Modeling New York in D-Flow FM

Delft, 30th of October 2013

Internship Tuinhof, T.J. (Taco)
August - October 2013

Supervisors:
Sander Post (Royal HaskoningDHV)
Maarten van Ormondt (Deltares)

Contents

Introduction

As part of my master Hydraulic Engineering at the TU Delft I worked as an intern for Royal HaskoningDHV and Deltares for nine weeks during the summer and fall of 2013. Goal of the internship was setting up a model for the New York region in the new software engine developed by Deltares: D-Flow Flexible Mesh. D-Flow FM makes it possible to do hydro dynamical simulations on unstructured grids. This means that in one mesh besides curvilinear also, for example, triangular structures can be present.

D-Flow FM makes more flexible gridding possible in, for instance, delta regions or harbors. When looking at the New York area could be an advantage. In the past, Deltares and Royal HaskoningDHV have developed a model for the region in Delft3D. This model was used to simulate water levels during storms like hurricane Sandy. This Delft3D model was set up using a fairly large grid size of 500x500 meters. This grid size is accurate enough for modeling the part of the North Atlantic Ocean included in the model. However, it is too large to capture the channels of the Hudson or Eastern River as accurate as wanted.

The purpose of this study is to model the region of New York again in D-Flow FM, and using the flexible mesh grid to more accurately model inland areas, which would need a smaller grid size. However, one of the requirements is a reasonable computation time, which means that it should not take more than a few hours to model a month of hydrodynamics. The goal of this internship will be to finally simulate hurricane Sandy in the developed D-Flow FM model. However, during the internship it was decided to first focus on a tide analysis before modeling the hurricane. After nine weeks the model was set up and a tide analysis was performed. In the last weeks hurricane Sandy was also simulated, but not much research was done on that, just to show that it works.

This report provides an overview of the work I did during the nine weeks at Deltares and Royal HaskoningDHV. It includes a description of the model set-up, the tide analysis which was done, the results of modeling hurricane Sandy and some recommendations for further development of the model. It also includes some personal experience of the software engine (D-Flow FM) itself.

Model set-up

Grid generation

At first instance it seemed reasonable to create the new D-Flow FM starting from the original Delft3D model grid. However, this grid was generated in the most simplified way, so starting from this grid would not save any time. Therefore a new grid was made. First was decided which parts to include in the model. The southern boundary was chosen to start somewhat south of the inlet to Lower Bay, in order to include some part of the North Atlantic Ocean. The boundary shape follows the land boundary when going west. The western boundary was chosen at such a position that the inlet to Long Island Sound was included. This was important since waves travel through the Long Island Sound and reach New York through the East River. At the Eastern boundary some of the larger rivers were included in the grid. To the North the Hudson River was added in the model until the city of Troy, just north of Albany. At Troy the Federal Dam is situation which makes it impossible for water travelling up the river to penetrate. Troy is the most upstream point on the Hudson River until which, for instance, the tidal wave can reach.

First a curvilinear grid, with a grid size of 1000x1000 meters, was created for the whole region. This grid size is small enough to capture changes in the bathymetry on the Ocean. However, at some locations finer meshes were wanted to capture the bathymetry better or follow the coast line more accurate. So, the next step was to remove or replace certain parts of the grid.

The grid size of 1000x1000 meters made it possible to follow the coast line quite accurate at some locations, however, in the Long Island Sound and in the Lower Bay a smaller grid size was needed near the coast. Therefore, the grid size was reduces a factor 2, to be able to follow the coastline more accurately. At some locations near the coast the grid should be even finer, however, seeing the purpose of this study it was decided not to do so.

All rivers and bays near Manhattan were captured with a finer grid, since that was the area of interest for this study. Where possible a curvilinear grid was used. All river sections were gridded with a curvilinear grid following the flow lines. An example is the Hudson River. Here a curvilinear grid was used of four grid cells wide. The number of grid cells used to capture the width of the river was chosen on basis of the wanted accuracy, available bottom data and computation time. The smallest river section, the East River, was also an important link in the grid. The smallest grid cell is of large influence on the computation time through the Courant number, and the smallest grid cells are present in this river. However, to capture the bottom level difference in the cross section of the river grid cells of 40-50 meters were used. If the computation time proves to be too slow, adjustments are possible in the future.

Unstructured meshes were also used at some locations. In bays or parts of the river where a curvilinear grid did not easily fit, an unstructured grid was used. An example is Jamaica Bay. Also, for connections between river sections triangles were used where needed. Figure 1 shows an overview of the final grid, used for computation.

