ABSTRACT

POISSON, AUTUMN CHAR. Identifying the Earliest Signs of Storm Impacts to Improve Hurricane Flooding Forecasts. (Under the direction of Casey Dietrich).

One of the most unpredictable and deadly parts of a coastal storm is the storm surge, which can cause devastating flooding of coastal regions, and can result in loss of property and life. Storm surge is a result of winds pushing water from the nearshore ocean to rise above regular tide levels. Storm surge can have a short duration; elevated water levels are limited to when the storm winds are strongest at the coast, typically for a few hours as the storm makes landfall. This short duration is a challenge for predictions of when storm surge will start, how long it will persist, and which regions will be flooded.

To predict storm surge and its associated flooding in coastal areas, numerical models must have a detailed representation of the impacted region, and thus be accurate, without having too much detail, which can limit efficiency. This research examines the use of a multi-resolution approach to improve both efficiency and accuracy. The key idea is that, as a storm travels across the ocean, it will affect different regions at different times. Early on, as the storm moves in open water and far from people, efficiency is more important in the model predictions. But as the storm moves toward the coast, it becomes necessary to have a higher accuracy near coastal communities and infrastructure. This research examines the use of multi-resolution simulations in which, as a storm travels along its track, the model 'switches' from lower resolution in open water to higher resolution as the storm moves closer to land. The main research question is to determine when is it most beneficial to switch resolutions by determining when storm effects are first seen at the coast.

This research will explore the arrival of storm effects for Florence, which made landfall along the North Carolina coast during September 2018. It is an ideal storm for this research as its track was shore-normal, and thus its coastal effects increased as it approached landfall. This will allow for investigating the most optimal switch by focusing on a single switch between a lower-resolution mesh to a higher-resolution mesh. The switches will be initiated by several triggers, including wind speeds and water levels at the coast and inland locations, and with several lead times, including near and several days before landfall. Model performance will be quantified via comparisons to observations of storm effects in the region, as well as to a single, high-resolution simulation for the full storm. It will be shown that switching from a coarse resolution mesh to a fine resolution mesh will lead to an increase in efficiency gains across all switching simulations with the most optimal

switch time resulting in the most accurate predictions of water levels as compared to our full high-resolution simulation.

The results of this research will provide valuable contributions to forecasters working tirelessly during hurricane season to produce accurate and efficient predictions of coastal flooding impacts. With this information, real-time forecasts can be delivered sooner to emergency managers for informing evacuation zones, thus saving lives.

Identifying the Earliest Signs of Storm Impacts to Improve Hurricane Flooding Forecasts

by
Autumn Char Poisson

A thesis submitted to the Graduate Faculty of
North Carolina State University
in partial fulfillment of the
requirements for the Degree of
Master of Science

Civil Engineering

Raleigh, North Carolina
2021

APPROVED BY:

Elizabeth Sciaudone	Laura Tateosian
Jason Fleming External Member	Casey Dietrich Chair of Advisory Committee

DEDICATION

To my parents Edward and Nancy Poisson, and sister Amanda Poisson
For supporting me and being there every step of the way on this seemingly endless academic journey.

BIOGRAPHY

Autumn Poisson was born in Grand Rapids, MI, in 1988 to Edward and Nancy Poisson. She graduated from Rockford High School in 2007. Autumn grew up with a love of the outdoors and graduated from the University of Michigan in Ann Arbor with a B.S. degree in environmental science and a minor in mathematics in 2011 and a M.S. degree in conservation ecology in 2014. From there, she worked in the fields of environmental education and outreach, natural resource management, and program management all while building her interests in the research process and developing her own research ideas. The most pivotal experience was spending 10 months on a small island in the Pacific Ocean, where she fell in love with the coastal environment and decided to go back to school one last time (I think) for a M.S. in civil engineering at North Carolina State University. After communicating with her now advisor, Dr. Casey Dietrich, for over a year, she began her studies in coastal engineering and hurricane storm surge modeling in Fall 2018. Beyond academia, Autumn loves reading, spending time in the outdoors both backpacking and hiking, and is very passionate about physical fitness and Olympic weightlifting.

ACKNOWLEDGEMENTS

I am extremely grateful to so many people over these past few years as I completed my Master's thesis and it is important I take time to thank everyone for all they have done for me. A lot happened during this time including a global pandemic and a great many thanks are in order for those who helped me make it through.

I would first like to thank my advisor, Dr. Casey Dietrich for taking me on as a graduate student in engineering when I did not have a background in engineering. Since the beginning he has been supportive of my goals and aspirations and I would not be where I am without his guidance. Working with him has given me a new perspective on work in this field and I am looking forward to implementing what I have learned in my future career.

I would next like to thank my committee members for their guidance and support. Dr. Jason Fleming for teaching me all I need to know about ADCIRC and the ADCIRC prediction system, as well as GitHub and so much more. I will always be eternally grateful for his willingness to take time out of his busy schedule and meet with me. To Dr. Elizabeth Sciaudone for her enthusiasm, practicality and openness about the processes involved in research, and being there when I needed extra encouragement and advice. To Dr. Laura Tateosian for building my knowledge in Python and ArcGIS which helped greatly in my research and will undoubtedly help me stand out in my future career.

I would also like to thank my lab members and office mates, including those prior to the COVID pandemic who made it through working in the sad windowless lab in Mann Hall. Thanks to Johnathan Woodruff for answering all of my silly questions about Python, ADCIRC and everything in between and for being there to commiserate on occasion when things got difficult. I also want to thank him and Ashley Woodruff for being great friends as well and for always being down to grab a drink or two. To Alireza Gharagozlou and Dylan Anderson for being Python experts and helping me when I had questions. To Ajimon Thomas who got me up and running on this research project and introduced me to *Adcirpolate*. And to Chloe Stokes for her help trying to solve the problems with the high-resolution mesh and sharing all of her knowledge on *Adcirpolate* with me. To Nicole Harvey a fellow engineering masters student who took part in this journey with me. Her help through the first math class I had taken since undergrad and introducing me to rock climbing were some of the most memorable times I've had at NCSU.

And last but certainly not least, to my family, I can't begin to express my gratitude for their love and support over the past few years and my prior time in school. To my sister, Amanda Poisson who was there with a listening ear and expert advice. To my parents Edward and Nancy Poisson whose love and support never faltered even when things were challenging and during the moments I wanted to give up. It is difficult to put into words how grateful I am to have them and I can't imagine having done this without them.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER

1

INTRODUCTION

Coastal regions are increasingly susceptible to hazards from tropical cyclones. These hazards are expected to increase due to anthropogenic global warming and accompanying sea level rise, as well as continued coastal urbanization. The global average intensity of tropical cyclones is increasing (Knutson et al. 2020), which results in stronger storms and more inundation of coastal regions. The 2020 Atlantic hurricane season set a record for the most activity with 30 named storms, 12 of which made landfall in the U.S. (National Oceanic and Atmospheric Administration 2020). Before this, the 2017 season had 17 named storms, 10 hurricanes, and six major hurricanes, and was the costliest season on record (Klotzbach et al. 2018). The increase in tropical cyclone intensity will affect millions of people across the globe. About 2.4 billion people worldwide live within 100 km of the coast, and more than 600 million people live in areas that are less than 10 meters above sea level (United Nations 2017). In the U.S., more than 123 million people, or nearly 40 percent of the U.S. population, live in coastal counties (NOAA and U.S. Census Bureau 2013).

With tropical cyclones come heavy rainfall, high winds, and dangerous storm surge that can damage both life and property. Storm surge is defined as a rise in water levels above normally occurring tides caused by storm forces, which can result in significant flooding. Storm surge is the most deadly component of a hurricane, causing about 50 percent of hurricane-related deaths (Erdman 2019). Beyond loss of life, a storm surge of 7 m has the ability to inundate 67 percent of interstates, 57 percent of arterials, almost half of rail miles, 29 airports, and virtually all ports along the U.S. Gulf coast (National Hurricane Center 2021). Additional damages caused by storm surge and coinciding waves can result in significant erosion and transformation of the coastal areas (Stockdon et al. 2012). Due to these hazards and vulnerabilities, predictions of the effects of tropical cyclones are of upmost importance.

There are many challenges for the prediction of storm surge, but two critical aspects are the timing and duration of flooding. With respect to duration, depending on a storm's approach direction, speed, intensity, coastline configuration, and other factors, the storm surge can be limited to 1 to 2 days during and after landfall. For fast-moving storms, such as Michael (2018), which made landfall on the Florida panhandle, the flooding was focused during the day of landfall (Beven et al. 2018). However, for slower moving storms like Harvey (2017) and Florence (2018), coastal wave and storm surge effects can occur over multiple days (National Hurricane Center 2017, 2018). With respect to timing, some storms may result in earlier coastal effects like a forerunner surge, such as Ike (2008), which raised water levels along the Texas coast for 12 to 24 hours before landfall (Kennedy et al. 2011). Model predictions will need to represent the relatively short window in which storm surge effects take place, and minimize the uncertainty around where and when the flooding occurs.

To better prepare for storm surge and coastal flooding, numerical models are designed to help predict which areas will experience hazards due to tropical cyclones. This information is useful for many people and is shared with emergency managers to aid in decision-making on which areas to evacuate and how to deploy resources (like water and food, medical supplies, rescue teams) for use immediately after the storm (Massey et al. 2007; Rucker et al. 2021). Model results can be shared to aid evacuation and to plan for mitigation and response before the storm makes landfall. Additionally, effective delivery of guidance on which regions will experience storm surge and flooding is of upmost importance (Coastal Emergency Risks Assessment 2019). Model predictions and forecasts are also used between storms for the design and planning of mitigation structures and coastal communities (Fleming et al. 2008; Massey et al. 2007).

Recent improvements in computational capacity have led to the development of numerical storm surge models that use high geospatial resolution and highly detailed domains called *meshes* (Westerink et al. 1994; Hope et al. 2013; Roberts et al. 2019). These meshes represent geographical features like coastal inlets, channels, estuaries, and nearshore bathymetry over a range of spatial scales, including the finest scales near critical infrastructure. The models then compute sea surface elevations and currents by forcing with tidal information and meteorological data (Luettich et al. 1992). Thus the high spatial resolution allows for better representation of the coastal features (Taeb and Weaver 2019), leading to improved predictions of water levels (Blain et al. 1998; Kerr et al. 2013).

While several models can represent coastal circulation and overland flooding (Chen et al. 2003; Zhang et al. 2016), two models have gained prominence for real-time forecasting of storm surge. The first is the Sea, Lake, and Overland Surges from Hurricanes (SLOSH) model developed by the National Weather Service (Jelesnianski et al. 1992). SLOSH uses curvilinear polar, hyperbolic, and elliptical structured grids (which they call basins) to represent the domain (Figure 1.1). This allows for resolution to increase from about 3 km in the open ocean to about 350 m near the shore. Each cell in the grid is assigned an elevation and may contain sub-grid features such as barriers, channels, and cuts to improve representation of these features within the structured grid (Massey et al. 2007). SLOSH uses the finite difference method to discretize its conservation equations, and its input data are based on the best available storm track data, which includes storm location and radius to maximum winds (Jelesnianski et al. 1992). The second model is the ADvanced CIRCulation (ADCIRC) model, which has been developed by academic researchers for more than 30 years (Luettich et al. 1992). ADCIRC can implement both 2D depth-integrated and 3D options using high-resolution, unstructured meshes to solve for tidal circulation and storm surge in complex coastal regions (Luettich et al. 1992; Westerink et al. 1994). The mesh resolution can vary over three to four orders of magnitude (e.g. Figure 1.2), and minimum element sizes in coastal regions can vary from 500 m down to 10 m (Riverside Technology and AECOM 2015). ADCIRC uses the finite element method to discretize its conservation equations, and its input data include tidal and atmospheric data, and additional wave data when coupled with the Simulating WAves Nearshore (SWAN) model (Dietrich et al. 2011).

There are many trade-offs between these two prominent models that make them better equipped for specific tasks. SLOSH is typically used for probabilistic forecasts with hundreds of low-resolution simulations with varying storm inputs. Each simulation is extremely fast, and the large number of simulations can allow for predictions of how the coastal water levels will respond to uncertainty in the storm forecasts. However, due to their low resolution, each SLOSH simulation has limited accuracy. In comparison, ADCIRC is used typically for deterministic forecasts with only one or a small number of high-resolution simulations using best-available storm inputs. Each simulation is relatively slow but provides an accurate prediction of how the coastal water levels will respond. Thus, there is a trade-off between accuracy and efficiency for storm surge simulations within both of these models.

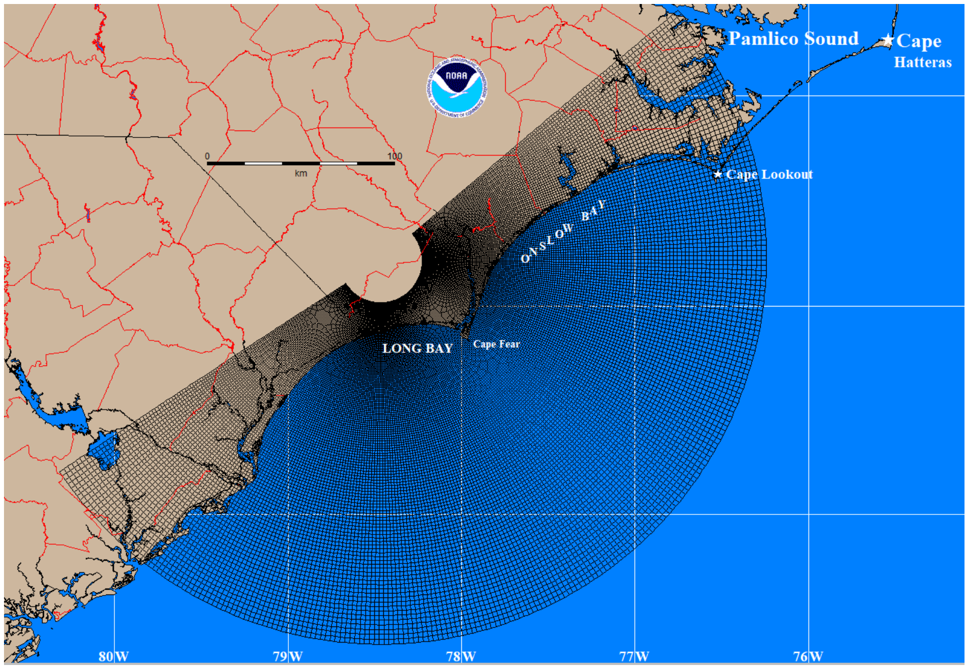

Figure 1.1: Example of a SLOSH Basin for North Carolina. Resolution varies from about 3 km in the open ocean to about 350 m near the shore.

To overcome this trade-off, a common technique in these and other models is to adapt resolution during the simulation, so that the highest levels of spatial resolution are used only when and where they are needed. This adaptivity can be implemented during a simulation, by increasing resolution of individual cells/elements, or between simulations, by mapping results onto a smaller, nested mesh. When applied during a simulation, the adaptive mesh refinement can follow the location of a tsunami wave front (Qin et al. 2019; Berger et al. 2011) or a coastal storm (Mandli and Dawson 2014; Beisiegel et al. 2020), but it requires a hierarchy of resolution at each cell/element as well as decision-making about when to refine/coarsen during the simulation. When applied between simulations, nesting can allow grids with unique coverages and higher resolution near coastal regions (Taeb and Weaver 2019), but it requires a separate simulation on a large domain to provide forcing at the boundaries of the nested grid. Thus, although these methods can reduce computational costs while providing higher resolution near a region of interest, they have limitations with respect to effort required during and between simulations.

This research will leverage a recent advancement, in which meshes are switched during a single simulation of a tropical cyclone (Thomas et al. 2021). This approach uses meshes with different resolutions during different parts of the storm. When the storm is in the open

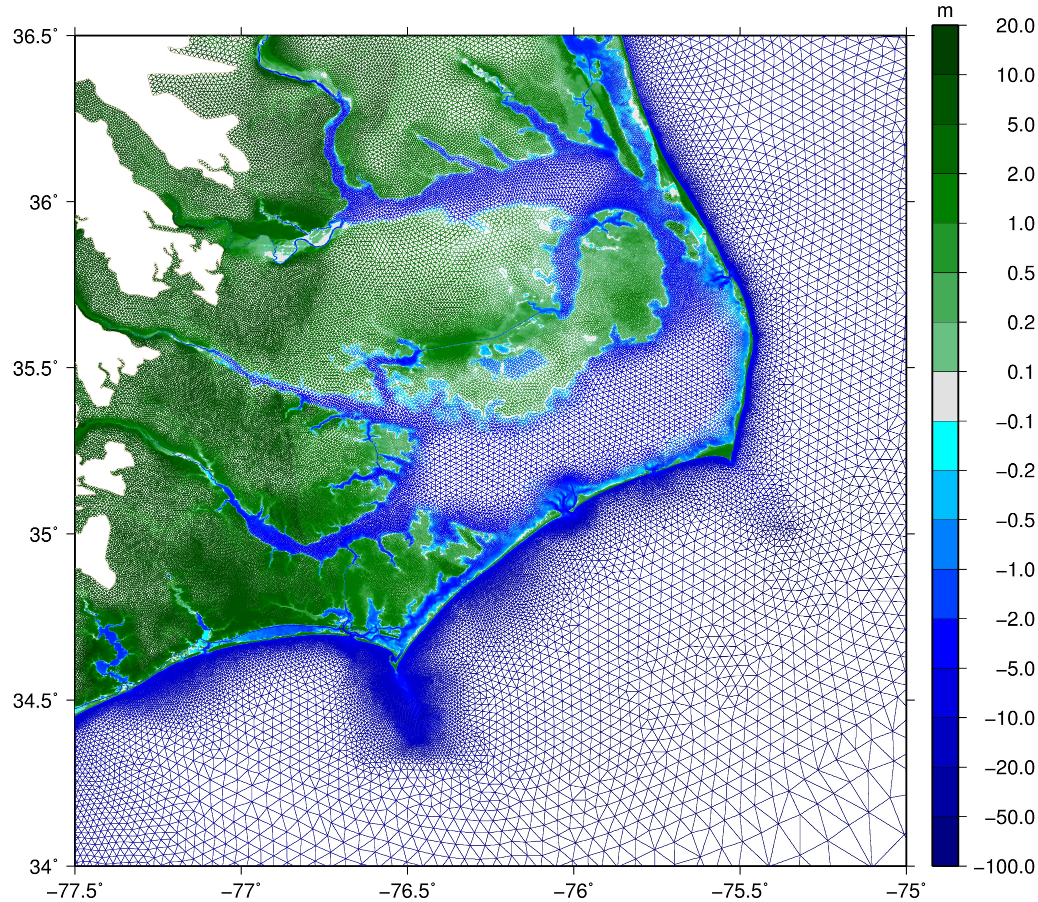

Figure 1.2: Example of a high-resolution ADCIRC mesh for North Carolina. Colors show the elevation of the ground surface in meters relative to NAVD88. The full mesh extent is not shown, but it includes the western North Atlantic Ocean, Gulf of Mexico, and Caribbean Sea. Resolution varies from tens of kilometers in the Atlantic Ocean to tens of meters along the NC coast.

ocean, the model uses a coarse-resolution mesh with elements of relatively large sizes and with a limited representation of the coastal region. Then, as the storm moves closer to the coast, the model switches to a mesh with smaller elements and higher resolution, allowing for better representation of the coastal region. The resulting combined simulation saves on computation time compared to using a single, high-resolution mesh. Results from (Thomas et al. 2021) show a comparable accuracy (Figure 1.3) and increased efficiencies between 38 percent and 53 percent for storm surge simulations during Matthew (2016) and Florence (2018), respectively. These results show the potential benefits of using this multi-resolution approach, but there are still remaining questions to be answered on how to best implement this technology during real-time forecasts.

