Carolina.

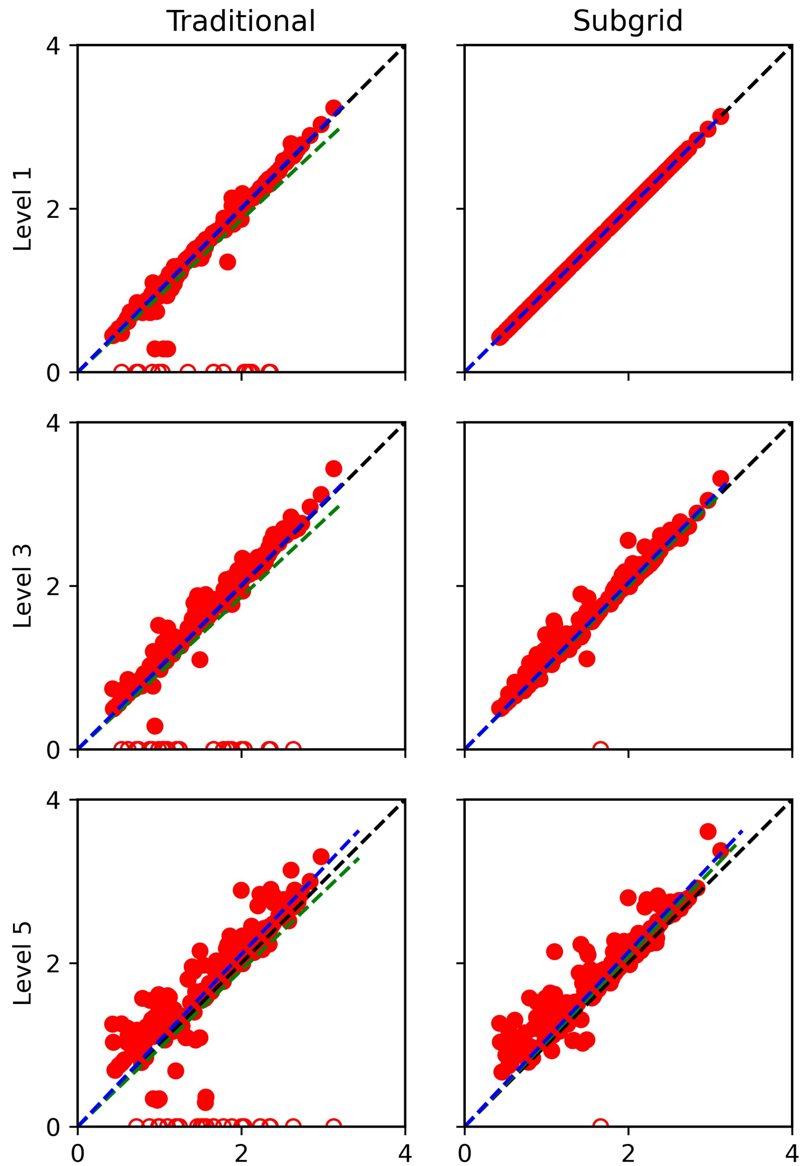

Figure 4.6: Comparison of max water levels of SABv3-60m, SABv3-200m, and SABv3-1000m simulations to the high resolution subgrid simulation for Matthew (2016).

Table 4.4: Statistics for maximum water levels of Matthew simulations compared to the reference SABv3-60m subgrid simulation. The All Stations stats used data from every station point where as the Wet Stations only used data at stations that were wet in all subgrid or conventional simulations.

All Stations	E_{RMS} (m)	R^2	B_{MN}	Slope (m/m)	# Dry
SABv3-60m Conventional	0.48	0.43	-0.0681	0.94	16
SABv3-100m Conventional	0.52	0.31	-0.0759	0.93	22
SABv3-200m Conventional	0.55	0.25	-0.0548	0.95	24
SABv3-400m Conventional	0.6	0.11	-0.0205	0.96	22
SABv3-1000m Conventional	0.67	-0.14	-0.0202	0.96	24
SABv3-60m Subgrid	0.0	1.0	0.0	1.0	0
SABv3-100m Subgrid	0.06	0.99	0.0035	1.0	0
SABv3-200m Subgrid	0.12	0.96	0.0418	1.03	0
SABv3-400m Subgrid	0.21	0.88	0.0855	1.06	0
SABv3-1000m Subgrid	0.32	0.74	0.132	1.1	0
Wet Stations	E_{RMS} (m)	R^2	B_{MN}	Slope (m/m)	# Dry
SABv3-60m Conventional	0.11	0.97	0.0033	1.01	–
SABv3-100m Conventional	0.19	0.91	0.0126	1.01	–
SABv3-200m Conventional	0.14	0.95	0.0488	1.04	–
SABv3-400m Conventional	0.21	0.89	0.0898	1.07	–
SABv3-1000m Conventional	0.34	0.72	0.1049	1.08	–
SABv3-60m Subgrid	0.0	1.0	0.0	1.0	–
SABv3-100m Subgrid	0.06	0.99	0.0035	1.0	–
SABv3-200m Subgrid	0.12	0.96	0.0418	1.03	–
SABv3-400m Subgrid	0.21	0.88	0.0855	1.06	–
SABv3-1000m Subgrid	0.32	0.74	0.132	1.1	–

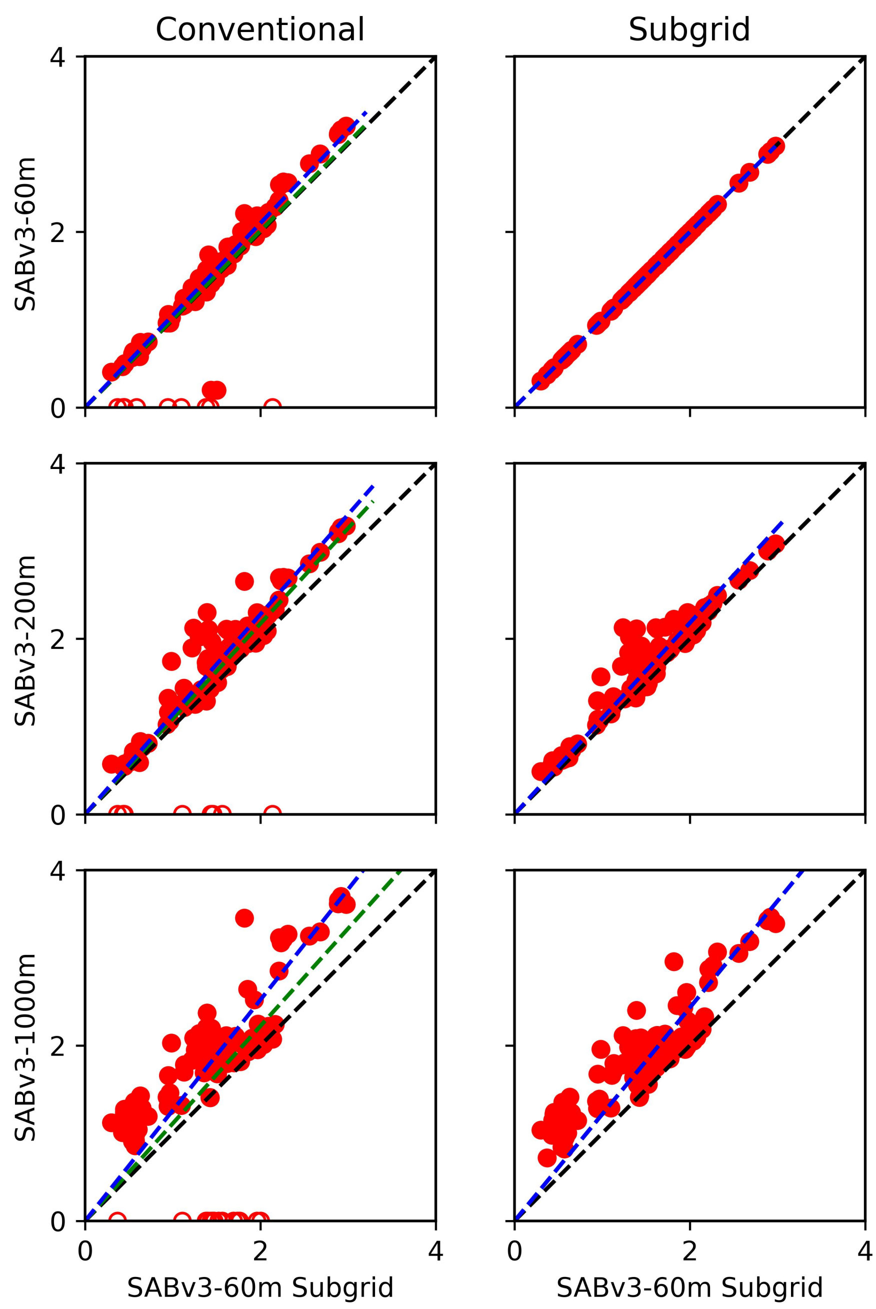

Figure 4.7: Comparison of max water levels of SABv3-60m, SABv3-200m, and SABv3-1000m simulations to the high resolution subgrid simulation for Florence (2018).

Table 4.5: Statistics for maximum water levels of Florence simulations compared to the reference SABv3-60m subgrid simulation. The All Stations stats used data from every station point where as the Wet Stations only used data at stations that were wet in all subgrid or conventional simulations.

All Stations	E_{RMS} (m)	R^2	B_{MN}	Slope (m/m)	# Dry
SABv3-60m Conventional	0.38	0.64	-0.0126	1.01	10
SABv3-100m Conventional	0.38	0.64	0.0142	1.03	10
SABv3-200m Conventional	0.45	0.51	0.0914	1.09	8
SABv3-400m Conventional	0.61	0.07	0.1206	1.09	11
SABv3-1000m Conventional	0.76	-0.42	0.1665	1.11	14
SABv3-60m Subgrid	0.0	1.0	0.0	1.0	0
SABv3-100m Subgrid	0.16	0.94	0.0081	1.01	2
SABv3-200m Subgrid	0.22	0.88	0.1075	1.09	0
SABv3-400m Subgrid	0.33	0.73	0.1895	1.15	0
SABv3-1000m Subgrid	0.48	0.43	0.285	1.22	0
Wet Stations	E_{RMS} (m)	R^2	B_{MN}	Slope (m/m)	# Dry
SABv3-60m Conventional	0.22	0.88	0.0469	1.05	–
SABv3-100m Conventional	0.25	0.84	0.0708	1.07	–
SABv3-200m Conventional	0.3	0.77	0.1548	1.14	–
SABv3-400m Conventional	0.42	0.59	0.2406	1.2	–
SABv3-1000m Conventional	0.57	0.26	0.3374	1.26	–
SABv3-60m Subgrid	0.0	1.0	0.0	1.0	–
SABv3-100m Subgrid	0.08	0.98	0.0208	1.02	–
SABv3-200m Subgrid	0.22	0.88	0.1075	1.09	–
SABv3-400m Subgrid	0.33	0.73	0.1895	1.15	–
SABv3-1000m Subgrid	0.48	0.43	0.285	1.22	–

4.4.3 Propagation of Peak Surges along River Thalwegs

One potential advantage of the subgrid ADCIRC should be its ability to represent flows through channels smaller than the model scale. To investigate this potential advantage, we quantified the propagation of peak surges along coastal rivers. These rivers become narrower at upland locations, and thus there is a potential for artificial blockages of surge in the model. The maximum water levels along the thalwegs of major waterways were taken at 6 locations for Matthew (2016) and 4 locations for Florence (2018). The locations for Matthew include the St. Johns River near Jacksonville, FL (JAX), the Savannah River near Savannah, GA (SAV), the Cooper River near Charleston, SC (CHA), the Cape Fear River near Wilmington, NC (CF), the New River near Jacksonville, NC (NR), and the Neuse River near New Bern, NC (NB). For Florence, only the Cooper River, Cape Fear River, New River, and Neuse River locations were used because storm effects were localized to the Carolinas. For each location, synthetic stations were placed 1 km apart starting at the open coast, and extending to the farthest reaches of the channel within the mesh bounds (Figure 4.8). Maximum water surface elevations that occurred over the entirety of each simulation are then plotted against the length of the channel (Figures 4.9, 4.10,4.11, 4.12).

The rivers were chosen for further analysis because they each have a unique set of attributes that affect storm surge propagation in different ways, thereby giving us a better understanding of how subgrid corrections will perform with varying coastal geometries. The St. Johns River in Florida is a massive waterway that extends several hundred kilometers from Jacksonville in the northeast part of the state to south of Melbourne in east central Florida. Along this stretch, the river connects a myriad of lakes, and goes from being highly channelized near Jacksonville to natural and meandering further upstream. The Savannah River represents a channelized waterway in a deltaic system that has frequent ship traffic and very deep water levels up to the Port of Savannah. In addition, there are numerous islands along the river that create multiple channels apart from the main shipping channel that could alter the movement of surge. Near Charleston, SC, the Cooper River lies on the northeast side of the city and is connected to numerous other water bodies and rivers including: Charleston Harbor, the Wando River, the Ashley River, and Lake Moultrie which is a dammed lake far upstream. In North Carolina, the Cape Fear River is a large estuarine system that is separated from the Atlantic Ocean on its eastern bank by barrier islands. This waterway contains a deep shipping channel that connects the ocean to the Port of Wilmington upstream. The New River location was chosen because of the unique, narrow inlet that connects the Atlantic to its numerous inland back bays. This channel has an additional inland choke point created by urban and military development near Jacksonville, NC. Finally, the Neuse River study area extends from the open ocean at Ocracoke Inlet, across the expansive Pamlico Sound, and up through the broad Neuse River estuary. Because of the openness and width of the nearby waterways, the Neuse River is susceptible to large storm surges caused by prolonged wind setup in the area.

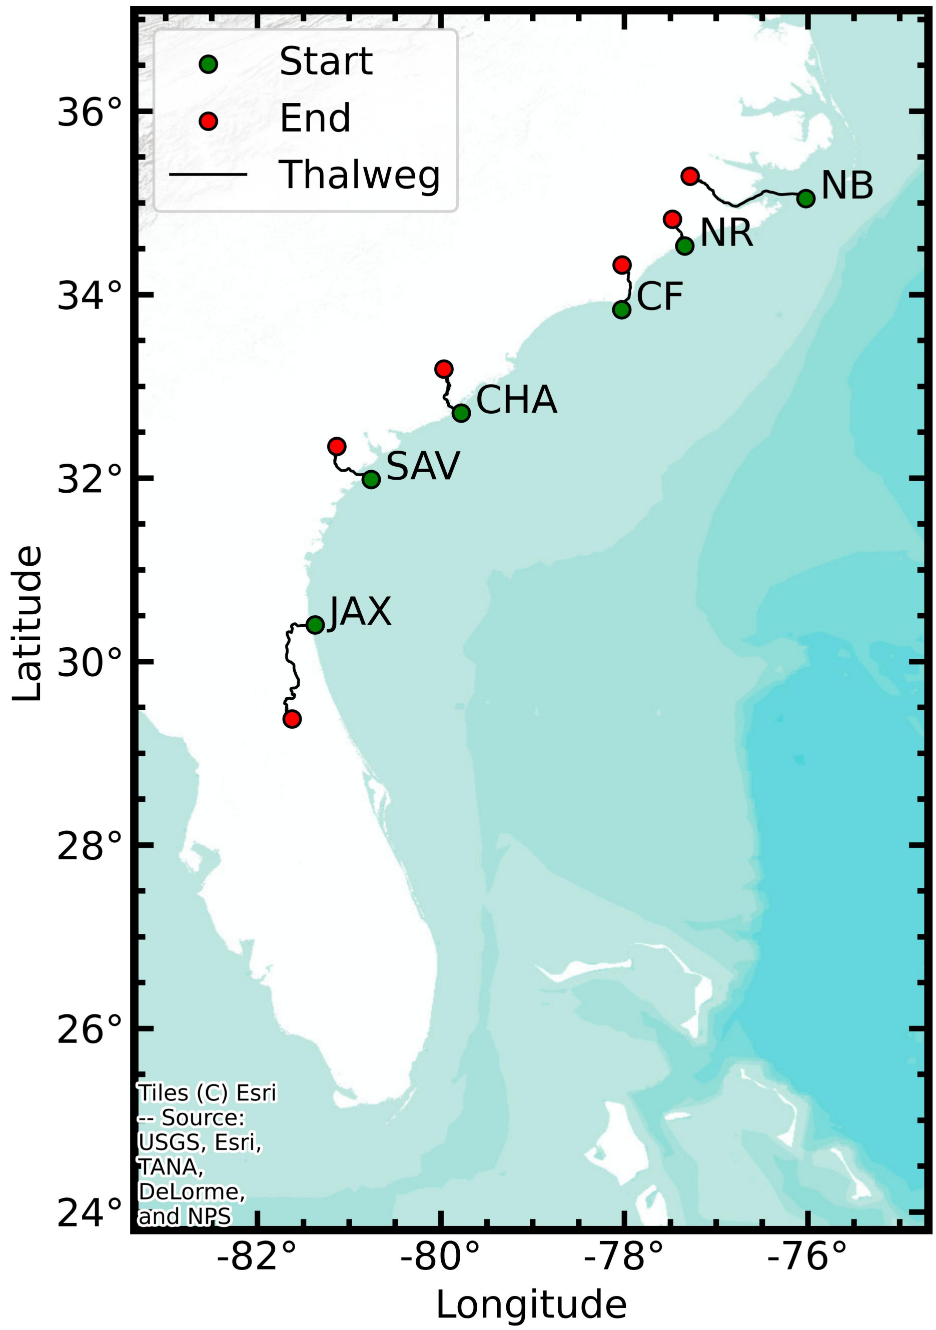

Figure 4.8: Thalweg locations for maximum water level analysis with starting and end points.

Matthew

For Matthew (2016) at the St. Johns River/Jacksonville location (JAX), maximum water levels were fairly consistent from mesh resolution SABv3-60m to SABv3-200m in the conventional simulations. However, for the SABv3-400m mesh, water levels tend to pile up at the channel constrictions located at the 42-km station and the 159-km station (shown as discontinuities at these locations in Figure 4.9, upper left), indicating that flow is not being effectively transmitted through these constrictions. The channel narrows to around 375 m and 370 m at the 42-km and 159-km stations respectively. In the SABv3-1000m mesh, after the first constriction starts near the station at 35 km, flow is unable to pass on to the rest of the river, and the only raised water levels past this station are likely caused by wind setup from kilometers 35 to 131 where flow is again stopped due to another channel constriction. In the subgrid simulations (Figure 4.10), all 5 of the mesh resolutions produced similar maximum water level results throughout the river from kilometer 0 to 197. The SABv3-1000m simulation produced slightly higher water levels from the station at 8 km to 33 km, likely due to incorrect flow connectivity across barrier islands and other flow blocking features in the area. However, after this section, the SABv3-1000m maximum water levels collapses back to that of the higher resolution meshes.

At the Savannah River location (SAV), only the SABv3-60m and SABv3-100m simulations maintain hydraulic connectivity over the full extent of station locations in the conventional simulations. While the SABv3-200m simulation maintains wet cells throughout most of the range, it reaches a choke point at the 48 km station where surge piles up and does not propagate further upstream. At this location, the river becomes significantly shallower since port activities are no longer present, and so the cell averaged interpolation technique used by the conventional model no longer defines a channel at this location in the coarsened simulations. The SABv3-400m and SABv3-1000m simulations lose consistent hydraulic connectivity around the 30 km station, after which maximum water level predictions deviate from the highest resolution simulations. The station lies very close to a strong bend in the river with a large island separating the main channel from a smaller, shallower channel to the south, and although the coarsened meshes can kind of resolve these two channels, it cannot do so in a way that is conducive to flow propagation. In the subgrid simulations, connectivity is maintained from the majority of the channel, but not past the 77 km station in the SABv3-1000m simulation, and the 81 km station in the SABv3-100m simulations. These upstream locations in the river do not have deep water depths represented in the DEM and contain numerous sharp bends and so they are hard to represent even in the subgrid simulations. Up to these points at the end of the channel, all simulation are in agreement.

Along the Cooper River at the Charleston, SC location (CHA), the SABv3-60m through SABv3-400m conventional simulations are able to maintain connectivity along the entire length of the channel. However, the SAVv3-1000m conventional simulation is stopped at the 32 km station due to insufficient model resolution near the channel. For the conventional simulations that did maintain connectivity, although the pattern of maximum water level was maintained across all meshes, the coarser meshes tended to over predict the maximum water level further up the channel. This is likely due to the aliasing of flow blocking features toward the open coast which allows more surge to propagate into Charleston Harbor and into the Cooper River. These flow blocking features include: large jetties at the entrance to Charleston Harbor, the surrounding barrier islands on the north and south side of the inlet, or the numerous islands inside the harbor. The subgrid simulations show a very similar pattern to the conventional, except the coarsest SABv3-1000m simulation maintains connectivity through the channel.

On the Cape Fear River (CF), the conventional SABv3-400m and SABv3-1000m simulations lose hydraulic connectivity at the 45 km station which is where the channel constricts just down stream of Wilmington, NC. The SABv3-60m, SABv3-100m, and SABv3-200m conventional simulation were able to maintain connectivity throughout this channel; however, similar to previous locations, as mesh resolution becomes coarser, flow blocking features that usually confine and stop the surge as it propagates inland are aliased and thus the coarser simulation have higher maximum water levels than the highest resolution simulation. In the Cape Fear River, the Cape Fear peninsula that separates the river from the Atlantic Ocean is fairly narrow in some areas. These narrow areas are aliased by the coarsest resolution simulations which appears as gaps in the conventional model which allow flow to pass through from the ocean into the waterway. The higher maximum water level at the inland extent of the channel in the coarser simulation transfers over to the subgrid simulations since many of the barrier islands and channel constrictions such as the peninsula, barrier islands, and islands within the river itself are aliased by the coarser model and thus have flow propagating past them.