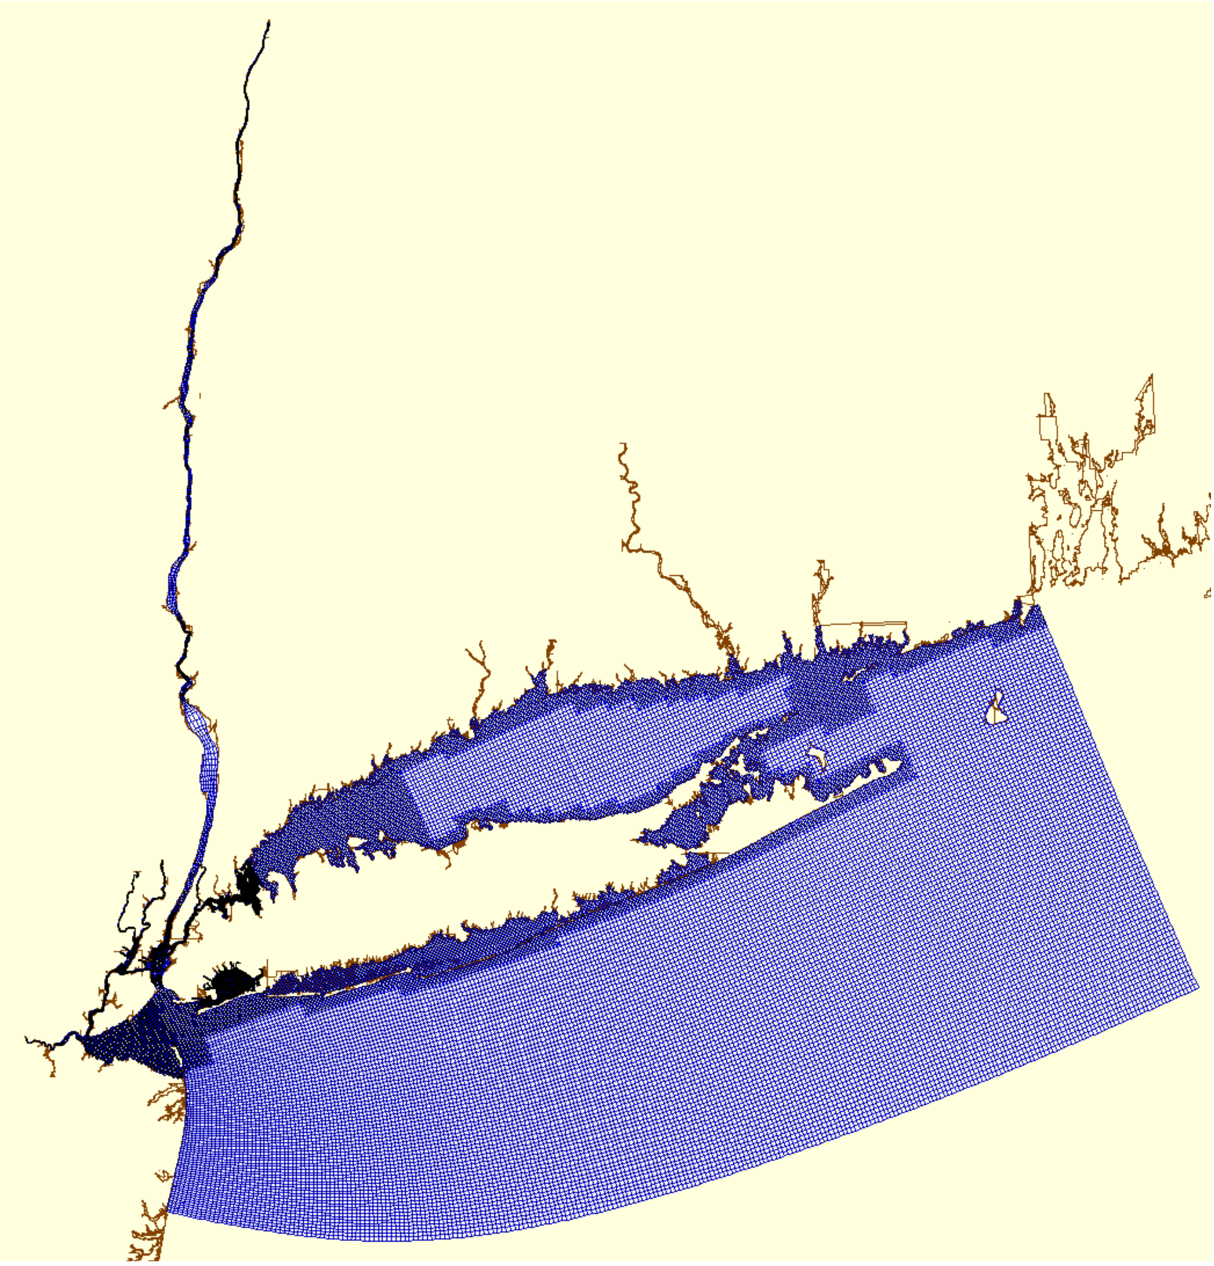

Figure 1 - Map showing land boundaries and generated grid

Bathymetry

For the bathymetry several datasets were available. For the Atlantic Ocean and large parts of the Long Island sound the NGDC Coastal Relief Model dataset was used. These data are from the USGS and the Shuttle Radar Topography Mission (SRTM) and cover the U.S. coasts. It has a data resolution of 3 arc-second, which corresponds to a cell size of approximately 90 meter. For the inland waterways (East River, Kill van Kull etc.) the NGDC Coastal Relief Model was lacking data, so other datasets had to be used. For these parts soundings and contour information was available from the National Oceanic and Atmospheric Administration. These data is generated from sea maps. The resolution of these data was varying within the project area.

In D-Flow the different datasets were loaded as samples and used to interpolate on the grid, creating a grid including depths. In the areas were the depth data was denser than the grid cells, the technique of grid cell averaging was applied. In areas with only a few sample points per grid cell, the technique of triangular interpolation was used. After the interpolation the depth was checked manually and several adjustments were made, especially in areas were not many sample points were present. During the tide analysis further adjustments were made to the grid. These adjustments will be mentioned in the next chapter.

Boundary conditions

For the tide analysis harmonic constituents are used as boundary condition at sea. In D-Flow a polyline is drawn along the boundary and in every point of this polyline harmonic constituents have to be determined as boundary conditions. In between these points D-Flow FM interpolates to obtain boundary conditions at every location. In this way harmonic constituents were applied on 20 points along the sea side boundary.

Around the boundary the M2 tide had the largest amplitude which varies between 0.4 and 0.6 meters. The harmonic constituents were obtained from the Atlantic Ocean 2008 model. In this model eleven constituents are taken into account (M2, S2, N2, K2, K1, O1, P1, Q1, M4, MS4, MN4). The resolution of the model is 1/12 degree (approximately 9 km).

At the river boundary no boundary conditions were necessary for the tide analysis since the Hudson is modeled until the Federal Dam near Troy. No water can travel upstream this dam, so it can be seen as a closed boundary.

For the model computation to simulate Sandy no harmonic constituents could be used as boundary conditions at sea. Time series of water levels were imposed at the sea side boundary now. These time series were generated from output at the specific locations from a larger Delft3D model covering a larger part of the Atlantic Ocean.

Initial parameter settings

For the bottom roughness a uniform Chézy coefficient is chosen. For the first calculation in the tide analysis the Chézy coefficient is set to 65 $m^{1/2}/s$, for the whole model. During the tide analysis the Chézy coefficient was made special variable and was adjusted in several areas to get better results. This is explained in the next chapter. The initial water level was set to 0 for both the tide analysis as for the simulation of hurricane Sandy.

For the calculations a reference date was chosen (01-01-2002) and seconds were used as unit for time. Important is the time step used. In D-Flow there are two ways of setting the time step. Either an automatic time step or a user defined time step can be chosen. When the automatic time step is chosen the time step is determined for every calculation using the CFL condition. A maximum Courant number can be selected and based on this number the time step is determined. The other possibility is just choosing a time step which is used throughout the whole calculation. At first instance the auto time step was chosen, however, after some trial and error a user defined time step of 30 seconds turned out to work fine as well and speeded up the computation.

For simulating hurricane Sandy the model had to be extended with wind and atmospheric pressure data. Therefore atmospheric pressure and wind files were created which specified wind and pressure development in space and time. The data used came from the North American Mesosclae Forecast System (NAM) from the National Centers for Environmental Prediction (NCEP). Wind and atmospheric pressure data was available on a grid of 20000 by 20000 meter and with an interval of three hours.

Tide Analysis

To test the performance of the model a tide analysis was done. The tide analysis will give a first insight in the behavior of the model. An advantage is the fact that data from many tidal stations in the area is available, which makes it possible to compare computed and observed data. In the tide analysis the harmonic constituents of the tidal stations as observed are compared to the computed harmonic constituents at that same location in the model. Seeing the purpose and duration of this study only a quick and dirty tide analysis was performed.

Tide

As mentioned above several tidal stations are present in the area. Figure A.1 in apendix A shows a map including all stations (each red dot represents a station). In total 24 stations are available to test the performance of the model.

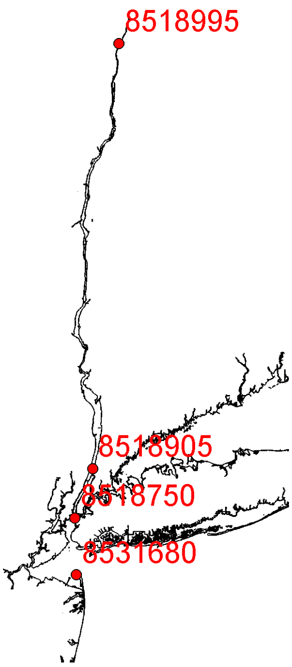

Figure 2 - Tidal stations Hudson River

To analyse the propagation of the tidal wave in the area first the the amplitude and phase of the main tidal constituents in certain tidal stations are analysed (see figure left). Analysing the observed data for these stations should say something about the propagation of the tide travelling up the Hudson River. In the two histograms in figure 3a and 3b the amplitudes and phases of the main tidal constituents for the four stations are given.
From the most south situated station Sandy Hook (8531680) until the station near Riverdale (8518905) the amplitude of all constituents decreases. This is caused by energy dissipation due to friction. This energy loss leads to a decrease of the tidal amplitude. The biggest loss is seen from The Battery (8518750) to Riverdale, which is the most downstream end of the Hudson River. From Riverdale to Albany (8518995) an increase in amplitude is visible for most constituents. This is due to the fact that the changes in geometry cause the amplitude to increase. While travelling upstream the river the crosssectional area decreases. A shallower river means a lower wave speed, because in shallow water the wave speeds depends on the waterdepth: wave speed=sqrt(g*waterdepth). To conserve energy the amplitude of the wave increases, clearly visible in figure 3a.

When looking at the phases for the different tide stations it is possible to determine the speed of the wave travelling up the Hudson River. For the M2 constituent there is a difference in phases between the Battery and Albany of about 260 degrees. The M2 tide has a period of 12.42 hours, so a difference of 260 degrees between the two stations corresponds to a time lag of approximately nine hours. This means that high water in Albany happens nine hours after high water at the Battery. The distance between the two stations is 250 kilometers. This leads to a wave speeds of 250 km/9 hours = 27.8 km/h. The speed of the tidal wave in the river can also be approximated using the fact that the river has a average depth of about 10 meters. This gives a wave speeds of 10 m/s. Friction slows the tide down, about 20% of its speed is lost[1], which leads to a wave speed of 8 m/s on average. This gives a wave speed of 28.8 km/h, close to the wave speed calculated using the phase difference.