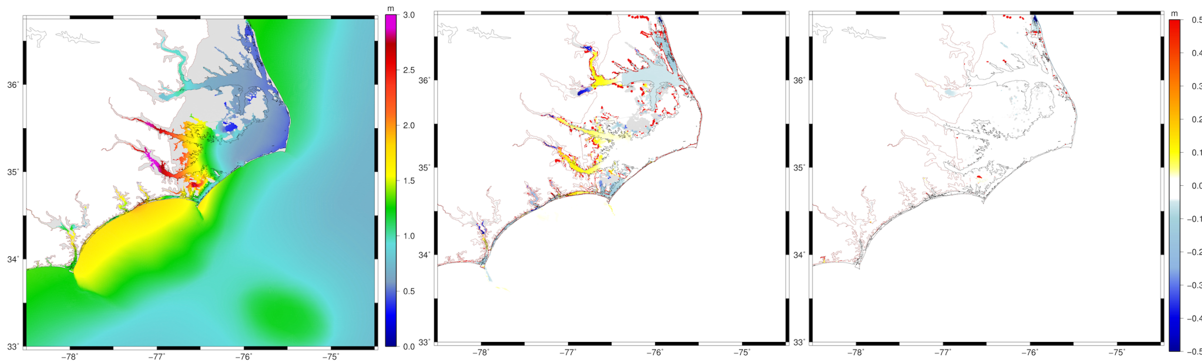

Figure 1.3: Example of the multi-resolution approach from Thomas et al. (2021). Panels are: (left) maximum water levels from a fine-mesh simulation of Florence (2018) in coastal NC, (center) difference in maximum water levels between simulations using coarse and fine meshes, and (right) difference in maximum water levels between simulations using mixed and fine meshes. By switching early in the storm, the mixed-mesh simulation maintains accuracy of coastal flooding predictions.

A key unknown is when to switch between meshes. For any storm simulation, there must be an ideal switch, which includes the coarse mesh for as long as possible to optimize efficiency, but which also includes the fine mesh for as long as necessary to optimize accuracy. To identify the ideal switch, we must consider a broader question of: when are storm effects first seen at the coast? As a storm approaches landfall, its high winds will extend into coastal regions, thus creating large waves and higher water levels. However, these effects will vary spatially, due both to proximity to the storm and configuration of the coast. Can we identify the onset of these effects as a storm approaches? And can we use these effects to design effective triggers for switching between meshes?

The goals of this research are to explore the onset of storm effects in coastal regions and then identify triggers for switching between meshes during storm surge forecasts. *It is hypothesized that, if the trigger is selected as the earliest sign of high winds or water levels, then the simulation will maintain its accuracy during the storm while receiving an efficiency boost before the storm.* This hypothesis will be tested via simulations of Florence (2018), which caused widespread flooding in coastal North Carolina (NC). By using observed time series of winds and water levels at coastal and inland stations, we will consider a range of possible triggers for switching meshes. Additionally, by varying the time of switching both near and several days before landfall, we will consider a range of trade-offs between accuracy and efficiency. The results of this research will better inform the forecasts of storm surge and coastal flooding in NC and elsewhere.

CHAPTER

2

BACKGROUND

This research will leverage and extend the state-of-the-art system for modeling storm surge and coastal flooding. As mentioned previously, the ADvanced CIRCulation (ADCIRC) model can predict coastal circulation and flooding with high accuracy, including for real-time forecasting. ADCIRC describes the coastal environment with unstructured, finite-element meshes, which can vary in their resolution of flow pathways and barriers along the coast. A new technology was developed recently to allow for switching between meshes during a single simulation, so the highest resolution is used only when the storm affects the coastal region of interest. This new technology has the potential to greatly improve the efficiency of the model, and speed up the delivery of forecast guidance provided to stakeholders.

In this chapter, these models and switching technologies are described, along with how they have been used recently to improve storm surge predictions. However, there are remaining questions about when storm effects are evident at the coast and thus when to switch meshes. These questions will then be expanded to motivate the research goals and objectives in this thesis.

2.1 Study Area and Storm

To explore effects of a real storm in a real setting, this research will consider the coastal flooding due to Florence (2018) in coastal North Carolina (NC). This region is highly susceptible to tropical cyclones, with an average of one landfalling storm every 2.5 years (State Climate Office of North Carolina 2017). Florence was especially devastating due to its strength and slow forward speed, which caused widespread flooding along the open coast of Onslow Bay and the inland river estuaries to the west of Pamlico Sound.

2.1.1 North Carolina

The onset of storm effects will be explored in coastal NC, which is well-suited for this purpose due to its varying shelf width, complex coastline, extensive inland water bodies, and low-lying floodplains. The continental shelf extends offshore by a distance of about 100 km in Onslow Bay near Wrightsville Beach, NC, but is much narrower to the north, where it extends only about 50 km offshore of Cape Hatteras. The coast is characterized by barrier islands with managed beaches and dunes; these islands are dotted with coastal communities. Along the southern coast, the lagoons are relatively narrow and include the managed intracoastal waterway, while along the northern coast, the Outer Banks protect the larger Pamlico and Albemarle Sounds. These sounds are very shallow, with typical depths less than 2 m. Several rivers, notably the Pamlico, Neuse, Roanoke, and Chowan Rivers, empty into the sounds; during storms, water can be pushed up the sounds and estuaries into these rivers. Throughout the region, the floodplains have a low gradient in their topography, and thus storm waters can be pushed inland by several kilometers.

Coastal NC has been affected by numerous storms over the years, including several notable storms in the past decade: Irene (2011), Sandy (2012), Arthur (2015), Matthew (2016), Florence (2018), Dorian (2019), and Isaias (2020). Depending on the storm track and intensity, these storms can have varying effects on coastal water levels (Cyriac et al. 2018). There have been extensive efforts to predict these effects, including via real-time forecasting (Fleming et al. 2008; Dresback et al. 2013) and subsequent dissemination of guidance to decision makers (Rucker et al. 2021).

2.1.2 Florence (2018)

For the analyses in this thesis, Florence (2018) will be used as forcing for simulations of coastal flooding in NC. Florence developed off the west coast of Africa on 31 August 2018

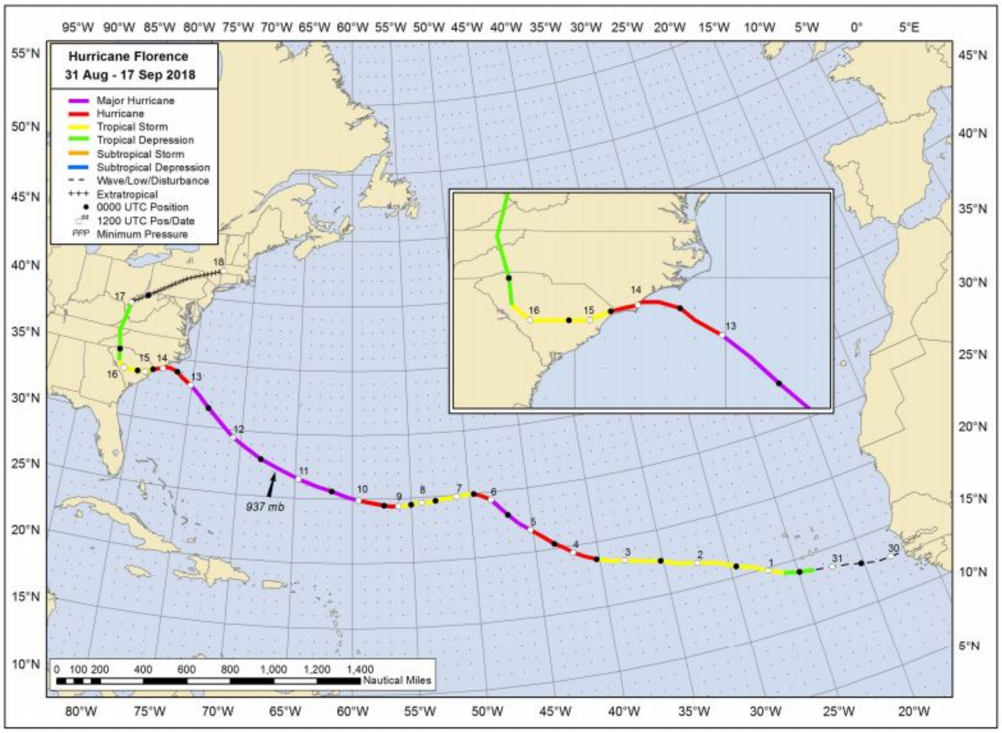

Figure 2.1: Best track for Hurricane Florence, based on analyses from the NOAA Weather Prediction Center (National Hurricane Center 2018).

and was upgraded to hurricane status 4 September 2018 when winds speeds reached 33 m/s. Within 30 hr, it intensified to a Category-4 hurricane with wind speeds of 59 m/s. However, in the next 12 hr, it weakened and slowed down and briefly became a tropical storm, before strengthening again. During this time, track estimates remained steady with landfall predictions leading to the U.S. Atlantic coast near the Carolinas. For the next five days, the storm continued along this track and maintained its status as a major hurricane until shortly before landfall, when it weakened as it approached the NC coast.

Florence made landfall near Wrightsville Beach, NC, at 1115 UTC on 14 September 2018. At landfall, it was a Category-1 storm with wind speeds recorded at 43 m/s (National Hurricane Center 2018). It moved slowly and remained near the coast for more than a day. Shortly after landfall, it reduced to a tropical storm and moved slowly over South Carolina; later, it moved farther inland and became extra-tropical about four days after landfall (Stewart and Berg 2019). Florence was a major rainfall event, as it produced large amounts of precipitation leading to more than 76 cm of rain in some locations. However, in this thesis, the focus will be on the storm's effects on coastal and inland water levels.

As it made landfall, Florence created devastating storm surge in coastal NC, both at the open coast and within the sounds and estuaries. On the open coast, surge was estimated about 1.5 to 2.5 m above ground level at Onslow Bay and surrounding counties (Stewart and Berg 2019). Due to the storm's slow forward speed, this storm surge interacted with several parts of the tidal cycle, thus increasing the maximum water elevations at many locations. Maximum storm surge inundation was estimated to be 2.5 to 3.5 m above ground level along the Neuse River and its tributaries. The highest levels of inundation were recorded near the community of New Bern along the Neuse and Trent Rivers. Within other parts of Pamlico, Beaufort, and Hyde counties, storm surge inundation levels were recorded about 1.5 to 2 m above ground level. Florence also produced significant amounts of damage to buildings, tress, and other structures throughout the area (National Hurricane Center 2018).

2.2 Storm Surge Prediction

To explore the prediction of storm surge and coastal flooding, this research will use the ADvanced CIRCulation (ADCIRC) modeling system. ADCIRC represents the coastal environment with unstructured, finite-element meshes that can vary resolution over several orders of magnitude, including down to 10 m in small-scale flow pathways and barriers in coastal regions. The ADCIRC Prediction System (APS) has been developed to provide forecasts of storm surge in a quasi-operational setting.

2.2.1 ADCIRC

For Florence or any tropical cyclone, the prediction of coastal water levels requires the understanding of complex interactions between and input data about tides, atmospheric wind speed and surface stresses, wave energy, and the detailed ocean floor bathymetry and coastal topography. All of these factors vary both spatially and temporally and are specific to a given tropical cyclone. To represent this variability, ADCIRC solves the two-dimensional, depth-integrated, shallow water equations for conservation of mass and momentum (Westerink et al. 2008; Blain et al. 1998). The equation for conservation of mass is restated as the generalized wave continuity equation (GWCE) (Kinnmark 1986; Luettich et al. 1992). ADCIRC solves these equations for water surface elevations and depth-averaged ocean current velocities at specific locations by discretizing the domain using the finite-element method. This model has been used both for forecasting water levels in real-time simulations during hurricane season (Fleming et al. 2008; Mattocks and Forbes 2008; Blanton et al. 2012), and also for hindcast simulations for investigative and research purposes such as evaluating the accuracy of predictions, or understanding pollutant and/or sediment transport (Riverside Technology and AECOM 2015; Akbar et al. 2017; Forbes et al. 2010). After Sandy (2012), a study was initiated for identifying flood risks, and to aid planning on how to reduce risk from future events which included the use of floodplains and evacuation planning, to nature-based and blended solutions (U.S. Army Corps of Engineers 2015).

Perhaps the most detailed and important component to ADCIRC is the unstructured, finite-element mesh that describes the region of interest at varying levels of spatial resolution. Within most meshes, elements located in the open ocean are typically larger and coarser in resolution with sizes on the scale of tens or hundreds of kilometers, whereas elements in the nearshore and those that represent coastal features are smaller or finer in resolution with sizes down to tens of meters. To properly represent these features in coastal regions, meshes often include important hydraulic features and controls along with small element sizes to allow for the model to represent overland flow, allowing for improved storm surge predictions.

Additionally, these meshes can include high-resolution coverage only in specific regions. Some meshes represent the open ocean with no coverage of coastal floodplains, such as meshes used for predictions of tides and development of tidal databases, with minimum element sizes ranging from 750 m down to 13 m (Mukai et al. 2002b; Szpilka et al. 2016). Other meshes provide focused, high levels of resolution in a single state, such as for North Carolina (NC9) (Blanton and Luettich 2008) with minimum element sizes of about 50 m, for coastal Louisiana (SL16) (Dietrich et al. 2008) with minimum element sizes of about 20 to 50 m, and for the coast of Texas (TX2008) (Kennedy et al. 2011) with minimum element sizes of about 30 m. Other meshes provide coverage of a larger region or coastline, such as the South Atlantic Bight (SABv1) (Thomas 2020) with minimum resolution of 20 m, the northern Gulf of Mexico (GOMEX) (Hagen et al. 2001; Bilskie et al. 2015) with minimum element sizes of 700 m, and the northeast Atlantic coast (NACCS) (Cialone et al. 2015) with element sizes down to 10 m.

Each mesh has its strengths and weaknesses. So-called 'open-water' meshes, which do not contain details about floodplains and inland regions, are highly efficient and useful for speeding up computation times and focusing on open water processes and circulation, but are not high-resolution enough to accurately represent the coastal floodplains. Higher-resolution meshes are highly accurate and useful for targeted research questions and design studies about a specific region and allow for high levels of detail and accuracy, but are computationally expensive during real-time forecasting settings. Some meshes attempt to bridge this gap by representing floodplains at coarse resolution; these meshes are useful for forecasts but limited in their accuracy. Thus, neither an open-water mesh or coarse-resolution floodplain coverage mesh nor a fine-resolution, state-based mesh are optimal for large-scale storm surge forecasts. This highlights a need for technologies that will bridge this gap between efficiency and accuracy for forecasting.

2.2.2 ADCIRC Prediction System

Over the years, ADCIRC has evolved to support real-time forecasting via the ADCIRC Prediction System (APS) (Fleming et al. 2008). This system is fully automated, and it detects when a new set of atmospheric forcing is available (either for a tropical cyclone or during normal conditions), downloads and processes the atmospheric forcing to be compatible with ADCIRC, submits and monitors simulations in a high-performance computing environment, and post-processes and disseminates the results. During a tropical cyclone, the APS uses forecast advisories from the National Hurricane Center (NHC) (Mattocks and Forbes 2008). These advisories are generated every 3 to 6 hr during an event and contain predictions about storm parameters like location of the storm center, maximum observed wind speed at an elevation of 10 m, radius to maximum winds, and central pressure (Fleming et al. 2008). These parameters are used inside ADCIRC to construct a parametric vortex model for the surface pressures and wind velocities (Gao 2018), which are then used as forcing to the ADCIRC simulations. Although these predictions are only as accurate as the NHC advisory parameters, the system has been applied for successful forecasts for every major storm in the last 10 to 15 years, including Gustav (2008) (Forbes et al. 2010), Irene (2011) (Dresback et al. 2013), and Isaac (2012) (Dietrich et al. 2013).

The APS uses results from the end of a previous simulation to 'hot-start' its next simulation, and continues the same simulation over the lifetime of the storm. This is done by using three types of simulations: (1) hindcasts, which are forced with known information from after the storm; (2) nowcasts, which are forced with currently known storm information between advisories; and (3) forecasts, which are forced with predicted information from the advisories. During a tropical cyclone, the APS will generate an initial hot-start file from a previous ADCIRC hindcast simulation. Then the storm simulations can be started. For the first advisory, the hindcast results are used to hot-start the forecast, which can run for 3 to 5 days. For subsequent advisories, nowcasts are needed to represent the 6 hr between each advisory.

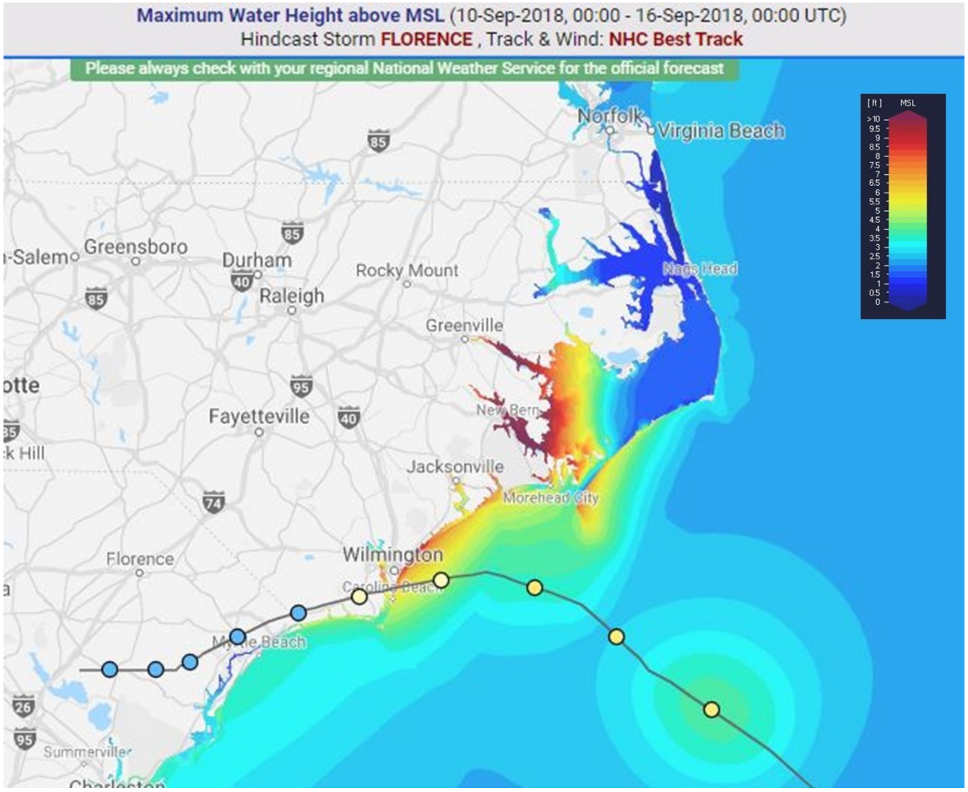

Figure 2.2: Water level predictions using the best track hindcast data from NHC for Florence provided by CERA (Coastal Emergency Risks Assessment 2019).