The New River location (NR) is characterized by and extremely narrow channel separating a larger, open river section from the Atlantic Ocean. A steep drop off in maximum water level height can be seen across all conventional simulations as flow moves through this narrow pass. This is especially evident in the SABv3-60m conventional simulation where channel banks and flow blocking barrier islands are properly represented and therefore constrict the flow. From there, the New River goes through a series of sharp turns and constrictions leading up to Jacksonville, NC. After the first constriction at the 6 km station location there is a pile-up and subsequent collapse in maximum water levels in the SABv3-1000m conventional simulation, indicating that this resolution mesh is not sufficient to resolve flow through the channel. Proceeding inland, the SABv3-100m, SABv3-200m, and SABv3-400m conventional simulations become unable to pass flow though the channel constriction at and around the 34 km station where the New River narrows near Jacksonville, NC. Past this point, only the SABv3-60m conventional simulation is able to resolve flow. In spite of this, all of the subgrid simulations were able to connect flow from the open coast all the way past the channel constriction near Jacksonville, NC. However, akin to the previous waterways, the coarsened mesh simulations both with and without subgrid corrections create increased maximum water levels along the main channel due to insufficiently resolving flow blocking features.

Moving to the New Bern and Neuse River thalweg analysis, only the SABv3-1000m conventional simulation was unable to resolve flow past the 135 km station. Apart from that, all of the conventional simulations give a similar representation for how maximum water level undulates from the open coast, across the Pamlico Sound, through the Neuse River Estuary, and up the Neuse River towards New Bern, especially across the SABv3-60m, SABv3-100m, SABv3-200m, and SABv3-400m simulations. These results were fairly similar in the subgrid simulations except that the SABv3-1000m subgrid simulation was able to push water along the majority of the channel. Again, due to the improper representation of the flow blocking Outer Banks Barrier Islands in the coarser simulations, the maximum water levels through the Pamlico Sound in these simulations is noticeably higher since surge is allowed to pass over the islands and into the sound.

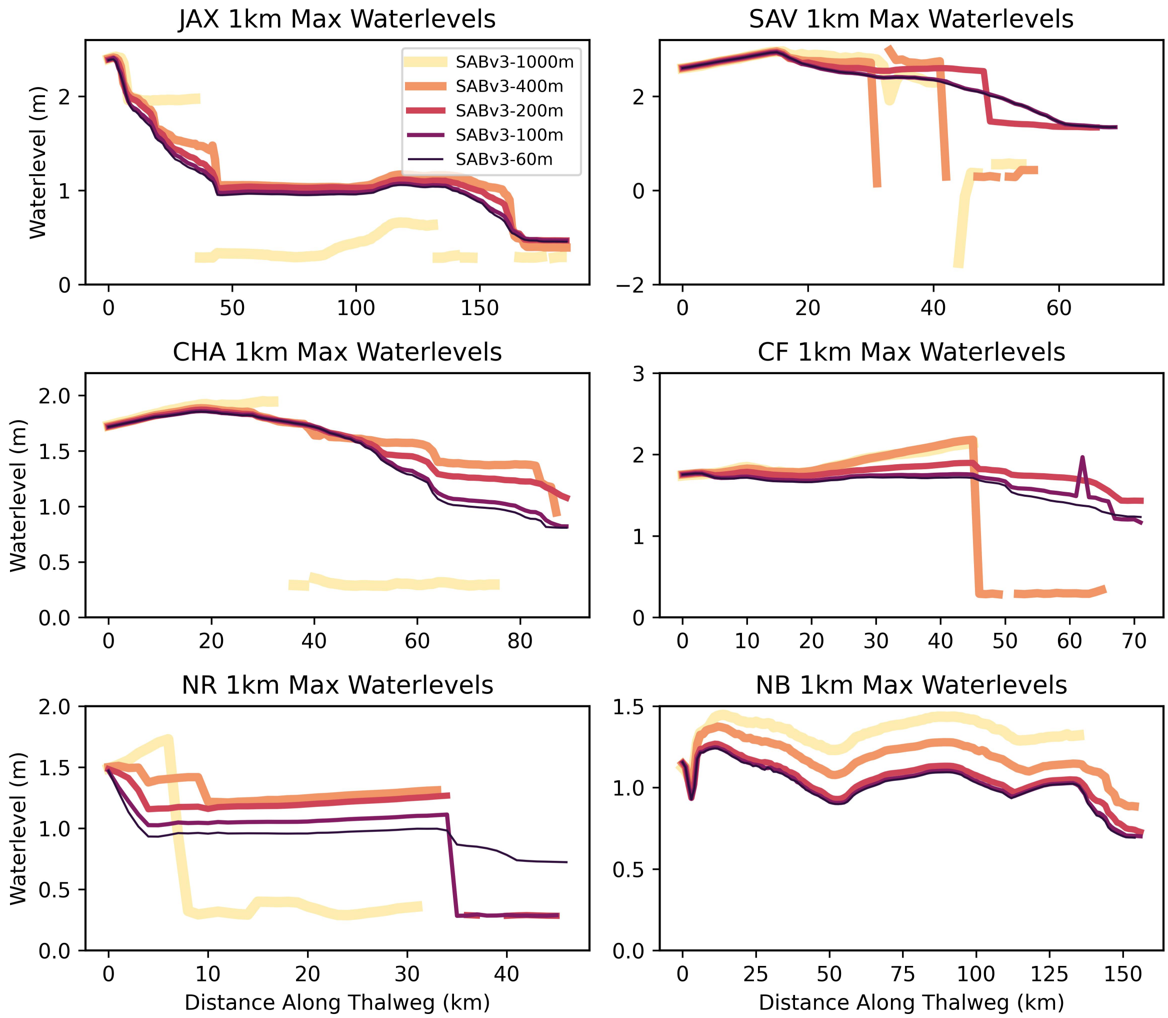

Figure 4.9: Maximum water levels along the thalweg taken from the conventional simulations along each of the 6 thalwegs during Matthew (2016).

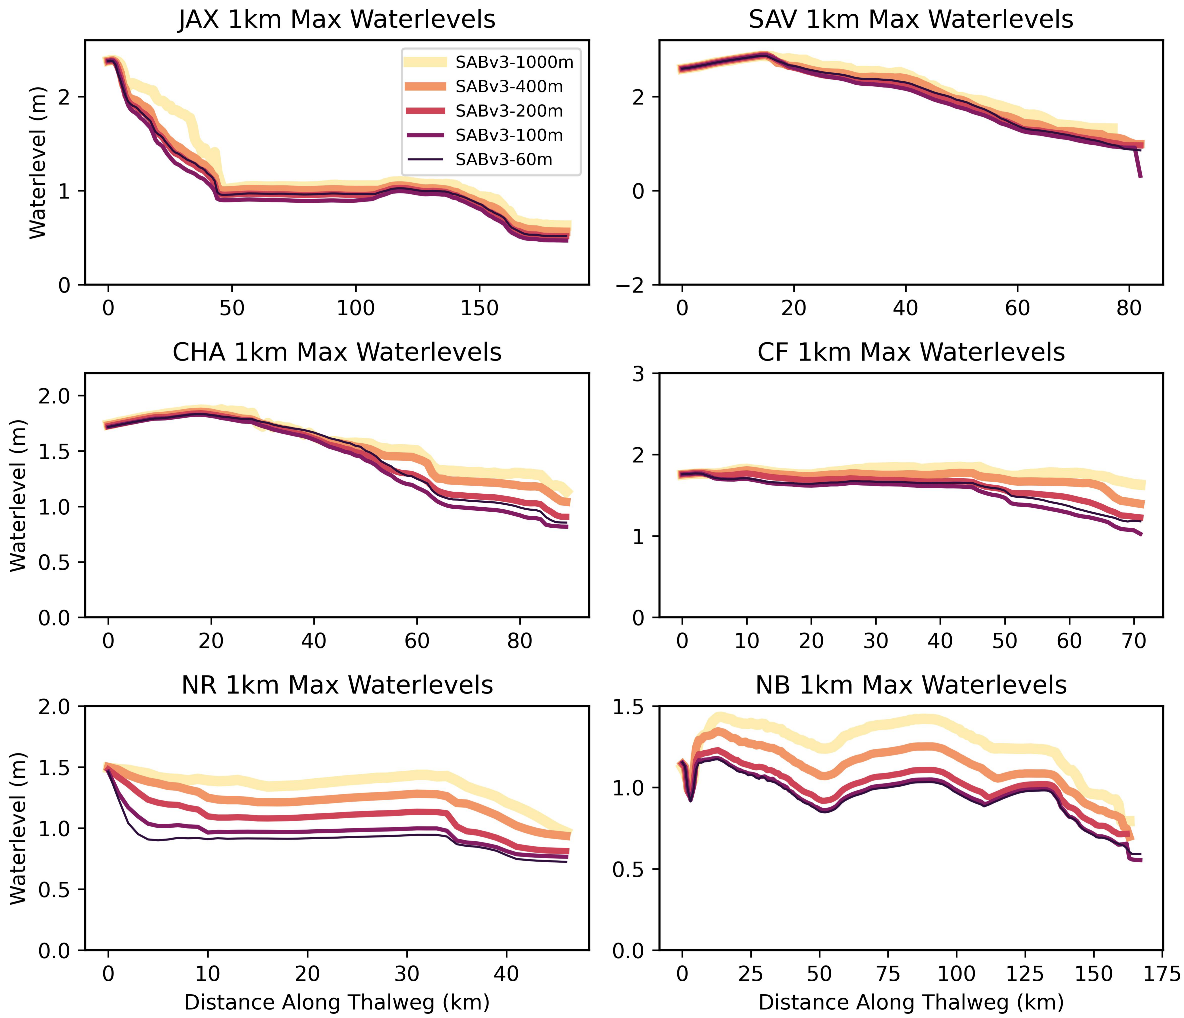

Figure 4.10: Maximum water levels along the thalweg taken from the subgrid simulations along each of the 6 thalwegs during Matthew (2016).

Florence

In the simulations of Florence (2018), at the Charleston, SC (CHA), Cooper River location in the conventional simulations, only the SABv3-60m was able to propagate surge up the entire river. The SABv3-100m conventional goes dry at the 86 km station, and although the SABv3-200m stays wet over the full extent of the channel, water piles up at various choke points where the river narrows or becomes more shallow causing unrealistically high water levels at the last station location. A similar phenomena occurs in the SABv3-400m simulation, and the SABv3-1000m simulation loses hydraulic connectivity at the

32 km point where there is a large bend in the river and the mesh aliases the majority of the channel. Similar to the results from Matthew above, all of the subgrid simulations are able to communicate storm surge from the open coast to the end of the channel. Only the SABv3-400m and SABv3-1000m simulations have water levels build up at around the 60 km station after a series of sharp bends in the river.

For the Cape Fear River (CF) case, only the SABv3-60m conventional simulation had consistent water levels from the open coast, up stream, to the end of the channel stations. The SABv3-100m conventional simulation loses hydraulic connectivity at 66 -km after a series of sharp bends in the river. The SABv3-200m simulation predicts a choke point at 45 -km causing unrealistically high water levels further downstream. At 45 -km there is a fork where the Brunswick River feeds into the Cape Fear. Here, Eagle Island confines the main channel and constricts storm surge propagation. The SABv3-400m and SABv3-1000m conventional are also hindered by this choke point; however, these simulations become completely disconnected from upstream stations. In contrast, all of the subgrid simulations maintained hydraulic connectivity along the full extent of the channel. Similarly to the conventional, the subgrid simulations predicted a slight water pile-up at the 45 km point caused by the channel narrowing closer to Wilmington and being confined by its eastern bank and Eagle Island to the west.

The entrance of the New River (NR) is approximately 65 -km northeast from where Florence made landfall, and so experienced a significant amount of surge during the storm. Fortunately, the entrance of the river is very narrow, so much of the incoming surge got hindered by the neighboring barrier islands on either side of the inlet. This flow constriction is observed in most of the subgrid and conventional simulations as a steep drop-off in predicted water level height as we move from the open coast, through the narrow inlet, and into the back bay. However, in the coarsened conventional simulations there is less of a drop-off in water levels, and in the SABv3-1000m simulation, storm surge piles up so much at the open coast that it bleeds through to the back bay. The pile-up is decreased as mesh resolution increases and water is allowed to pass through the narrow channel resolved by the meshes. Of all the conventional simulations, only the SABv3-60m simulation connects water from the ocean to the end of the station locations. At the 35 -km mark, the New River contracts considerably around Jacksonville, NC, cutting off flow in the SABv3-100m to SABv3-400m simulations. The SABv3-1000m simulation loses most hydraulic connectivity after the first river contraction around the 10 -km station. Although the subgrid simulations experience flow contractions at the 10 -km and 34 -km stations, hydraulic connectivity is maintained to the last channel station in all simulations.

For the New Bern (NB), Neuse River channel, both the subgrid and conventional simulations had very similar water level results until around the 135 -km station. After passing through Ocracoke Inlet, this thalweg traverses the wide open Pamlico Sound and enters the Neuse River Estuary which is several kilometers wide even at its narrowest points. Thus, subgrid corrections are not necessary to resolve flow in this area. However, after the 135 -km station, which lies upstream of New Bern, NC, the channel narrows significantly, and so there is a steep drop off in minimum water levels across all simulations. In the conventional, the SABv3-1000m simulation goes dry at 155 -km, but all other simulations maintain connectivity. It should be noted that the SABv3-60m to SABv3-400m conventional simulations seem to lose some flow connectivity at the very end of the channel indicated by slight upticks in water levels. This means the flood waters may have begun to pile up ahead of choke points past the end of the channel. In the subgrid, all simulations maintained hydraulic connectivity; however, only the SABv3-100m simulation mimicked the results of the SABv3-60m. The coarsened simulations had higher water levels along the entire length of the channel likely due to the aliasing of flow blocking features such as the North Core Banks and Ocracoke Island.

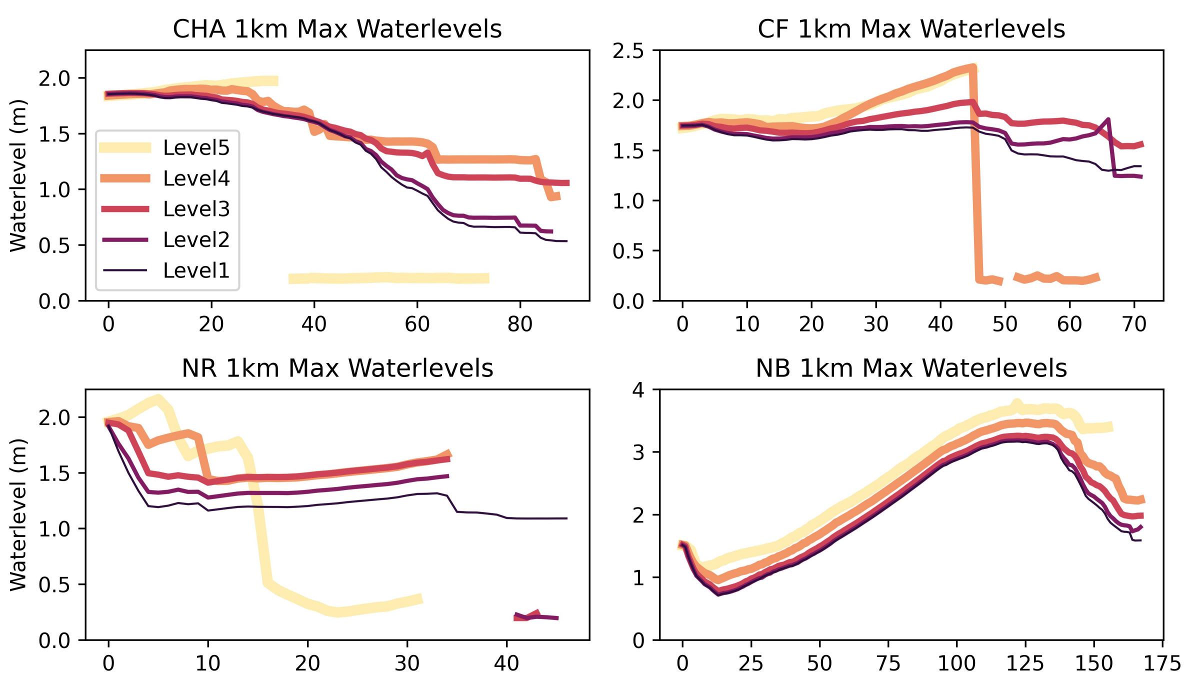

Figure 4.11: Maximum water levels along the thalweg taken from the conventional simulations along each of the 6 thalwegs during Florence (2018).

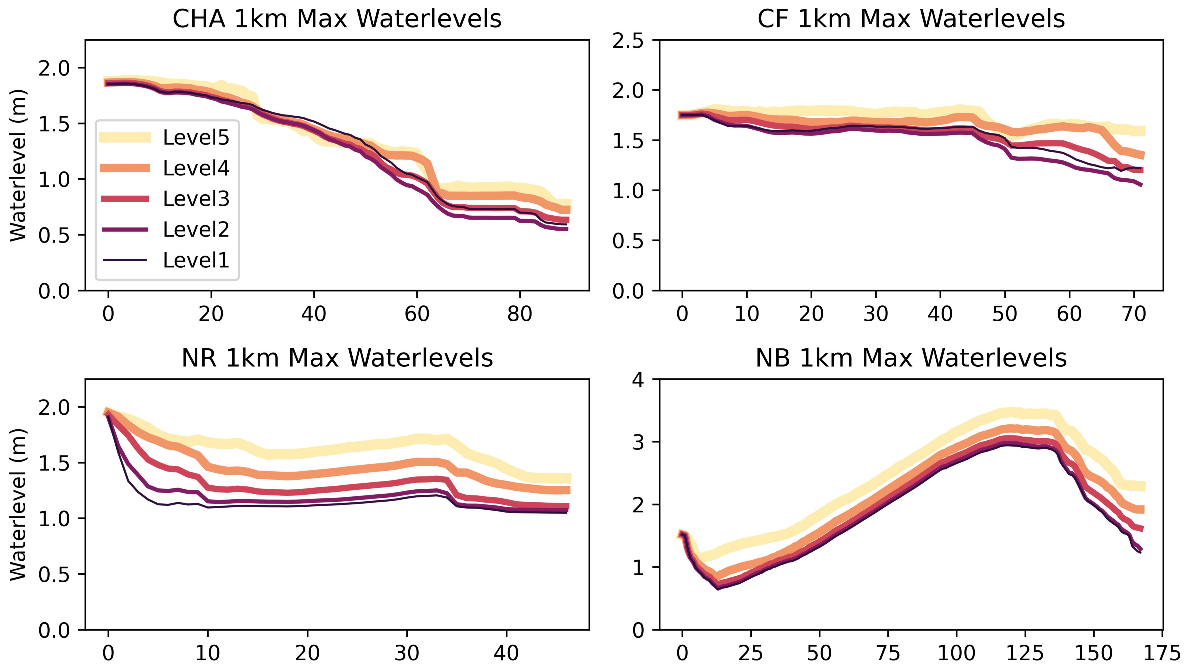

Figure 4.12: Maximum water levels along the thalweg taken from the subgrid simulations along each of the 6 thalwegs during Florence (2018).

4.4.4 Magnitudes and Phases of Flows from Coast to Inland

To investigate how the magnitude and phase of storm surge and tidal propagation changed from the open coast to inland water ways based on mesh resolution, water level hydrographs were extracted from the Beginning (coast), Middle, and End (inland) of three thalwegs used in the previous analysis.

Matthew

For Matthew, water level times series were extracted from the beginning, middle, and end of the St. Johns River, Jacksonville, FL (JAX) thalweg, the Cooper River, Charleston, SC (CHA) thalweg, and the New River, NC (NR) thalweg. The St. Johns River thalweg is by far the longest analyzed. The Beginning and End stations are nearly 200 km apart and there are several large open water bodies connected with relatively narrow channels. Due to its length and geometry, the water level hydrographs produced along this waterway should give a good indication of how subgrid corrections and changes in resolution affect the timing and magnitude of surge propagation over long distances. The Cooper River location begins right outside of Charleston Harbor, which is a wide, relatively shallow estuarine system with many connected waterways. Moving upstream in the Cooper River, the channel gradually narrows and river bends become more radical similar to hair pin turns on roadways. This location also has a fairly significant tidal range and thus will provide a good representation of how well the conventional and subgrid simulations can match phase and amplitude of combined tides and surge over the length of the thalweg. Lastly, the New River location has an extremely narrow, natural inlet opening which is different from the other channels which are deeply dredged and maintained for shipping traffic. This will show us how narrow constrictions at the beginning of the the channel can affect upstream surge propagation in the simulation. These three locations were also chosen because they provide a representative look at how Matthew affected water levels all along the SAB. Conventional and subgrid time series results were extracted for a time period between 0000 UTC 6 October 2016 to 2300 UTC 10 October 2016.

To start, all 'Beginning' station locations were placed at the open coast. Water level time series at these locations should line up very well since only very coarse resolution is required to resolve flow in this region. As such, all predictions line up almost perfectly across all mesh resolutions in both the conventional and subgrid simulations (Figure 4.13 and 4.14). However, at the middle stations, things start to differ between subgrid and conventional and coarse and fine resolution simulations. Although the higher resolution simulations like the SABv3-100m and SABv3-200m behave similarly to the SABv3-60m, in the conventional model the coarsened simulations (SABv3-400m and 1000m) have trouble representing flow, and in the NR location are unable to resolve any water. In the JAX and CHA locations the SABv3-1000m conventional simulation is able to resolve water, but no flow. This is due to constrictions upstream from this location that prevent hydraulic connectivity. The SABv3-1000m subgrid simulation does not have this problem and is able to resolve phase and amplitude fairly well at both locations.