1. W. Rockwell Geyer and Robert Chant, The Physical Oceanography Processes in the Hudson River Estuary, Chapter 3 of The Hudson River Estuary, Cambridge University Press, 2006

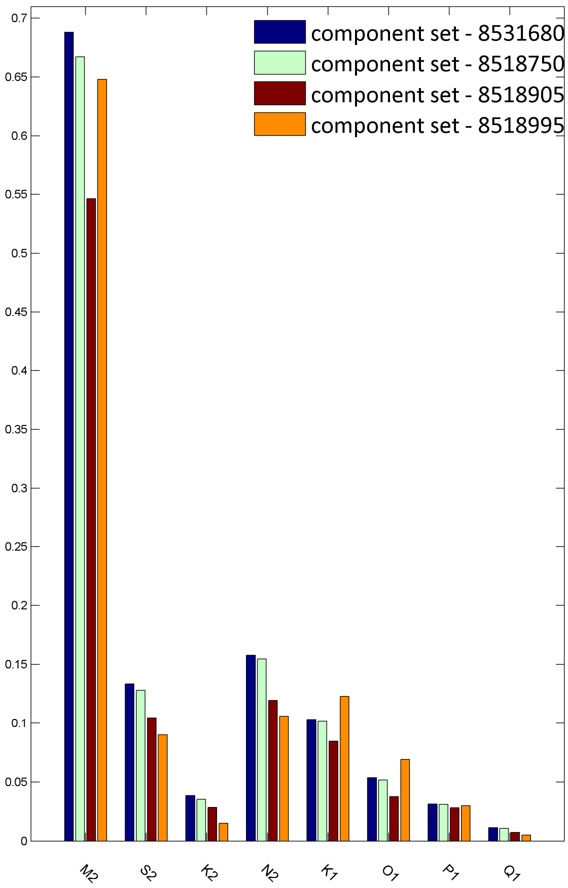

Figure 3a – Amplitude [m] of harmonic constituents

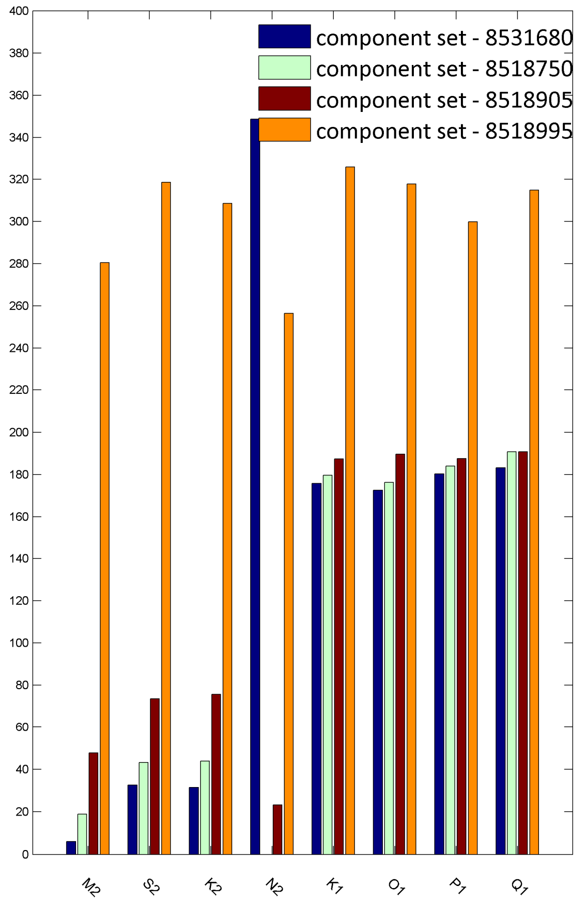

Figure 3b - Phase [deg] of harmonic constituents

Another interesting area to have a closer look at is the Long Island Sound. See figure 4 for the tidal stations included and figure 5 for the amplitudes of the constituents at the different stations.

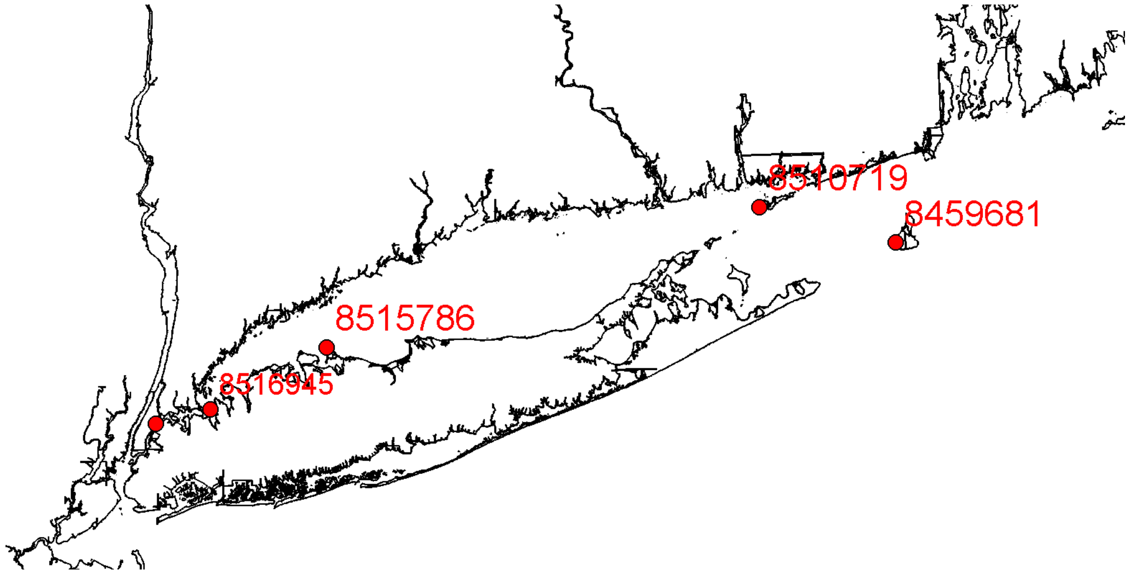

Figure 4 - Tidal stations Long Island Sound

At Block Island (8459681) the amplitude of the tidal constituents is still close to the amplitude at the boundary, which is convenient since this station is situated outside the entrance of the Long Island Sound. At Silver Eal Pond (8510719), near the entrance of the Long Island Sound, the amplitudes for most constituents are even decreased slightly. However, within the Long Island Sound the tidal amplitudes are strongly increased from east to west. At King's Point (8516945), the far most west tidal station, the tidal amplitude of the M2 tide is almost 1.25 meters. This is caused by the size and shape of the Sound (long, narrow and shallow), which amplifies the oceanic tide.

Going even more westward, into the East River, the tidal amplitude is damped again, which is shown in figure 5 by the smaller tidal amplitude for the station at Horns Hook (8518668). In the East River the tide is influenced by both the tides in the Long Island Sound as well as the tides in New York Harbor. The difference in phase and amplitude between these two is large, resulting in strong tidal currents in the East River.

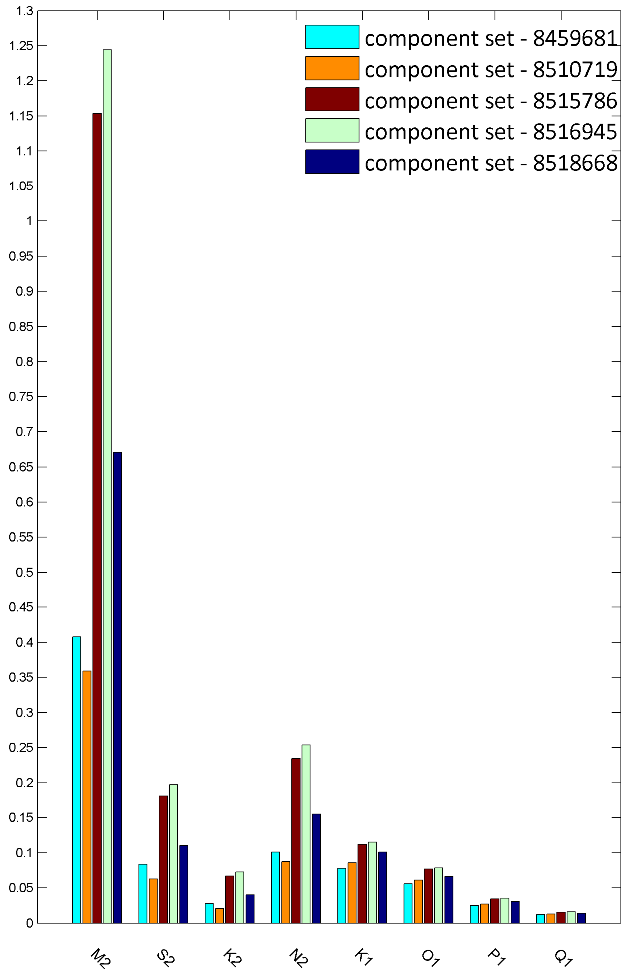

Figure 5 - Amplitude of harmonic constituents

Harmonic analysis

During a computation the output delivered by D-Flow are time series of water levels for each observation point. There are also observed time series available for each tidal station. However, comparing the computed time series to the observed data by eye would be both time consuming and inaccurate. To perform a proper tide analysis the harmonic constituents of the observations and the computations are compared. For the observed data the harmonic constituents can be obtained from the National Oceanic and Atmospheric Administration (NOAA). For the computed data harmonic constituents have to be generated from the time series. This can be done by performing a harmonic analysis on the time series. In this study the Matlab toolbox T_tide was used to do this.