For each advisory, more than one forecast can be simulated with perturbations to the storm, including changes in predicted track directions by veering them to the left or right of current predictions. For example, during Florence, the APS ran simulations from NHC advisory 46 through its final advisory 68. Several of the forecasts included ensemble members in which the track was veered left by 50 percent or right by 50 percent. The results from these forecasts, as well as the post-storm hindcast, can be shared with emergency managers, other stakeholders, and the general public, e.g. via a web service (Figure 2.2).

When a tropical cyclone track becomes more understood, with less uncertainty in the region of landfall, different meshes can be used that allow for better representation of the coastal landscape. Previously, switching from one mesh to another required starting over from scratch for a given forecast, because each component of the simulation needs to be hot-started from information on the same mesh. Notably, the hindcast simulation can represent several weeks or months, and thus it can be computationally costly to repeat on a new mesh. This highlights a strong motivation for switching between meshes in the middle of a forecast, so APS operators can divert the current forecast to a new mesh without restarting from scratch with a new hindcast.

2.3 Multi-Resolution Approach

To bridge the gap between accuracy and efficiency, and enable forecasts with higher resolution of coastal floodplains and inland regions, a new technology was recently developed (Thomas 2020). This new technology was developed around the idea of changing mesh resolution during a storm event. When the storm is far away from coastal regions and its coastal effects are minimal, the simulation can be started on a mesh with a relatively coarse resolution of the coastal environment. Then, as the storm approaches the coast, the simulation can be continued on a mesh with a relatively fine resolution. This so-called 'multi-resolution approach' requires a switching of meshes during the ADCIRC simulation.

This mesh switching is performed with a new tool called *Adcirpolate*, and its benefits have been demonstrated for storm simulations along the U.S. Gulf and Atlantic coasts (Thomas et al. 2021). However, its usage for forecasting is still limited, due to gaps in our understanding about when storm effects will be first observed in coastal regions. These gaps will motivate the research in this thesis.

2.3.1 *Adcirpolate*

As noted above, there is an existing capability to continue an ADCIRC simulation, e.g. by linking together the nowcast and forecast simulations in the APS. Solution data are written into a hot-start file (called `fort.67` or `fort.68` in the ADCIRC convention). These files contain information from the last time-steps of the simulation about surface water elevations, depth averaged velocities, and wet and dry states of nodes and elements, among other details. These files can then be used to restart the simulation on the same mesh.

Adcirpolate enables the mapping of solution data into a hot-start file for a new mesh. It is implemented using the Earth System Modeling Framework (ESMF) (Hill et al. 2004), which allows for efficient mapping between meshes. For the research in this thesis, the first mesh is denoted alternately as the source or coarse mesh, while the second mesh is denoted alternately as the destination or fine mesh.

The mapping is done in two steps. For points within both meshes, the solution data are mapped using a bilinear interpolation. For destination points outside the source mesh, the solution data are mapped using a nearest-neighbor extrapolation. Then a new hot-start file is written for the destination mesh. *Adcirpolate* is highly efficient (Thomas 2020).

Thus, the general procedure is to (a) start the hurricane simulation, i.e. a hindcast or nowcast in the APS; (b) write a hot-start file from the source/coarse mesh at the time of switching; (c) use *Adcirpolate* to map the solution data and write a hot-state file for the destination/fine mesh; and (d) continue the hurricane simulation, i.e. as a new forecast.

2.3.2 Previous Findings

The benefits of the multi-resolution approach have been demonstrated previously, notably by Ajimon Thomas in his dissertation (Thomas 2020) and a subsequent journal article (Thomas et al. 2021). The feasibility of *Adcirpolate* was investigated for switching between meshes during a single simulation. For simulations of Matthew (2016) and Florence (2018), simulations were switched between a lower-resolution, source mesh called HSOFS (Hurricane Surge On-Demand Forecasting System) and a higher-resolution, target mesh called SABv1 (South Atlantic Bight). HSOFS provides coverage in the nearshore at relatively coarse resolutions, with SABv1 having more detailed resolution of the nearshore and coastal floodplains. The switching time was determined based on water level time series at stations along the coastline, with elevated water levels being the only trigger (Thomas et al. 2021).

Those studies showed that the multi-resolution simulations did not sacrifice accuracy and showed gains in efficiency. Water levels were compared to 'true' predictions from a fine-mesh simulation, with correlation coefficients of $R^2 = 0.91$ and $R^2 = 0.96$ for Matthew coarse and mixed simulations, respectively, and $R^2 = 0.86$ and $R^2 = 0.90$ for Florence coarse and mixed simulations, respectively. Mixed simulations refer to simulations that ran on both the coarse and fine mesh using *Adcirpolate* to switch between meshes. Meanwhile, the efficiency gains were significant, with a gain for Matthew of 38 percent and a gain for Florence of 53 percent. The difference in the storms' approaches was the main result for the difference in efficiency savings, as Matthew remained close to the coast for a longer period of time and thus required more time on the higher-resolution mesh. However, results demonstrated the benefits of using *Adcirpolate* when running full storm surge simulations.

Further work in this area investigated the same approach using multiple switches between meshes, rather than a single switch, by splitting the high resolution SABv1 mesh using watershed boundaries as delineations. The work explored which combination of meshes to use and which storm parameters to use for determining when to switch between meshes that would result in the most gains in efficiency and the most accurate simulation results. Water levels from the forecast simulations were used to determine the trigger switching criteria. Switching from the coarse mesh to the fine mesh or subsequent fine mesh occurred when water levels became elevated above normal tidal heights at sub-mesh boundaries. Three simulations were performed on the SABv1 sub-meshes with a focus on either efficiency, accuracy or both with each simulation showing valuable gains in efficiency (61%, 29%, and 48% respectively). The simulation focused on accuracy had the best error metrics when compared to a full, fine-resolution simulation.

2.4 Motivation

Although these previous studies demonstrated the potential benefits of the multi-resolution approach, there are remaining questions about its utility in a forecast setting. These questions will motivate the research objectives and tasks in the rest of this thesis.

2.4.1 Remaining Questions

The first set of questions are about triggers for switching. Previous studies used triggers based on elevated water levels, by comparing to expected water levels from normal tide conditions. When the water levels increased to a set percentage (e.g. 10 percent, 20 percent, etc.) above the normal peak tide, then it was assumed that the storm was affecting the coast, and the simulation was switched onto the fine-resolution mesh with better representation of the coastal floodplains. Although this method worked well, there are other possible triggers that may also be suitable. These possible triggers include wind speeds and non-tidal residuals, and they may be activated at different times during the storm and at different locations along the coast. Thus, there is a need for a systematic analysis of the potential benefits of a range of triggers for switching meshes.

By addressing this first set of questions, this research will also address a larger question of: When are storm effects first seen along the coast? This question has implications beyond the numerical modeling of storm surge, notably for disaster preparedness and evacuation. If decision-makers can understand and watch for the earliest signs of storm effects, then they can better prepare and act quickly during evacuations.

The second set of questions are about the difference in mesh resolution between the coarse and fine simulations. Previous studies used HSOFS as the coarse mesh, which worked well because HSOFS represents the coastal floodplains, albeit at coarser resolution. For instance, it includes rivers, estuaries, and other channels to convey flows to inland locations.

Thus, it was a good source for the mapping of solution data. However, HSOFS is relatively expensive to run operationally, and so there is a motivation to explore the use of coarser meshes with less representation of coastal floodplains. A so-called 'open-water' mesh would be extremely efficient, but it would not allow surge into rivers, estuaries, and inland channels. This may have implications for the mapped solution on the destination/fine mesh. Thus, there is a need for studies with a larger gap in resolution and coverage, to better understand the limits of the switching.

By addressing this second set of questions, this research will also address a larger question of: What are the critical flow pathways during a coastal storm? This question also has implications beyond the numerical modeling of storm surge, notably for mitigation via built structures. In recent years, there has been momentum toward the construction of large seawalls and gates to protect large coastal communities, e.g. the proposed 'Ike Dike' in Texas (Davlasheridze and Fan 2019) and the proposed sea wall near New York City (U.S. Army Corps of Engineers 2016). If decision-makers can understand how surge propagates from the open coast to inland regions, then they can better prepare.

2.4.2 Objectives

By identifying the most appropriate trigger for a single switch from a coarse-resolution mesh (about 250,000 vertices) to a fine-resolution mesh (about 620,000 vertices) using storm parameters from Florence (2018) for wind speeds and coastal and inland water levels of values above max tide and above non-tidal residual, the optimal switching simulation will show highest gains in efficiency by switching at the latest possible time while maintaining the most accuracy when compared to a single high resolution simulation and being validated with nearby observations. To investigate the hypotheses articulated above, research will have the following objectives:

1. For predictions of Florence (2018) and its storm surge and flooding in coastal North Carolina (NC), consider simulations on a coarse mesh (EC2001), on a fine mesh (NC9), and on mixed meshes (switching from EC2001 to NC9).

2. As the storm approaches NC, evaluate the storm's effects on parameters of wind speeds, coastal water levels, and inland water levels. For each storm parameter, consider variations in value of trigger and timing of switch. For example, test simulations will be conducted with triggers of wind speeds for values of 8 m/s, 10 m/s, and 15 m/s,

and with switches immediately at the trigger, 12 hr before the trigger, and 24 hr before the trigger.

3. For each test simulation (storm parameter, trigger value, switch timing), quantify the efficiency gains by comparing to wall-clock times for a full simulation on the high-resolution mesh.

4. For each test simulation (storm parameter, trigger value, switch timing), quantify the accuracy losses by comparing to observations and model predictions throughout coastal NC.

By the end of this research, we will have a better understanding of when a storm's impacts begin being felt near the coast. Although the focus is on Florence, these ideas and methods can be applied to almost any storm in any location. Additionally, results will provide insight into improving real-time prediction simulations.

CHAPTER

3

METHODOLOGY

This chapter describes the data and methods used to explore the effects of different triggers and switching times on flooding predictions. Effects during Florence (2018) on coastal water levels will be explored via ADvanced CIRCulation (ADCIRC) model simulations on two meshes: a coarse mesh with coverage only in open water, and a fine mesh with high-resolution of the inland water bodies and floodplains of coastal North Carolina (NC). A single simulation on the fine mesh will be treated as 'truth' for later comparisons, but the focus will be on exploring the performance of simulations using a single switch between the coarse and fine meshes. We will examine triggers from winds speeds, coastal water levels, and inland water levels, which will be identified from results from the 'truth' simulation. Because Florence was a shore-normal and relatively slow moving storm, investigations will also include buffer times from given triggers. All simulations will be analyzed for how well they improve upon the 'truth' or fine simulation, in terms of both accuracy and efficiency.

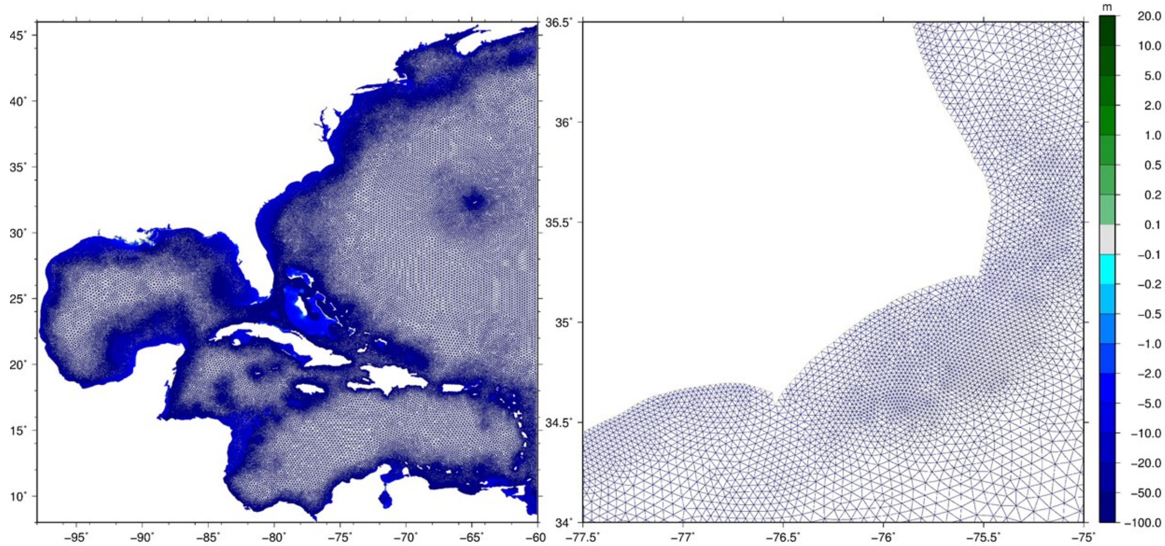

Figure 3.1: EC2001, with elements colored to show their ground surface elevations (m relative to NAVD88). The panels are (left) full mesh and (right) zoom of coastal NC.

3.1 Model Simulations

3.1.1 Unstructured Meshes

To analyze the accuracy and efficiency of mixed-mesh simulations of Florence (2018), we use two meshes. The coarse mesh is an 'open-water' mesh, so it has a relatively coarse representation of the offshore regions, and it does not include any coastal floodplains. The fine mesh was developed for flood risk mapping studies in North Carolina (NC), and has a relatively fine representation of the bathymetry and topography throughout the region. This combination of meshes will allow for efficiency gains when the storm is offshore, while maintaining accuracy as the storm makes landfall.

The coarse mesh is the Eastcoast 2001 (EC2001) mesh developed for the U.S. Army Corps of Engineers (Mukai et al. 2002a). It includes the western North Atlantic Ocean, Gulf of Mexico, and Caribbean Sea (Figure 3.1). It has an open ocean boundary on its eastern edge at the 60°W longitude. This mesh contains 254,629 vertices and 492,182 elements. It has a minimum element size of 1 to 4 km along the coast, and a maximum element size of about 25 km in the open ocean (Mukai et al. 2002a). The shoreline information was provided

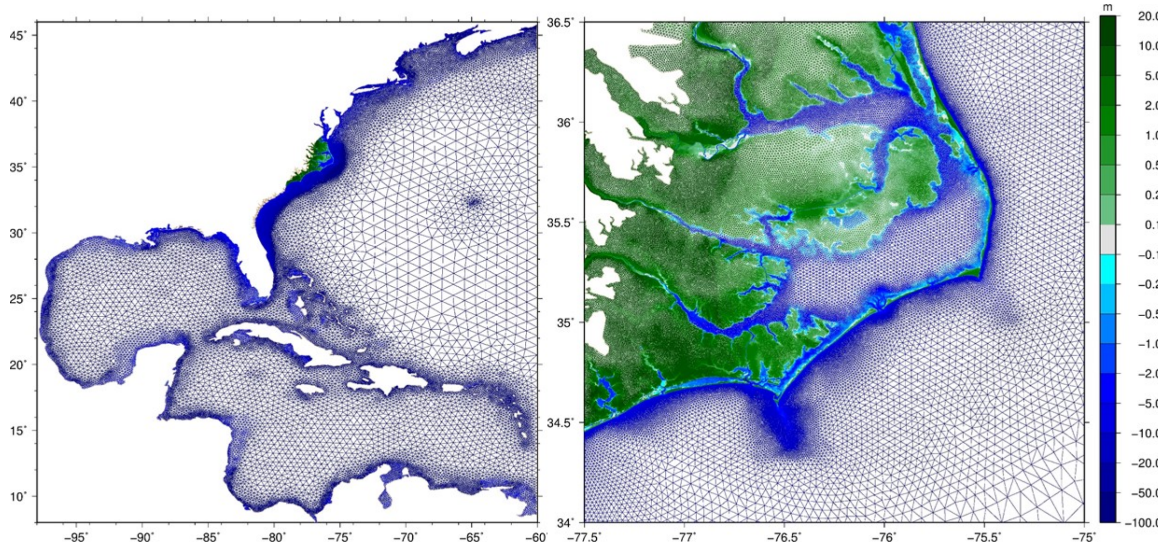

Figure 3.2: NC9, with elements colored to show their ground surface elevations (m relative to NAVD88). The panels are (left) full mesh and (right) zoom of coastal NC.

by the Defense Mapping Agency's World Vector Shoreline database. Bathymetric data were provided from ETOPO5, the Digital Nautical Charts (DNC) (National Imagery and Mapping Agency), and NOS raw sounding bathymetric database (NOS Hydrographic Survey Data) (Mukai et al. 2002a). Rivers, estuaries, and other inland features beyond the shoreline are not represented within the mesh. EC2001 was selected as the coarse mesh for this study because: (a) its resolution is relatively constant along the entire U.S. coast, which is appropriate for simulations of the early part of a storm when the eventual landfall is uncertain; and (b) its relatively small size and coarse resolution will allow for fast simulations.

The fine mesh is the North Carolina (NC9) mesh developed by researchers at the Renaissance Computing Institute (RENCI) (Blanton and Luettich 2008). It is a high-resolution mesh covering the NC coastal region. It extends from the 60°W longitude to the U.S. mainland, including into NC to the 15-m contour, allowing for storm surge flooding predictions (Figure 3.2). It resolves major bathymetric and topographic features relative to NAVD88 (North American Vertical Datum 1988), such as inlets, dunes, and rivers as identified through satellite images, NOAA charts, and several digital elevation models or DEMs. Conversion from NAVD 88 to mean sea level (MSL) was done for use in the ADCIRC model. The mesh contains 622,946 vertices and 1,230,430 elements within the v9.15 mesh. About 50 percent of the elements have spacings of 50 to 500 m, with the smallest elements in the NC coastal region (Blanton and Luettich 2008). NC9 has been used successfully for studies of storm surge and flooding during Irene (2011) (Dresback et al. 2013), Arthur (2015) (Cyriac et al. 2018), and other storms. NC9 was selected as the fine mesh for this study because its resolution is relatively high in coastal NC, and thus it will allow for accurate predictions of coastal circulation and flooding during Florence.