Some of the largest differences between the subgrid and conventional simulations can be seen at the End stations. This makes sense because these stations were located in narrow channels far upstream from the coast with many flow obstacles in between. Looking at the Charleston and New River locations, only the highest resolution simulation resolved realistic flow, while all of the subgrid simulations produced time series results. However, the water levels created by the coarser models lagged in phase with the high resolution simulations and had non-realistic tidal amplitudes for such inland locations. At the JAX End station, the results were largely the same in the conventional and subgrid models since very little flow could make it so far upstream.

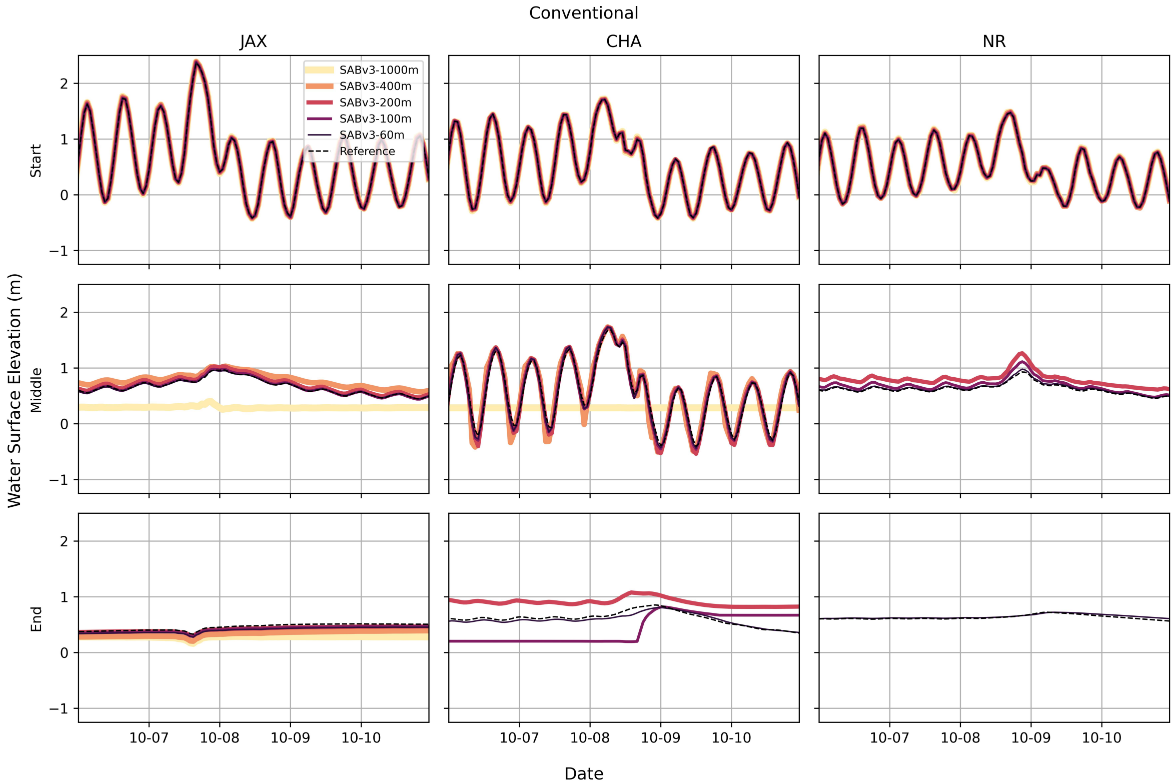

Figure 4.13: Conventional ADCIRC water level hydrographs for synthetic water level stations along the (left) JAX, (center) CHA, and (right) NR thalwegs taken at the (top) start, (middle) middle, and (bottom) end of the thalweg for Matthew (2016).

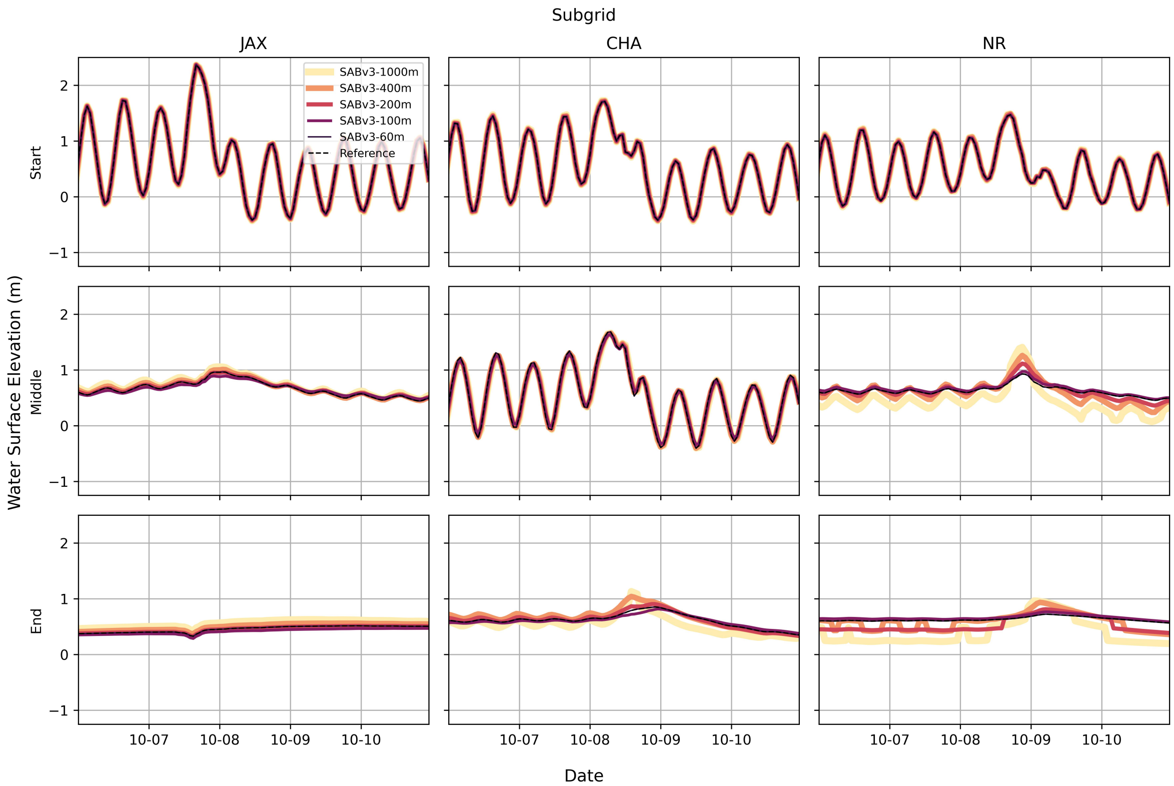

Figure 4.14: Subgrid ADCIRC water level hydrographs for synthetic water level stations along the (left) JAX, (center) CHA, and (right) NR thalwegs taken at the (top) start, (middle) middle, and (bottom) end of the thalweg for Matthew (2016).

Florence

For Florence, water level time series were taken at the beginning, middle, and end of the Cape Fear River (CF) thalweg near Wilmington, NC, the New River (NR) thalweg near Jacksonville, NC, and the Neuse River thalweg, near New Bern, NC (NB). The Cape Fear location was chosen because it has several unique channel characteristics including a deep shipping channel, a wide estuarine system toward the outlet of the river, a narrow barrier island separating the river from the ocean (Cape Fear Peninsula), and a channel constriction upstream near the City of Wilmington, NC. The New River location serves a similar purpose as it did for the Matthew test case in that it has a very narrow inlet separating the ocean from a back bay, and a severe channel constriction towards the end of the thalweg. The New River was also located very close to where Florence made landfall and so experienced very high water levels during the storm. The City of New Bern, NC is located along the

Neuse River estuary and is connected to the open coast by a series of expansive sounds. When Florence stalled over southeastern North Carolina, its winds blew over these open bodies of water and generated extremely high storm surge in the areas near New Bern. The North Carolina Outer Banks separates the inland sounds and estuaries from the open coast, and so it will be interesting to see how the mesh resolution and subgrid corrections affect the ability of these barrier islands to hinder surge propagation. Additionally, all of these thalwegs are located in close proximity to where Florence made landfall and will give a good representative look at how its surge varied along the coast. Water level time series for both conventional and subgrid simulations was predicted from 0000 UTC 12 September 2018 to 0000 UTC 18 September 2018.

Similar to the Beginning stations in the Matthew test case, all of the Florence hydrographs located at the open coast lined up perfectly with the reference solution both in the conventional and subgrid models. Where things start to differ is at the Middle and End stations. These locations are tens of kilometers inland and so simulations need to resolve various flow barriers and hydraulically influencing features along the way. These unresolved barriers and changes in flow appear as flow discontinuities in the coarsened conventional simulations. For example, along the CF thalweg, the SABv3-1000m conventional simulation does not show any water and the SABv3-400m simulation cannot resolve tidal propagation. A similar result can be seen in the NR hydrographs where only the SABv3-60m conventional simulation effectively mimics the flow pattern of the reference solution. However, with an overwhelmingly large storm surge, even the coarsest resolution models can resolve flow at the inland stations, albeit with some flow discontinuity, timing, and amplitude issues. In contrast, we see fairly consistent replication of the reference solution in all of the subgrid simulations, even when the overall storm surge height is not incredibly significant. There are however still issues in the amplitude in the coarsened simulations where the model over and under predicts water levels. This most evident at the End stations where tidal signals are amplified in the CF hydrographs and storm surge is over-predicted by nearly 1 m at the NB End station.

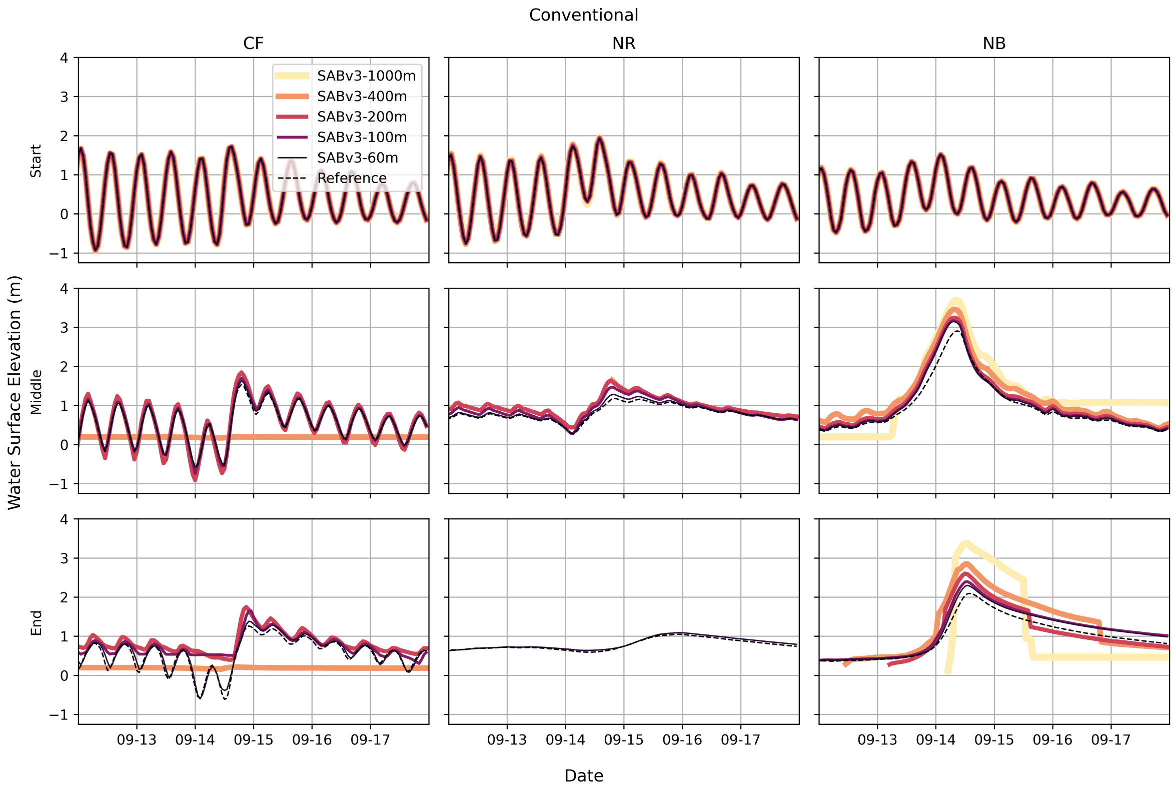

Figure 4.15: Conventional water level hydrographs during Florence for synthetic water level stations along the (left) CF, (center) NR, and (right) NB thalwegs taken at the (top) start, (middle) middle, and (right) end of the thalweg for Florence (2018).

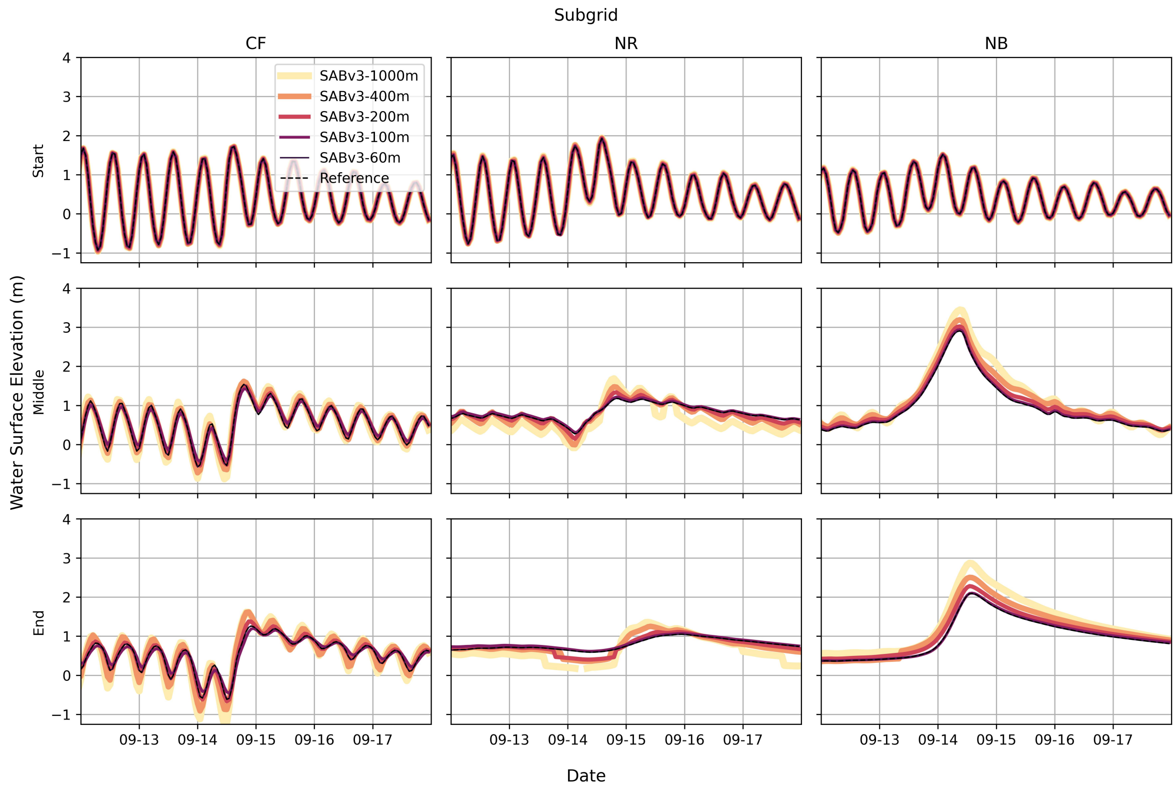

Figure 4.16: Subgrid water level hydrographs during Florence for synthetic water level stations along the (left) CF, (center) NR, and (right) NB thalwegs taken at the (top) start, (middle) middle, and (right) end of the thalweg for Florence (2018).

4.4.5 Maximum Wet Areas in Select Regions

In addition to time series comparisons, the maximum wet area contained within several $2,256.25$ km^2 bounding boxes was calculated for each simulation. For the Matthew simulations, three locations were used surrounding Jacksonville, FL (JAX), Charleston, SC (CHA), and Carteret County, NC (CC). These three separate areas were selected because they both give a good representation of how Matthew affected water levels along the entire SAB, and because each location has a distinct set of features that will affect how subgrid corrections assist model performance. The JAX location has barrier islands, a extensive riverine system, and a deepened port and shipping channel. The CHA location is at the focal point of a large deltaic system with expansive tidal marshes and extremely low lying topgraphy. In North Carolina, the CC location includes a set of narrow barrier islands separating a massive inland lagoon and estuarine system from the Atlantic Ocean. For the Florence simulations, only the Carteret Count, NC location was used because this storm, for the most part, only affected eastern North Carolina. The wet area was calculated because it is a good indication of improved hydraulic connectivity in the model. Wet area was calculated by first downscaling the results onto a DEM with 3.4 m resolution. The total area of the wet cells in each region was then calculated for all subgrid and conventional simulations. It should be noted that although wet area is a good indicator of increased hydraulic connectivity, special attention needs to be paid to make sure the increases in wet areas are realistic.

Matthew

The results of this process for Matthew show an increase in wet area and thereby hydraulic connectivity between the conventional and subgrid simulations, especially at the coarsest resolutions. For the region near Jacksonville, FL (Figure 4.17), an increase in wet area was seen across all mesh levels especially in the coarsest simulation, where the subgrid results give a 26 percent increase in wet area when compared to the conventional simulation. In all simulations, this increase came through improvements in hydraulic connectivity through the St. Johns River, intracoastal water waterways, and smaller tributaries. It should be noted that, although the SABv3-1000m simulation showed the largest increase in wet area, some of these increases are not realistic. For example, this simulation shows intracoastal water ways connected to the ocean via flow across narrow barrier islands (Figure 4.18). This phenomenon is present in the SABv3-60m and other higher resolution simulations, but it is exacerbated in the coarsest simulations since large elements are able to span between hydraulically disconnected areas. This leads to more storm surge propagating inland and thus higher water levels along inland water ways.

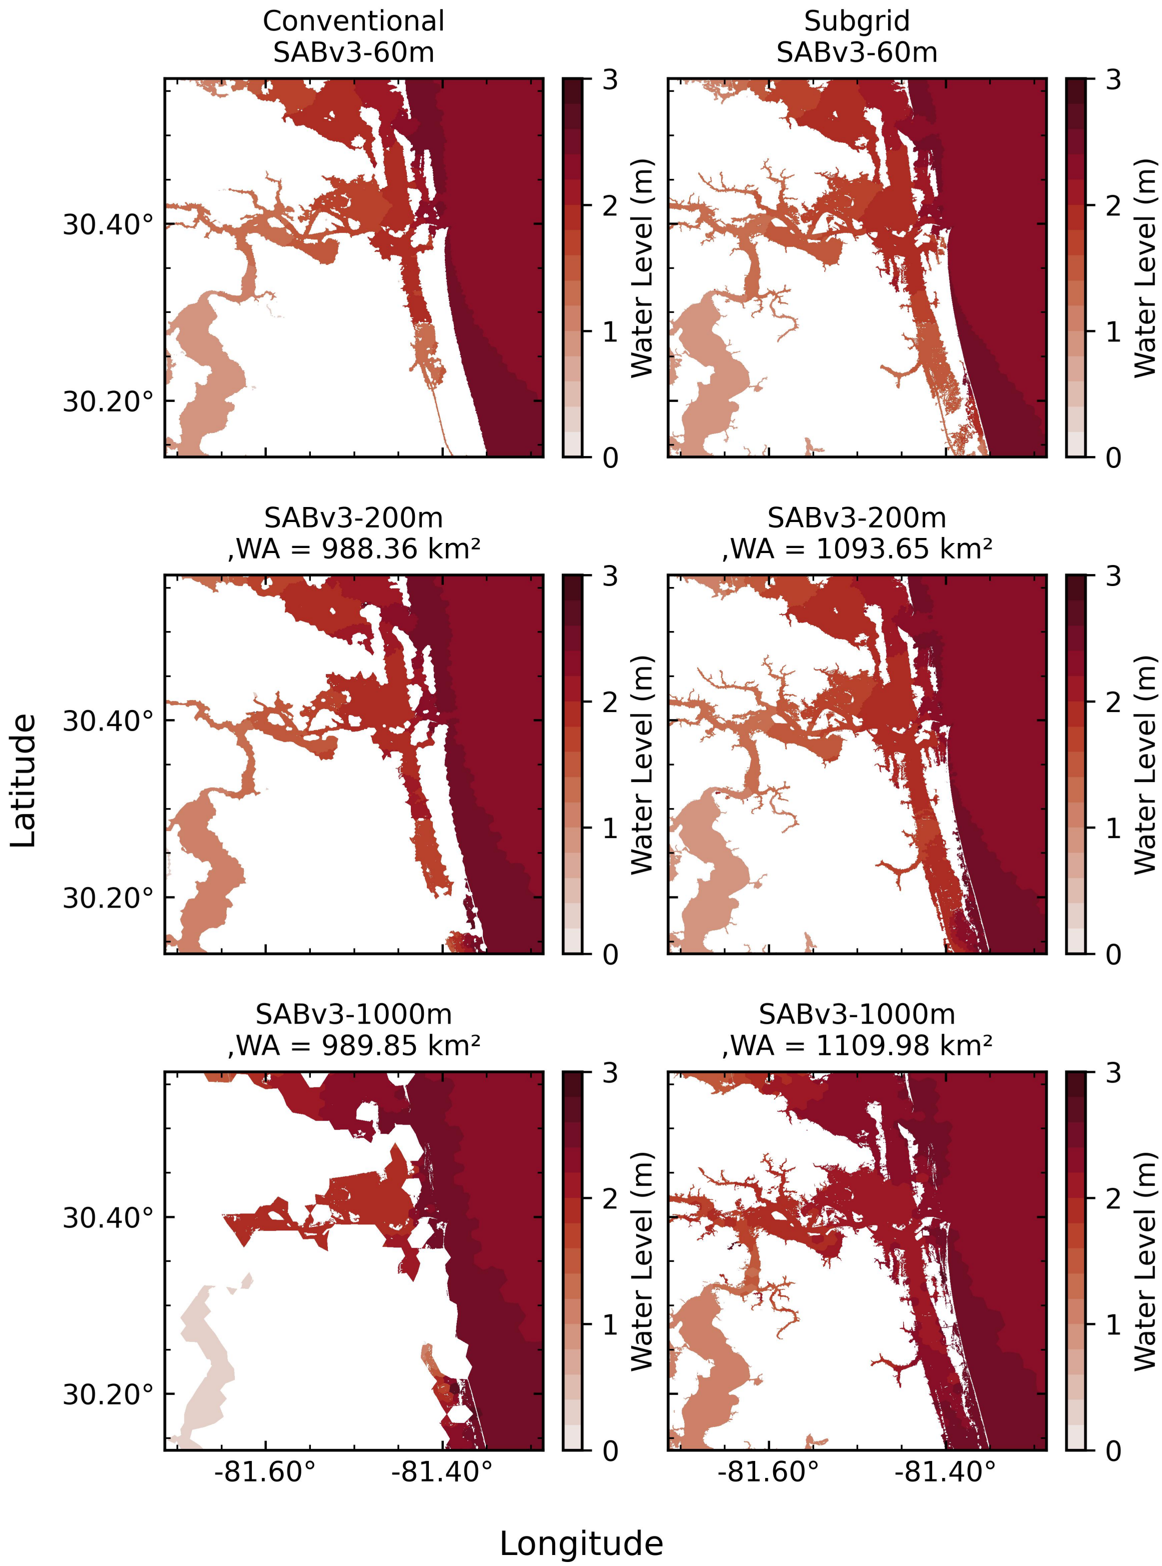

Figure 4.17: Maximum water level simulation results for (left) conventional and (right) subgrid simulations downscaled to high resolution DEM of Jacksonville, FL for Matthew (2016).