At first instance the observed water levels and harmonic constituents for one tidal station were compared, to show that T_tide was able to correctly calculate harmonic constituents from a data set. Tidal station 8531680 (Sandy Hook) was chosen to perform this check. Several sets of time series for this station were downloaded: one year, six months, three months, two months and one month. In this way it is possible to see what duration of time series is needed to get the correct harmonic constituents during a harmonic analysis.

Amplitude [m]	NOAA	1 Year	6 Months	3 Months	2 Months	1 Months
M2	0,688	0,688	0,683	0,684	0,682	0,681
S2	0,134	0,134	0,145	0,145	0,129	0,114
N2	0,158	0,158	0,168	0,174	0,168	0,155
K1	0,103	0,103	0,103	0,105	0,117	0,129
O1	0,054	0,054	0,055	0,055	0,054	0,054
Q1	0,011	0,011	0,012	0,013	0,012	0,012
K2	0,038	0,038				
P1	0,031	0,031				
Phase [deg]	**NOAA**	**1 Year**	**6 Months**	**3 Months**	**2 Months**	**1 Months**
M2	6	6,0	5,9	5,8	5,8	5,8
S2	32,6	32,6	36,6	45,0	47,4	44,8
N2	348,6	348,6	348,9	353,5	356,9	358,4
K1	175,7	175,7	176,2	184,2	185,5	182,7
O1	172,5	172,6	172,6	173,1	173,7	173,9
Q1	183,1	183,1	182,2	185,4	192,4	194,0
K2	31,5	31,5				
P1	180,2	180,2				

Table 1 - Amplitude and phase of harmonic constituents for NOAA observations and different time series' durations

Table 1 shows the results for the main eight constituents; both the amplitude and the phase for each constituent. The first column shows the data obtained from NOAA and the other columns the values calculated for the several time series durations. Using the one year time series it is possible to simulate the harmonic constituents most accurately. Already for the six months' time series two constituents disappear in the harmonic analysis, K2 and P1. Furthermore it can be concluded that the amplitude and phase for the other constituents show differences with the NOAA data for the shorter time series. The differences in amplitude (percentages) and phase (degrees) for the constituents are drawn in figure A.2 and figure A.3 in appendix A.

For most constituents the amplitude and phase difference between the NOAA and the calculated amplitude and phases become bigger when the time series' duration becomes smaller. This problem is caused by the resolution of certain constituents. In short time series durations it is hard to distinguish constituents that are close to others in frequency. For example, S2 has a frequency of 2 cycles per day and K2 of 2.0055 cycles a day. When the time series is not long enough, there is not enough data to accurately determine the phase and amplitude of these constituents. In this example K2 even disappears completely when time series of several months are used.

However, what also becomes clear from the table and graphs is that for M2 and O1 the phase and amplitude error do not grow extremely when shorter time series are used. So, if someone is interested in the characteristics of the M2 constituent only, a tidal analysis of one month would be sufficient.

Which time series duration to choose is thus determined by the accuracy wanted and the computation time of the model. The computation time of this model is about 24 hours for one year of simulation (on my core i7 laptop). Therefore it was chosen to run the model for one year and then compare the harmonic constituents to get the most accurate results.

Model

To test the performance of the model simulation of the tide, the computed water levels and corresponding tidal constituents are compared to the observed amplitudes and phases. This is done by comparing both the absolute differences in phase and amplitude of the main eight tidal constituents. Furthermore, the vector difference between the observed and computed constituent are calculated. This vector difference says something about the phase difference and amplitude difference in combination with the amplitude itself. An amplitude difference of 25% and phase difference of 20 degrees may seem unacceptable; however, if the amplitude of the constituent is only 1 centimeter, the relative influence of this difference is not that large. On the other hand, small amplitude and phase differences do have influence on larger amplitudes. Calculating the vector difference makes it possible to compare the amplitude and phase difference of the several constituents and determine the relative influence on the tidal signal for the specific tidal stations. Furthermore, the sum of these vector differences per tidal station says something about the quality of the computed tidal singal for these tidal stations. The smaller the sum of the vector difference, the smaller the differences in phase and amplitude.

Several simulations were performed to calibrate the model to compute the tidal wave as accurate as possible. Here the results of the first and final simulation are presented. In between the first and final simulation the bathymetry and roughness parameter were adjusted to get a better performance of the model.

For the first simulation all model parameters as described before, were used. The roughness was set to a constant Chézy friction coefficient value of 65 $m^{1/2}/s$ and the bathymetry as obtained from interpolation was used.

In appendix A (figures A.4-A.13) the phase and amplitudes are plotted in a histogram for all tidal stations. Furthermore, in figures A.14-A.21 maps are shown in which for several stations the amplitude, phase and vector difference are added (for all eight constituents). In this was the development of the errors through the system can be evaluated for each constituent separately. Figure 6 shows this same map with the sum of the vector differences for the different stations.

From the results can be concluded that at Sandy Hook and Block Island the tide is computed quite well, but when travelling more inland bigger differences occur at the tidal stations. The output as presented here was generated for each simulation after every change in roughness or bathymetry.

First the boundary conditions were checked. This was done by looking at the phase and amplitude differences for computed and observed water levels for each tidal station. A constant difference for each tidal station could be solved by changing the boundary conditions. This was only done for the phase of the N2 constituent. For each station in the Lower Bay and up the Hudson the computed N2 phase was overestimated 10 degrees or more. Therefore the phase of the N2 at the boundary was lowered for the southern boundary. Other constant differences over the tidal stations were not observed.

Secondly several adjustments were made to the bathymetry. For example, for tidal station New Haven Harbor (8465705) in the Long Island Sound the amplitudes and phases computed by this first simulation of the model differ from the observed values. A closer look at the bathymetry showed that the depth of the surrounding cells was incorrect, which made proper calculations of the tidal signal impossible. This was corrected by hand.

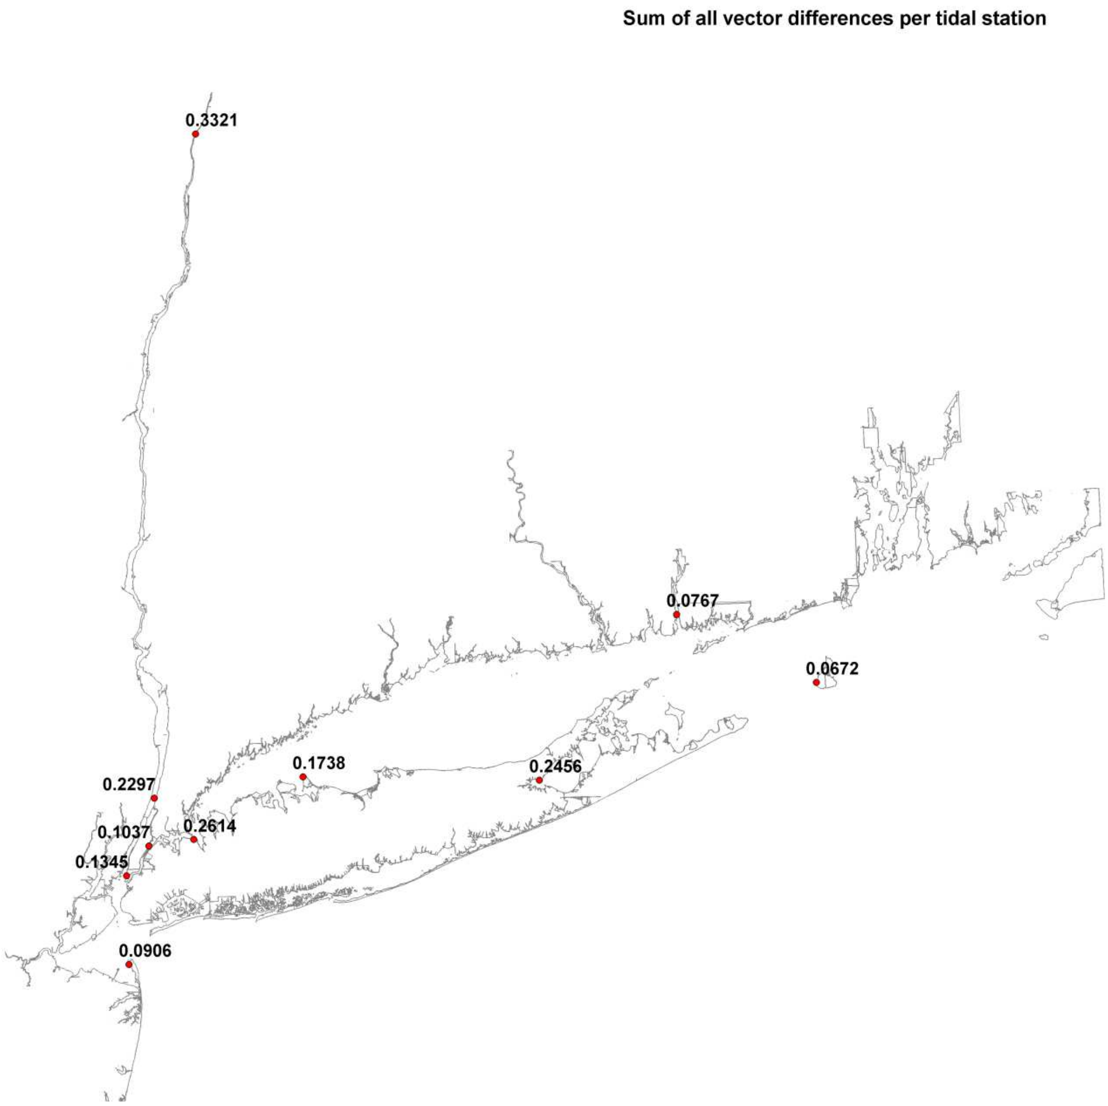

Figure 6 - Sum of vector differences of all harmonic constituents for several tidal stations for the first simulation