It is emphasized that, although NC9 is about 2.5 times larger than EC2001 in terms of numbers of vertices and elements, its resolution is at least 20 times higher in the NC coastal region. Furthermore, NC9 represents the inland water bodies and floodplains behind the coastline (comparing Figures 3.1 and 3.2). Thus it is expected that NC9 will be a good target/destination mesh for switching.

3.1.2 ADCIRC Settings

ADCIRC is used to simulate Florence's effects on coastal water levels. These simulations were conducted in two stages: a spin-up simulation with forcing from only tides, to allow tides to ramp to a dynamic equilibrium; and a storm simulation with forcing from tides, surface atmospheric pressures, and wind stresses. The spin-up simulations covered 15 days from 23 August to 07 September 2018, while the storm simulations covered 9 days from 07 September to 16 September 2018. For the 'truth' simulation, both stages (spin-up and storm) used NC9. For the mixed-mesh simulations, EC2001 was used for the spin-up and the first few days of the storm, and then the simulation was switched onto NC9 for the last few days of the storm.

Most ADCIRC parameters were set the same for the simulations on the two meshes. Eight tidal constituents were used as forcing for the time period of 23 August to 16 September 2018. These tidal constituents included the M_2 principal lunar semi-diurnal, S_2 principal solar semi-diurnal, N_2 larger lunar elliptic semi-diurnal, K_2 luni-solar semi-diurnal, K_1 and O_1 lunar diurnal, P_1 solar diurnal, and Q_1 Larger lunar elliptic diurnal (National Oceanic and Atmospheric Administration 2018b). The time step was 1 s, the finite amplitude terms were enabled in the governing equations, wetting and drying were enabled, and bottom drag was represented with a depth-dependent quadratic friction law based on Manning's n values. Before each simulation, the water levels were increased by 0.18 m throughout the domain to adjust to the NAVD88 vertical datum.

However, some ADCIRC parameters were set differently. For simulations on EC2001, the advective terms were disabled in the governing equations, and constant values for

Manning's $n = 0.022$ and horizontal eddy viscosity of 20 m^2/s were applied everywhere in the domain. For simulations on NC9, the advective terms were enabled, the horizontal eddy viscosities used a two-stage scheme with values of 10 m^2/s in open water and 20 m^2/s in the nearshore and overland, and the Manning's n values varied spatially as derived from land-use/land-cover data.

For atmospheric forcing, the storm simulations used a data-assimilated product from Oceanweather Inc. (OWI). Surface atmospheric pressures and wind velocities (reported with 15-min averaging and at 10 m elevation) are provided on regular grids and regular intervals during the storm. The basin-scale grid has coverage from 5°N to 47°N and from 99°W to 55°W with a spatial resolution of 0.20°. Within this broader coverage, a nested, region-scale grid has coverage from 31°N to 37°N and from 82°W to 74°W with a spatial resolution of 0.05°. Both grids have data from 0000 UTC 07 September 2018 until 0000 UTC 18 September 2018 (Thomas et al. 2021). These data are interpolated onto the unstructured meshes described above. The wind velocities are converted to stresses by using the drag law from Garratt (1977) with an upper limit of $C_D \leq 0.002$.

3.2 Wind Speeds and Water Levels

Florence was described by a wealth of observations at buoys, temporary gauges, and permanent stations in both the open ocean and along the coast of NC. These observations describe the evolution of the storm's effects on wind speeds and water levels as it made its approach and then made landfall. For this research, these observations will be used for two purposes: (1) the data will be used to quantify the performance of the mixed-mesh simulations, and (2) the locations will be used to identify triggers for switching meshes. For the first purpose, it was necessary to remove any observations that had obvious errors in vertical datums or included processes that will not be modeled. For the second purpose, it was necessary to use the locations of the observation stations to interpolate results from the 'truth' simulation, and then then classify the locations based on proximity to the coast. The tasks for these purposes are described in the following subsections.

3.2.1 Observations as Comparison for Model Performance

The coastal effects of Florence are described by observations from several sources. These observations describe the wind speeds and water levels in coastal NC during the storm. Observations of wind speeds are available at six buoys operated by the National Oceanic and Atmospheric Administration (NOAA) National Data Buoy Center (NDBC) (National Oceanic and Atmospheric Administration 2017) and six long-term stations operated by NOAA National Ocean Service (NOS) (National Oceanic and Atmospheric Administration 2018a). These stations are distributed throughout coastal NC (Figure 3.3). The wind speeds are reported with an averaging period of 10 min and elevation of 10 m. At the buoys, the wind speeds are reported every 1 hr, whereas at the stations, the wind speeds are reported every 6 min. In this research, while wind speeds from observations were not used as comparisons for the model performance (because the OWI products are already data-assimilated), their locations will be useful later when identifying triggers.

Observations of water levels are available at six long-term stations operated by the NOAA NOS (National Oceanic and Atmospheric Administration 2018a), 120 permanent stations operated by the U.S. Geological Survey (USGS), and 25 rapidly deployed gauges operated by the USGS (U.S. Geological Survey 2020b,a). These stations are also distributed throughout coastal NC (Figure 3.3), both at the open coast and at locations in river estuaries and other inland regions. These water levels are reported relative to a vertical datum of NAVD88 and at frequencies of every 6 min for the NOS stations, every 30 sec for the USGS stations, and every six to 15 min for the USGS gauges. These observations were interpolated in time to correspond to the model predictions, before any other analyses were conducted.

Further preparation was taken with the water level observations for quality control and quality assurance. If a station had water level observations with a large vertical offset from the model predictions, then its vertical installation was assumed to be mis-surveyed, and it was removed from our analyses. Additionally, as mentioned previously, Florence was a slow-moving storm, which led to significant amounts of precipitation in the landfall area. As a result, some observation time series showed elevated water levels during the later portion of the storm. Water level observations were ignored if they did not record the storm peak or they showed elevated water levels due to wave run-up or freshwater run-off, neither of which will be modeled herein. Thus, of the initial 157 stations identified with observations of water levels, 54 stations were removed (including data from wind-only stations), and our error comparisons will consider observed water levels from a total of 103 stations.

3.2.2 Classifying Locations of Observations

To identify triggers (in the next section), it will be necessary to examine the wind speeds and water levels at specific locations throughout coastal NC. Although the observations could be used for this purpose, they are imperfect because (a) they have different vertical datums and

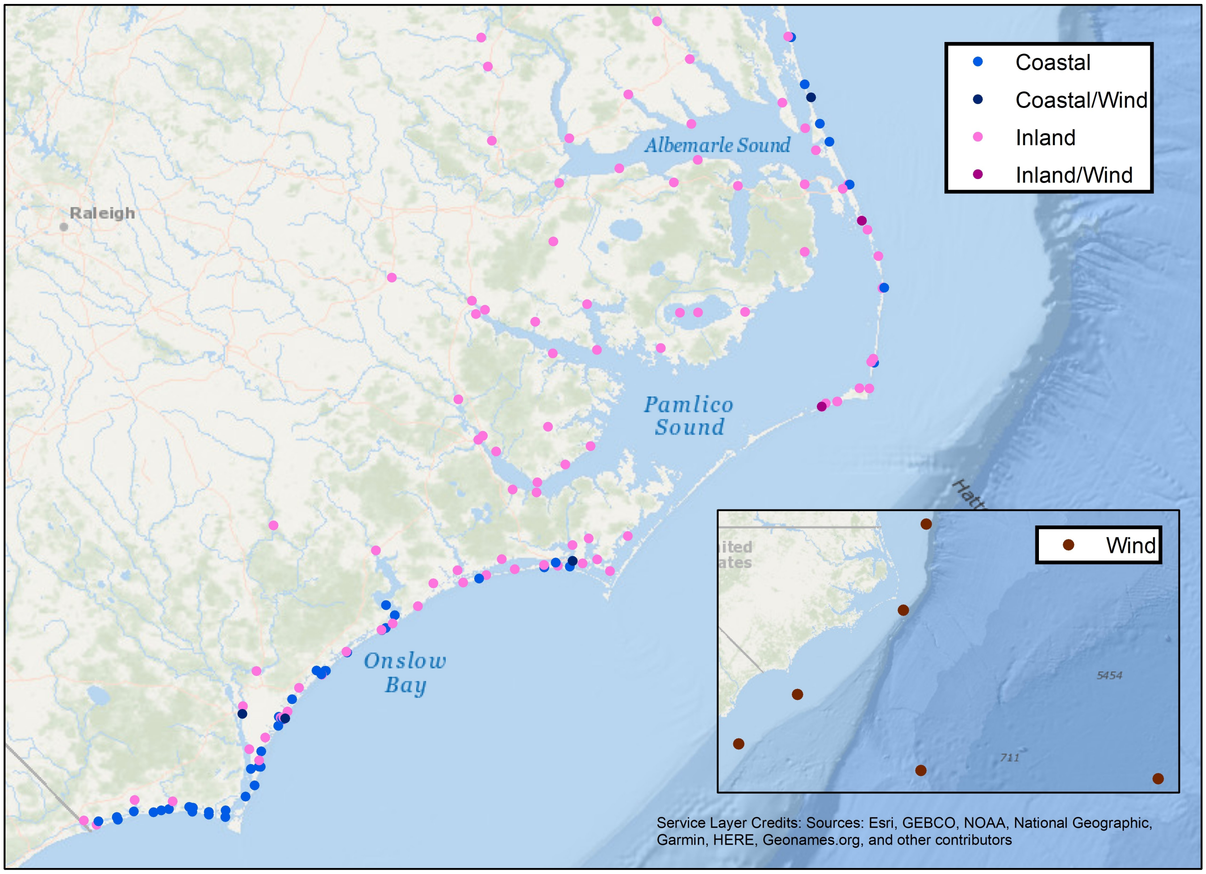

Figure 3.3: Locations of observation stations, categorized as either wind, coastal or inland based stations.

errors, due to being collected by different agencies; and (b) they are not distributed evenly throughout the region affected by the storm. Thus, to examine wind speeds and water levels from a consistent source and a larger data set, we interpolated results from the ADCIRC 'truth' simulation at the observation locations. This has the benefit of expanding the data used to identify triggers; instead of having only a wind speed or water level at an observation station, we can interpolate both parameters from the ADCIRC 'truth' simulation.

For the following water level analyses, it was desired to classify the observation stations based on their location (coastal or inland). This classification was based on time series of tides, as well as the geographic location of the station relative to the open ocean. Stations were classified at the coast if they were on the open coast, on the back side of a barrier island in close proximity to a tidal inlet, or slightly upstream in an estuary with minimal barriers to the ocean, and if they had a tidal magnitude greater than 1 m. Stations were classified as inland if they were located on land and had no tidal signal (even if they were near to the coast), or if they were located higher in estuaries, lakes, or inland streams, regardless of their tidal signal presence and magnitude. There were 56 'coast stations' and 89 'inland stations' (Figure 3.3).

3.3 Storm Parameters for Switching

Several storm-related parameters can be used as triggers for switching meshes during simulations. These parameters can include: time, date and location of storm center; storm direction and forward speed; maximum sustained wind speeds; storm eye diameter; and the radius to maximum winds in four quadrants (NE, NW, SE, SW). These information are available in the forecast advisories during the storm. Then, several days or months after the storm, hindcasts are made available with more detailed and accurate information including: storm track, size, intensity on the Saffir-Simpson scale, forward speed. Similarly to the storm itself, observations of waves and water levels may be available during or after the storm. At permanent observation locations, the observations may be available in real-time during the storm, whereas at rapidly deployed gauges, the observations may be released several months after the storm.

For the following analyses, we will consider two of these parameters, specifically the wind speeds and water levels at the coast and at inland stations. These parameters are selected partly because they can be available in real-time during future storms, and partly because they can be generated from the ADCIRC predictions – wind speeds and water levels can be interpolated at a synthetic station from the model results for a previous forecast advisory, and then used to inform switching of meshes before the next forecast advisory. In this study, as described above, we will leverage a large suite of observation station locations that include data from several sources and data types, resulting in 157 station locations within the NC9 mesh. These observations will then be used to identify triggers for wind speeds and water levels.

Each test with these triggers will additionally have buffers of 0 hr, 12 hr, or 24 hr before each trigger is activated. For example, the first trigger for wind speeds will be identified, and then tests will be designed to switch as the trigger is activated (with a 0-hr buffer), a half-day before the trigger is activated (with a 12-hr buffer), and a full day before the trigger is activated (with a 24-hr buffer). Buffers are implemented based on the uncertainty of using these triggers, as well as due to the shore-normal approach of Florence to coastal NC.

It is possible that, when these selected triggers are hit, the effects of Florence may already have been experienced at other non-station locations within our region of interest.

3.3.1 Wind Speeds

There are a few methods to suggest values for the wind speed trigger. An obvious method is to use the wind speeds from the Saffir-Simpson scale for tropical cyclones, specifically the wind speeds that correspond to Category 1, 2, etc. However, this method is imperfect because many storms are weakened as they make landfall. Their wind speeds may be reduced to the status of a tropical storm or minor hurricane, even as they push and high surges along the coast. This was true during Florence; at observation locations along the NC coast, the wind speeds did not reach the lowest category on the Saffir-Simpson scale.

Instead, we examined the time series from all 157 station locations and used our engineering judgment as to when wind speeds became elevated above a 'base' level. It was determined that, overall, a majority of stations had base wind speeds at or below 5 m/s. Higher values indicated winds from Florence were beginning to be felt in the area. As wind speeds became elevated above these base levels and storm presence became recognized within the time series plots, wind speed triggers were selected. Based on the wind speeds from the high-resolution, 'truth' simulation, the following wind triggers were selected:

Trigger	Wind Speed
WT1	$W > 8$ m/s
WT2	$W > 10$ m/s
WT3	$W > 15$ m/s

These triggers will describe different portions of the storm's effects on the coast, ranging from wind speeds slightly above the base value (WT1 with 8 m/s) to wind speeds closer to the level of a tropical storm (WT3 with 15 m/s).

3.3.2 Coastal and Inland Water Levels

The water level triggers were determined differently for coastal and inland stations, to reflect the storm's varying effects at these locations. In addition, for each location, one set of triggers was determined as total water levels increased above the normal tidal range, and another set of triggers was determined based on the storm surge.

The first set of water level triggers corresponds to elevated values above the base peak tidal value at each station. A separate simulation was performed on the fine, high-resolution mesh with no hurricane winds, which allowed for estimation of tidal ranges at all stations. Peak values were identified for each station, and triggers were identified when the total water levels increased during the storm to exceed this peak value by 10, 20 and 30 percent. Thus, the trigger values are different at every station, because are relative to the local tidal maximum. Thus the first set of water level trigger values are:

Triggers	Water Levels vs. Tidal Maxima
CWLT1 & IWLT1	$\eta > 1.1 \cdot \eta_{\mathrm{max}}$
CWLT2 & IWLT2	$\eta > 1.2 \cdot \eta_{\mathrm{max}}$
CWLT3 & IWLT3	$\eta > 1.3 \cdot \eta_{\mathrm{max}}$

in which the trigger acronyms refer to coastal water level triggers (CWLT) and inland water level triggers (IWLT), and the local tidal maxima is denoted by η_{max}.

The second group of triggers for water levels were determined from the storm surge, or specifically the non-tidal residual. This quantity was obtained by subtracting the water levels from the tides-only NC9 simulation from the with-storm ('truth') NC9 simulation, thus leaving only the storm effects on the water levels.

Triggers	Non-Tidal Residuals
CWLT10 & IWLT10	$\left\|\eta_{\mathrm{NTR}}\right\| > 0.3$ m
CWLT20 & IWLT20	$\left\|\eta_{\mathrm{NTR}}\right\| > 0.4$ m
CWLT30 & IWLT30	$\left\|\eta_{\mathrm{NTR}}\right\| > 0.5$ m

in which the trigger acronyms are now appended with zeros, the non-tidal residuals are denoted as η_{NTR}, and the absolute values are compared with the specific trigger values. Thus, for example, CWLT10 can be activated if the non-tidal residual increases above 0.30 m or decreases below −0.30 m. This allows for large drawdowns to also be identified.

3.3.3 Summary

These triggers lead to 45 possible simulations scenarios for switching from the coarse to fine mesh 3.1. The three wind triggers each have a set of three simulations with the offsets, whereas the coastal and inland have 6 triggers, each with three offsets. The time at which each of these triggers is activated at corresponding stations will signify when we switch between meshes.

Table 3.1: Summary of 45 possible mixed-mesh simulations, based on wind and water-level triggers and offsets.

Trigger	Value	Offsets		
WT1	$W > 8$ m/s	0 hr	12 hr	24 hr
WT2	$W > 10$ m/s	0 hr	12 hr	24 hr
WT3	$W > 15$ m/s	0 hr	12 hr	24 hr
CWLT1	$\eta > 1.1 \cdot \eta_{\text{max}}$	0 hr	12 hr	24 hr
CWLT2	$\eta > 1.2 \cdot \eta_{\text{max}}$	0 hr	12 hr	24 hr
CWLT3	$\eta > 1.3 \cdot \eta_{\text{max}}$	0 hr	12 hr	24 hr
IWLT1	$\eta > 1.1 \cdot \eta_{\text{max}}$	0 hr	12 hr	24 hr
IWLT2	$\eta > 1.2 \cdot \eta_{\text{max}}$	0 hr	12 hr	24 hr
IWLT3	$\eta > 1.3 \cdot \eta_{\text{max}}$	0 hr	12 hr	24 hr
CWLT10	$\lvert\eta_{\text{NTR}}\rvert > 0.3$ m	0 hr	12 hr	24 hr
CWLT20	$\lvert\eta_{\text{NTR}}\rvert > 0.4$ m	0 hr	12 hr	24 hr
CWLT30	$\lvert\eta_{\text{NTR}}\rvert > 0.5$ m	0 hr	12 hr	24 hr
IWLT10	$\lvert\eta_{\text{NTR}}\rvert > 0.3$ m	0 hr	12 hr	24 hr
IWLT20	$\lvert\eta_{\text{NTR}}\rvert > 0.4$ m	0 hr	12 hr	24 hr
IWLT30	$\lvert\eta_{\text{NTR}}\rvert > 0.5$ m	0 hr	12 hr	24 hr

3.4 Error Analysis

These triggers will be used to identify the appropriate times to switch meshes during simulations of Florence. The performance of these mixed-mesh simulations will be judged by both accuracy relative to observed water levels and efficiency relative to wall-clock times.