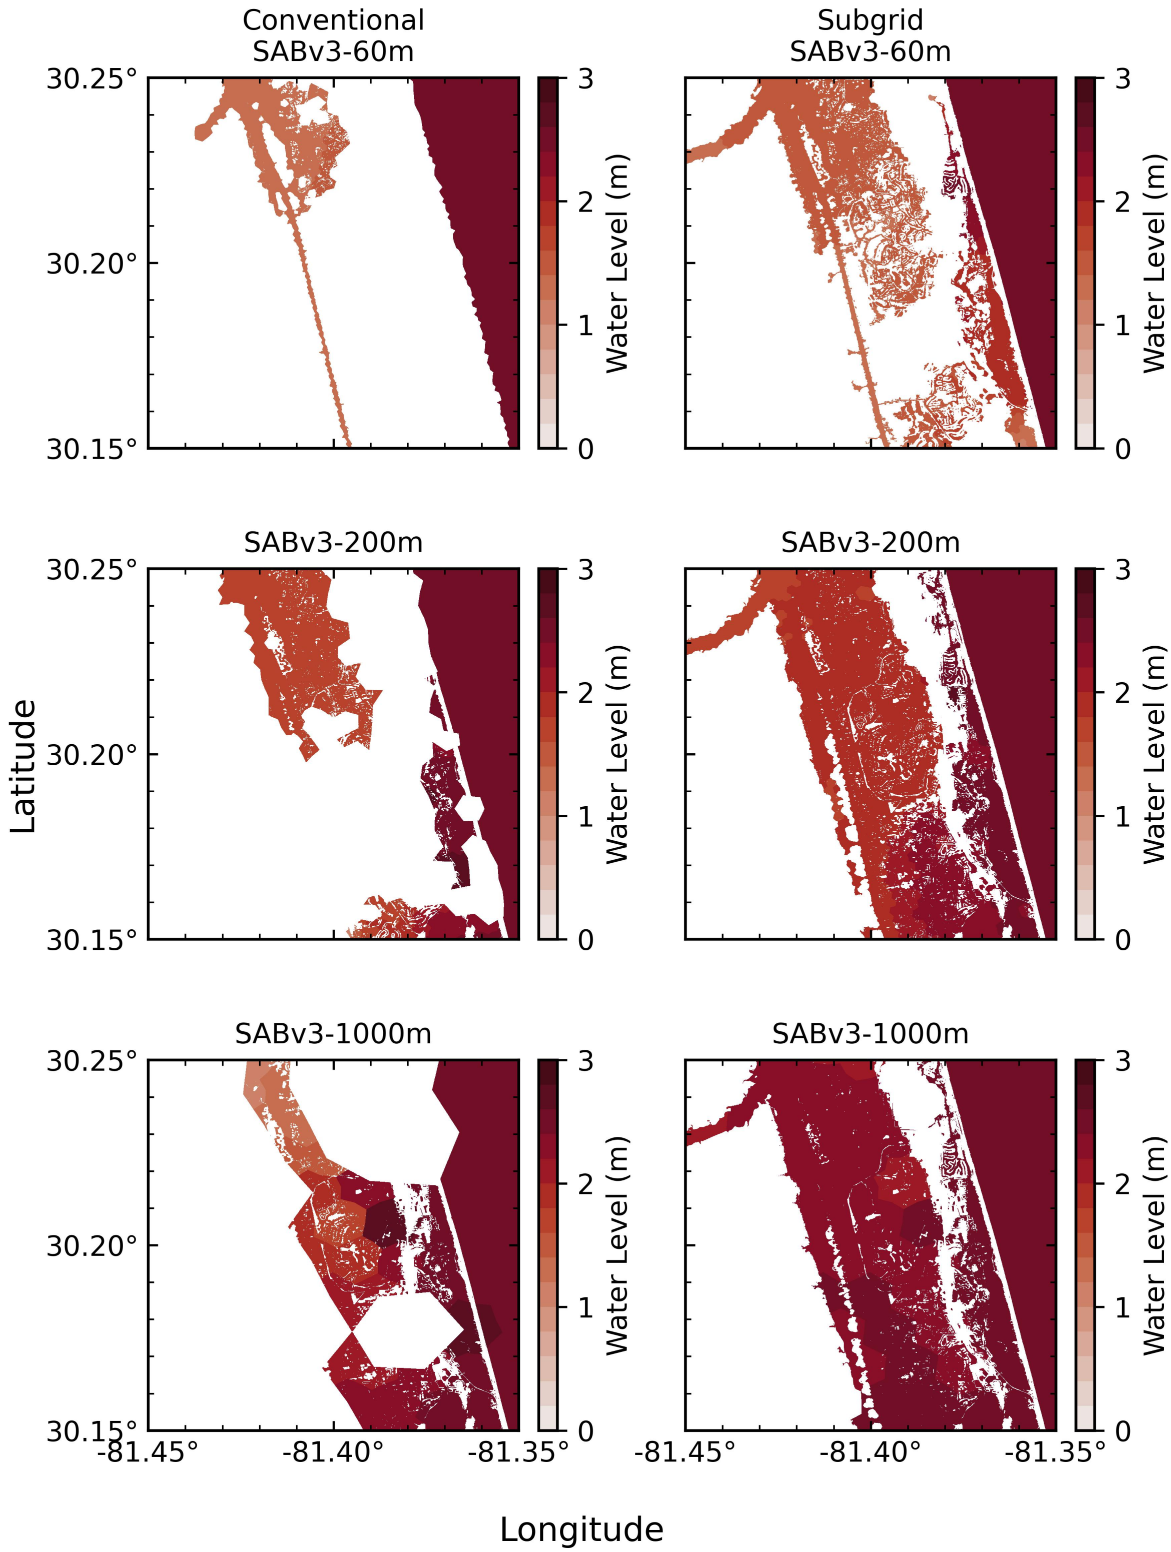

Figure 4.18: Maximum water level simulation results for (left) conventional and (right) subgrid with emphasis on flow connectivity across flow blocking features near Jacksonville, FL for Matthew (2016).

Florence

During Florence there was significant flooding throughout Carteret County, NC, especially along the Neuse River Estuary. This is reflected in the downscaled results of the region (Figure 4.19) where large swaths of the county are inundated even in the coarse resolution conventional simulations. In this test case, something very interesting happened in the subgrid simulations. For the most part the subgrid simulations showed increases in hydraulic connectivity and wet area across most of Carteret County (Table 4.6) including Adams Creek and many of the intricate tidal bays and channels that exist in this area. However, in the eastern section of Carteret County, the subgrid simulations show a reduced wet area when compared to the conventional simuations (Figure 4.20), especially in the higher resolution cases. In this area of Carteret County there is a large farming operation that has drastically changed the natural landscape with elevated areas for planting and irrigation canals in between. This farm covers $57,000$ acres or about one-fifth of Carteret County's total area. The subgrid results suggest that because of these landscape changes, this area does not flood during the storm. Unfortunately, there are no water level gauges located within this farm, and no post-storm imagery to confirm if flooding did or did not occur during the storm.

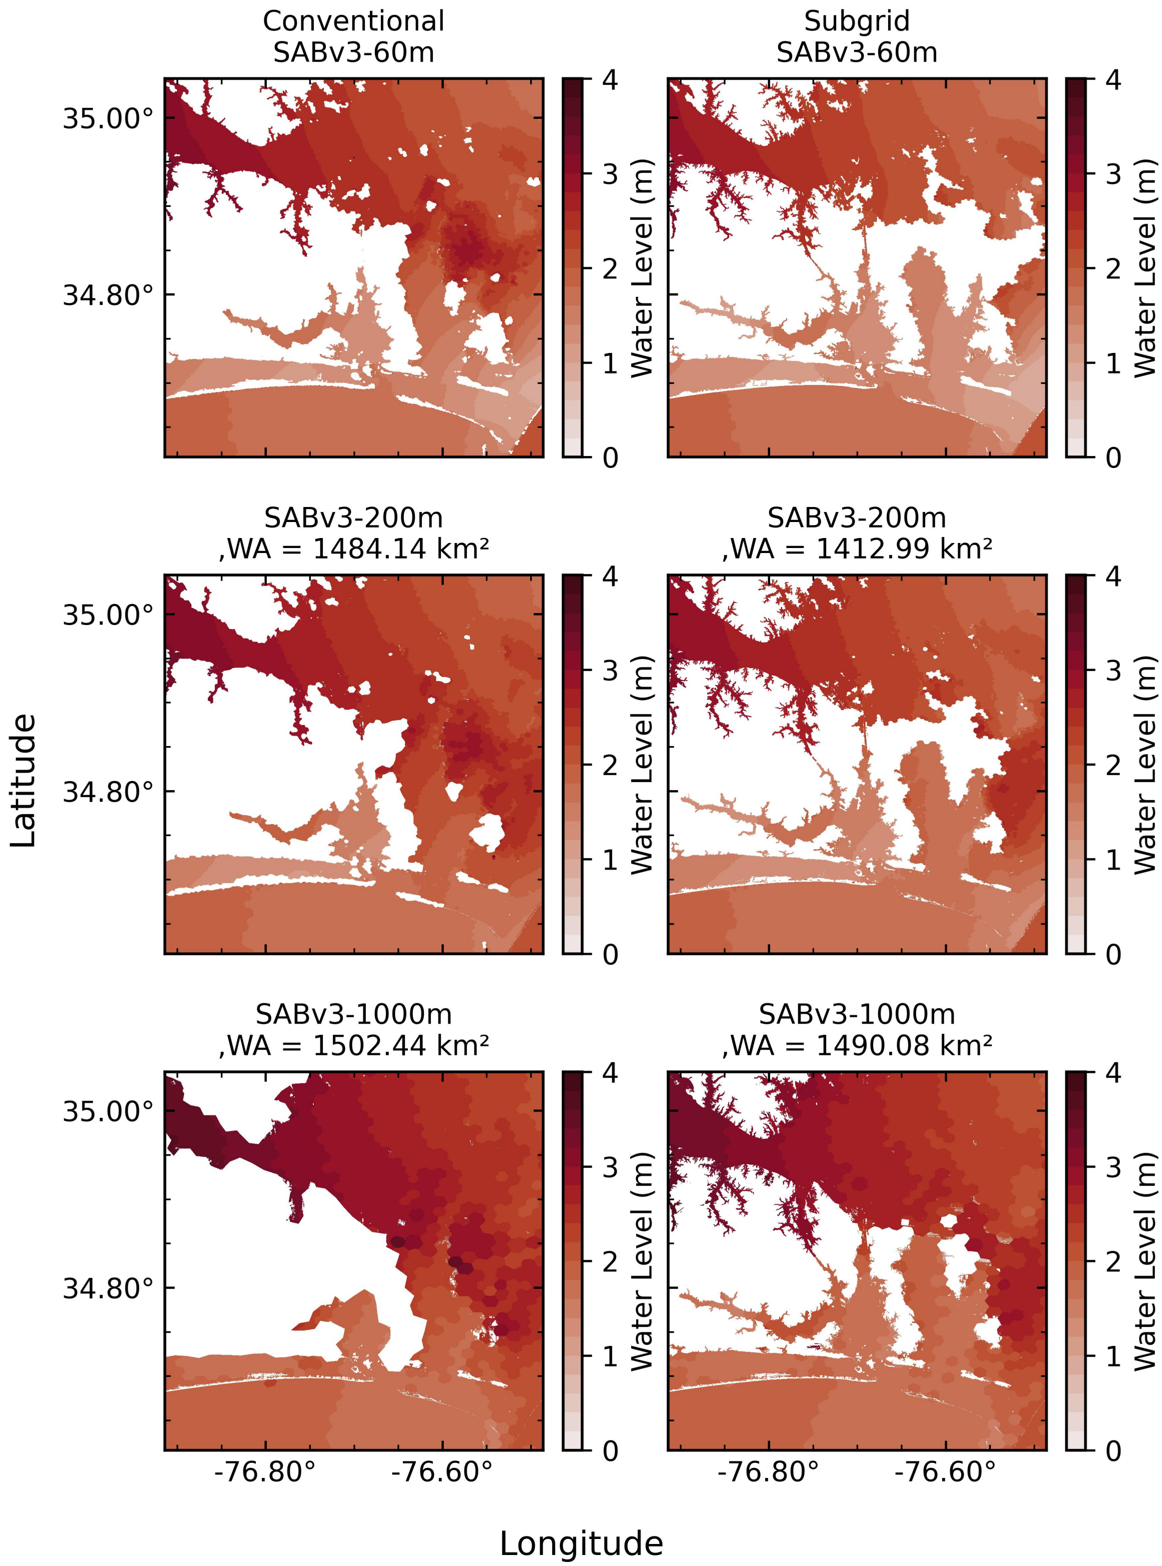

Figure 4.19: Maximum water level simulation results for (left) conventional and (right) subgrid simulations downscaled to high resolution DEM of Carteret County, NC for Florence (2018).

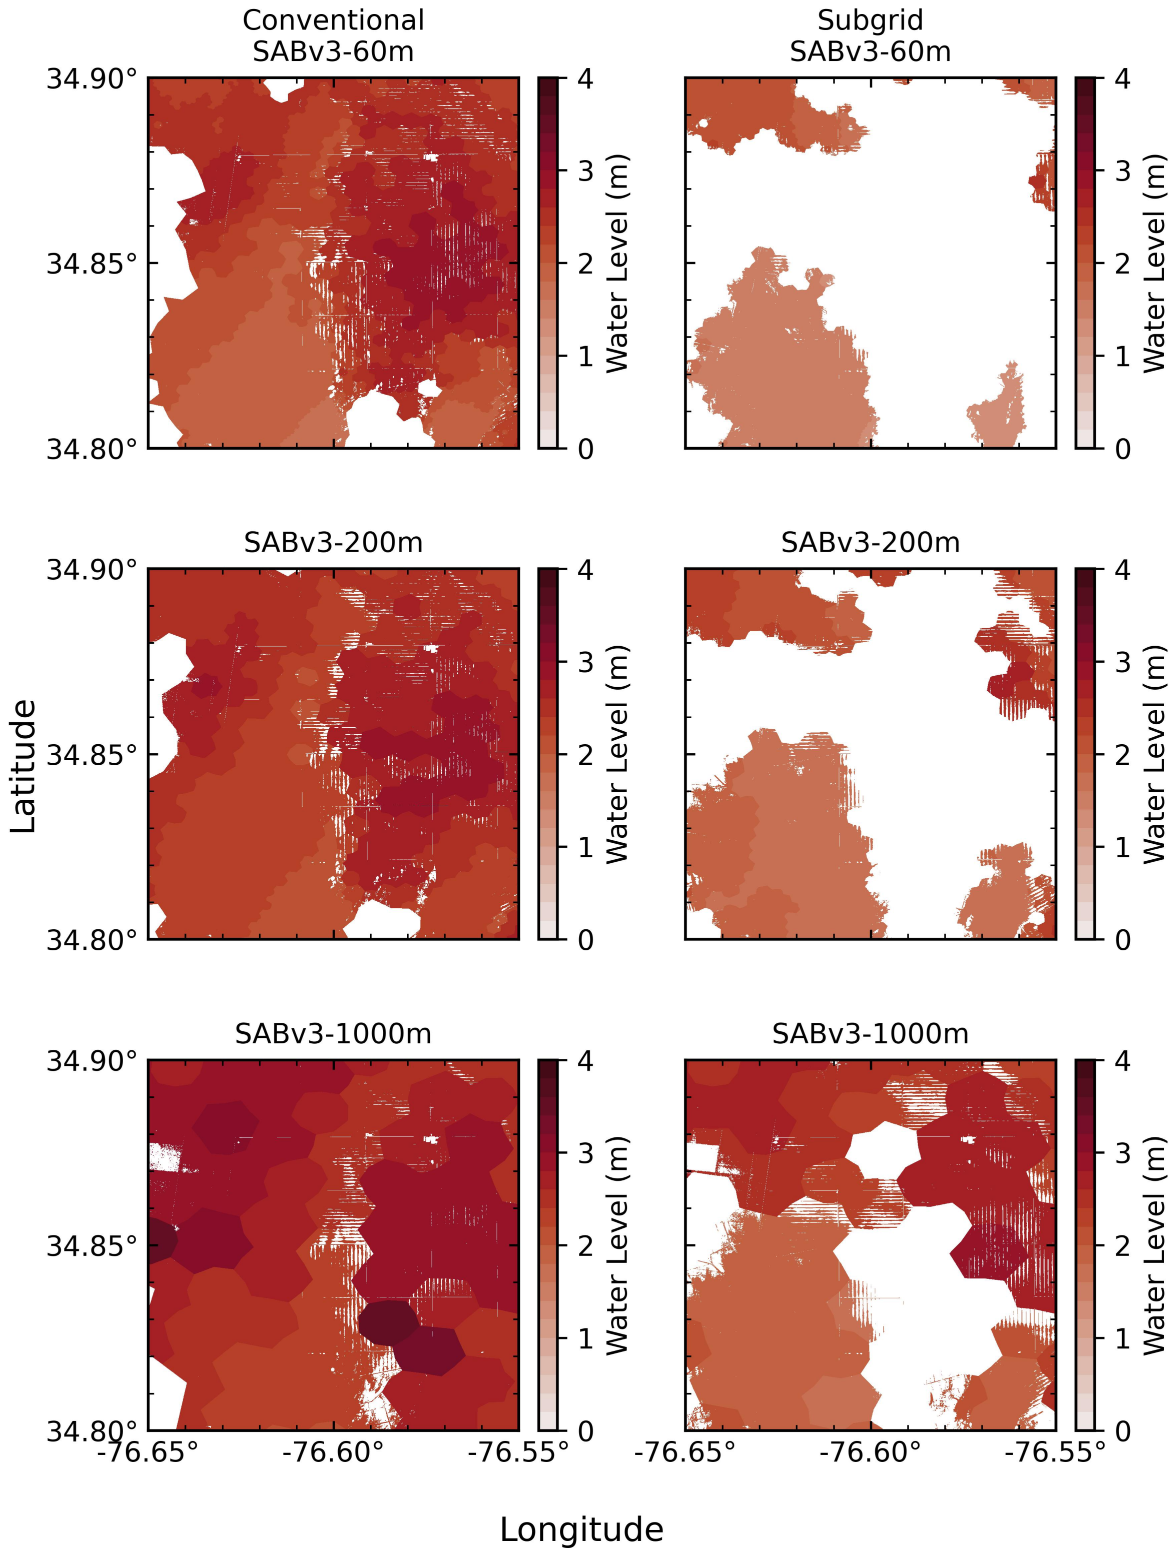

Figure 4.20: Maximum water level simulation results for (left) conventional and (right) subgrid with emphasis on storm surge propagation across low lying farmland within Carteret County, NC for Florence (2018).

Table 4.6: The wet area in km^2 calculated using the maximum water levels during Matthew (2016) and Florence (2018).