When looking at the results of the first simulation for the tidal stations on the Hudson River, it can be concluded that the difference between the observed and the computed data is growing while the tide travels up the river. Part of this was due to the fact that the data used to generate the bathymetry for the river was lacking. Some research on the internet learned that the river had a minimum depth of 9 meters throughout the whole river sections, to guarantee transport by vessels. When looking at the interpolated bathymetry large sections did not reach this depth. So large parts of the bathymetry of the Hudson River were checked and adjusted by hand.

When looking at the development of, for instance, M2 in the first simulation (figure A.14), several conclusions can be drawn. The computed amplitude is too large for all tidal stations in the Long Island Sound. As explained earlier the tidal amplitude is expected to grow in the Long Island Sound, however the model is over estimating this growth. More friction was added to the Sound to compute the tide more accurate. Also the roughness in the Hudson River as adjusted, since amplitude and phase were still not correct after adjusting the bathymetry. Figure 7 shows the results after the last simulation. Here again the sum of all vector differences for the same stations is added to the map. Also, in appendix A, histograms and map illustrations per constituent are added for this last simulation.

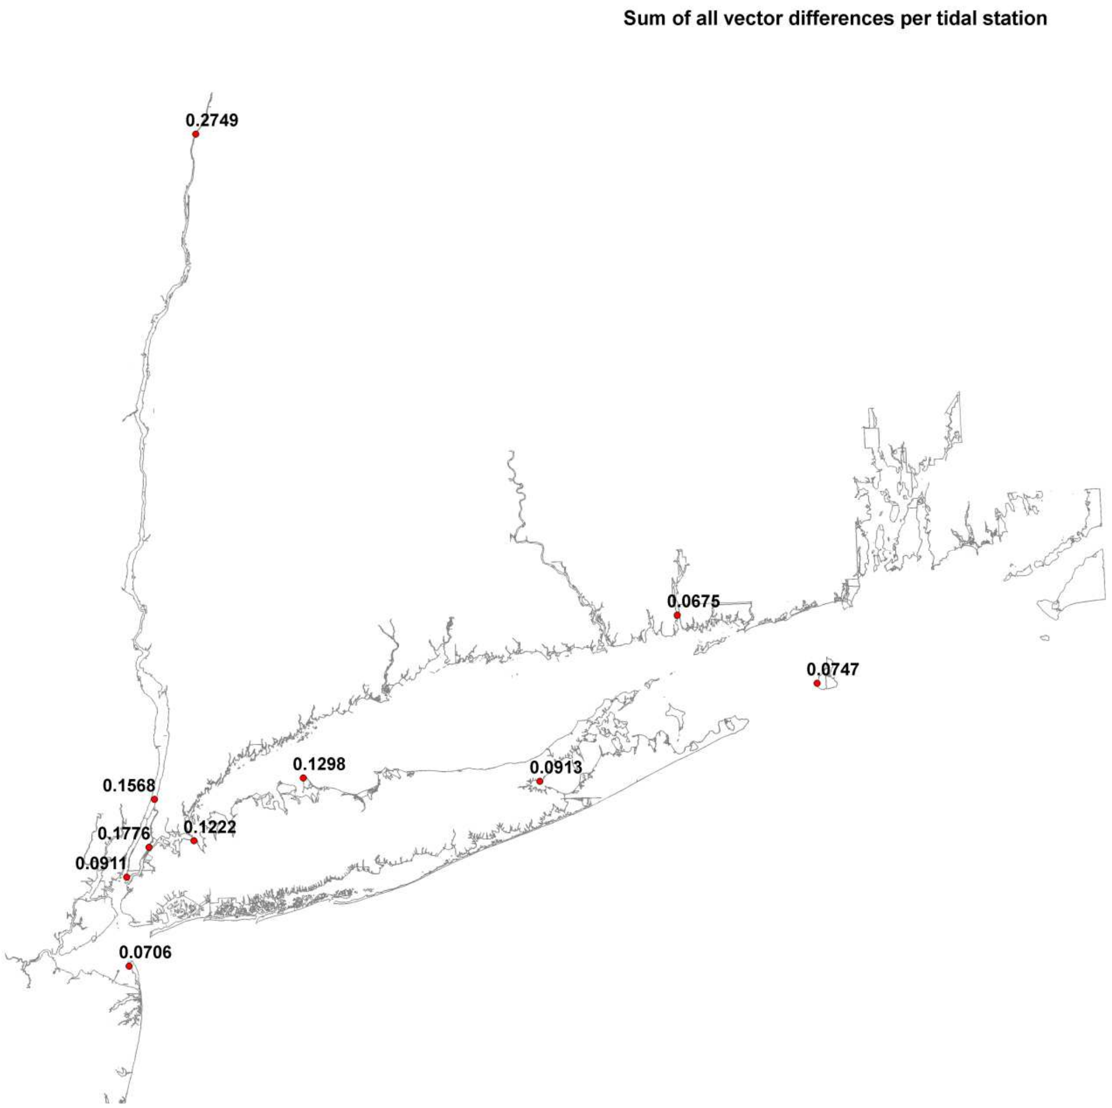

Figure 7 - Sum of vector differences of all harmonic constituents for several tidal stations for the final simulation

During the calibration process the amplitude and phase errors for most stations were lowered. As can be concluded from figure 6 and 7, where for all tidal stations, except Horns Hook, the sum of vector difference is lower for the final simulation. This tidal analysis focuses most on the three largest constituents (M2, N2 and S2). In the histograms and figures in appendix A can be seen that for most stations the amplitude and phase errors are within reasonable limits. Take for example figure A14 where the errors for the first and final simulation are shown for the M2 constituent. During the process some errors have slightly risen, but overall the amplitude and phase errors have gone down.

Sandy

Sandy was modeled using the model input described in the model setup chapter. Time series of water levels at the boundary were available from a model covering a larger part of the North Atlantic Ocean. Also wind and atmospheric pressure data was added. The model simulates the hydrodynamics from the 25th of October 2012, which is several days before the hurricane strikes New York, until the 30th of October. In the appendix several graphs are added (figures A.22-A.25) showing both the observed and computed water levels before and during hurricane Sandy. New London, New Haven, Kings Point, Montauk and Bridgeport are situated in the Long Island Sound. Sandy Hook, The Battery and Bergen Point are situated near the city of New York. Beneath the results are drawn for Sandy Hook and Kings Point for the observed, D-Flow FM computed and Delft3D computed water levels. It is hard to compare the Delft3D and the D-Flow FM results, and also not really relevant. From the two figures beneath we can conclude that both models simulate the storm quite well.