3.4.1 Accuracy

The mixed-mesh simulations will be analyzed for accuracy by comparing them against both observational data points and the fine simulation ran on the high-resolution mesh. Observation data and data from the model simulations are in the form of time series at hourly increments for water levels. All points of comparison are based on physical gauges stationed within the region of interest. The analysis of comparing our mixed simulation predictions to the fine simulation and to the observations is quantified using to statistical error metrics. The first is root-mean-square-error or E_{RMS}:

$$E_{RMS} = \sqrt{\frac{1}{N}\sum_{i=1}^{N} E_i^2}$$

With the above equation, we will have two E_{RMS} values. The first value is for the water levels on the mixed simulation compared to the fine simulation, so E_i will be fine compared to mixed. The second value will be for the mixed results compared to observations, so E_i will be observations compared to mixed.

We also have mean normalized bias, or B_{MN}:

$$B_{MN} = \frac{\frac{1}{N}\sum_{i=1}^{N} E_i}{\frac{1}{N}\sum_{i=1}^{N} |O_i|}$$

For this equation, E_i will be the same as above, depending on the comparison, with mixed results minus the fine or observations. The O_i is the absolute value of the water level for the comparison, so when comparing mixed to the fine results, it would be the absolute value of the fine water level.

In addition to these time-series-based error metrics, we will also consider a metric based on the predictions of the peak water levels. High water mark (HWM) observations will be compared to the maximum water levels produced from the mixed simulations. These will be analyzed using the coefficient of determination, R^2, to see how well observations and the fine simulation compared to the mixed simulations.

3.4.2 Efficiency

Efficiency will be measured by using wall-clock times from the component simulations. Each simulation was ran on the North Carolina State University High Performance Computing (HPC) cluster using 256 cores. It is noted that this cluster has a heterogeneous hardware, and thus some simulations were ran on older, slower cores while some were ran on newer, faster cores. Resulting wall-clock times for each simulation were reported in seconds. For the fine simulation, the full time for the simulation will be its wall-clock time, whereas for the mixed-mesh simulations, the wall-clock time will be the combined time for the coarse, *Adcirpolate* and the fine components of the simulation. Efficiency gains will be quantified by using two metrics: actual speedup, S_{actual}:

$$S_{actual} = \frac{T_{\text{fine}}}{T_{\text{mixed}}}$$

and theoretical speedup, $S_{\text{theoretical}}$:

$$S_{\text{theoretical}} = \frac{NT}{\sum_{i=1}^{n} N_i T_i}$$

in which T_{fine} is the total wall-clock time for the full high-resolution truth simulation; T_{mixed} is the total wall-clock time for the full mixed simulation including coarse, *Adcirpolate*, and fine simulations all in seconds; N is the number of vertices in the high-resolution mesh, T is the number of days (indicated in run files, for example 9 days for a typical ADCIRC simulation) on the high-resolution simulation; n is the number of component meshes used, here would only be 2; and N_i and T_i are the number of vertices and number of days respectively for simulation on each component mesh.

CHAPTER

4

RESULTS AND DISCUSSION

In this chapter, the triggers for wind speeds, and coastal and inland water levels will be used to initiate switching of meshes in ADCIRC simulations of Florence. Results will be analyzed to understand when and where each trigger was activated along the coast of NC, and then to quantify the accuracy and efficiency of the mixed simulations based on each trigger. Results will be discussed in the context of how flow pathways can affect extrapolation onto the fine mesh, and how these results can inform real-time forecasting.

4.1 Results

The performance of the mixed-mesh simulations will be compared to the full, high-resolution, fine-mesh simulation ('truth'), as well as to observations for water levels. Performance will be quantified via accuracy evaluations using error metrics of root-mean-squared error, bias, and coefficient of determination, and via efficiency gains using speed-up time calculations. Comparisons between max water level elevations on the mixed simulations and full fine-resolution simulation will also be reported and discussed.

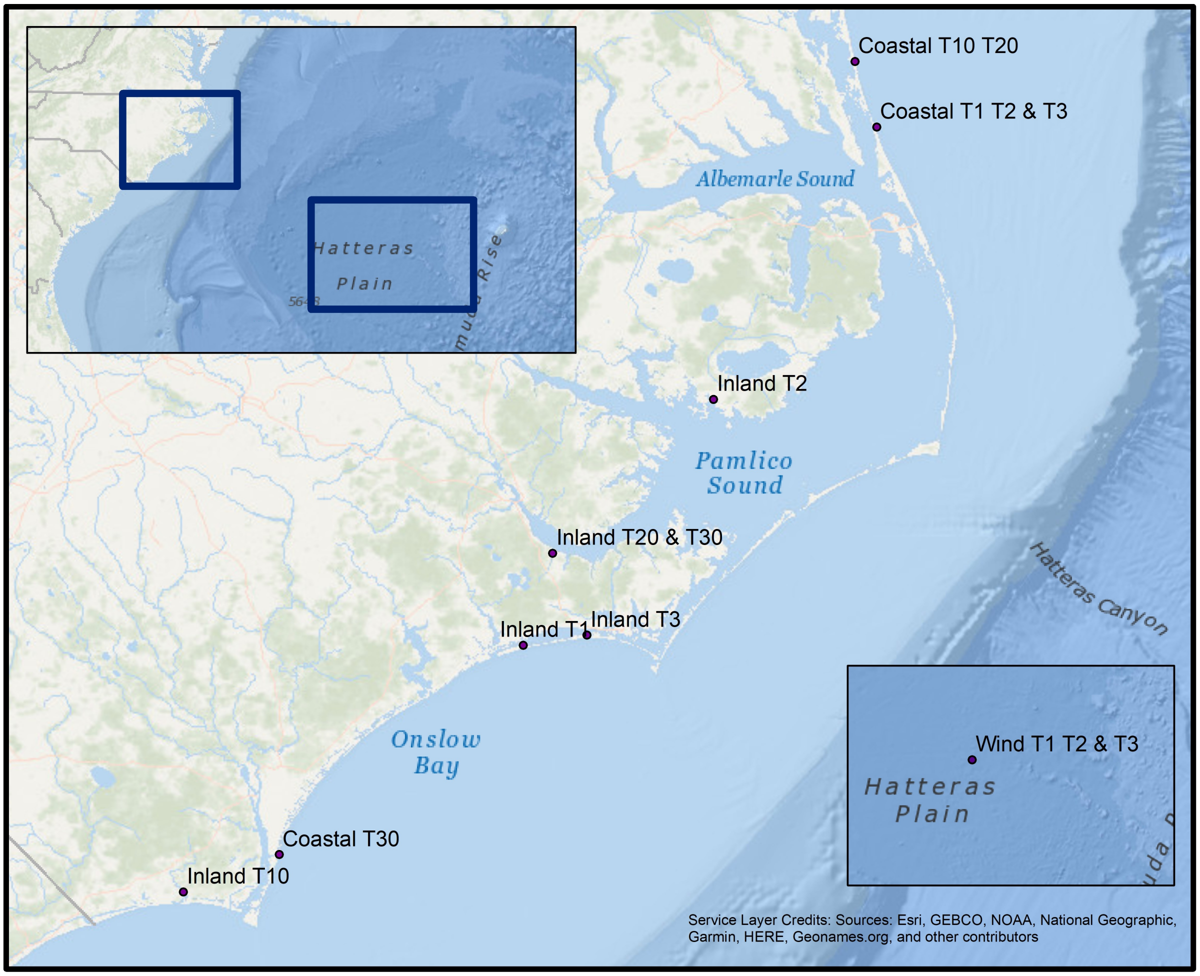

Figure 4.1: Locations of stations where each of the 15 triggers were first exceeded. Note the stations are distributed along the NC coast, not clustered near the storm's landfall location.

4.1.1 Simulations

Based on the triggers identified in the previous chapter, a set of 45 possible mixed-mesh simulations were identified (Table 3.1). (Remember the number is 45 simulations because each of the 15 triggers will have three simulations with varying offsets of 0 hr, 12 hr, and 24 hr.) The first step was to apply these triggers to the wind speeds and water levels during Florence, as computed from a single high-resolution 'truth' simulation.

From the time series results, we determined when wind speeds and water levels exceeded the identified triggers at each station. Figure 4.1 shows the location of the station where each trigger was activated during Florence. For each station in the simulation, we determined when each of our triggers were activated during the storm. Then, for each of

Table 4.1: Summary of the earliest time when each trigger was activated (date of switch), and the resulting duration in days of the coarse and fine components of the mixed-mesh simulations. It is reiterated that each mixed-mesh simulation was also run with offsets of 12 hr and 24 hr, e.g. the three simulations for WT1 would switch at day 4.75 (offset of 0 hr), at day 4.25 (offset of 12 hr), and at day 3.75 (offset of 24 hr).

Trigger	Date of Switch	Time (days)	
		Coarse (Source)	Fine (Target)
WT1	9/11/18 18:00	4.75	4.25
WT2	9/12/18 3:00	5.125	3.875
WT3	9/12/18 9:00	5.375	3.625
CWLT1	9/10/18 0:00	3	6
CWLT2	9/10/18 0:00	3	6
CWLT3	9/10/18 0:00	3	6
IWLT1	9/10/18 3:00	3.125	5.875
IWLT2	9/10/18 18:00	3.75	5.25
IWLT3	9/11/18 3:00	4.125	4.875
CWLT10	9/10/18 0:00	3	6
CWLT20	9/10/18 12:00	3.5	5.5
CWLT30	9/13/18 15:00	6.625	2.375
IWLT10	9/10/18 21:00	3.875	5.125
IWLT20	9/13/18 3:00	6.125	2.875
IWLT30	9/13/18 6:00	6.25	2.75

these timings, we identified the earliest occurrence of when the trigger was exceeded. This timing was then used to determine when we would switch from the source/coarse mesh to the target/fine mesh in each of our mixed simulations.

Then, from the earliest exceedance timings at any station, we determined the actual simulations needed to evaluate each trigger (Table 4.1). Some triggers were activated at the same times during the storm, whereas other triggers showed days of separation. For example, the possible triggers for CWLT1, CWLT2, CWLT3, and CWLT10 were activated at the same time (00:00 UTC 10 September 2018), and thus only one simulation was needed to evaluate these triggers. For future discussion, these triggers will be referred with the combined acronym of CWLT123. Thus, from the initial set of 45 possible mixed-mesh simulations, this analysis identified 36 actual simulations.

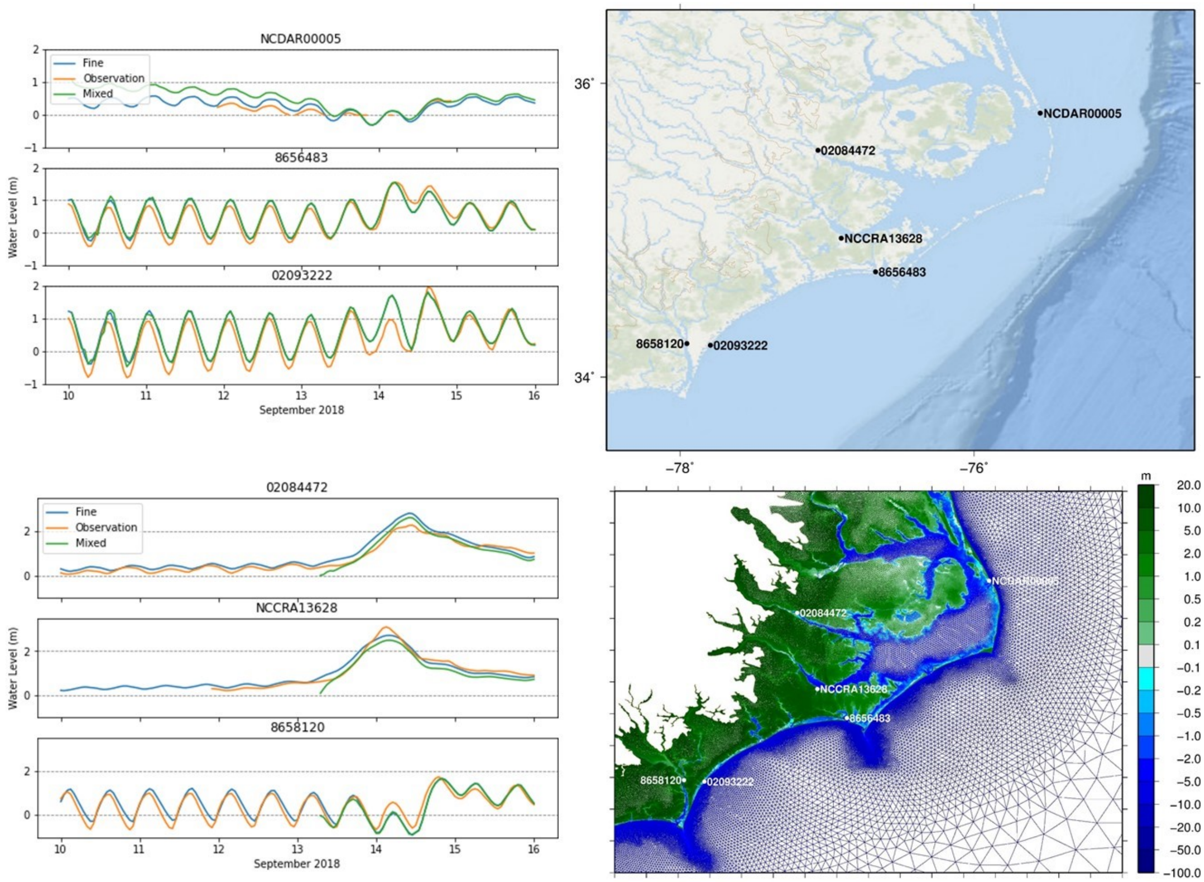

Figure 4.2: Hydrographs of water levels (m): (top) from CWLT123 coastal stations and from IWLT30 (bottom) inland stations, with corresponding locations along the NC coast and bathymetry of NC9. Note the variability in the timing of peaks of each hydrograph individually and then within each coastal and inland group.

4.1.2 Earliest Signs of Storm Effects

Using only the information about when and where the triggers were activated (Table 4.1 and Figure 4.1), we can learn a lot about how Florence's effects were first seen in coastal NC. There are wide variations in both time, as some triggers were activated several days before others, and space, as some triggers were activated in different parts of the coast.

Some triggers were activated at similar times (e.g. coastal and inland water elevated above tidal maximum), while others vary in space and time (e.g. inland and coastal water level greater than non-tidal residual absolute values). For these analyses and the implications of this research, it is important to consider what is physically happening near the coast of NC as Florence is approaching landfall. It is important to consider observations both near the coast and further inland to capture the difference in the timing of water levels rising above base levels. There is a significant temporal variation between water levels at the coast and inland stations for the non-tidal residual trigger as compared to the tidal maximum trigger. Based on the trigger values, there is a larger temporal variation from when the initial trigger for non-tidal residuals, which is 0.3 m, was activated to when the third trigger, 0.5 m, was activated for both coastal and inland stations (e.g. CWLT10 and CWLT20 were activated on 10 September versus CWLT30 activated on 13 September, and then IWLT10 activated on 10 September versus IWLT20 and IWLT30 activated on 13 September). For the other set of triggers, a day or less separates their activation.

Expanding this, both spatial and temporal variability in water level responses can be considered (Figure 4.2) for time series of water levels at selected stations for a single coastal trigger, 10 percent above the tidal maximum (CWLT1). The USGS station NCDAR00005 (northernmost in Figure 4.2) does not show a peak water level, but rather a drawdown to below sea level due to its location relative to the coast and barrier island. Moving south, NOAA station 8656483 is located at the open coast. Here, we see peak water levels of about 1.5 m at the start of 14 September, which is approximately when Florence made landfall at Wilmington (at 07:15 EDT 14 September). Moving to the NOAA station 02093222 (southernmost in Figure 4.2) shows a peak of closer to 2 m later on 14 September, which is slightly after Florence made landfall. Spatially, there are a lot of variations in what is happening along the coast as shown by these few coastal station results. Not all locations will show elevations in water levels, as seen with USGS station NCDAR00005, while some closest to landfall may show delayed peaks, such as NOAA station 02093222.

Similar behavior is seen at inland water level stations (Figure 4.2). While the hydrographs for two of these stations lack tidal signals, there are slight temporal variations in the peak water level for each station based on its location. The northernmost NOAA station 02084472 has the latest peak at around mid-day of 14 September at higher than 2 m. This later peak is likely due to the distance that the station is located inland. Moving south, USGS station NCCRA13628 had a peak of near 2.5 m early on 14 September, nearest to Florence's landfall time. Because it is closer to the sound side of the barrier islands, water forced inland by the storm reached this location prior to the northernmost station in this figure. Looking at the last NOAA station 8658120, which is closest to the location of landfall, we see a drawdown below sea level followed by a peak below 2 m later on 14 September. This rapid increase is likely due to changes in wind direction as Florence traveled over this location during 14-15 September. Again, similarly to what we saw in the coastal stations, there are subtle differences between the locations of these stations and how water levels at each location

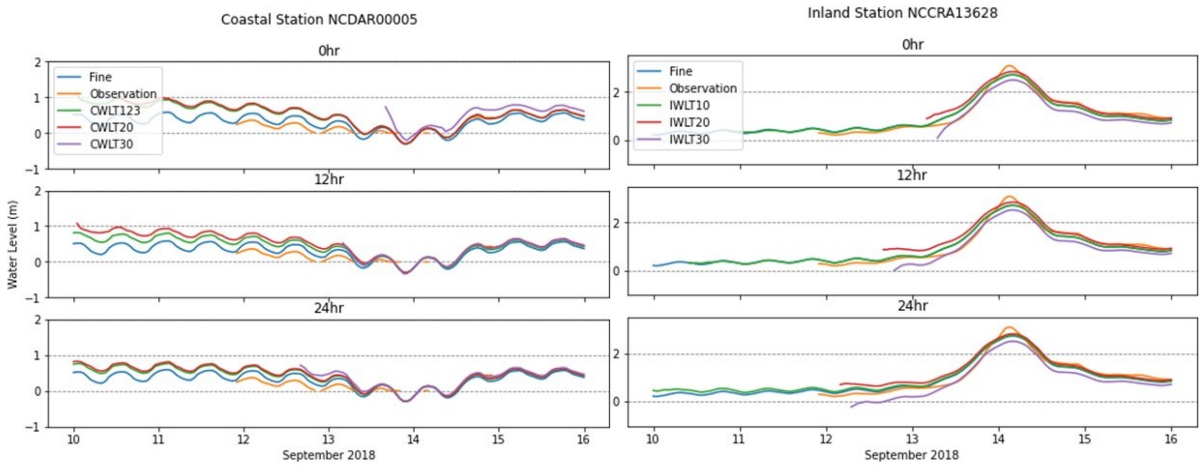

Figure 4.3: Hydrographs for coastal station NCDAR00005 for all coastal triggers (CWLT123, CWLT20 and CWLT30) and inland station NCCRA13628 for all inland triggers (IWLT10, IWLT20, IWLT30) with each offset timing, 0 hr (top), 12 hr (middle), and 24 hr (bottom).

responded as Florence approached.