Matthew	SABv3-60m	SABv3-100m	SABv3-200m	SABv3-400m	SABv3-1000m
JAX Conventional	999.75 km^2	988.36 km^2	989.85 km^2	972.20 km^2	906.82 km^2
JAX Subgrid	1079.04 km^2	1093.65 km^2	1109.98 km^2	1123.32 km^2	1145.64 km^2
CHA Conventional	870.6 km^2	858.60 km^2	841.27 km^2	791.57 km^2	652.10 km^2
CHA Subgrid	958.86 km^2	952.35 km^2	966.30 km^2	983.53 km^2	999.78 km^2
CC Conventional	1084.56 km^2	1071.80 km^2	1058.84 km^2	1063.18 km^2	1042.31 km^2
CC Subgrid	1150.43 km^2	1153.88 km^2	1185.15 km^2	1231.00 km^2	1268.27 km^2
Florence	**SABv3-60m**	**SABv3-100m**	**SABv3-200m**	**SABv3-400m**	**SABv3-1000m**
CC Conventional	1483.88 km^2	1484.14 km^2	1502.44 km^2	1503.83 km^2	1513.88 km^2
CC Subgrid	1390.39 km^2	1412.99 km^2	1490.08 km^2	1551.58 km^2	1599.82 km^2

The accuracy of results is not the only factor to consider when designing a mesh to predict storm surge. The time it takes to run the model can significantly influence how much resolution can be included in the mesh. To analyze the difference in run times between the conventional and subgrid simulations, each simulation was run in triplicate on 128 cores contained on 2 AMD Epyc "Milan" processors, each processor has 64 cores, and each node has 2 processors with 256 GB of memory and a clock speed of 2.45 GHz. The SABv3-60m simulations were run on the same hardware, but had 1 TB of memory available since the lookup table for this size of mesh was very large. The processors are connected via an Infiniband switch in the Anvil high-performance computing cluster at Purdue University. The minimum wall-clock time was used for timing comparisons.

It was found that the subgrid additions add anywhere between 14% and 80% to the model when compared to the conventional model run on the same mesh. However, simulations with subgrid additions can achieve results with accuracy as good or better on a mesh with nearly 20 times fewer computational cells. And so, the slight loss in efficiency in the subgrid model is more than made up for. That being said, the correct combination of accuracy and computational efficiency needs to be found for the particular modeling scenario being considered and what may be suitable for a design study may not work for a forecasting scenario or vice versa. Interestingly, the SABv3-60m simulation shows a smaller efficiency difference between the conventional and subgrid model than the SABv3-100m, suggesting that the subgrid model scales differently that conventional model when implemented on high-resolution meshes.

Table 4.7: Wall-clock times (sec) for ADCIRC simulations on 128 processors, and ratios of wall-clock times. The minimum time of three simulations was reported.

	Conventional	Subgrid	Runtime Difference (%)
SABv3-60m	58,588	91,657	56%
SABv3-100m	23,984	43,181	80%
SABv3-200m	8,514	15,270	79%
SABv3-400m	4,030	5,921	47%
SABv3-1000m	2,234	2,539	14%

4.5 Discussion

We quantified how varying mesh resolution on ocean-scale domains affected water level predictions for Matthew (2016) and Florence (2018) using maximum water elevation, hydrographs, and wet area summations. It was found that simulations with subgrid corrections can improve model results when running on coarsened meshes by increasing hydraulic connectivity throughout the domain, and by accounting for subgrid-scale variations in bottom roughness and bathymetry. In this section, we will try to understand and quantify just how much the mesh can be coarsened (while still balancing accuracy and efficiency), as well as discuss the implications of our results and how they may provide guidance for future studies.

4.5.1 How coarse is too coarse?

Simulations with subgrid corrections can represent flow through channel constrictions, even when those constrictions were not resolved by the mesh. As long as the underlying DEM resolved the constrictions, flow could pass through the coarsened mesh. But there must be an upper limit – if the elements become too large, then even the subgrid model cannot represent the full complexity and connectivity in the flow. For example, in this study, the coarsened subgrid models often aliased many flow blocking features that separate hydraulically disconnected regions, such as in the area surrounding Jacksonville, FL (Figure 4.18) where flow passed over much of the barrier island connected in the ocean to the intercoastal waterway when in reality there are only a few channels that connect through this area. Other studies found similar issues in their subgrid model including one from Kennedy et al. (2019) where their coarsest simulation could not account for non-continuous flow pathways in the domain.

We attempt to develop a relationship between mesh size and subgrid model skill. This relationship will provide guidance for future subgrid studies by creating a formula to design meshes based on the size and geometry of important bathymetric features. The goal being to describe the feature with maximum element for computational efficiency while also properly describing necessary flow processes to accurate predictions. As an example, the Cooper River near Charleston, SC, is nearly 1000-m wide but, as it extends upstream over several tens of kilometers, it narrows to a width less than 200 m. As it narrows, the channel bends and constricts, which can cause flow connectivity issues in coarsened conventional models. However, since these changes in channel orientation and size are included in the subgrid information, subgrid ADCIRC is able to maintain hydraulic connectivity. To quantify this ability and further show how subgrid ADCIRC can maintain water levels when running on coarsened meshes, we analyzed maximum water depth and the difference in maximum water depth from the reference solution with respect to channel width along the Cooper River near Charleston, SC.

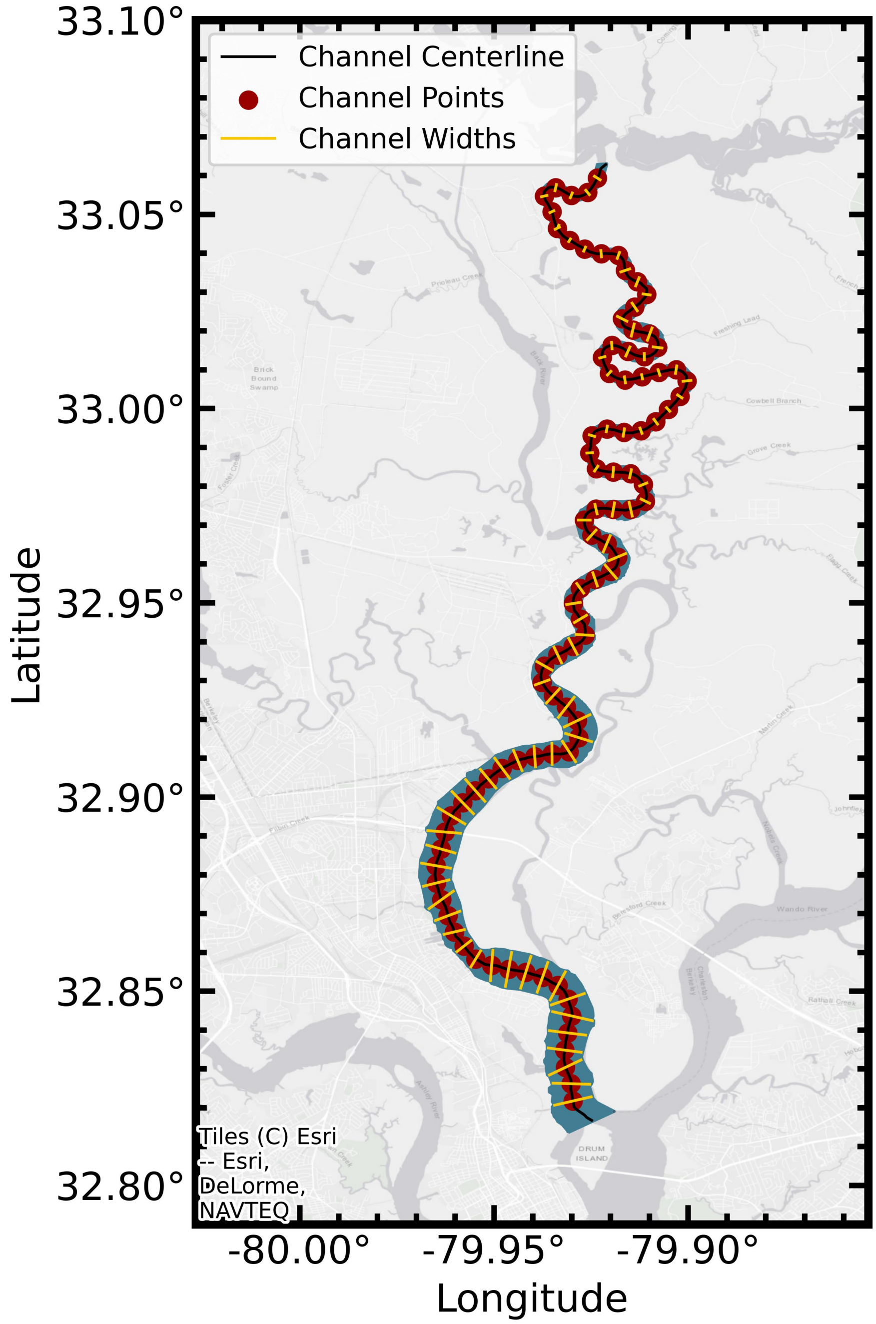

Figure 4.21: Isolated section of the Cooper River, SC, that was used in the channel width analysis.

We first isolated and simplified a section of the Cooper River. Then we created a centerline, placed points at a 0.5-km increment, and calculated the channel width at each of the points (Figure 4.21). We then found the maximum water levels predicted during Matthew at each point along the channel and compared them with channel width (Figure 4.22). In the conventional simulations, channel connectivity is lost by the SABv3-1000m simulation after a sharp channel contraction before the 10-km point indicated by the darker grey color on the left side of the plot. The other simulations begin to diverge from the reference solution as the channel narrows moving inland. Although the coarsened conventional simulations are able to maintain connectivity and generally mimic the maximum water levels of the reference solution, they do not adequately resolve flow processes in the narrower portions of the channel. As the channel narrows to widths smaller than the minimum resolution of the mesh, the chance that those coarsened elements will not have all (or a reduced amount) of their vertices contained within the bounds of the channel increases. In the conventional ADCIRC model, this can lead to a flow disconnect or at the very least reduce the ability of the model to adequately resolve channel geometry, thereby leading to unresolved flow processes as mentioned previously. Along inland portions of the Cooper River, the unresolved flows have an increased maximum water level of 0.32 m and 0.20 m in the SABv3-400m and SABv3-200m simulations at the end of the thalweg.

The subgrid model is able to maintain hydraulic connectivity on all meshes, and its predictions of maximum water levels more closely align with the higher resolution results. However, the subgrid model is not totally immune to the channel contractions and other features that affect flow. This can be seen in increases and decreases in maximum water levels in the model when compared to the reference solution. This is especially apparent at the first channel contraction mentioned previously where the SABv3-1000m shows a noticeable drop-off in water level. Also, in the SABv3-1000m to SABv3-200m there is a pile-up of water towards the last channel constriction located around the 40 km point. The subgrid model uses an averaged representation of bathymetry and water depth to represent flow; therefore, as a channel constricts, the deepest parts of the channel get averaged with shallower parts and surrounding topography. The extent of this averaging is dependent on element size (i.e. larger elements average more extensively). And so, what could be a very narrow but deep channel that is capable of conveying a lot of flow, might be represented as a wider, shallower channel in the model where surge may pile up or fail to proceed upstream. This is seen as an over-prediction of 0.27 m, 0.18 m, and 0.02 m in the SABv3-1000m, -400m, and -200m simulations when the channel width is around 200 to 300 m wide near the 40 km station.

The relationship between maximum water level and channel width shows us that (1) subgrid corrections are necessary in coarsened simulations to hydraulically connect the river from start to end, and (2) in areas further inland where the channel narrows, the subgrid model may need smaller element spacing, not to maintain hydraulic connectivity, but to improve predictions. In this example, the extent to which the subgrid model needs to be refined is proportional to the width of the channel, with the best accuracy to efficiency trade-off coming when mesh resolution is equivalent or slightly less than the channel width as seen in the SABv3-200m simulation.

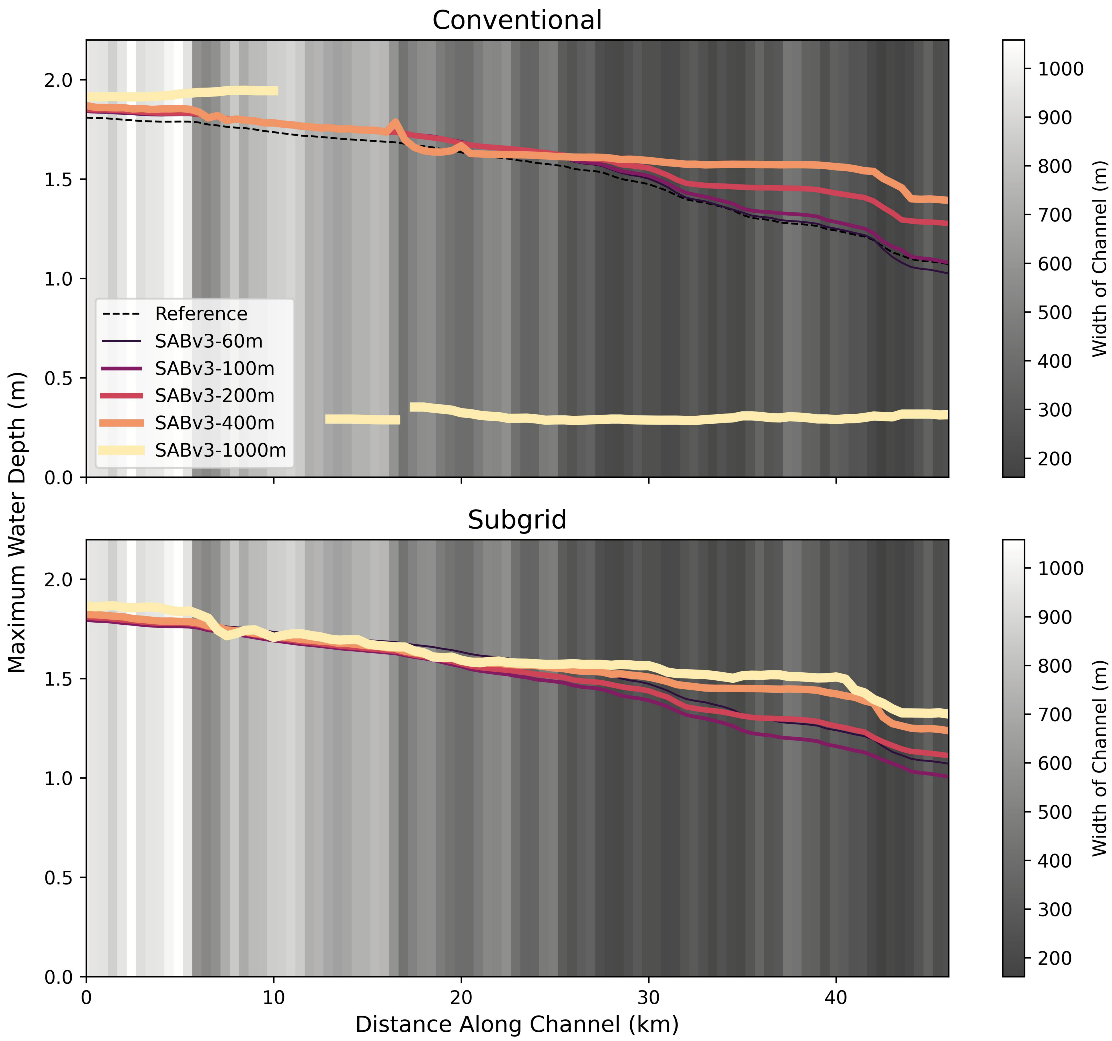

Figure 4.22: Maximum water level along channel centerline for the (top) conventional and (bottom) subgrid simulations compared to channel width. Here, the background colors represent the channel width from (light gray) wide to (dark gray) narrow.

To further investigate how mesh resolution affects flow propagation in narrow channels, the absolute difference in maximum water level compared to the reference solution were plotted against channel width for the conventional and subgrid simulations (Figure 4.23). In the conventional simulations, after the channel width drops below about 800 m, the SABv3-1000m simulation deviates significantly (more than 1 m) from the reference solution. This pattern repeats in the SABv3-400m and SABv3-200m simulations when channel contractions dropped below 350 m. At locations where the channel is about 200 m wide, the difference in maximum water level from the reference solution can deviate by as much as 1.3 m in the SABv3-1000m simulation and 0.25 m in the SABv3-200m simulation. Again, when the width of the channel drops near or below the minimum resolution of the mesh, the conventional model cannot resolve the channel because not as many if any mesh vertices will fall within the channel. In the subgrid simulations, the deviations from the reference solution are much smaller, with very few being greater than 0.25 m, and begin at smaller channel widths. The deviations between coarseness levels are also smaller indicating that the accuracy of results is not as dependent on minimum mesh resolution. These results demonstrate that although the coarsened subgrid simulations do not perfectly replicate the reference solution, they do a much better job than the conventional model, especially in the SABv3-1000m simulation where much of the channel was aliased by the conventional model.

Figure 4.23 helps illustrate that the subgrid model needs significantly less mesh resolution to achieve results of much higher resolution simulations when compared to the conventional model, but the level of coarsening has a limit depending on the hydraulic feature being described. In this particular example, increasing element side lengths to larger than channel widths show a clear drop off in accuracy when compared to the reference. Therefore, for this channel geometry, a minimum element side length size would be around 200 m in this area. This information is very useful when designing meshes, because it helps eliminate unnecessary resolution from being used in areas that do not need it, thereby making the mesh more computationally efficient.

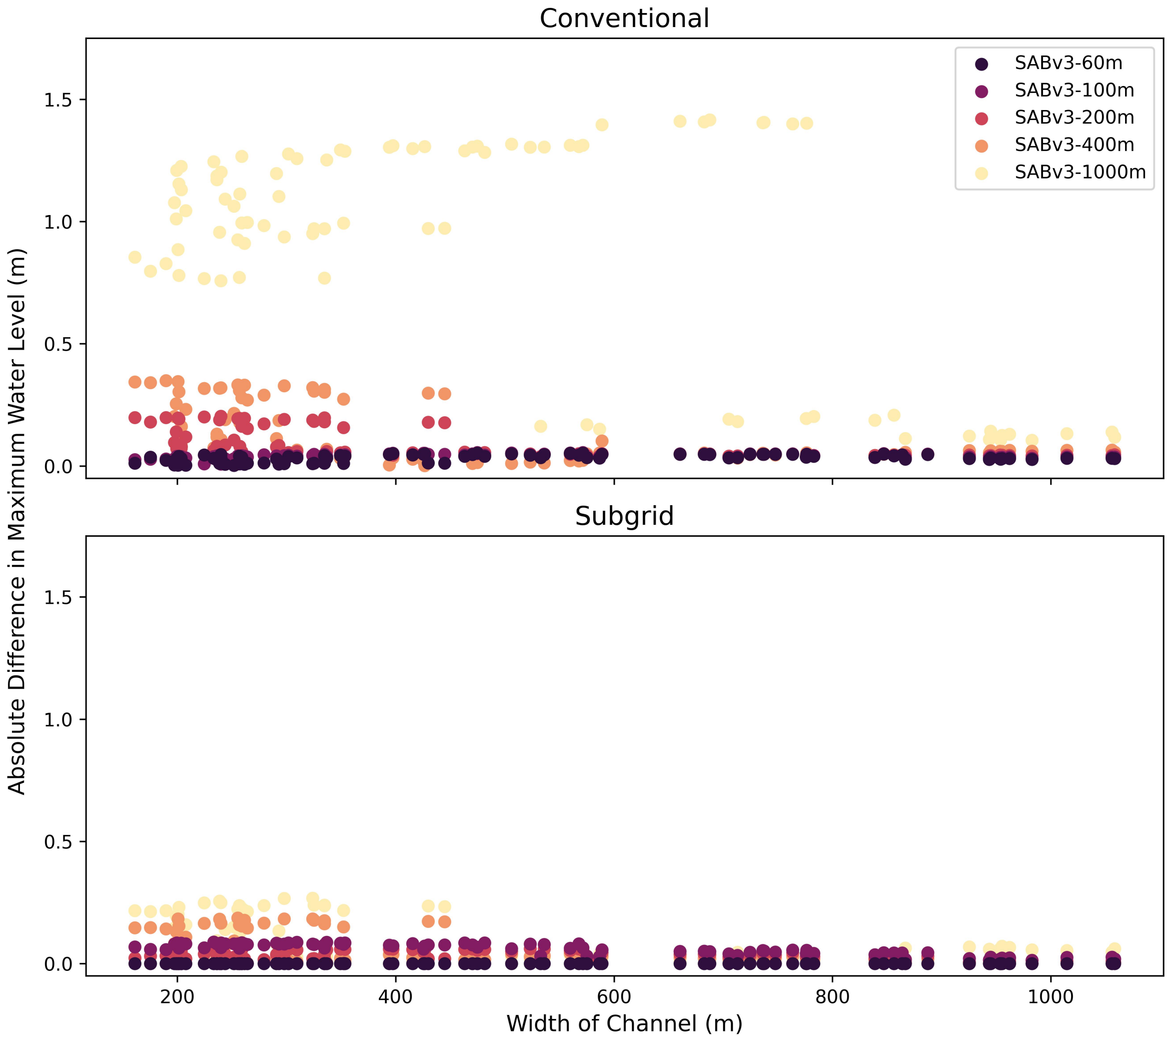

Figure 4.23: Difference in maximum water levels relative to reference solution, for (top) conventional and (bottom) subgrid.

4.5.2 Guidance for Coastal Flooding Applications of Subgrid Models

There is no doubt that subgrid corrections can enhance shallow water flow models when compared to conventional methodologies. This has been found numerous times in previous subgrid studies (Defina 2000; Casulli 2009; Kennedy et al. 2019). However, these previous publications fail to offer guidance to future users of their subgrid model. This includes guidance for the level of coarsening that is possible, time constraints of the model, and how to avoid other pitfalls common in subgrid modeling like aliasing flow blocking features. For example, in (Begmohammadi et al. 2022) the authors mention that implementing higher resolution grids with subgrid corrections would improve results, but they do not explicitly say why or how much the grid would need to be refined. This both limits the applicability of the model and also leads to less than ideal results for other users. To reduce confusion, and preemptively answer future questions about subgrid ADCIRC (as well as other common subgrid models), the authors offer the following guidance.