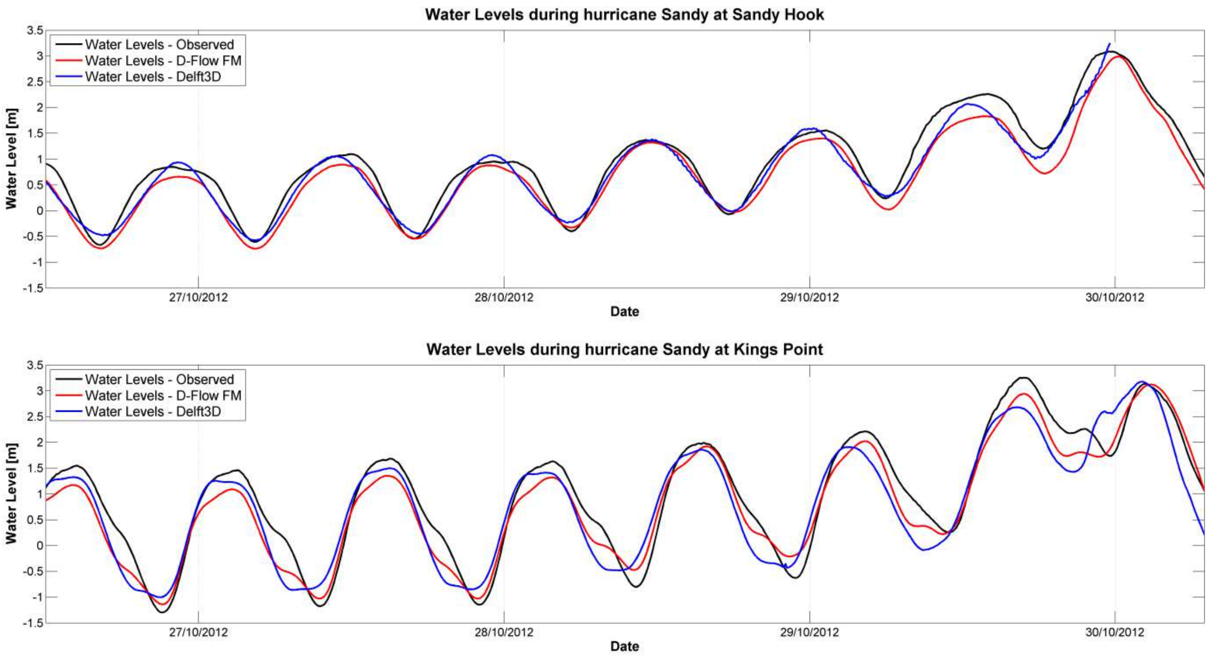

Figure 8 – Water levels during hurricane Sandy at Kings Point and Sandy Hook

It is more interesting to compare the observed levels to the D-Flow FM computed water levels. The figures in the appendix show that the water levels in the Long Island Sound are computed accurately. But for the other stations the results are reasonable as well. The maximum difference between the computed and observed water levels is about 50 centimeters, but in most cases it is less. It should get better, but this first simulation at least shows that D-Flow FM is able to compute the surge levels during the hurricane. And during the days before the hurricane we see the tide being simulated rather well on most locations.

Several remarks have to be made regarding this model. Some of these have negative influence on the results and fixing these problems could lead to more accurate modeling.

Firstly, the initial water level in the D-Flow FM model is not correct (see figures in appendix). The initial water level should be zero (initial water level is set to zero as input), but for all stations the initial water level is lowered to a level of about -0.23 meter. This might be a bug in D-Flow FM and has to do with the atmospheric pressure field. When running this model without applying this atmospheric pressure, the simulation starts at zero meter water level. However, the results throughout the rest of the calculation do not differ that much (from the results including the atmospheric pressure field), so it seems like a problem only occurring at the initial state.

Another remark is the fact that the bathymetry data used was measured in meters below MLLW, while MSL should be used. This was something which is not that hard to correct. However, for this study this could not be done anymore since it was found out too late. In practice MSL is about 50 centimeters higher than MLLW. So water depths are not completely correct in the model, which does have its effects on the propagations of the tidal wave.

Furthermore, the boundary conditions should be improved as well. The larger model on which the boundary conditions are based needs more calibration to get better results. Through these boundary conditions this influences the results of this D-Flow FM model as well.

The last thing is the fact that several phenomena cannot be simulated by the model, while they can play a role. For example, wind set up due to a wind blowing from one direction persistently (high wind speed is not needed) can cause a setup which cannot be calculated in this model and should be included in the boundary conditions. Also larger scale oceanographic processes are known to play a role in this region, but are not included in the models. These processes might cause a setup which is present in the observed results, but not in the computed water levels. This can explain some of the differences, since at most locations the computed water levels are lower than the observed levels.

Improvements

The duration and focus of this study did not allow a thorough calibration of the model. It was decided to do only a quick and dirty tide analysis to be able to show that the model works correctly. The goal of this internship was to finally simulate hurricane Sandy with the model, even if it is not completely calibrated correctly. So, after the tide analysis hurricane Sandy was simulated to show that the model works. The results of the tide analysis and the model run on hurricane Sandy have been presented in this report. These results show that the model is capable of modeling Sandy; however, improvements to the model are certainly needed to make its performance even better. Of course, whether better results are needed or not depends on the available time and the purpose of the study (the wanted accuracy of the results).

The bathymetry data used for this model was not sufficient at all places. Improvements can be made on the bathymetry in the Hudson River, as already mentioned. Another example is Jamaica Bay, for which not enough sufficient data was available to capture all small islands and wetlands. Overall, it would be best to update most of the bathymetry data of inland waterways. The Ocean is captured quite well, but more inland, when rivers get smaller, data with a higher resolution could improve the results significantly. Also, a correction should be made to the bathymetry data used at this moment. As already mentioned, the data is a depth below MLLW while a depth below MSL should be used.

The grid itself can also be improved in certain areas. Jamaica Bay is such a location; the grid could be adjusted just as the bathymetry such that it captures the islands and wetlands in the area more accurate. Another improvement could be the use of 1D channels for some of the smaller rivers which are included in this model with a grid with only one or two grid cells over the river width. A third improvement could be including low lying areas into the grid. These are the areas which are flooded when water levels rise. Including these areas in the grid makes it possible to see which areas get inundated and at which water levels. Running a hurricane model this could give insight in the risk of low lying areas. Also this makes it possible to generate animations in which you can clearly indicate which areas are at risk and visualize these risks.

Furthermore, for some tidal stations (tide analysis) the difference between the observed and computed tidal constituents is still too large, especially on the Hudson River. This could be improved by adjusting the above mentioned, but at some locations local adjustments might still be needed.

When it comes to modeling hurricane Sandy several improvements were already mentioned in the chapter before. Overall it can be concluded that the model works quite well. The computation time is not too long, and the results are for most stations close to the observed water levels. However, the initial water level difference, the calibration of the large scale model, the large scale process which are not taken into account and therefore the errors in boundary conditions are things which could be improved to achieve even better results.

D-Flow FM experiences

After working with D-Flow Flexible Mesh for several weeks I have some experience on the usability of the software at this stage. I started off knowing nothing about setting up a model in D-Flow FM, but with some experience in modeling in Delft3D. In Delft3D the grid is made in RGFGRID. I had a new version of RGFGRID in which it was possible to make Flexible Mesh grids, however, this is still being developed so it was not working correctly. Fortunately it is possible to generate grids in D-Flow FM itself as well, so I decided to make it in D-Flow.

Generating the grid works quite well in D-Flow FM. Using splines it is easy to make a curvilinear grid, which afterwards can be converted to a flexible mesh grid. Also it is possible to import curvilinear grids and then convert to FM grids. For some parts a curvilinear grid was not suitable and a triangular grid was needed. This was the case for connections between parts of the grid and for areas with a very irregular shape. The simplest way to get an irregular grid is to click a polygon in which the program can generate a triangular grid. This option works fine. It is also possible to click a grid by hand, which can be the easiest way to connect two grids.

Interpolation of the sample files on the grid worked fine as well. The only disadvantage compared to RGRGRID was the fact that it was not possible to load several sample files at the same time. But this could be easily done in other programs and then saved as one sample file. Viewing of and zooming in on the created network is possible. After you get used to the keys it is easy to zoom in or out. However, when many things have to be displayed (for example many sample points) loading the new window can take quite a long time. To conclude, building a grid in D-Flow FM is not difficult and with the current software this works fine.

Applying boundary conditions is done in a different way than it was done in Delft3d. Since the grid is not numbered anymore with N and M values, it is not possible anymore to define a boundary conditions to a line between two points (N,M). In D-Flow a polygon is drawn consisting of several points. At these points boundary conditions can be specified, which are then interpolated over the polygon. Applying those conditions is not possible in the GUI yet, so this have to be done in the polygon files itself.

The initial conditions can be changed in D-Flow itself, however, the MDU file (which is like the MDF file in Delft3D) is the easiest way to change roughness, initial water level etc. Furthermore, things wind data, atmospheric data or local roughness factors also have to be added by adding data files and couple these to the model by the external forcing file (.ext). This is not possible to do in the GUI at the moment. However, Deltares is working on Delta Shell, which is the new GUI and would make this all possible. I have worked with it before (only with the 1D module) and it looked good and was easy to work with. However, until then, the current user interface is not very sophisticated, but it does work. And this way of working (with the .ext and .mdu file) works fine as well.

What is not working correct at the moment is the way the program handles with wrong input. Sometimes the software just shuts down without warning or error after starting a model run. During the tide analysis I had this problem several times. It can be very annoying that the program does not give an error, but just shuts down. I also had the experience where it happened the other way around. The model does not stop running, even though some input is not correct. After the computation you find out that you did something wrong and you have to start over again. However, for the more experienced user these issues are of less importance, since then you exactly know how to deal with the software.