The offsets for these simulations can illustrate how the timing of the switch can affect water levels on the fine portion of the mixed simulations match with the full fine simulation. Figure 4.3 shows two of the stations discussed previously, one coastal (NCDAR00005) and one inland (NCCRA13628), but now with several of the simulations and corresponding offsets. Both the coastal and inland stations show separation between water level estimates for most of the mixed simulations. The coastal station (Figure 4.3, left) with the 0-hr offset has the CWLT123 and CWLT20 overestimating to start by around 0.4 m, and the CWLT30 overestimating by closer to 0.5 m, with each catching up very quickly with the fine and observations around the time of landfall and then showing a separation again later on 14 September, with CWLT30 continuing with the largest overestimation while the others showing very minimal differences. Extending out to the 12 hr and 24 hr offset, we see similar trends with variations in overestimation to start, catch up to match the fine and observations, with minimal separation afterwards.

A similar result is shown for the inland station (Figure 4.3, right). The IWLT30 simulation appears to underestimate for all offset values with largest differences around 0.5 m prior to mid-day on 13 September, while IWLT20 shows slight over estimations by 0.2 to 0.3 m prior to the peak. IWLT20 minimizes its difference from the full fine simulation after peak water levels, while IWLT30 still remains a slight underestimation. One other interesting thing to note is that the difference between IWLT20 and the full fine result does not appear to decrease at mid-day on 13 September, as the offset is increased from 0 hr to 12 hr, but there is a decrease moving to the 24 hr offset timing, similar to CWLT20 for the coastal station.

This variation in water levels, both spatially and temporally and based on switching timing illustrates a few key points. First, Florence did not immediately cause a change in water levels everywhere along the coast, and instead resulted in spatial variations in water levels including both rising of water levels and draw downs in certain locations. Second, it is important to have a large coverage of observation data available for gathering data to perform the switching-based analysis, to understand what is happening over a large spatial scale in the region of predicted storm landfall. And lastly, the timing of the switching does matter, with some simulations needing a larger offset to allow for ample time to 'catch-up' to the fine simulation.

4.1.3 Accuracy

The accuracy of the mixed simulations will be quantified in this section, beginning with a brief discussion of comparisons of the fine high-resolution simulation with observations, and then moving on to the trigger based simulations starting with winds, then moving into water levels for coastal and then inland. First, we examine the time series at the specific stations discussed previously, and we compute error metrics via comparisons to both the full fine simulation or 'truth' and observations over the full time series. Then, we examine the maximum water levels as predicted in the full region, and we compute differences relative to the full fine simulation. We look at both because we want understand how accurate our predictions are at various spatial scales both at station locations through time series and regionally along the coast of North Carolina. The error metrics calculated include root-mean-squared error (RMSE), Bias (B_{MN}), and coefficient of determination (R^2). Smaller values closer to zero represent good values for both RMSE and Bias, with those greater than 0.5 reflecting poor ability of the model to predict the data. Whereas values closer to unity are good for R^2, with less than 0.3 showing weak or no ability of mixed simulations to predict variation in the observations or fine results.

Fine vs. Observations

To better compare later accuracy metrics it is important to have a base line of what the full high-resolution simulation results show compared to the observations. In Table 4.2, a RMSE of 0.23 m and Bias of 0.21 m with a R^2 of 0.93 for comparing the fine results to observations will provide a valuable reference to how well each mixed simulations performs. We see that

Table 4.2: Accuracy metrics of fine simulation on high-resolution NC9 mesh and coarse simulation on EC2001, compared to observation data. Metrics for the coarse simulation were computed from 4 stations that were within the coverage of the coarse mesh, whereas metrics for the fine simulation used 103 stations. No HWM (High Water Mark) data were available within the bounds of the coarse mesh, so no R^2 metric was computed.

Simulation	Observation		
	RMSE	B_{MN}	R^2
Fine	0.23	0.21	0.93
Coarse	0.46	-0.15	NA

the fine simulation over estimates water levels by 0.21 m as compared to observation data at stations used in this analysis.

Wind Triggers

Starting with the wind triggers, the water level predictions from the three sets of mixed simulations were compared to both the fine simulation and observations using time series for all stations in the analysis (Table 4.3). Comparisons of the mixed simulations to the fine simulation are very favorable, with R^2 values near unity, RMSE below 0.15 m for all switches, and bias below 0.23 m. Comparisons to observation data are not as favorable, which is to be expected, with large RMSE for all simulations with values above 0.21 m and bias varying from 0.59 m to 0.20 m. The bias results for WT2 are least favorable with values over 0.50 m. However, R^2 values were all above 0.90 for each simulation. Based on the results in this table, the best-performing simulation would be WT3 with the 24 hr offset. Comparing to the values shown in 4.2, WT1 and WT3 RMSE and Bias, along with R^2 have results more favorable to the full fine simulation with WT3 showing better comparisons to the observations.

Trends show the error metrics for wind triggers typically improve as we increase the offset timing when compared to the fine simulation, with similar trends compared to observations for both RMSE and Bias, while R^2 tends to remain fairly constant for all mixed simulations. Increase in accuracy with increasing offset confirms the assumption that shifting back the switch time using increasing offset times allows more time for water levels in the fine portion of the mixed simulations to adjust from water levels on the coarse portion. This is particularly important within the sound side of the Outer Banks, as this region is not

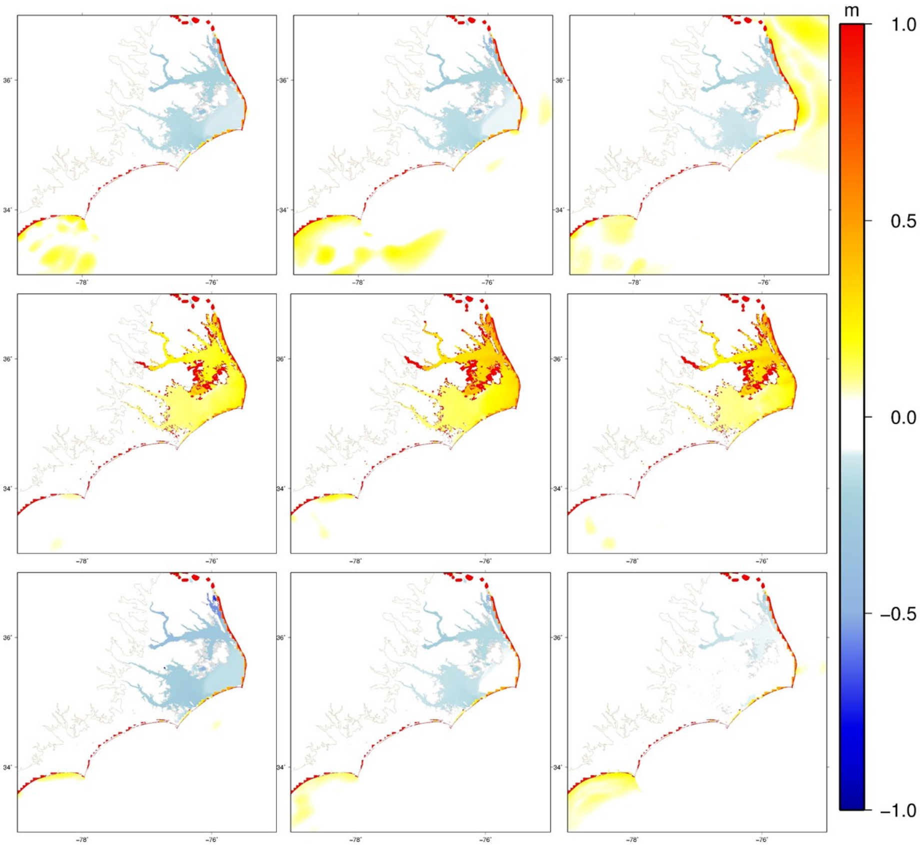

Figure 4.4: Wind based trigger differences in maximum water levels, computed as the maximum water levels of mixed minus fine. Each trigger simulation is represented per row (WT1 top, WT2 middle, WT3 bottom) with increasing offset timing across columns (left 0 hr, middle 12 hr, and right 24 hr).

represented within the coarse portion of the simulation.

In addition to these time-series-based analyses, we can also consider differences in maximum water levels between the mixed and fine simulation for coastal NC (Figure 4.4). For locations in the mesh on the ocean side of the Outer Banks, results show good matches of maximum water levels between the mixed and fine simulations. Whereas on the sound-side of the Outer Banks, the differences are much greater. A trend that is noticeable here is that WT1 and WT3 show the mixed simulations underestimating water levels and WT2 overestimating. There is a 9-hr difference between WT1 and WT2, and a 6-hr difference between WT2 and WT3. While it is possible the timing of the tides could play a role in this variation, time series plots for the wind based triggers show the WT2 mixed simulation overestimating water levels for these switches (Figure 4.5).

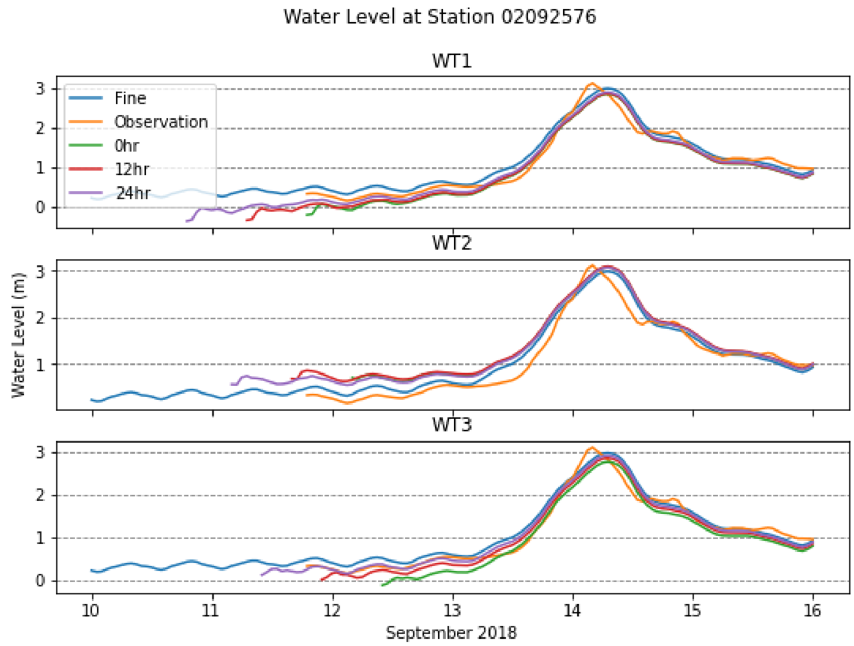

Figure 4.5: Example time series plot for station 02092576 showing over estimations in the water levels at the time of switching for WT2 compared to WT1 and WT3.

Coastal Water Level Triggers

Similarly to the wind triggers, we have both error metrics and difference plots for coastal and inland water level trigger simulations. For the coastal water level based triggers, the accuracy metrics for the coastal triggers are slightly underperforming 4.3, particularly compared to wind triggers. RMSE reaches 0.20 m for CWLT30 0 hr and is larger than 0.10 m for all mixed simulations when looking at comparisons to the fine simulation. When looking at comparisons to observations, RMSE is above 0.28 m for all simulations, which is higher than several wind-trigger-based mixed simulations. All mixed simulations are overpredicting water levels based on the positive bias values compared to both observations and the fine simulation. Bias values are larger than 0.15 m and 0.46 m compared to the fine simulation and observations, respectively. Stand out simulations are coastal triggers CWLT1, CWLT2, and CWLT3, with 24 hr offset, and CWLT30 with 24 hr offset. Values for the coastal based triggers show larger error metrics and slightly lower R^2 values than the values for the fine simulation compared to the observations 4.2.

Trends for the coastal water level triggers with regards to RMSE and Bias are similar to those for winds, with increasingly favorable values as the offset timing increases, due to similar reasons as discussed previously. However, bias values are large for comparisons

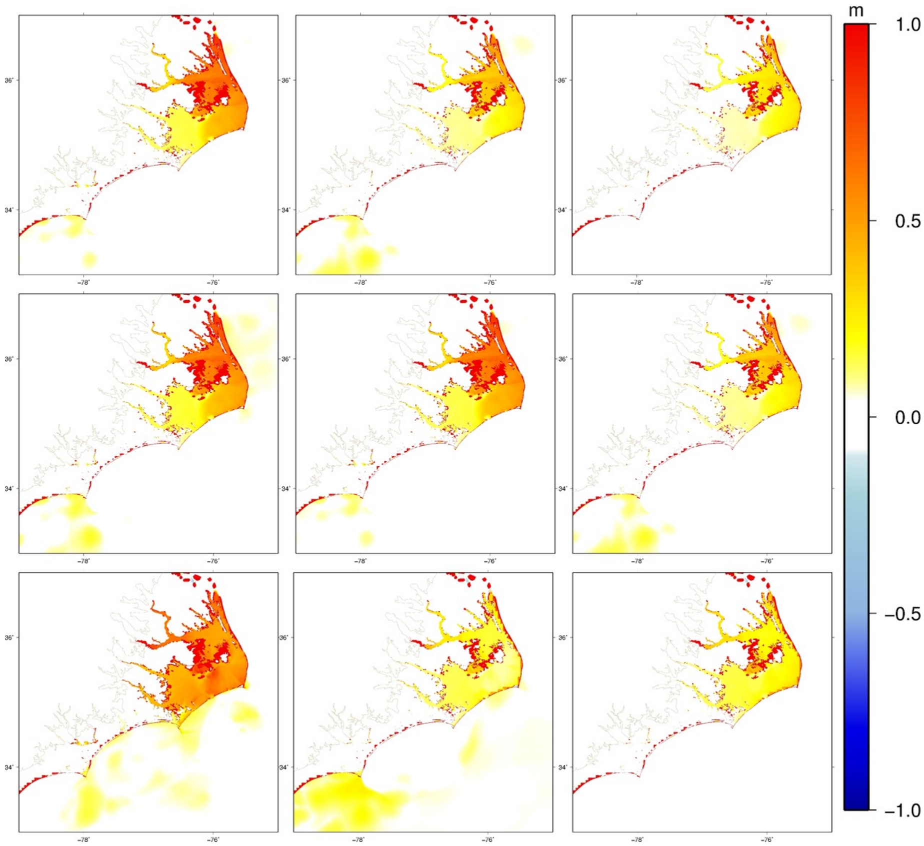

Figure 4.6: Coastal based trigger difference plots, maximum water levels of fine simulation subtracted from the mixed simulation. Each trigger simulation is represented per row (CWLT1/CWLT2/CWLT3 top, CWLT20 middle, CWLT30 bottom) with increasing offset timing moving to the right (left column 0 hr, middle 12 hr, and right 24 hr).

to observations at all triggers and offsets, meaning the mixed simulations significantly overpredict water levels compared to the observations. The timing of these trigger switches compared to wind switches has switches for CWLT123 and CWLT20 occurring prior to all wind switches and CWLT30 occurring after all wind switches. Switching post all wind switches (CWLT30) may result in larger errors due to the understanding of needing time for water levels on the fine mesh in the mixed simulation to 'catch-up' to the full fine simulation water levels.

For maximum water levels in the difference plots (Figure 4.6) the mixed results over-predict water levels overall, which corresponds with bias results (Table 4.3), particularly within the sound side of the Outer Banks. As the offset time increases, the results improve for each trigger. On the ocean side of the Outer Banks, we see more favorable comparisons between max water levels on the mixed versus fine, with the exception of trigger CWLT30 (bottom), particularly the for 0 hr and 12 hr offsets.

Inland Water Level Triggers

Lastly, for the inland-trigger-based mixed simulations, the accuracy is within the range of both wind and coastal trigger results, with comparisons to the fine simulation faring better than comparisons to the observations 4.3. The IWLT10 trigger, which corresponds to non-tidal residuals higher than 0.30 m above or below sea level, has the best error metrics of all simulations, with trigger IWLT3 coming in close second. Overall, R^2 values are very favorable for mixed results compared to the fine simulation, with values greater than 0.90 and all RMSE are under 0.24 m with several well under 0.10 m and Bias less than 0.25 m with most less than 0.10 m. For observation comparisons, results are fair for RMSE with all within the 0.20 to 0.40 m and Bias no greater than 0.56 m in either direction. All of the mixed simulations are overpredicting water levels compared to observations with the exception of IWLT30. Several of the inland simulations have comparable values to the results from the fine mesh compared to the observations such as IWLT1 12 hr and IWLT10 0 hr and 12 hr with a majority showing larger error metrics for RMSE and Bias but comparable values for R^2.

Trends interestingly do not show consistent increases in accuracy metrics as offset timing increases particularly for IWLT1, IWLT10, and IWLT30 for both RMSE and Bias. The switch time between IWLT1 and IWLT10 is only 18 hours, with the IWLT30 switch time being nearly two days later than IWLT1. A possible explanation for this is likely due to the location of the inland water level trigger stations being located farther inland. These stations would be the last to experience impacts from a coastal storm, thus pushing back 24 hr might not be early enough as it might be too close to landfall time. However, the timing of these switches is not significantly different from several of the coastal-based switches. It is possible that the configuration of the stations along the coast for both inland and coastal stations does not provide enough coverage to adequately represent the full arrival of storm effects and either a more broad coverage or the use of combination triggers would prove beneficial to investigating storm impacts.