Before creating a mesh or grid to be used in a subgrid model, time, accuracy, and computational constraints for future simulations need to be taken into consideration. Even with subgrid models, there is still a balance between accuracy and efficiency. This accuracy over efficiency curve is significantly improved in subgrid models when compared to conventional, but still needs to be considered. Ideally, we would use the highest resolution subgrid model to obtain our best possible results; however, this is not always practical and a coarsened grid may offer the best balance. While in this study the SABv3-1000m simulations used a mesh with 20-fold fewer computational nodes, it only maintained an accuracy to within 50 cm of the reference solution. This may be fine for some circumstances, but in others a slightly finer resolution would be required like the SABv3-200m, which overestimated water levels by about 20 cm and still offers a significant decrease in run time. Thus, although there must always be a balance between accuracy and efficiency, our results suggest that mesh resolution can be coarsened by a factor of ≈ 4.5 to achieve a speed-up of 6 times with a minimal effect on accuracy. Although this is not quite as large of an improvement to similar subgrid studies (e.g. a 20-fold decrease in run time, Sehili et al. 2014), many other studies were on much smaller domains and used models developed originally for subgrid corrections.

Another question that should be asked when creating a mesh or grid for a subgrid model is if there are important flow blocking features in the domain that must be resolved. If a feature such as a raised roadway, levee, or narrow barrier island exists in the area of interest and must be resolved to properly represent flow, the authors advise that this feature not have elements that span it entire width. Otherwise, incorrect hydraulic connectivity could occur across this flow barrier. However, this increase in resolution near important structures can easily be offset by dramatic decreases in resolution around areas that still need flow connectivity through subgrid features, but do not have barriers that must be resolved such is a large deltaic marsh system with interconnected tidal channels. For example, if there is a 1000-m wide barrier island (commonly found in areas along the North Carolina Outer Banks) a maximum resolution in these areas should be several times less than the width of the island. Otherwise, it is likely that the subgrid model would allow flow through this flow blocking feature which would in turn create unrealistic high water levels on the sound side of the island. In relation to this study, this would mean using the SABv3-200m mesh instead of the SABv3-400m or -1000m meshes.

Once a suitable mesh resolution is settled on for a subgrid simulation, the subgrid data used to describe important bathymetric features and flow pathways needs to be quality controlled. Subgrid models are only as good as the data used to inform them. Thus, it is essential that accurate bathymetric and landcover data sets are curated. This is especially true for large-scale hurricane storm surge simulations like the ones performed in this study. High-quality data sets are more available now than they have ever been; however, there can be gaps in this high-quality data over a large area. These gaps are often filled with low resolution, low quality data. Poor data set quality can lead to poor performance in subgrid models even with a high-resolution numerical grid since the model relies solely on this data to solve for flow variables (water levels and current velocities). For example, in some estuaries along the SAB where ship traffic in uncommon (meaning there is less frequent dredging and bathymetric surveys) it is common for data sets to have a constant elevation value like 0 or -1 m for all bathymetric cells. This then makes the subgrid think there is an extremely flat, shallow sea floor which is an inaccurate representation of reality, leading to poor model accuracy when compared to *in-situ* observations. So, it is important to check subgrid data sets for areas like these, and fix the data either manually or by merging it with another, more accurate data set.

4.6 Conclusion

It has been shown that mesh resolution does not limit the subgrid model's ability to represent hydraulic connectivity through unresolved subgrid features. Instead, the extent to which subgrid corrections to ADCIRC can improve results compared to conventional methodologies is only limited by the quality and resolution of the DEMs used to calculate the corrections, and the size of the flow blocking features that need to be represented to properly constrict flow.

This dissertation chapter explored the methods used to test the efficacy of subgrid corrections on sequentially coarsened ocean-scale domains. The main contributions of the chapter to the field of subgrid corrections and storm surge modeling are:

1. *Improvements to subgrid lookup table size on ocean-scale domains and implementation of a new wetting and drying scheme allowed for more robust and efficient*

simulations. To accommodate subgrid corrections on high resolution meshes, an additional reduction in the vertex tables was made. This reduced the lookup table to a manageable size for large meshes. The wetting and drying scheme of the subgrid model was also modified to rely on grid-averaged water depth ($\langle H \rangle_G$) instead of the wet area fraction ϕ, which improved the model stability and moved it towards a quasi-volumetric criteria for wetness. This contribution will allow for further domain size extension in future subgrid studies including the movement to global scale water level and flood predictions.

2. *Testing of subgrid corrections on incrementally coarsened ocean-scale domains established guidelines for future subgrid studies.* This study was the first to test an ocean-scale hurricane storm surge subgrid model on successively coarsened model grids. This enabled the authors to give a general guidance on the resolution required by a subgrid model to effectively resolve storm surge in a particular area. It was found that subgrid models still need to represent flow blocking features to prevent unrealistic connectivity, and special care needs to be taken when collecting subgrid data sets to ensure they represent important features in the area of interest.

3. *There is still a balancing between accuracy and efficiency in ocean-scale subgrid models.* The complexity of coastal geometry does not allow for drastic coarsening of mesh resolution (past around a factor of 20) in Subgrid ADCIRC depending on the area. Past this level of coarsening, it is likely that some important feature such as barrier islands could be aliased by the model. Nevertheless, subgrid ADCIRC offers a significant improvement to the accuracy versus efficiency curve over the conventional model.

The systematic analysis of subgrid performance in ocean-scale hurricane storm surge modeling is essential to understanding the proper resolution required when designing a mesh for use in a subgrid model. This knowledge would be critical for a surge forecasting mesh since the coarser a mesh can be made, the faster it can be run. This reduces the wait time between tropical cyclone track and intensity updates and the delivering of flood level results to emergency managers and decision makers in communities that may be affected by the storm.

Future work for this study would be to incorporate cell clones into the model to properly resolve flow blocking structures (Casulli 2019; Begmohammadi et al. 2021), adding other sources of flooding like rainfall, and finely tuning the automated mesh production to have a higher variability in resolution based on model results for a particular area.

CHAPTER

5

SUMMARY

Storm surge is the principal cause of loss of life and damages to property and infrastructure during a tropical cyclone. As such, it is essential that storm surge and coastal flooding are represented accurately in numerical models, both in forecasts prior to a storm's arrival and in long-term planning and design studies. For a numerical model to predict the hydrodynamic processes during a flooding event, it must resolve important bathymetric and topographical features that influence flow such as rivers, tidal channels, raised roadways, marsh platforms, and other built infrastructure. However, not all of these features can be resolved using conventional methods, because they exist on spatial scales that would be computationally expensive to model. In recent years, the use of subgrid corrections has become increasingly popular in numerical models to bridge the gap between computational expense and model accuracy by allowing for subgrid-scale features to be resolved on coarsened numerical grids.

In this dissertation, we expand the use of subgrid corrections into the widely use ADvanced CIRCulation (ADCIRC) hydrodynamic model. We first implemented corrections to wetting and drying on a small, regional-scale study site. Then, we expanded our subgrid storm surge model to ocean-scale domains and included higher-level corrections to bottom friction and advection. And finally, we systematically tested the ocean-scale subgrid model across an array of different meshes with varying degrees of coarseness. The primary contributions of this work to the numerical modeling community are as follows:

1. The theoretical development and implementation of subgrid corrections to wetting and drying in a widely used storm surge and ocean circulation model.

2. The expansion of subgrid corrections to the ocean-scale, and the incorporation of higher-level corrections to bottom friction and advection.

3. The evaluation of subgrid corrections on the ocean-scale using numerical grids designed to resolve flow processes on varying scales.

In Chapter 2, we laid out and discussed the theoretical development and implementation of subgrid corrections into ADCIRC. Here, the governing shallow water equations were averaged and the technique used for defining the nodal and elementary subgrid areas was explained. The ability for the code to allow for partially wet areas required a completely new wetting and drying scheme. The new subgrid ADCIRC was first tested on a synthetic winding channel test case and a small tidal simulation on Buttermilk Bay, Massachusetts. The model was then tested on a regional domain in southwestern Louisiana which encompassed Calcasieu Lake and connected Bayou Contraband. The model was then forced with water levels and winds generated by Rita (2005). Results from a coarse-resolution subgrid simulation were compared to results run with the conventional ADCIRC run on the same mesh and a high resolution mesh. Water levels from pressure gauge sensors were used to validate the results. The implementation work completed in Chapter 2 was the first time subgrid corrections had been added to a widely used hurricane storm surge model with hurricane strength forcing. This study showed that subgrid corrections could allow for accurate results on a significantly coarsened computational mesh, thereby saving a significant amount of computational expense.

Chapter 3 added higher-level subgrid corrections to bottom friction and advection, and extended the implementation of subgrid ADCIRC for use on ocean-scale domains. The higher level corrections for bottom friction and advection allowed the model to better account for small-scale changes to bottom roughness and elevation, which can significantly affect flow. This is especially true during a tropical cyclone where bottom friction is often the dominant force controlling how far storm surge propagates inland. The expansion of subgrid corrections to ocean-scale numerical grids is essential to properly resolving the majority of flow processes involved in a tropical cyclone flooding event, including astronomical tides and the evolution of storm surge from its origins in the open ocean to its build up on the continental shelf and finally its progression into bays and estuaries. Increasing the size of the domain also allowed us to move the boundaries far away from the area of interest, thereby decreasing the chance of boundary effects influencing water level results. This study was the first to develop an ocean-scale subgrid simulation and required careful memory management and subgrid lookup table design.

Although the work done in Chapters 2 and 3 was a leap forward towards applying subgrid modeling capability to tropical cyclone storm surge models, there are still several gaps left in the research. In the previous chapters, only two mesh resolutions were used to show how subgrid corrections could increase model accuracy while running on coarsened meshes. This is suitable for proof-of-concept studies; however, to understand how subgrid model results are affected by decreases in resolution a more systematic approach needs to be used. Therefore, in Chapter 4 we created five ocean-scale meshes of different resolution to analyze how the results of the subgrid model changed with incremental decreases in mesh resolution. It was found that, although there were decreases in accuracy in the subgrid model as resolution was decreased, hydraulic connectivity was maintained as long as the underlying DEM used to create the subgrid lookup tables resolved the necessary hydraulic features. In addition to this finding, the work done in Chapter 4 emphasized the importance of resolving flow blocking features to obtain accurate and realistic model results. This chapter left further work to be done in improving unrealistic hydraulic connectivity in the subgrid model caused by elements spanning flow blocking features such as barrier islands, roadways, and raised land masses separating water bodies.

The work completed in this dissertation made contributions to the fields of subgrid modeling and storm surge modeling. The implementation of subgrid corrections into a widely used, ocean-scale, hurricane storm surge model had never been done before, and required innovative numerical techniques and memory management. Introducing subgrid corrections into the field of hurricane storm surge modeling will contribute to improvements in wall clock times and water level prediction accuracy in future flood hazard studies and forecasting efforts.

REFERENCES

Akbar, M. K., Kanjanda, S., and Musinguzi, A. (2017). Effect of Bottom Friction, Wind Drag Coefficient, and Meteorological Forcing in Hindcast of Hurricane Rita Storm Surge Using SWAN + ADCIRC Model. *Journal of Marine Science and Engineering*, 5.

Asher, T. G., Luettich, R. A., Fleming, J. G., and Blanton, B. O. (2019). Low frequency water level correction in storm surge models using data assimilation. *Ocean Modelling*, 144.

Atkinson, L. P. and Menzel, D. W. (1985). *Introduction: Oceanography of the Southeast United States Continental Shelf*, volume 2. American Geophysical Union.

Bacopoulos, P. and Hagen, S. C. (2017). The intertidal zones of the South Atlantic Bight and their local and regional influence on astronomical tides. *Ocean Modelling*, 119:13–34.

Bates, P. D. and Hervouet, J. M. (1999). A new method for moving-boundary hydrodynamic problems in shallow water. *Proceedings of the Royal Society A*, 455:3107–3128.

Begmohammadi, A., Wirasaet, D., Poisson, A., Woodruff, J. L., Dietrich, J. C., Bolster, D., and Kennedy, A. B. (2022). Numerical extensions to incorporate subgrid corrections in an established storm surge model. *Coastal Engineering Journal*.

Begmohammadi, A., Wirasaet, D., Silver, Z., Bolster, D., Kennedy, A. B., and Dietrich, J. C. (2021). Subgrid surface connectivity for storm surge modelling. *Advances in Water Resources*, 153.

Bender, C. (2013). Georgia and Northeast Florida Coastal Storm Surge and Mapping Study. `https://fsbpa.com/2013TechPresentations/Bender.pdf`. [Retrieved 23 April 2020].

Bender, C. (2014). East Coast Central Florida Coastal Storm Surge and Mapping Study. `http://www.floods.org/Files/Conf2014_ppts/E7_Bender.pdf`. [Retrieved 23 April 2020].

Bender, C. (2015). South Florida Coastal Storm Surge and Mapping Study. `http://www.floods.org/Files/Conf2015_ppts/A7_Bender.pdf`. [Retrieved 23 April 2020].

Berenbrock, C., Mason, R., and Blanchard, S. (2008). Mapping hurricane rita inland storm tide. In *4th International Symposium on Flood Defence*.

Bilskie, M. V., Coggin, D., Hagen, S. C., and Medeiros, S. C. (2015). Terrain-driven unstructured mesh development through semi-automatic vertical feature extraction. *Advances in Water Resources*, 86:102–118.

Bilskie, M. V. and Hagen, S. C. (2013). Topographic accuracy assessment of bare earth lidar-derived unstructured meshes. *Advances in Water Resources*, 52:165–177.

Bilskie, M. V., Hagen, S. C., and Medeiros, S. C. (2020). Unstructured finite element mesh decimation for real-time hurricane storm surge forecasting. *Coastal Engineering*, 156:103622.

Blain, C. A., Cambazoglu, M. K., Linzell, R. S., Dresback, K. M., and Kolar, R. L. (2012). The predictability of near-coastal currents using a baroclinic unstructured grid model. *Ocean Dynamics*, 62(3):411–437.

Blain, C. A., Westerink, J. J., and Luettich, R. A. (1994). The influence of domain size on the response characteristics of a hurricane storm surge model. *Journal of Geophysical Research*, 99:467–479.

Blain, C. A., Westerink, J. J., and Luettich, R. A. (1998). Grid convergence studies for the prediction of hurricane storm surge. *International Journal for Numerical Methods in Fluids*, 26:369–401.

Blanton, B. O. and Luettich, R. A. (2008). North Carolina Coastal Flood Analysis System: Model Grid Generation. Technical Report TR-08-05, Renaissance Computing Institute.

Blanton, B. O., McGee, J., Fleming, J. G., Kaiser, C., Kaiser, H., Lander, H., Luettich, R. A., Dresback, K. M., and Kolar, R. L. (2012). Urgent computing of storm surge for North Carolina's coast. In *Proceedings of the International Conference on Computational Science*, volume 9, pages 1677–1686.

Blanton, B. O., Werner, F. E., Seim, H. E., Luettich, R. A., Lynch, D. R., Smith, K. W., Voulgaris, G., Bingham, F. M., and Way, F. (2004). Barotropic tides in the South Atlantic Bight. *Journal of Geophysical Research*, 109:1–17.

Blumberg, A. F. and Mellor, G. L. (1983). Diagnostic and Prognostic Numerical Circulation Studies of the South Atlantic Bight. *Journal of Geophysical Research Oceans*, 88:4579–4592.

Borns, P. (2022). Did Lee County's 'Just-in-Time' evacuation protocol for Hurricane Ian cost lives? `https://wusfnews.wusf.usf.edu/weather/2022-10-25/did-lee-county-just-in-time-evacuation-protocol-hurricane-ian-cost-lives`. [Retrieved 17 April 2023].

Bradford, S. F. and Sanders, B. F. (2002). Finite-Volume Model for Shallow-Water Flooding of Arbitrary Topography. *Journal of Hydraulic Engineering*, 128.

Bucci, L., Alaka, L., Hagen, A., Delgado, S., and Beven, J. (2023). Tropical Cyclone Report for Hurricane Ian, 23 September – 30 September 2022. Technical report, National Hurricane Center.

Bunya, S., Dietrich, J. C., Westerink, J. J., Ebersole, B. A., Smith, J. M., Atkinson, J. H., Jensen, R. E., Resio, D. T., Luettich, R. A., Dawson, C. N., Cardone, V. J., Cox, A. T., Powell, M. D., Westerink, H. J., and Roberts, H. J. (2010). A high-resolution coupled riverine flow, tide, wind, wind wave and storm surge model for southern Louisiana and Mississippi: Part I – Model development and validation. *Monthly Weather Review*, 138:345–377.

Burrough, P. A. and McDonnell, R. A. (1998). *Principles of geographical information systems*. Oxford, New York: Oxford University Press.

Casulli, V. (2009). A high-resolution wetting and drying algorithm for free-surface hydrodynamics. *Numerical Methods in Fluids*, 60:391–408.

Casulli, V. (2019). Computational grid, subgrid, and pixels. *International Journal for Numerical Methods in Fluids*, 90:140–155.

Casulli, V. and Stelling, G. S. (2011). Semi-implicit subgrid modelling of three-dimensional free-surface flows. *Numerical Methods in Fluids*, 67:441–449.

Chen, C., Zheng, L., and Blanton, J. O. (1999). Physical processes controlling the formation, evolution, and perturbation of the low-salinity front in the inner shelf off the southeastern United States: A modeling study. *Journal of Geophysical Research Oceans*, 104:1259–1288.

Cheung, K. F., Phadke, A. C., Wei, Y., Rojas, R., Douyere, Y. J. M., Martino, C. D., Houston, S. H., Liu, P. L. F., Lynett, P. J., Dodd, N., Liao, S., and Nakazaki, E. (2003). Modeling of storm-induced coastal flooding for emergency management. *Ocean Engineering*, 30:1353–1386.

Cho, H. (2020). A recursive algorithm for calculating the longest flow path and its iterative implementation. *Environmental Modelling & Software*, 131.

Cialone, M. A., Grzegorzewski, A. S., Mark, D. J., Bryant, M. A., and Massey, T. C. (2017). Coastal-storm model development and water-level validation for the North Atlantic Coast Comprehensive Study. *Journal of Waterway, Port, Coastal, and Ocean Engineering*, 143.

CIRES (2014). Continuously Updated Digital Elevation Model (CUDEM) - 1/9 Arc-Second Resolution Bathymetric-Topographic Tiles. `https://coast.noaa.gov/htdata/raster2/elevation/NCEI_ninth_Topobathy_2014_8483/`. Accessed 15 July 2021.

Coastal Emergency Risks Assessment (2019). `https://cera.coastalrisk.live/`. [Retrieved 28 May 2019].

Cobell, Z. (2020). ADCIRCModules: a C++ and Python interface for manipulation of ADCIRC model data. `https://github.com/zcobell/ADCIRCModules`.

Contreras, M. T., Woods, B., Blakely, C., Wirasaet, D., and Westerink, J. (2021). A channel to basin-scale unstructured grids for the US East and Gulf of Mexcico coasts. Noaa technical memorandum nos cs 41, University of Notre Dame.

CoreLogic (2020). 2020 storm surge report. `https://www.corelogic.com/downloadable-docs/storm-surge-report-20200528-screen-2.pdf`. [Retrieved 20 October 2020].

Cyriac, R., Dietrich, J. C., Blain, C. A., Dawson, C. N., Dresback, K. M., Fathi, A., Bilskie, M. V., Graber, H. C., Hagen, S. C., and Kolar, R. L. (2020). Wind and tide effects on the Choctawhatchee Bay plume and implications for surface transport at Destin Inlet. *Regional Studies in Marine Science*, 35:101131.

Danielson, J. J., Poppenga, S. K., Tyler, D. J., Palaseanu-Lovejoy, M., and Gesch, D. B. (2018). Coastal National Elevation Database: U.S. Geological Survey Fact Sheet 2018–3037.

Deb, M. and Ferreira, C. M. (2016). Simulation of cyclone-induced storm surges in the low-lying delta of bangladesh using coupled hydrodynamic and wave model (swan + adcirc). *Journal of Flood Risk Management*, 11:S750–S765.

Deb, M., Kirby, J. T., Abdolali, A., and Shi, F. (2023). A surface porosity approach for eliminating artificial ponding in coastal salt marsh simulations. *Coastal Engineering*, 179.

Defina, A. (2000). Two-dimensional shallow flow equations for partially dry areas. *Water Resources Research*, 36:3251–3264.

Dick, C. G., Tromble, E. M., Dresback, K. M., and Kolar, R. L. (2013). Implementation and analysis of a partial-element wetting and drying framework for generalized wave continuity equation-based hydrodynamic models. *International Journal for Numerical Methods in Fluids*, 72:1015–1033.

Dietrich, J. C., Bunya, S., Westerink, J. J., Ebersole, B. A., Smith, J. M., Atkinson, J. H., Jensen, R. E., Resio, D. T., Luettich, R. A., Dawson, C. N., Cardone, V. J., Cox, A. T., Powell, M. D., Westerink, H. J., and Roberts, H. J. (2010). A high-resolution coupled riverine flow, tide, wind, wind wave and storm surge model for southern Louisiana and Mississippi: Part II – Synoptic description and analysis of Hurricanes Katrina and Rita. *Monthly Weather Review*, 138:378–404.