Appendix A - Figures

Figure A.1 – Used tidal stations (red dots) in project area

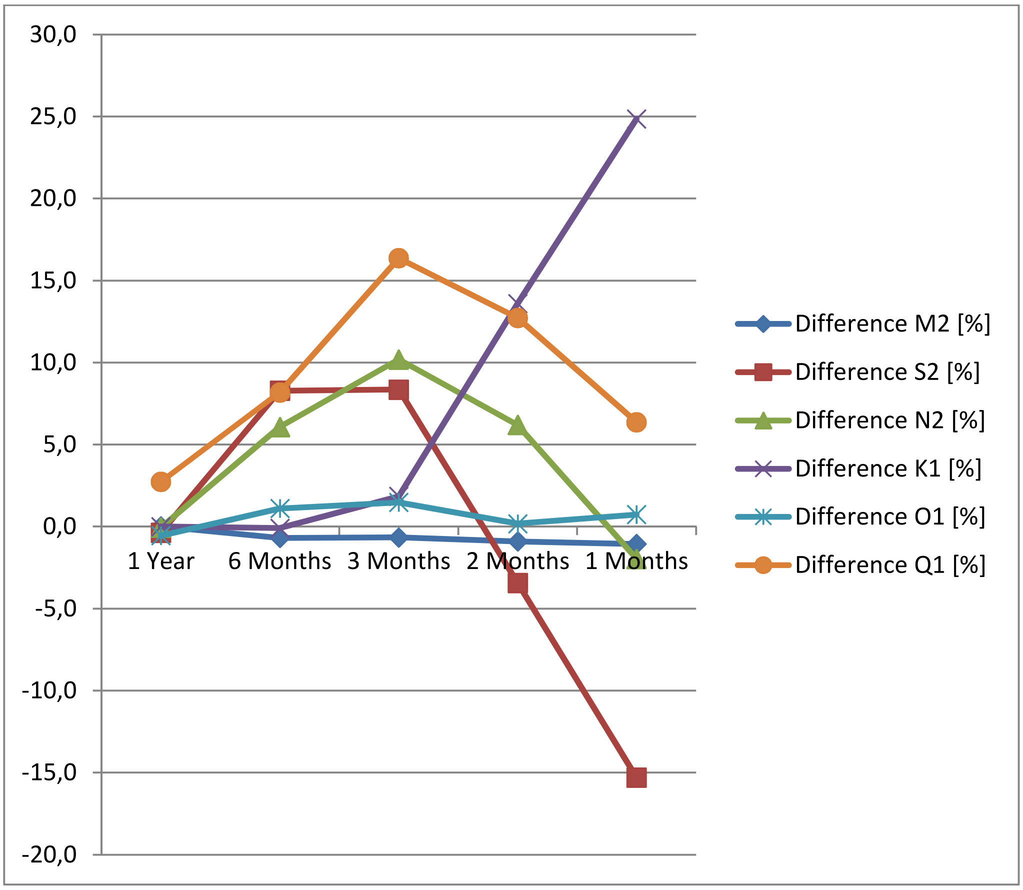

Figure A.2 – Amplitude difference between NOAA data and computed data [percentages] for each time series duration

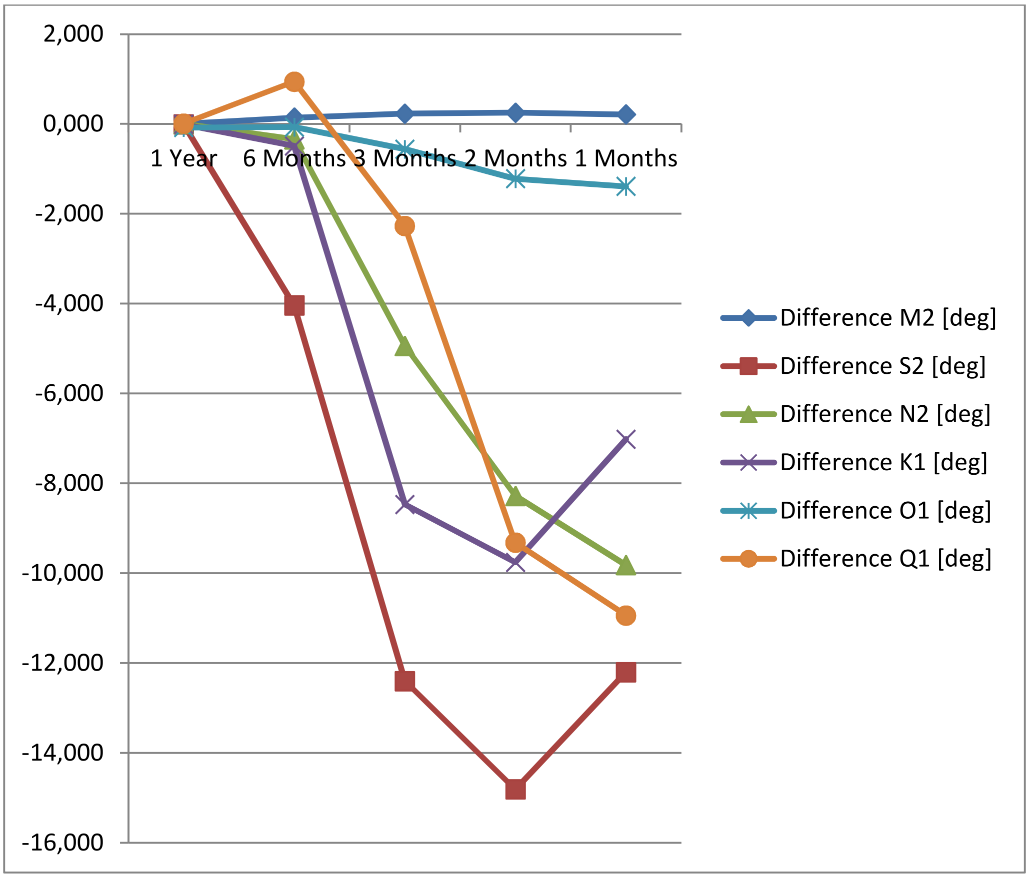

Figure A.3 – Phase difference between NOAA data and computed data [degrees] for each time series duration