At the broader spatial scale, for differences of maximum water levels (Figures 4.7 and 4.8), a few of the mixed simulations performed well (e.g. IWLT1 0 hr and 12 hr offsets, IWLT3 24 hr offset, and IWLT10 0, 12 and 24 hr offset). We do see similar results as for winds, with some triggers showing overpredictions (IWLT20 and IWLT3) in the mixed simulations and others showing underpredictions (IWLT1 24 hr, IWLT2, and IWLT30). Overall trends for

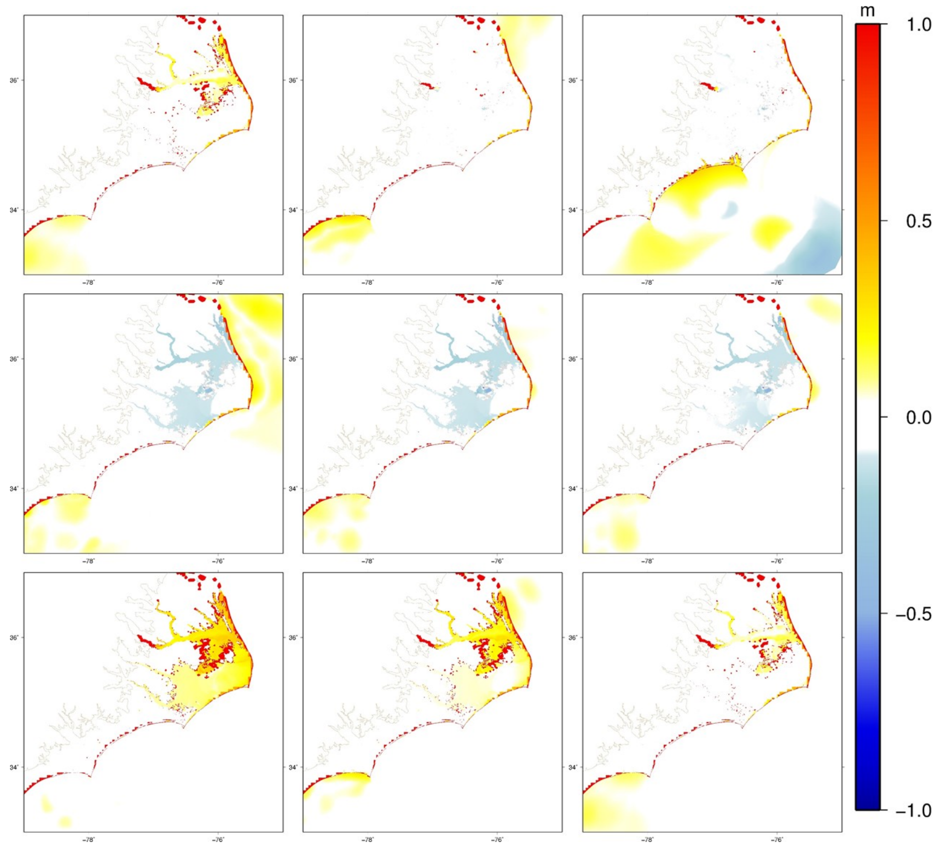

Figure 4.7: Inland trigger differences of maximum water levels (m) for mixed-mesh simulations using the inland-based triggers for tidal maximum, with differences computed as mixed minus fine. Each trigger simulation is represented per row (IWLT1 top, IWLT2 middle, IWLT3 bottom) with increasing offset timing across columns (left 0 hr, middle 12 hr, and right 24 hr).

increasing performance as offset increases are not apparent for all switches with IWLT1 24 hr showing larger differences on the ocean side of the Outer Banks as compared to later switches (0 hr and 12 hr).

4.1.4 Efficiency

Efficiency gains were quantified via wall-clock times and speed-up metrics for all mixed-mesh simulations. Before considering these results, it is emphasized that the wall-clock times are sensitive to the hardware on the high-performance computing system at NCSU. There are several hundred compute nodes, each with varying numbers of cores, on which to run simulations. These compute nodes have varying levels of efficiency. This is often

Table 4.3: Accuracy metrics of all trigger simulation relative to the fine simulation ran on the high-resolution NC9 mesh and observation data.

Simulation	Offset	Fine			Observation		
		RMSE	B_{MN}	R^2	RMSE	B_{MN}	R^2
WT1	0 hr	0.12	-0.18	0.99	0.24	-0.10	0.93
	12 hr	0.13	-0.19	0.97	0.23	-0.11	0.90
	24 hr	0.11	-0.17	0.99	0.23	-0.07	0.93
WT2	0 hr	0.10	0.14	0.99	0.29	0.51	0.91
	12 hr	0.10	0.18	0.99	0.30	0.59	0.91
	24 hr	0.10	0.14	0.99	0.28	0.52	0.91
WT3	0 hr	0.14	-0.23	0.99	0.25	-0.20	0.93
	12 hr	0.10	-0.16	0.99	0.22	-0.04	0.92
	24 hr	0.07	-0.10	0.99	0.21	0.07	0.92
CWLT123	0 hr	0.15	0.28	0.98	0.33	0.72	0.90
	12 hr	0.11	0.20	0.98	0.29	0.56	0.91
	24 hr	0.10	0.17	0.99	0.28	0.53	0.91
CWLT20	0 hr	0.15	0.28	0.98	0.34	0.74	0.90
	12 hr	0.15	0.28	0.98	0.33	0.72	0.90
	24 hr	0.11	0.20	0.98	0.29	0.56	0.91
CWLT30	0 hr	0.20	0.39	0.96	0.40	0.90	0.88
	12 hr	0.11	0.15	0.99	0.30	0.46	0.91
	24 hr	0.10	0.17	0.99	0.31	0.56	0.91
IWLT1	0 hr	0.06	0.04	0.99	0.25	0.30	0.92
	12 hr	0.04	-0.03	0.99	0.22	0.16	0.92
	24 hr	0.24	-0.06	0.96	0.40	0.12	0.88
IWLT2	0 hr	0.11	-0.18	0.99	0.30	0.44	0.93
	12 hr	0.12	-0.20	0.99	0.29	0.44	0.93
	24 hr	0.11	-0.18	0.99	0.29	0.43	0.92
IWLT3	0 hr	0.06	-0.02	0.99	0.28	0.52	0.91
	12 hr	0.03	-0.01	0.99	0.26	0.40	0.92
	24 hr	0.05	-0.01	0.99	0.25	0.30	0.92
IWLT10	0 hr	0.03	-0.01	0.99	0.22	0.21	0.92
	12 hr	0.02	-0.001	0.99	0.23	0.22	0.92
	24 hr	0.05	0.06	0.99	0.25	0.33	0.91
IWLT20	0 hr	0.11	0.15	0.99	0.30	0.46	0.91
	12 hr	0.10	0.17	0.99	0.31	0.56	0.91
	24 hr	0.10	0.13	0.99	0.29	0.51	0.91
IWLT30	0 hr	0.12	-0.18	0.99	0.24	-0.11	0.92
	12 hr	0.12	-0.20	0.99	0.24	-0.12	0.92
	24 hr	0.15	-0.25	0.99	0.25	-0.20	0.93

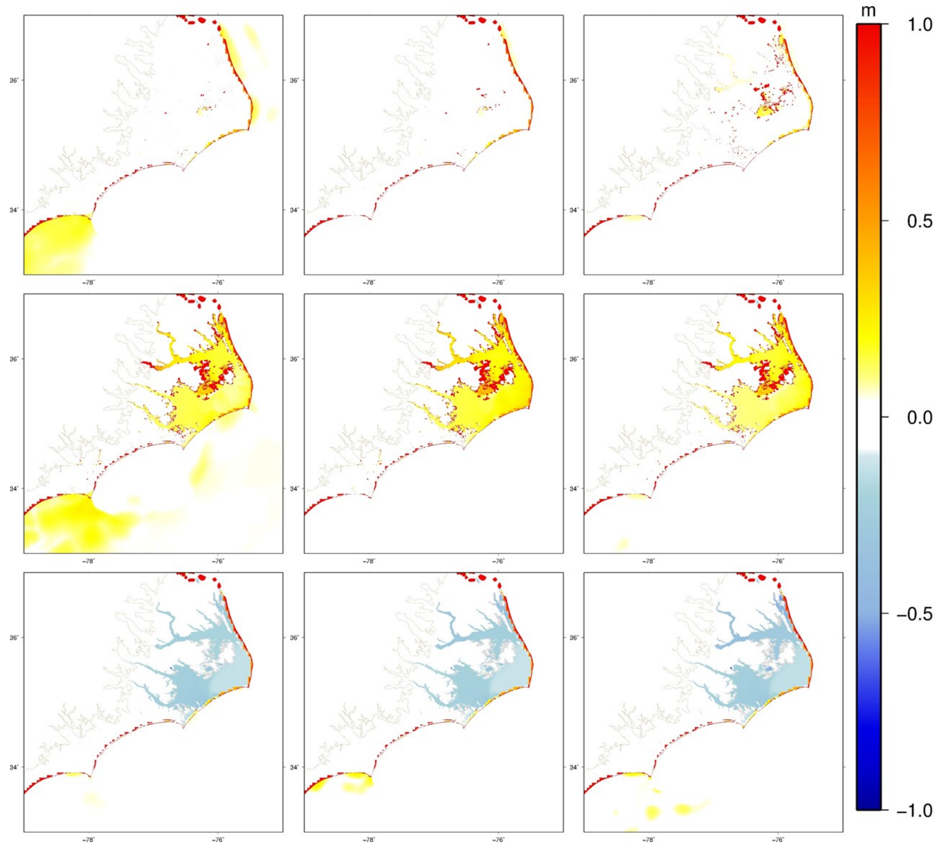

Figure 4.8: Inland trigger differences of maximum water levels (m) for mixed-mesh simulation using the inland-based triggers for non-tidal residual, with differences computed as mixed minus fine. Each trigger simulation is represented per row (IWLT10 top, IWLT20 middle, IWLT30 bottom) with increasing offset timing across columns (left 0 hr, middle 12 hr, and right 24 hr).

because nodes are purchased by individual researchers on an ongoing basis, with some nodes being newer than others. When submitting a simulation or job, the cores running that job can vary from one job to the next. While users can specify the cores used for a specific job, this does not allow for batch job submission. Therefore, the simulations ran for this work did not specify the cores used, and thus some simulations were ran on older and likely slower than average cores. If these simulations would be submitted several more times, they may by chance run on faster cores for one of those submissions, and times would be more similar to other triggers in this set. This difference in hardware speeds will explain why a handful of simulations do not share the same results, and there is no apparent trend to why this may occur based on the change in time spent running on the coarse and fine meshes.

Table 4.4: Efficiency data of all trigger simulations with wall-clock times and resulting speedups (actual and theoretical).

Trigger	Offset	Mixed (time in seconds)				Fine	Actual	Theoretical
		Coarse	*Adcirpolate*	Fine	Total			
WT1	0 hr	2688	149	1532	4369	7870	1.80	1.45
	12 hr	1549	167	2747	4463	7870	1.76	1.39
	24 hr	1264	162	3295	4721	7870	1.67	1.33
WT2	0 hr	2369	158	2410	4937	7870	1.59	1.51
	12 hr	3132	171	3136	6439	7870	1.22	1.44
	24 hr	1610	202	1553	3365	7870	2.34	1.37
WT3	0 hr	2065	281	2571	4917	7870	1.60	1.55
	12 hr	3219	188	7423	10830	7870	0.73	1.47
	24 hr	2676	193	5802	8671	7870	0.91	1.40
CWLT123	0 hr	1180	80	2307	3567	7870	2.21	1.25
	12 hr	1257	85	2005	3347	7870	2.35	1.20
	24 hr	3399	94	5397	8890	7870	0.89	1.15
CWLT20	0 hr	4833	181	3355	8369	7870	0.94	1.30
	12 hr	643	182	2587	3412	7870	2.31	1.25
	24 hr	955	154	9935	11044	7870	0.71	1.20
CWLT30	0 hr	1642	203	2462	4307	7870	1.83	1.77
	12 hr	4471	202	1221	5894	7870	1.34	1.67
	24 hr	2786	175	1560	4521	7870	1.74	1.59
IWLT1	0 hr	1347	103	3112	4562	7870	1.73	1.26
	12 hr	1564	173	3379	5116	7870	1.54	1.21
	24 hr	477	96	5157	5730	7870	1.37	1.16
IWLT2	0 hr	843	178	2807	3828	7870	2.06	1.33
	12 hr	650	123	3159	3932	7870	2.00	1.27
	24 hr	614	90	4016	4720	7870	1.67	1.22
IWLT3	0 hr	1487	127	3386	5000	7870	1.57	1.37
	12 hr	716	88	1994	2798	7870	2.81	1.31
	24 hr	739	90	2187	3016	7870	2.61	1.26
IWLT10	0 hr	777	88	5176	6041	7870	1.30	1.34
	12 hr	671	89	3240	4000	7870	1.97	1.28
	24 hr	1028	94	2718	3840	7870	2.05	1.23
IWLT20	0 hr	2991	95	1683	4769	7870	1.65	1.67
	12 hr	1107	87	1057	2251	7870	3.50	1.59
	24 hr	1010	85	1491	2586	7870	3.04	1.51
IWLT30	0 hr	1628	85	2778	4491	7870	1.75	1.70
	12 hr	1111	88	2678	3877	7870	2.02	1.61
	24 hr	2422	137	1988	4547	7870	1.73	1.53

For wind-based triggers, the actual speed up times were significantly higher than theoretical speed up times, with the exception of WT2 12 hr offset and WT3 12 and 24 hr offsets (Table 4.4). Coastal efficiency gains show more variation with several under performing (CWLT123 24 hr offset, CWLT20 0 hr and 24 hr offset, and CWLT30 12 hr offset) while others showed larger gains in efficiency (Table 4.4). Lastly, for inland efficiency gains for actual speed up were typically over 1.5 with an average of 1.73 with the slowest at 1.30 and the fastest at 3.50 (Table 4.4).

Overall, there are potential gains in efficiency, with a majority of simulations showing gains over 1.50. Ideal speed-up time would show actual times matching up with the theoretical time in the table. Of the 36 simulations, only 9 had actual speed up times below the theoretical speed up time, meaning 75 percent of simulations ran faster than would have theoretically been expected. Additionally, based on the offset timing increase, we would expect simulations to lose efficiency as we go from 0 hr to 24 hr offsets due to longer time spent on the fine mesh. In the wind triggers, this is the case for WT1, and IWLT1 and IWLT2 but not the case for WT2 or WT3 or other triggers. For coastal triggers, there appears to be no trend in any switches. The reason for this as discussed previously is the simulations with greatest speed up times likely ran on newer and faster cores, whereas the simulations with smaller speed up times ran on slower and older cores.

4.2 Discussion

The irregularities in the results mentioned previously will be further discussed in this section with a focus on the meshes used in these mixed simulations as well as the general recommendations made to key stakeholders and real-time forecast operators who have the most to gain from the results of this work.

4.2.1 Mesh Selection

A key finding from these results is the effect of different floodplain coverage between the coarse and fine meshes. Previous research by Thomas et al. (2021) used meshes that provided the same domain coverage for each part of the simulation. In those studies, the coastal floodplains of the Atlantic coast were included in both the coarse HSOFS mesh and fine SABv1 mesh. This allowed *Adcirpolate* to interpolate water levels in most locations when switching meshes. There was little need for extrapolation, because there were relatively few inland water bodies in the fine mesh that were not also in the coarse mesh.

However, in the research presented here, EC2001 does not provide coverage in the coastal region. Its land boundary is located at the ocean side of the Outer Banks in NC, and thus water in the coarse simulation is not able to propagate into coastal inlets and estuaries. When switching to NC9, *Adcirpolate* had to extrapolate water levels from the edges of the coarse mesh onto the fine mesh. This extrapolation led to differences in water levels that were larger than expected for the mixed simulation compared to the full fine simulation. To help understand the impacts this lack of coverage on the coarse mesh, two stations were arbitrarily selected to highlight the results from extrapolating water levels using *Adcirpolate* (Figure 4.9). The first station is located on the ocean side of the Outer Banks and is represented within both the coarse and fine mesh. The second station is located on the sound side and is only represented within the fine portion of the mesh. The two stations are approximately 12 km apart. The nearest node in the coarse mesh is 4 km to the point on the sound side of the Outer Banks.

At the end of the coarse simulation, water levels on the ocean side of the barrier islands of the Outer Banks in EC2001 are extrapolated onto the elements on the sound side within NC9. This extrapolation is imperfect, because there is no direct hydraulic connectivity between these locations, leading to larger differences in water levels within the sound side of the Outer Banks compared to the full fine simulation. When looking at the time series plots of select stations (Figure 4.9), depending on the location of the station, there can be significant variation in the starting water levels for the fine portion of the mixed simulations. The station on the ocean side shows that both simulation, mixed and fine result in the same values for water levels. However, at the sound side station the mixed simulation starts with water levels below those of the full fine simulation by close to 0.5 m. This is reflected in the spatial plots in Figure 4.9 with water levels in the full fine simulation or 'truth' being higher than those in the beginning fine portion of the mixed simulation. For further context, the water level value at the beginning of the fine part of the mixed simulation on the sound side would correspond to the water levels at last point in the time series for the coarse simulation on the ocean side of the mesh.

The extrapolation from the end of the coarse simulation to the beginning of the fine can lead to predictions with larger than expected errors in important coastal locations. Although the plots shown here are for one of the trigger switches, the difference plots (e.g. Figures 4.6, 4.7, 4.8, 4.4) show almost all simulations show errors on the sound side of the Outer Banks. The inclusion of offset timings that push back on the switching timing do lead to improvements in some cases. The simulation with the smallest errors would be IWLT10, which corresponds to 0.3 m non-tidal residual, which switched at 10 September

2018 at 21:00. This is approximately 3.5 days prior to when Florence made landfall. Similar resulting simulations switched earlier on the 10 September 2018 (03:00 for IWLT1 0 hr, and IWLT3 at 24 hr which would be the same time as IWLT1 0 hr).

When it comes to implementing *Adcirpolate* in real-time forecasting scenarios, operators should note variations in coverage of the meshes used during simulations. If a coarse mesh is used without coverage in the coastal floodplains such as EC2001, then it is important to switch with ample time before the storm is predicted to make landfall. If operators use a mesh with floodplain coverage, such as HSOFS, then it is possible that an offset will not be necessary, which could lead to further gains in efficiency. The optimal mesh choice would be to use a source mesh that contains floodplain coverage at a coarse scale in the the region where landfall is predicted.

4.2.2 Recommendations

By using the earliest timing of storm impacts such as wind speeds and water levels at coastal and inland locations as a trigger for switching from a coarse to a fine mesh, the resulting mixed simulations can achieve increases in efficiency and a maintenance accuracy. This leads to a recommended use of these triggers along with *Adcirpolate* during real-time storm forecasting scenarios. As more information about a storm becomes available and a greater confidence in landfall location is available, APS operators can implement *Adcirpolate* and switch to an alternate mesh with high resolution of the newly predicted landfall location based on the trigger selections used here. Furthermore, the most significant recommendation is to use inland based triggers as switching signals in combination with other storm parameters such as winds and coastal water levels, with possible inclusion of other parameters that are readily available during forecasting scenarios such as wave heights and atmospheric pressure values. For efficiency gains, consideration also needs to be made during real-time predictions on the time saved by switching meshes on the fly using *Adcirpolate* as opposed to the current requirement of running a 30 day tidal spin up on each mesh used during a simulation.

Specific recommendations could apply to the inland based triggers as they showed the most favorable error metrics for a select few trigger options including water levels exceeding 10 and 30 percent above tidal maximum (IWLT1 and IWLT3) and a non-tidal residual greater than 0.3 m above or below zero (IWLT10). While the non-tidal residual of greater than 0.3 m above or below zero showed small errors for all offsets, the two offsets for tidal maximum of 10 percent above maximum and the last offset for the 30 percent above maximum show greatest accuracy gains. The switch time for all inland triggers occurs prior to all wind based triggers and at a similar time to coastal triggers with inland having slightly later switching times due to their inland locations. This is interesting because one would expect water levels at inland locations to result in switches occurring well after both coastal and wind triggers. However, this is not what we are seeing. There appears to be an optimal time for switching based on the results from inland based triggers, which is around four days prior to land fall on early morning of the 10 September 2018. Possible reasons for this could be due to the meshes used as well as the approach angle and strength of Florence.