Dietrich, J. C., Kolar, R. L., and Luettich, R. A. (2004). Assessment of ADCIRC's wetting and drying algorithm. In Miller, C. T., Farthing, M. W., Gray, W. G., and Pinder, G. F., editors, *Proceedings of Computational Methods in Water Resources*, volume 2, pages 1767–1778.

Dietrich, J. C., Tanaka, S., Westerink, J. J., Dawson, C. N., Luettich, R. A., Zijlema, M., Holthuijsen, L. H., Smith, J. M., Westerink, L. G., and Westerink, H. J. (2012). Performance of the unstructured-mesh, SWAN+ADCIRC model in computing hurricane waves and surge. *Journal of Scientific Computing*, 52:468–497.

Dietrich, J. C., Zijlema, M., Westerink, J. J., Holthuijsen, L. H., Dawson, C. N., Luettich, R. A., Jensen, R. E., Smith, J. M., Stelling, G. S., and Stone, G. W. (2011). Modeling hurricane waves and storm surge using integrally-coupled, scalable computations. *Coastal Engineering*, 58:45–65.

Dresback, K. M., Fleming, J. G., Blanton, B. O., Kaiser, C., Gourley, J. J., Tromble, E. M., Luettich, R. A., Kolar, R. L., Hong, Y., Van Cooten, S., Vergara, H. J., Flamig, Z., Lander,

H. M., Kelleher, K. E., and Nemunaitis-Monroe, K. L. (2013). Skill assessment of a real-time forecast system utilizing a coupled hydrologic and coastal hydrodynamic model during Hurricane Irene (2011). *Continental Shelf Research*, 71:78–94.

Dresback, K. M., Kolar, R. L., Blain, C. A., Szpilka, C. M., Szpilka, A. M., Luettich, R. A., and Shay, T. (2010). Development and application of the coupled HYCOM and ADCIRC system. In *Proceedings of the 11th International Conference*, volume 388, pages 259–277. Estuarine and Coastal Modeling.

Egbert, G. D. and Erofeeva, S. Y. (2002). Efficient Inverse Modeling of Barotropic Ocean Tides. *Journal of Atmospheric and Oceanic Technology*, 19:183–204.

Emanuel, K. (2020). Evidence that hurricanes are getting stronger. *Proceedings of the National Academy of Sciences of the United States of America*, 117:13194–13195.

Federal Emergency Management Agency (2019). Guidance for flood risk analysis and mapping: Determination of wave characteristics. Technical report, Federal Emergency Management Agency.

Ferry, N. and Reverdin, G. (2000). Seasonal sea surface height variability in the North Atlantic Ocean. *Journal of Geophysical Research*, 105:6307–6326.

Fleming, J., Fulcher, C., Luettich, R. A., Estrade, B., Allen, G., and Winer, H. (2008). A real time storm surge forecasting system using ADCIRC. In *Estuarine and Coastal Modeling (2007)*, pages 893–912.

Fritz, H. M., Blount, C., Sokoloski, R., Singleton, J., Fuggle, A., McAdoo, B. G., Moore, A., Grass, C., and Tate, B. (2007). Hurricane Katrina storm surge distribution and field observations on the Mississippi barrier islands. *Estuarine, Coastal, and Shelf Science*, 74:12–20.

Gao, J., Luettich, R. A., and Fleming, J. G. (2017). Development and evaluation of a generalized asymmetric tropical cyclone vortex model in ADCIRC. In Preparation.

Garratt, J. R. (1977). Review of drag coefficients over oceans and continents. *Monthly Weather Review*, 105:915–929.

Gill, A. E. and Niiler, P. P. (1973). The theory of the seasonal variability in the ocean. *Deep-Sea Research*, 20:141–177.

Hagen, S. C., Westerink, J. J., Kolar, R. L., and Horstmann, O. (2001). Two-dimensional, unstructured mesh generation for tidal models. *International Journal for Numerical Methods in Fluids*, 35:669–686.

Harris, D. L. (1963). Characteristics of the hurricane storm surge. Technical Report Technical Paper No. 48, U.S. Department of Commerce and Weather Bureau, Washington D.C.

Helderop, E. and Grubesic, T. H. (2019). Hurricane storm surge in volusia county, florida: evidence of a tipping point for infrastructure damage. *Disasters*, 43:157–180.

Herdman, L., Erikson, L., and Barnard, P. (2018). Storm Surge Propagation and Flooding in Small Tidal Rivers during Events of Mixed Coastal and Fluvial Influence. *Journal of Marine Science and Engineering*, 6.

Hope, M. E., Westerink, J. J., Kennedy, A. B., Kerr, P. C., Dietrich, J. C., Dawson, C. N., Bender, C. J., Smith, J. M., Jensen, R. E., Zijlema, M., Holthuijsen, L. H., Luettich Jr, R. A., Powell, M. D., Cardone, V. J., Cox, A. T., Pourtaheri, H., Roberts, H. J., Atkinson, J. H., Tanaka, S., Westerink, H. J., and Westerink, L. G. (2013). Hindcast and Validation of Hurricane Ike (2008) Waves, Forerunner, and Storm Surge. *Journal of Geophysical Research: Oceans*, 118:4424–4460.

Jelesnianski, C. P., Chen, J., and Shafer, W. A. (1992). SLOSH: Sea, Lake, and Overland Surges from Hurricanes. Technical report, U.S. Dep. of Commer., Natl. Oceanic and Atmos. Admin., Natl. Weather Serv., Silver Spring, Maryland.

Kennedy, A. B., Wirasaet, D., Begmohammadi, A., Sherman, T., Bolster, D., and Dietich, J. C. (2019). Subgrid theory for storm surge modeling. *Ocean Modelling*, 144.

Kerr, P. C., Martyr, R. C., Donahue, A. S., Hope, M. E., Westerink, J. J., Luettich Jr, R. A., Kennedy, A. B., Dietrich, J. C., Dawson, C. N., and Westerink, H. J. (2013). US IOOS coastal and ocean modeling testbed: Evaluation of tide, wave, and hurricane surge response sensitivities to mesh resolution and friction in the Gulf of Mexico. *Journal of Geophysical Research: Oceans*, 118:4633–4661.

King, I. (2001). Modeling of marshes and wetlands. *Journal of Coastal Research*, 27:76 – 87.

Kinnmark, I. (1986). The shallow water wave equations: Formulation, analysis and application. In Brebbia, C. A. and Orszag, S. A., editors, *Lecture Notes in Engineering*, volume 15, pages 12–26. Springer-Verlag.

Knabb, R. D., Brown, D. P., and Rhome, J. R. (2005). Tropical Cyclone Report, Hurricane Rita, 18-26 September 2005. Technical report, National Hurricane Center.

Kyprioti, A. P., Adeli, E., Taflanidis, A. A., Westerink, J. J., and Hendrik, H. L. T. (2021). Probabilistic Storm Surge Estimation for Landfalling Hurricanes: Advancements in Computational Efficiency Using Quasi-Monte Carlo Techniques. *Journal of Marine Science and Engineering*, 9(12):1322.

Lee, J.-W., Irish, J. L., Bensi, M. T., and Marcy, D. C. (2021). Rapid prediction of peak storm surge from tropical cyclone track time series using machine learning. *Coastal Engineering*, 170:104024.

Leijnse, T., van Ormondt, M., Nederhoff, K., and van Dongeren, A. (2021). Modeling compound flooding in coastal systems using a computationally efficient reduced-physics solver: Including fluvial, pluvial, tidal, wind- and wave-driven processes. *Coastal Engineering*, 163.

Lin, N. and Chavas, D. (2012). On hurricane parametric wind and applications in storm surge modeling. *Journal of Geophysical Research*, 117.

Lin, N., Emanuel, K. A., Smith, J. A., and Vanmarcke, E. (2010a). Risk assessment of hurricane storm surge for New York City. *Journal of Geophysical Research Atmospheres*, 115.

Lin, N., Smith, J. A., Villarini, G., Marchok, T. P., and Baeck, M. L. (2010b). Modeling extreme rainfall, winds, and surge from Hurricane Isabel (2003). *Weather and Forecasting*, 25:1342–1361.

Lockwood, J. W., Lin, N., Oppenheimer, M., and Lai, C. Y. (2022). Using Neural Networks to Predict Hurricane Storm Surge and to Assess the Sensitivity of Surge to Storm characteristics. *Journal of Geophysical Research: Atmospheres*, 127.

Luettich, R. A. and Westerink, J. J. (1995a). Continental shelf scale convergence studies with a barotropic tidal model. In Lynch, D. R. and Davies, A. M., editors, *Coastal and Estuarine Studies*, pages 349–372.

Luettich, R. A. and Westerink, J. J. (1995b). Implementation and testing of elemental flooding and drying in the ADCIRC hydrodynamic model. Technical report, U.S. Army Corps of Engineers, Vicksburg MS.

Luettich, R. A. and Westerink, J. J. (2004). Formulation and numerical implementation of the 2D/3D ADCIRC finite element model version 44.XX. `https://adcirc.org/files/2018/11/adcirc_theory_2004_12_08.pdf`.

Luettich, R. A., Westerink, J. J., and Scheffner, N. W. (1992). ADCIRC: An advanced three-dimensional circulation model for shelves coasts and estuaries, report 1: Theory and methodology of ADCIRC-2DDI and ADCIRC-3DL. Technical report, United States Army Corps of Engineers.

Medeiros, S. C. and Hagen, S. C. (2013). Review of wetting and drying algorithms for numerical tidal flow models. *International Journal for Numerical Methods in Fluids*, 71:473–487.

National Centers for Environmental Information (2023). U.S. Billion-Dollar Weather and Climate Disasters (2023). `https://www.ncei.noaa.gov/access/billions/`. [Retrieved 6 April 2023].

National Hurricane Center (2020). Sea, lake, and overland surges from hurricanes (slosh). `https://www.nhc.noaa.gov/surge/slosh.php`. [Retrieved 22 December 2020].

National Hurricane Center (2021). Storm surge overview. `https://www.nhc.noaa.gov/surge/`. [Retrieved 10 May 2021].

National Hurricane Center (2023). Tropical Cyclone Forecast/Advisory (TCM). `https://www.nhc.noaa.gov/aboutnhcprod.shtml#TCM`. [Retrieved 7 April 2023].

National Oceanic and Atmospheric Administration (2020). Data access viewer. https://coast.noaa.gov/dataviewer/#/lidar/search/. [Retrieved 20 October 2020].

National Oceanic and Atmospheric Administration (2021a). Data access viewer. https://coast.noaa.gov/dataviewer/#/landcover/search/. [Retrieved 26 January 2021].

National Oceanic and Atmospheric Administration (2021b). Data access viewer. https://coast.noaa.gov/dataviewer/#/landcover/search/. [Retrieved 11 October 2021].

National Oceanic and Atmospheric Administration (2021c). NOAA Medium Resolution Shoreline. https://shoreline.noaa.gov/data/datasheets/medres.html. [Retrieved 6 October 2021].

Needham, H. F., Keim, B. D., and Sathiaraj, D. (2015). A review of tropical cyclone-generated stormsurges: Global data sources, observations,and impacts. *Reviews of Geophysics*, 53:545–591.

Oceanweather Inc. (2018). Tropical modelling. https://www.oceanweather.com/research/TropicalCyclone.html. [Retrieved 31 January 2022].

Owensby, M., Bryant, M., Provost, L., Hesser, T., Massey, T. C., Tritinger, A., and Ding, Y. (2020). Calibration and Validation of the South Atlantic Domain Model Setup for the South Atlantic Coastal Study (sacs). Technical report, US Army Corps of Engineers.

Park, K., Federico, I., Lorenzo, E. D., Ezer, T., Cobb, K. M., Pinardi, N., and Coppini, G. (2022). The contribution of hurricane remote ocean forcing to storm surge along the southeastern u.s. coast. *Coastal Engineering*, 173.

Pringle, W. J., Wirasaet, D., Roberts, K. J., and Westerink, J. J. (2021). Global storm tide modeling with ADCIRC v55: unstructured mesh design and performance. *Geoscientific Model Development*, 14:1125–1145.

Pringle, W. J., Wirasaet, D., Suhardjo, A., Westerink, J. J., Kennedy, A. B., and Nong, S. (2018). Finite-element barotropic model for the indian and western pacific oceans: Tidal model-data comparisons and sensitivities. *Ocean Modelling*, 129:13–38.

Ramirez, J. A., Lichter, M., Coulthard, T. J., and Skinner, C. (2016). Hyper-resolution mapping of regional storm surge and tide flooding: Comparison of static and dynamic models. *Natural Hazards*, 82:571–590.

Ratcliff, J. D. (2022). Analysis of wind and storm surge of hurricane florence. Master's thesis, University of North Carolina at Chapel Hill.

Redfield, A. C. (1958). The influence of the continental shelf on the tides of the Atlantic coast of the United States. *Journal of Marine Resources*, 17:432–448.

Rego, J. L. and Li, C. (2010). Nonlinear terms in storm surge predictions: Effect of tide and shelf geometry with case study from Hurricane Rita. *Journal of Geophysical Research*, 115:C06020.

Resio, D. T. and Westerink, J. J. (2008). Modeling of the physics of storm surges. *Physics Today*, 69:33–38.

Riverside Technology and AECOM (2015). Mesh development, tidal validation, and hindcast skill assessment of an ADCIRC model for the Hurricane Storm Surge Operational Forecast System on the US Gulf-Atlantic Coast.

Roberts, K. J., Dietrich, J. C., Wirasaet, D., Pringle, W. J., and Westerink, J. J. (2021). Dynamic load balancing for predictions of storm surge and coastal flooding. *Environmental Modelling and Software*, 140:105045.

Roberts, K. J., Pringle, W. J., and Westerink, J. J. (2019a). OceanMesh2D 1.0: MATLAB-based software for two-dimensional unstructured mesh generation in coastal ocean modeling. *Geoscientific Model Development*, 12:1847–1868.

Roberts, K. J., Pringle, W. J., and Westerink, J. J. (2019b). On the automatic and a priori design of unstructured mesh resolution for coastal ocean circulation models. *Ocean Modelling*, 144:101509.

Roig, L. C. (1994). Hydrodynamic modeling of flows in tidal wetlands.

Rucker, C. A., Tull, N., Dietrich, J. C., Langan, T. E., Mitasova, H., Flemming, J. G., Blanton, B. O., and Luettich, R. A. (2021). Downscaling of real-time coastal flooding predictions for decision support. *Natural Hazards*, 107:1341–1369.

Ryosuke, O., Yuya, U., Daisuke, N., Shohei, H., and Crissman, L. (2017). Cupy: A numpy-compatible library for nvidia gpu calculations. In *Proceedings of Workshop on Machine Learning Systems (LearningSys) in The Thirty-first Annual Conference on Neural Information Processing Systems (NIPS)*.

Sebastian, A. G., Proft, J. M., Dietrich, J. C., Du, W., Bedient, P. B., and Dawson, C. N. (2014). Characterizing hurricane storm surge behavior in Galveston Bay using the SWAN+ADCIRC model. *Coastal Engineering*, 88:171–181.

Sehili, A., Lang, G., and Lippert, C. (2014). High-resolution subgrid models: background, grid generation, and implementation. *Ocean Dynamics*, 64:519–535.

Shin, S. I. and Newman, M. (2021). Seasonal Predictability of Global and North American Coastal Sea Surface Temperature and Height Anomalies. *Geophysical Research Letters*, 48.

Stark, J., Oyen, T. V., Meire, P., and Temmerman, S. (2015). Observations of tidal and storm surge attenuation in a large tidal marsh. *Limnology and Oceanography*, 60:1371–1381.

Stewart, S. R. (2017). Tropical Cyclone Report for Hurricane Matthew, 28 September - 9 October 2016. Technical report, National Hurricane Center.

Stewart, S. R. and Berg, R. (2019). Tropical Cyclone Report for Hurricane Florence, 31 August – 17 September 2018. Technical report, National Hurricane Center.

Tanaka, S., Bunya, S., Westerink, J. J., Dawson, C. N., and Luettich, R. A. (2011). Scalability of an Unstructured Grid Continuous Galerkin Based Hurricane Storm Surge Model. *Journal of Scientific Computing*, 46:329–358.

Taylor, A. A. and Glahn, B. (2008). Probabalistic guidance for hurricane storm surge. Technical report, Meteorological Development Laboratory Office of Science and Technology National Weather Service.

Thomas, A., Dietrich, J. C., Asher, T. G., Blanton, B. O., Cox, A. T., Dawson, C. N., Fleming, J. G., and Luettich, R. A. (2019). Influence of storm timing and forward speed on tide-surge interactions during Hurricane Matthew. *Ocean Modelling*, 137:1–19.

Thomas, A., Dietrich, J. C., Dawson, C. N., and Luettich, R. A. (2022). Effects of Model Resolution and Coverage of Storm-Driven Coastal Flooding Predictions. *Journal of Waterway, Port, Coastal, and Ocean Engineering*, 148.

Thomas, A., Dietrich, J. C., Loveland, M., Samii, A., and Dawson, C. N. (2021). Improving coastal flooding predictions by switching meshes during a simulation. *Ocean Modelling*, 164:101820.

Tomiczek, T., Kennedy, A., and Rogers, S. (2014). Collapse Limit State Fragilities of Wood-Framed Residences from Storm Surge and Waves during Hurricane Ike. *Journal of Waterway, Port, Coastal, and Ocean Engineering*, 140.

URS Corporation (2009). South Carolina storm surge project deliverable 1: Grid development report.

U.S. Army Corps of Engineers (2015). North atlantic coast comprehensive study: Resilient adaption to increasing risk. Technical report, U.S. Army Corps of Engineers.

U.S. Army Corps of Engineers (2021). South Atlantic Coastasl Study (SACS) Main Report. `https://www.sad.usace.army.mil/SACS/`. [Retrieved 8 July 2022].

U.S. Geological Survey (2020a). Coastal National Elevation Database (CoNED) Applications Project. `https://topotools.cr.usgs.gov/topobathy_viewer/`. [Retrieved 20 October 2020].

U.S. Geological Survey (2020b). Flood event viewer. `https://stn.wim.usgs.gov/FEV/#Rita`. [Retrieved 29 December 2020].

Viero, D. P. (2019). Modelling urban floods using a finite element staggered scheme with an anisotropic dual porosity model. *Journal of Hydrology*, 568:247–259.

Volp, N. D., van Prooijen, B. C., and Stelling, G. S. (2013). A finite volume approach for shallow water flow accounting for high-resolution bathymetry and roughness data. *Water Resources Research*, 49:4126–4135.

Wang, H. V., Loftis, J. D., Liu, Z., Forrest, D., and Zhang, J. (2014). The storm surge and sub-grid inundation modeling in New York City during Hurricane Sandy. *Journal of Marine Science and Engineering*, 2:226–246.

Weaver, R. J. and Luettich, R. A. (2010). 2D vs 3D storm surge sensitivity in ADCIRC: Case study of Hurricane Isabel. In *Estuarine and Coastal Modeling 2009*, pages 762–779.

Westerink, J. J., Luettich, R. A., and Muccino, J. C. (1994). Modeling tides in the western North Atlantic using unstructured graded grids. *Tellus A*, 46A:178–199.

Westerink, J. J., Luettich Jr, R. A., Feyen, J. C., Atkinson, J. H., Dawson, C. N., Roberts, H. J., Powell, M. D., Dunion, J. P., Kubatko, E. J., and Pourtaheri, H. (2008). A Basin to Channel Scale Unstructured Grid Hurricane Storm Surge Model Applied to Southern Louisiana. *Monthly Weather Review*, 136:833–864.

Westerink, J. J., Muccino, J. C., and Luettich, R. A. (1992). Resolution requirements for a tidal model of the Western North Atlantic and Gulf of Mexico. In Russell, T. F., editor, *Computational Methods in Water Resources*, pages 669–674.

Whitacker, S. (1985). A Simple Geometrical Derivation of the Spatial Averaging Theorem. *Chemical Engineering Education*, 19.

Whitacker, S. (1999). *The Method of Volume Averaging*. Kluwer Academic Publishers.

Whitehead, J. C. (2003). One million dollars per mile? The opportunity costs of Hurricane evacuation. *Ocean & Coastal Management*, 46:1069–1083.

Woodruff, J. L., Dietrich, J., Wirasaet, D., Kennedy, A., Bolster, D., Silver, Z., Medlin, S., and Kolar, R. (2021). Subgrid corrections in finite-element modeling of storm-driven coastal flooding. *Ocean Modelling*, 167.

Woodruff, J. L., Dietrich, J. C., Wirasaet, D., Kennedy, A. B., and Bolster, D. (2023). Storm surge predictions from ocean to subgrid scales. *Natural Hazards*.

Wu, G., Shi, F., Kirby, J. T., Mieras, R., Liang, B., Li, H., and Shi, J. (2016). A pre-storage, subgrid model for simulating flooding and draining processes in salt marshes. *Coastal Engineering*, 108:65–78.

Xian, S., Lin, N., and Hatzikyriakou, A. (2015). Storm surge damage to residential areas: a quantitative analysis for Hurricane Sandy in comparison with FEMA flood map. *Natural Hazards*, 79:1867–1888.

Zarillo, K. A., Zarillo, G. A., McGinnis, D. I., Belanger, T. V., Adams, D. H., Paperno, R., Cox, A. C., and Irlandi, E. A. (2012). An Assessment of Coastal Water Resources and Watershed Conditions In and Adjacent to Canaveral National Seashore. Technical report, National Park Service.

Zhang, K., Xiao, C., and Shen, J. (2008). Comparison of the CEST and SLOSH Models for Storm Surge Flooding. *Journal of Coastal Research*, 24:489–499.

APPENDIX

APPENDIX

A

SUBGRID THEORY

A.1 Averaged Governing Equations for ADCIRC

The upscaled governing equations stated in Section 2.3 are derived by applying a formal averaging technique (Whitacker 1999) to the standard 2D shallow water equations written in the conservative form. By following such a technique, we define a mesh-scale average of any flow quantity Q as:

$$\langle Q \rangle_G = \frac{1}{A_G} \iint_{A_W} Q \, \mathrm{d}A, \tag{A.1}$$

where A_G denotes the mesh area and A_W the wet area within A_G (note that A_G and A_W are related through Equation 2.2). In addition, an alternative average use in the wet average (commonly known as intrinsic phase average) defined by:

$$\langle Q \rangle_W = \frac{1}{A_W} \iint_{A_W} Q \, \mathrm{d}A. \tag{A.2}$$

In addition, the following rules (Whitacker 1985) are used to interchange differentiation with respect to time and space and time-dependent spatial integration. In the formula below, $\boldsymbol{U}_B$ denotes the velocity of the potentially moving boundary, $\mathbf{n}_s = (n_{s,x}, n_{s,y})$ is an outward-pointing unit vector normal to the wet/dry boundary, Γ_W is the wet/dry boundary, and the subscript r denotes a dummy notation for the x or y coordinates.

$$\left\langle \frac{\partial Q}{\partial t} \right\rangle_G = \frac{\partial \langle Q \rangle_G}{\partial t} - \frac{1}{A_G} \int_{\Gamma_W} Q \boldsymbol{U}_B \cdot \mathbf{n}_s \, \mathrm{d}S, \tag{A.3}$$

and:

$$\left\langle \frac{\partial Q}{\partial r} \right\rangle_G = \frac{\partial \langle Q \rangle_G}{\partial r} + \frac{1}{A_G} \int_{\Gamma_W} \mathbf{n}_{s,r} Q \, \mathrm{d}S. \tag{A.4}$$

The development of subgrid equations involves roughly applying A.2 to the mass and momentum equations, making use of A.3 and A.4, and determining closures for terms that are not uniquely defined by the coarsened mesh-scale variables. The following subsections describe the development of the averaged mass equation, the averaged momentum equations, and the reformulation of the averaged continuity equation into the GWCE form.

A.1.1 Averaged Primitive Continuity Equation

The primitive continuity equation is:

$$\frac{\partial H}{\partial t} + \frac{\partial UH}{\partial x} + \frac{\partial VH}{\partial y} = 0, \tag{A.5}$$

in which $H = h + \zeta$ is the total water depth, h is the bathymetric depth measured positive downwards from a reference datum, ζ is the water surface elevation measured positive upwards from the datum, U and V are the depth-averaged horizontal velocity components in the x- and y-directions respectively. The mesh-scale averaging of each term is described below.

First, for the local rate of change in time, we use A.3 to pull the time derivative out of the integral, more specifically,

$$\begin{aligned} \left\langle \frac{\partial H}{\partial t} \right\rangle_G &= \frac{1}{A_G} \iint_{A_W} \frac{\partial H}{\partial t} \, \mathrm{d}A \\ &= \frac{1}{A_G} \frac{\partial}{\partial t} \iint_{A_W} H \, \mathrm{d}A - \frac{1}{A_G} \int_{\Gamma_W} H (\boldsymbol{U}_B \cdot \mathbf{n}_s) \, \mathrm{d}S. \end{aligned} \tag{A.6}$$

Because $H = 0$ at the wet/dry front, we eliminate the boundary integral and obtain

$$\left\langle \frac{\partial H}{\partial t} \right\rangle_G = \frac{1}{A_G} \iint_{A_W} \frac{\partial H}{\partial t} \, \mathrm{d}A = \frac{1}{A_G} \frac{\partial}{\partial t} \iint_{A_W} H \, dA = \frac{\partial \langle H \rangle_G}{\partial t}, \tag{A.7}$$

which is now temporal rate of change of the averaged total water depth.

Next, for the volume flux in the x-direction, we apply the spatial averaging theorem A.4 to pull the spatial derivative out of the integral, more precisely,

$$\begin{aligned} \left\langle \frac{\partial UH}{\partial x} \right\rangle_G &= \frac{1}{A_G} \iint_{A_W} \frac{\partial UH}{\partial x} \, \mathrm{d}A \\ &= \frac{1}{A_G} \frac{\partial}{\partial x} \iint_{A_W} UH \, \mathrm{d}A + \frac{1}{A_G} \int_{\Gamma_W} UH \mathbf{n}_{s,x} \, \mathrm{d}S. \end{aligned} \tag{A.8}$$

Again $H = 0$ at the wet/dry boundary, we eliminate the boundary integral and have:

$$\left\langle \frac{\partial UH}{\partial x} \right\rangle_G = \frac{1}{A_G} \iint_{A_W} \frac{\partial UH}{\partial x} \, \mathrm{d}A = \frac{1}{A_G} \frac{\partial}{\partial x} \iint_{A_W} UH \, \mathrm{d}A = \frac{\partial \langle UH \rangle_G}{\partial x}. \tag{A.9}$$

The interchange between differentiation and integration in the averaging of the last term, the volume flux in the y-direction, can be done in an analogous way. After above manipulation, the averaged primitive continuity equation becomes

$$\frac{\partial \langle H \rangle_G}{\partial t} + \frac{\partial \langle UH \rangle_G}{\partial x} + \frac{\partial \langle VH \rangle_G}{\partial y} = 0. \tag{A.10}$$

By postulating that ζ varies very slowly within A_W, one has $\langle H \rangle_G = \frac{1}{A_G} \int_{z=-\inf}^{z=\langle \zeta \rangle_W} \iint_{A_G} max(0, b + z) \, \mathrm{d}A \, \mathrm{d}z$. As a consequence, we can rewrite Equation A.10 as

$$\phi \frac{\partial \langle \zeta \rangle_W}{\partial t} + \frac{\partial \langle UH \rangle_G}{\partial x} + \frac{\partial \langle VH \rangle_G}{\partial y} = 0, \tag{A.11}$$

which is the final averaged form of the primitive continuity equation to be considered in the reformulation into the the GWCE described below in A.1.3.

Note that in this study, we consider $\langle \boldsymbol{U} H \rangle_G$ as the variable to be solved for. Intead of using A.2, the velocity when required is computed from the following formula

$$\langle \boldsymbol{U} \rangle = \frac{\iint_{A_W} H \boldsymbol{U} \, \mathrm{d}A}{\iint_{A_W} H \, \mathrm{d}A} = \frac{\langle H \boldsymbol{U} \rangle_G}{\langle H \rangle_G}, \tag{A.12}$$

or equivalently

$$\langle \boldsymbol{U} \rangle \langle H \rangle_G = \langle \boldsymbol{U} H \rangle_G. \tag{A.13}$$

It is worth mentioning that the so-called volume-averaged velocity defined above has an advantage over an averaged velocity defined by A.2 in that it permits a substitution of $\langle \boldsymbol{U} H \rangle_G$ by $\langle \boldsymbol{U} \rangle \langle H \rangle_G$ in the governing equation without the need to resort to a more complicated closure. From this point forward, unless otherwise indicated, the notation $\langle \boldsymbol{U} \rangle$ is understood as the volume averaged velocity. Note that various forms of governing equations presented (Kennedy et al. 2019) are obtained from making use of A.13; they are intended for the solution where $\langle \boldsymbol{U} \rangle$ is chosen as an unknown variable.

A.1.2 Averaged Conservative Momentum Equations

We now average to the mesh scale the conservative momentum equations, including terms for the barotropic pressure gradient and lateral momentum-mixing stress terms. Consider the momentum equation in the x-direction:

$$\frac{\partial UH}{\partial t} + \frac{\partial UUH}{\partial x} + \frac{\partial UVH}{\partial y} - fVH = -gH \frac{\partial [\zeta + P_A]}{\partial x} + \frac{\tau_{sx}}{\rho_0} - \frac{\tau_{bx}}{\rho_0} + M_x. \tag{A.14}$$

It can be verified through the use of A.3 and A.4 and $H = 0$ at the wet/dry boundary that the mesh scale averaging of A.14 is equivalent to:

$$\begin{aligned} &\frac{\partial \langle UH \rangle_G}{\partial t} + \frac{\partial \langle UUH \rangle_G}{\partial x} + \frac{\partial \langle VUH \rangle_G}{\partial y} - f \langle VH \rangle_G = -\left\langle gH \frac{\partial \zeta}{\partial x} \right\rangle_G \\ &- g \langle H \rangle_G \frac{\partial P_A}{\partial x} + \left\langle \frac{\tau_{sx}}{\rho_0} \right\rangle_G - \left\langle \frac{\tau_{bx}}{\rho_0} \right\rangle_G + \langle M_x \rangle. \end{aligned} \tag{A.15}$$

In the above equation, the Coriolis parameter f and the atmospheric pressure P_A are assumed to vary at a spatial scale much larger than the grid scale and hence can be moved out of their respective integral terms. There is no unique way to define the averaging of convective momentum, bottom friction, surface gradients, and lateral mixing stresses in terms of the mesh-scale quantities $\langle H \rangle_G$, $\langle \boldsymbol{U} H \rangle_G$, $\langle \boldsymbol{U} \rangle$; further assumptions to be described below are therefore required to close the system.

For the convective accelerations, we chose the closure of the form written below:

$$\langle UUH \rangle_G = C_{UU} \langle U \rangle \langle UH \rangle_G, \langle UVH \rangle_G = C_{UV} \langle U \rangle \langle VH \rangle_G,$$

which resemble the particular forms of the convective momentum considered in ADCIRC (see Equation (2.2) on p.15 of the ADCIRC theory report (Luettich and Westerink 2004)) with additional correction coefficients C_{UU} and C_{UV}.

For the surface gradient pressure term, we consider the following closure:

$$\begin{aligned} g\left\langle H\frac{\partial P_A}{\partial x}\right\rangle_G &= g\langle H\rangle_G\frac{\partial P_A}{\partial x} = g\phi\left\langle H\frac{\partial \zeta}{\partial x}\right\rangle_W \\ &= gC_\zeta\phi\langle H\rangle_W\frac{\partial\langle\zeta\rangle_W}{\partial x} = gC_\zeta\langle H\rangle_G\frac{\partial\langle\zeta\rangle_W}{\partial x}, \end{aligned} \tag{A.16}$$

where C_ζ is an additional correction coefficient. Although counterintuitive, numerical evidences demonstrate in Kennedy et al. (2019) indicated that C_ζ is clearly needed in some cases.

For the surface stress term, we consider the quadratic drag law for the surface stress caused by wind:

$$\left\langle\frac{\tau_{sx}}{\rho_0}\right\rangle_G = \phi\left\langle\frac{\tau_{sx}}{\rho_0}\right\rangle_W = \phi\frac{\rho_a}{\rho_0}C_D\,|\boldsymbol{W}_{10}|\,W_{10,x}, \tag{A.17}$$

where ρ_a denotes the air density and $\boldsymbol{W}_{10} = (W_{10,x}, W_{10,x})$ denotes the 10 m wind velocity assumed to be known (wind data comes typically from a numerical model with a spatial scale greater than the grid scale considered in the surge model).

The bottom stress τ_{bx} is assumed to obey a quadratic bottom friction law and the closure below is considered:

$$\left\langle\frac{\tau_{bx}}{\rho_0}\right\rangle_G = \phi\left\langle\frac{\tau_{bx}}{\rho_0}\right\rangle_W = \phi\left\langle\frac{C_f|\mathbf{U}|UH}{H}\right\rangle_W = \phi C_{M,f}\frac{|\mathbf{U}|\langle UH\rangle_W}{\langle H\rangle_W}, \tag{A.18}$$

where $C_{M,f}$ is to-be-determined equivalent frictional coefficients that may depend on water surface elevations. In this work, for simplicity, $C_{M,f}$ is taken to be:

$$C_{M,f} = \frac{g\langle n\rangle_W^2}{\langle H\rangle_W^{1/3}}, \tag{A.19}$$

where $\langle n\rangle_W$ is a value characterizing the Manning's roughness coefficient of the wet area.

Finally, consider the average of the lateral mixing term:

$$
\begin{aligned}
\langle M_x \rangle_G &= \left\langle \frac{\partial H\tau_{xx}}{\partial x} + \frac{\partial H\tau_{yx}}{\partial y} \right\rangle_G \\
&= \frac{1}{A_G}\frac{\partial}{\partial x}\iint_{A_W} H\tau_{xx}\,\mathrm{d}A + \frac{1}{A_G}\int_{\Gamma_W} H\tau_{xx}\mathbf{n}_{s,x}\,\mathrm{d}S \\
&+ \frac{1}{A_G}\frac{\partial}{\partial y}\iint_{A_W} H\tau_{yx}\,\mathrm{d}A + \frac{1}{A_G}\int_{\Gamma_W} H\tau_{yx}\mathbf{n}_{s,y}\,\mathrm{d}S \\
&= \frac{\partial \langle H\tau_{xx}\rangle_G}{\partial x} + \frac{\partial \langle H\tau_{yx}\rangle_G}{\partial y}.
\end{aligned} \tag{A.20}
$$

Boundary integrals go to zero because $H = 0$ at the wet/dry boundary. Indeed, the vertically-integrated lateral terms $H\tau_{xx}$ and $H\tau_{yx}$ by itself require a closure assumption. ADCIRC supports several lateral closures. Here, we consider one specific form of such closures, more precisely:

$$
H\tau_{xx} = \widetilde{E}_h \frac{\partial UH}{\partial x}, H\tau_{yx} = \widetilde{E}_h \frac{\partial UH}{\partial y}.
$$

The grid-average of these lateral closures are approximated as:

$$
\langle H\tau_{xx}\rangle_G = \widetilde{E}_h \frac{\partial \langle UH\rangle_G}{\partial x}, \langle H\tau_{yx}\rangle_G = \widetilde{E}_h \frac{\partial \langle UH\rangle_G}{\partial y}, \tag{A.21}
$$

where $\tilde{E}_h$ is a grid scale eddy viscosity (potentially of different value than that used in the high-resolution calculation).

With the closure terms given above, the averaged momentum equation in the x-direction becomes:

$$
\begin{aligned}
&\frac{\partial \langle UH\rangle_G}{\partial t} + gC_\zeta \langle H\rangle_G \frac{\partial \langle \zeta\rangle_W}{\partial x} = -\frac{\partial C_{UU}\langle U\rangle\langle UH\rangle_G}{\partial x} - \frac{\partial C_{VU}\langle V\rangle\langle UH\rangle_G}{\partial y} \\
&+ f\langle VH\rangle_G - g\langle H\rangle_G \frac{\partial P_A}{\partial x} + \phi \left\langle \frac{\tau_{sx}}{\rho_0} \right\rangle_W - \frac{g\langle n\rangle_W^2 |\langle \mathbf{U}\rangle| \langle UH\rangle_G}{\langle H\rangle_W^{4/3}} \\
&+ \frac{\partial}{\partial x}\widetilde{E}_h \frac{\partial \langle UH\rangle_G}{\partial x} + \frac{\partial}{\partial y}\widetilde{E}_h \frac{\partial \langle UH\rangle_G}{\partial y}.
\end{aligned} \tag{A.22}
$$

Similarly, the averaged momentum equation in the y-direction with closure terms is:

$$
\begin{aligned}
&\frac{\partial \langle VH\rangle_G}{\partial t} + g C_\zeta \langle H\rangle_G \frac{\partial \langle \zeta\rangle_W}{\partial y} = -\frac{\partial C_{UV}\langle U\rangle\langle VH\rangle_G}{\partial x} - \frac{\partial C_{VV}\langle V\rangle\langle VH\rangle_G}{\partial y} \\
&- f\langle UH\rangle_G - g\langle H\rangle_G \frac{\partial P_A}{\partial y} + \phi\left\langle \frac{\tau_{sy}}{\rho_0}\right\rangle_W - \frac{g\langle n\rangle_W^2 |\langle \mathbf{U}\rangle| \langle VH\rangle_G}{\langle H\rangle_W^{4/3}} \\
&+ \frac{\partial}{\partial x} E_h \frac{\partial \langle VH\rangle_G}{\partial x} + \frac{\partial}{\partial y} E_h \frac{\partial \langle VH\rangle_G}{\partial y}.
\end{aligned} \tag{A.23}
$$

The final step is to select the correction coefficients. In this work, we consider a so-called 'Level 0' closure (Kennedy et al. 2019), in which: $C_{UU} = C_{UV} = C_{VU} = C_{VV} = 1$, $C_\zeta = 1$. Then the only non-unity closure is the wet-area fraction, as shown in the final Equations 2.5 and 2.6.

A.1.3 Averaged Generalized Wave Continuity Equation

Then the GWCE is formed by differentiating Equation A.11 with respect to time, adding to this A.11 mutiplied by a positive spatially-varying numerical parameter τ_0. This leads to:

$$
\begin{aligned}
\frac{\partial}{\partial t}\left(\phi \frac{\partial \langle \zeta\rangle_W}{\partial t}\right) + \tau_0 \phi \frac{\partial \langle \zeta\rangle_W}{\partial t} + \frac{\partial \langle \tilde{J}_x\rangle_G}{\partial x} + \frac{\partial \langle \tilde{J}_y\rangle_G}{\partial y} \\
- \langle UH\rangle_G \frac{\partial \tau_0}{\partial x} - \langle VH\rangle_G \frac{\partial \tau_0}{\partial y} = 0,
\end{aligned} \tag{A.24}
$$

where:

$$
\langle \tilde{J}_x\rangle_G = \frac{\partial \langle UH\rangle_G}{\partial t} + \tau_0 \langle UH\rangle_G, \tag{A.25}
$$

and:

$$
\langle \tilde{J}_y\rangle_G = \frac{\partial \langle VH\rangle_G}{\partial t} + \tau_0 \langle VH\rangle_G. \tag{A.26}
$$

The time derivative terms $\frac{\partial \langle UH\rangle_G}{\partial t}$ and $\frac{\partial \langle VH\rangle_G}{\partial t}$ in the above equation are further eliminated by means of the momentum equation A.22 and A.23. With the Level 0 closure we obtain the final form of the GWCE as it appears in Equation 2.7 repeated below:

$$
\begin{aligned}
\phi \frac{\partial^2 \langle\zeta\rangle_W}{\partial t^2} + \frac{\partial \phi}{\partial t}\frac{\partial \langle\zeta\rangle_W}{\partial t} + \tau_0 \phi \frac{\partial \langle\zeta\rangle_W}{\partial t} \\
- \frac{\partial}{\partial x}\left(g \langle H\rangle_G \frac{\partial \langle\zeta\rangle_W}{\partial x}\right) - \frac{\partial}{\partial y}\left(g \langle H\rangle_G \frac{\partial \langle\zeta\rangle_W}{\partial y}\right) \\
+ \frac{\partial \langle\tilde{J}_x\rangle_G}{\partial x} + \frac{\partial \langle\tilde{J}_y\rangle_G}{\partial y} - \langle UH\rangle_G \frac{\partial \tau_0}{\partial x} - \langle VH\rangle_G \frac{\partial \tau_0}{\partial y} = 0,
\end{aligned}
$$

where:

$$
\langle\tilde{J}_x\rangle_G = (\text{RHS of A.22}) + \tau_0 \langle UH\rangle_G,
$$

and:

$$
\langle\tilde{J}_y\rangle_G = (\text{RHS of A.23}) + \tau_0 \langle VH\rangle_G.
$$

Note that for $\langle H\rangle_G > 0$ (i.e. in fully wet or partial wet areas) the GWCE is a second order wave equation.

A.1.4 Finite Element Discretization

In this study, the ADCIRC solvers were kept largely the same, the GWCE is solved implicitly via the use of a global mass matrix, while the momentum equations are solved semi-implicitly. Both element- and vertex- based quantities are using in these solutions. On each time marching step, the GWCE (Equation 2.7) uses elementally-averaged quantities (Figure 2.1) to find a vertex-averaged water surface elevation $\langle\zeta\rangle_W$. This quantity is then used to look up the corresponding vertex-averaged total water depth $\langle H\rangle_W$, wet area fraction ϕ, and wet averaged Manning's n $\langle n\rangle_W$, which are used along with elementally-averaged quantities to solve Equations 2.5 and 2.6 for the vertex-averaged water velocities. Because we are solving averaged equations, the solutions for $\langle\zeta\rangle_W$, $\langle U\rangle$, and $\langle V\rangle$ are appropriately averaged. Therefore, no further manipulation is required.

The only change was the addition of the $\frac{\partial \phi}{\partial t}\frac{\partial \langle\zeta\rangle_W}{\partial t}$ which was discretized in the following way:

$$
\frac{\partial \phi}{\partial t}\frac{\partial \langle\zeta\rangle_W}{\partial t} = \sum_{n=1}^{NE_j} \frac{A_n}{12} \frac{\partial \overline{\phi}_n}{\partial t} \sum_{i=1}^{3} \Phi_{i,j} \frac{\partial \langle\zeta\rangle_{W_i}}{\partial t},
$$

where:

$$\frac{\partial \overline{\phi}_n}{\partial t} = \frac{\overline{\phi}_n^s - \overline{\phi}_n^{s-1}}{\Delta t} \quad \text{and} \quad \frac{\partial \langle \zeta \rangle_{W_i}}{\partial t} = \frac{\langle \zeta \rangle_i^{s+1} - \langle \zeta \rangle_i^{s-1}}{2\Delta t}.$$

Here, A_n is the area of element n, NE_j is the number of elements containing node j, $\overline{\phi}_n$ is the average wet area fraction over element n, $\Phi_{i,j}$ is the weighting function, and s is the current timestep.