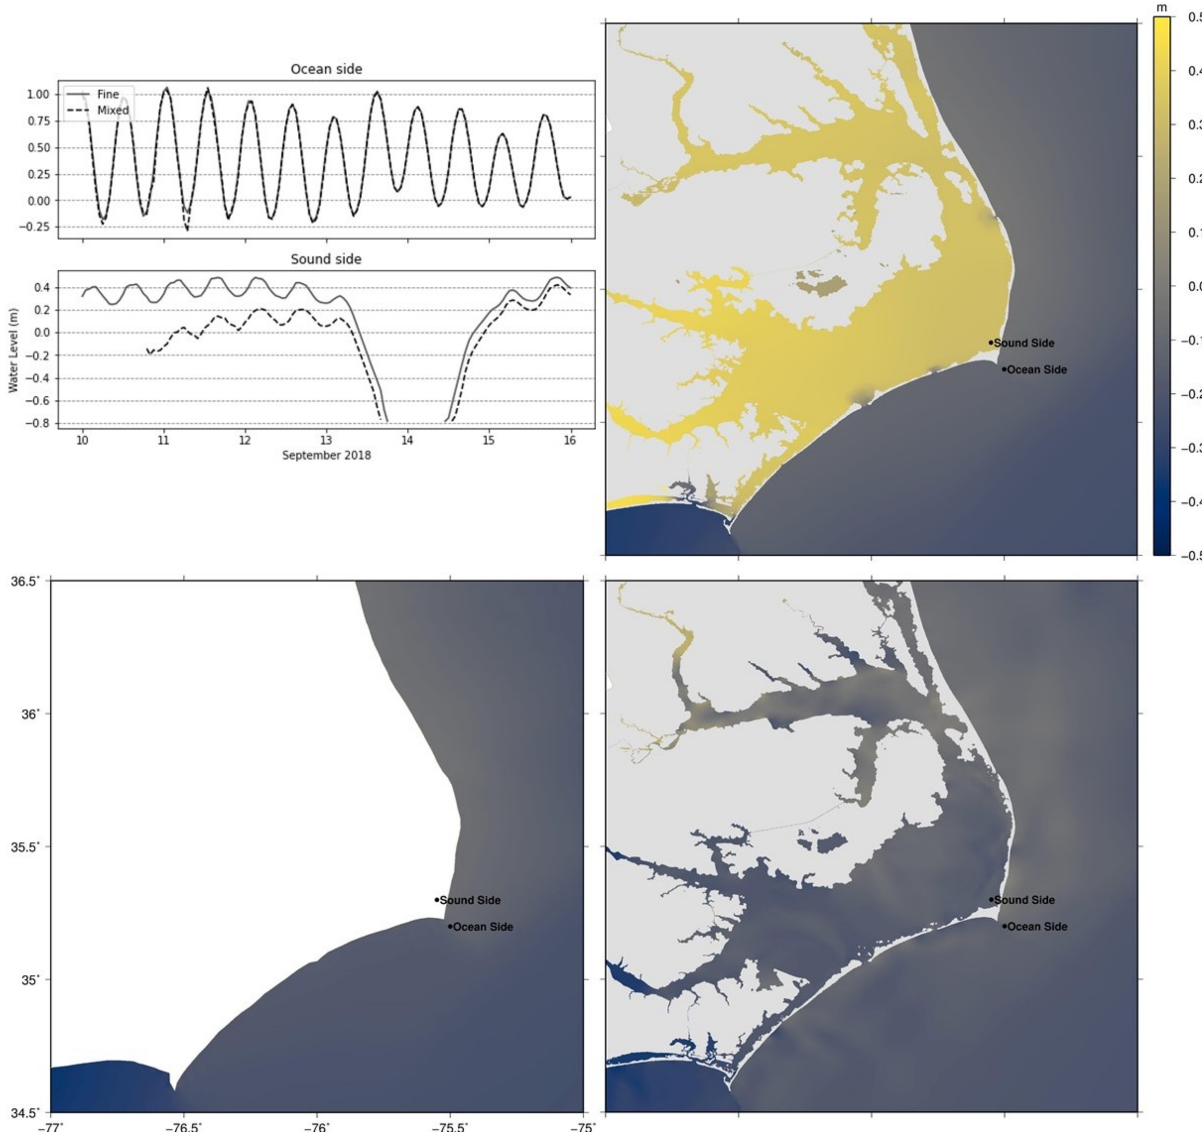

Figure 4.9: Differences in results for coarse and fine simulations for select stations near Hatteras Island. Top left shows time series plots for select stations from the inland T2 (IWLT2) mixed simulation, top right shows the fine simulation at the same time snap as the beginning of fine portion of the mixed-mesh simulation (bottom right) which corresponds to 10 September 2018 at 19:00, then bottom left shows end of the coarse portion of the mixed-mesh simulation time snap prior to the switch, which corresponds to 10 September 2018 at 18:00.

CHAPTER

5

CONCLUSIONS AND FUTURE WORK

During a coastal storm, emergency managers and decision makers need to know when a storm's effects will be felt, when and where flooding will occur, and how to improve predictions for coastal residents. In this thesis research, these questions were addressed via predictions of Florence (2018) with a multi-resolution approach, in which simulations were switched between two meshes of varying coverage and resolution based on the earliest storm impacts using wind speeds, and water levels at both the coast and inland locations. A suite of triggers was developed for wind speeds at varying levels of magnitude, and for water level exceedance above tidal maximum and non-tidal residual values at coastal and inland locations. The overwhelming majority of mixed-mesh simulations had faster computational times compared to a single fine-mesh simulation, with minimal losses in accuracy. Our validation with observation data also proves the usefulness of these triggers as a technique for real-time predictions. This research aids our understanding of when a storm may affect the coast, and will improve real-time predictions. The major conclusions from this study are:

- *Efficiency gains were substantial for simulations on mixed meshes.* By using *Adcirpolate*, efficiency was increased by an average of 37 percent across all mixed simulations.

While it is hard to evaluate specific trends based on the variability of which cores the simulations were ran on, we do expect increases in efficiency as time spent running on the fine mesh is decreased.

- *Accuracy losses were minimal, even with triggers switching from the coarse/source to fine/target mesh late in the simulation.* Accuracy was maintained, with an average RMSE of 0.11 m, Bias of 0.02 m and R^2 of 0.99 as compared to a full, high-resolution simulation. Triggers which switched well before landfall time showed better error metrics than those that switched closer to landfall time.

- *All of these triggers are viable for mixed-mesh simulation using Adcirpolate.* One particular group of triggers for wind speed or water levels at coastal or inland locations was not a stand out among those selected. However, by using individual triggers for inland water levels exceeding tidal maximums by 10 to 30 percent and a non-tidal residual of greater than 0.30 m above or below zero resulted in the highest levels of accuracy maintenance being achieved. Thus these are viable triggers that need to be included when deciding when to switch between meshes during a real-time forecasting scenario.

- *The timing of the switch between meshes should be informed by their floodplain coverage's.* It is not essential for both the coarse and fine meshes to contain the coastal floodplains. However, if the coarse mesh does not provide coverage in this region, then switching should occur earlier in the storm simulation, to maintain the highest accuracy.

Future work should include an expansion of the work discussed here. An expansion of triggers from a larger suite of parameters including wave heights, distances to isotachs, water velocities, radius to maximum winds, wetting of nodes along high and low tides, among others variables available with hindcast data as well as forecast data when transitioning to real-time implementation. However, applying this to real-time simulations would limit which parameters to use based on the data available at observation gauges within the region of the storm impact. A combination of these and other parameters can also be considered as triggers, such as applying a switch when both winds and water levels are elevated above a specified value. Another consideration could include different trigger values. As seen with the use of inland based triggers, it is possible that the trigger values of non-tidal residual absolute value above 0.30 m, 0.40 m, and 0.50 m are close to common variability seen at inland locations and separate values need to be used for these locations.

Lastly, an expansion of station coverage would improve understanding of earliest storm impacts and these additional stations could be used for all triggers to determine switching times.

The technique of trigger timing and using it as a guidance for switching can also be implemented outside of hurricane situations and be used for daily, non-hurricane conditions such as during Nor'easters and other drawdown events that can affect the ferry service in some locations. Additionally, this work could lead into future work within artificial intelligence, machine learning, and big data projects that build upon the suggested triggers highlighted here.

For efficiency work, all simulations should be ran on the same set of cores within the NCSU computing system, including the full high-resolution simulation. This would allow for more easy comparisons of efficiency gains across simulations. Furthermore, several sets of simulations would be submitted for each individual mixed simulation to get averages for the efficiency values. Further efficiency gains can be provided from advancements in more recent ADCIRC versions, such as v55 which allows for the implementation of larger time steps from 50 to 120 seconds compared to the 1 second used for the simulations here. This could further improve our efficiency in switching simulations by running the coarse part of the simulation with a time step of 50 seconds potentially leading to a simulation completing in 1/10 th of the time

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.

Beisiegel, N., Vater, S., Behrens, J., and Dias, F. (2020). An adaptive discontinuous galerkin method for the simulation of hurricane storm surge. *Ocean Dynamics*, 70:641–666.

Berger, M. J., George, D. L., Leveque, R. J., and Mandli, K. T. (2011). The GeoClaw software for depth-averaged flows with adaptive refinement. *Advances in Water Resources*, 34:1195–1206.

Beven, J. L., Berg, R., and Hagen, A. (2018). National hurricane center tropical cyclone report, hurricane michael (al142018) 7-11 october 2018. Technical report, National Hurricane Center. [Retrieved 10 May 2021].

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.

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.

Chen, C., Liu, H., and Beardsley, R. C. (2003). An unstructured grid, finite-volume, three-dimensional, primitive equations ocean model: Application to coastal ocean and estuaries. *J. Atmos. Oceanic Technol.*, 20:159–186.

Cialone, M. A., Massey, T. C., Anderson, M. E., Grzegorzewski, A. S., Jensen, R. E., Cialone, A., Mark, D. J., Pevey, K. C., Gunkel, B. L., McAlpin, T. O., Nadal-Caraballo, N. C., Melby, J. A., and Ratcliff, J. J. (2015). North atlantic coast comprehensive study (naccs) coastal storm model simulations: Waves and water levels. Technical report, US Army Engineer Research andDevelopment Center.

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

Cyriac, R., Dietrich, J. C., Fleming, J. G., Blanton, B. O., Kaiser, C., Dawson, C. N., and Luettich, R. A. (2018). Variability in coastal flooding predictions due to forecast errors during Hurricane Arthur (2014). *Coastal Engineering*, 137:59–78.

Davlasheridze, M. and Fan, Q. (2019). Valuing seawall protection in the wake of hurricane ike. *Economics of Disasters and Climate Change*, 3:257–279.

Dietrich, J. C., Dawson, C. N., Proft, J., Howard, M. T., Wells, G., Fleming, J. G., Luettich, R. A., Westerink, J. J., Cobell, Z., and Vitse, M. (2013). Real-time forecasting and visualization of hurricane waves and storm surge using SWAN+ADCIRC and FigureGen. In Dawson, C. N. and Gerritsen, M., editors, *Computational Challenges in the Geosciences*, pages 49–70.

Dietrich, J. C., Westerink, J. J., Kennedy, A. B., Smith, J. M., Jensen, R. E., Zijlema, M., Holthuijsen, L. H., Dawson, C. N., Luettich, R. A., Powell, M. D., Cardone, V. J., Cox, A. T., Stone, G. W., Pourtaheri, H., Hope, M. E., Tanaka, S., Westerink, L. G., Westerink, H. J., and Cobell, Z. (2008). Hurricane gustav (2008) waves and storm surge: Hindcast, synoptic analysis, and validation in southern louisiana. *Monthly Weather Review*, 139:2488–2522.

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.

Erdman, J. (2019). 88% of u.s. deaths from hurricanes, tropical storms are from water, not wind. `https://weather.com/safety/hurricane/news/hurricanes-tropical-storms-us-deaths-surge-flooding`. [Retrieved 10 May 2021].

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.

Forbes, C., Luettich, R. A., Mattocks, C. A., and Westerink, J. J. (2010). A retrospective evaluation of the storm surge produced by Hurricane Gustav (2008): Forecast and hindcast results. *Wea. Forecasting*, 25:1577–1602.

Gao, J. (2018). *On the surface wind stress for storm surge modeling*. PhD thesis, University of North Carolina at Chapel Hill.

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

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.

Hill, C., DeLuca, C., Balaji, V., Suarez, M., and da Silva, A. (2004). Architecture of the Earth System Modeling Framework. *Computing in 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., Gravois, U., Zachry, B. C., Westerink, J. J., Hope, M. E., Dietrich, J. C., Powell, M. D., Cox, A. T., Luettich, R. A., and Dean, R. G. (2011). Origin of the Hurricane Ike forerunner surge. *Geophysical Research Letters*, 38(8).

Kerr, P. C., Donahue, A. S., Westerink, J. J., Luettich, R. A., Zheng, L. Y., Weisburg, R. H., Huang, Y., Wang, H. V., Teng, Y., Forrest, D. R., Roland, A., Haase, A. T., Kramer, A. W., Taylor, A. A., Rhome, J. R., Feyen, J. C., Signell, R. P., Hanson, J. L., Hope, M. E., Estes, R. M., Dominguez, R. A., Dunbar, R. P., Semeraro, L. N., Westerink, H. J., Kennedy, A. B., Smith, J. M., Powell, M. D., Cardone, V. J., and Cox, A. T. (2013). U.S. IOOS coastal and ocean modeling testbed: Inter-model evaluation of tides, waves, and hurricane surge in the Gulf of Mexico. *Journal of Geophysical Research: Oceans*, 118:5129–5172.

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.

Klotzbach, P. J., Bowen, S. G., Pielke, R., and Bell, M. (2018). Continental u.s. hurricane landfall frequency and associated damage: Observations and future risks. *Bulletin of the American Meteorological Society*, 99:1359–1376.

Knutson, T., Camargo, S. J., Chan, J. C. L., Ho, K. E. C. H., Kossin, J., Mohapatra, M., Satoh, M., Sugi, M., Walsh, K., and Wu, L. (2020). Tropical cyclones and climate change assessment: Part ii: Projected response to anthropogenic warming. *Bulletin of the American Meteorological Society*, 101:E303–E322.

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.

Mandli, K. T. and Dawson, C. N. (2014). Adaptive mesh refinement for storm surge. *Ocean Modelling*, 75:36–50.

Massey, W. G., Gangai, J. W., Drei-Horgan, E., and Slover, K. J. (2007). History of coastal inundation models. *Marine Technology Society Journal*, 4.

Mattocks, C. and Forbes, C. (2008). A real-time, event-triggered storm surge forecasting system for the state of North Carolina. *Ocean Modelling*, 25:95–119.

Mukai, A., Westerink, J. J., and Luettich, R. A. (2002a). Guidelines for using the eastcoast 2001 database of tidal constituents within the western north atlantic ocean, gulf of mexico and caribbean sea. Technical Report ERDC/CHL CHETN-IV-40, U.S. Army Engineer Research and Development Center, Vicksburg MS.

Mukai, A., Westerink, J. J., Luettich, R. A., and Mark, D. J. (2002b). A tidal constituent database for the Western North Atlantic Ocean, Gulf of Mexico and Caribbean Sea. Technical Report ERDC/CHL TR-02-24, U.S. Army Engineer Research and Development Center, Vicksburg MS.

National Hurricane Center (2017). Hurricane Harvey & its impacts on Southeast Texas. `https://www.weather.gov/hgx/hurricaneharvey`. [Retrieved 10 May 2021].

National Hurricane Center (2018). Hurricane florence: September 14, 2018. `https://www.weather.gov/ilm/HurricaneFlorence`. [Retrieved 10 May 2021].

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

National Oceanic and Atmospheric Administration (2020). Record-breaking atlantic hurricane season draws to an end. `https://www.noaa.gov/media-release/record-breaking-atlantic-hurricane-season-draws-to-end`. [Retrieved 10 May 2021].

National Oceanic and Atmospheric Administration (2017). National Data Buoy Center. `http://www.ndbc.noaa.gov`.

National Oceanic and Atmospheric Administration (2018a). National Ocean Service. `https://oceanservice.noaa.gov/`. [Retrieved 10 March 2018].

National Oceanic and Atmospheric Administration (2018b). Sea level trends. `https://tidesandcurrents.noaa.gov/sltrends/sltrends.html`. [Retrieved 09 March 2018].

NOAA and U.S. Census Bureau (2013). Noaa coastal population report: Population trends from 1970 to 2020. Technical report.

Qin, X., LeVeque, R. J., and Motley, M. R. (2019). Accelerating an adaptive mesh refinement code for depth-averaged flows using gpus. *Journal of Advances in Modeling Earth Systems*, 11:2606–2628.

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. Technical report.

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

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.

State Climate Office of North Carolina (2017). `http://climate.ncsu.edu/climate/hurricanes/statistics`. [Retrieved 21 June 2017].

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

Stockdon, H. F., Doran, K. J., Thompson, D. M., Sopkin, K. L., Plant, N. G., and Sallenger, A. H. (2012). National assessment of hurricane-induced coastal erosion hazards: Gulf of Mexico. Technical Report 2012-1084, U.S. Geological Survey.

Szpilka, C., Dresback, K., Kolar, R., Feyen, J., and Wang, J. (2016). Improvements for the Western North Atlantic, Caribbean and Gulf of Mexico ADCIRC tidal database (EC2015). *Journal of Marine Science and Engineering*, 4:72.

Taeb, P. and Weaver, R. J. (2019). An operational coastal forecasting tool for performing ensemble modeling. *Estuarine, Coastal and Shelf Science*, 217:237–249.

Thomas, A. (2020). *Using a Multi-Resolution Approach to Improve the Accuracy and Efficiency of Flooding Predictions*. PhD thesis, North Carolina State University, Raleigh, NC, USA.

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*.

United Nations (2017). The ocean conference factsheet. `https://www.un.org/sustainabledevelopment/wp-content/uploads/2017/05/Ocean-fact-sheet-package.pdf`. [Retrieved 12 December 2020].

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 (2016). South shore of staten island, new york coastal storm risk management interim feasibility study for fort wadsworth to oakwood beach. Technical report, U.S. Army Corps of Engineers.

U.S. Geological Survey (2020a). Flood event viewer. `https://stn.wim.usgs.gov/FEV/#FlorenceSep2018`. [Retrieved 16 June 2020].

U.S. Geological Survey (2020b). USGS rapid deployment gages (rdgs). `https://www.usgs.gov/mission-areas/water-resources/science/usgs-rapid-deployment-gages-rdgs?qt-science_center_objects=0#qt-science_center_objects`. [Retrieved 23 May 2020].

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.

Zhang, Y., Ye, F., Stanev, E. V., and Grashorn, S. (2016). Seamless cross-scale modeling with SCHISM. *Ocean Modelling*, 102:64–81.