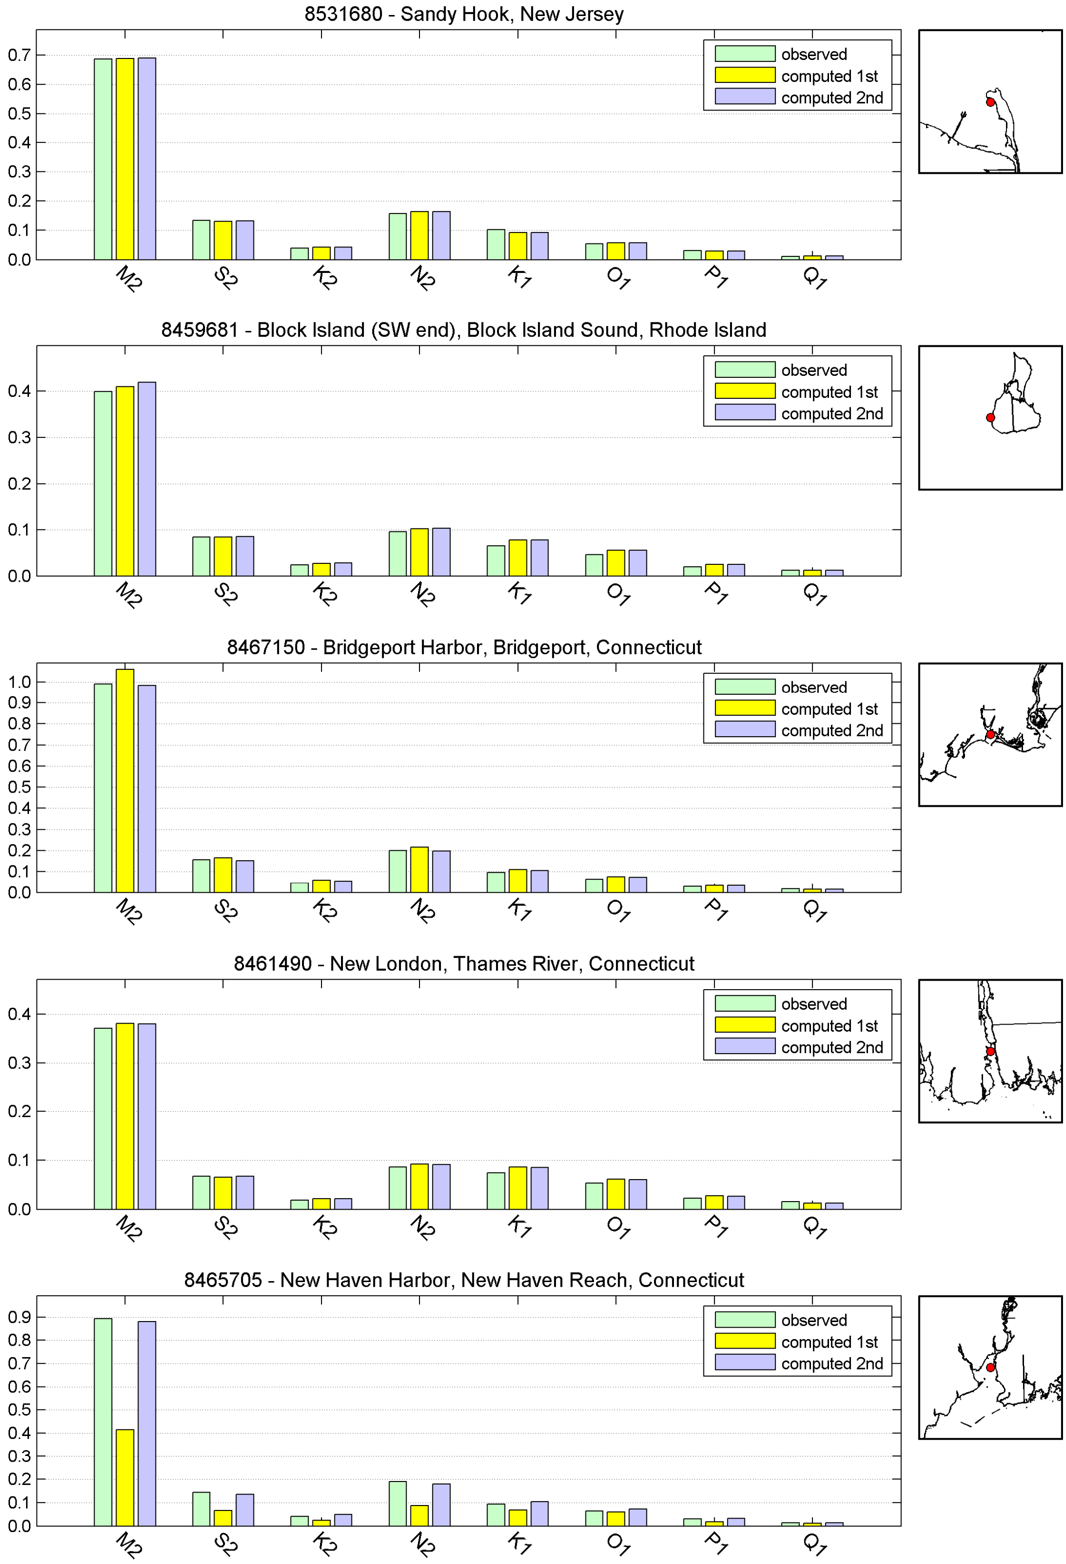

Figure A.4 – Amplitudes for main eight constituents of observed and simulated data (first and final simulation) per tidal station

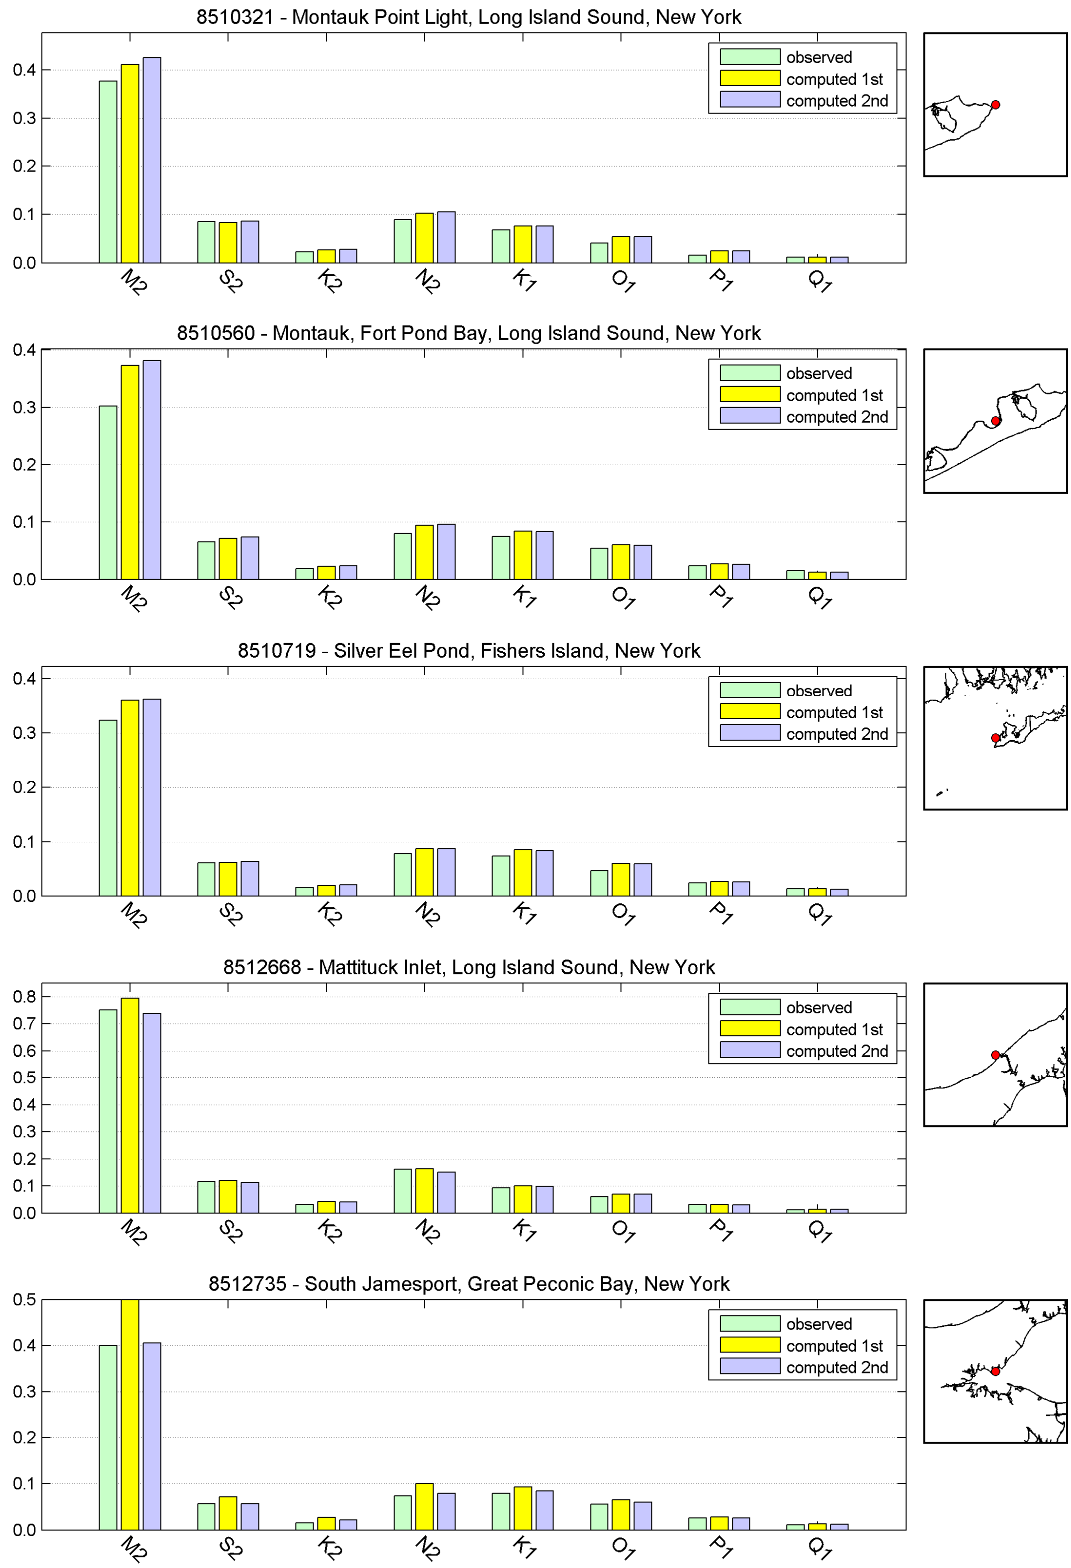

Figure A.5 – Amplitudes for main eight constituents of observed and simulated data (first and final simulation) per tidal station

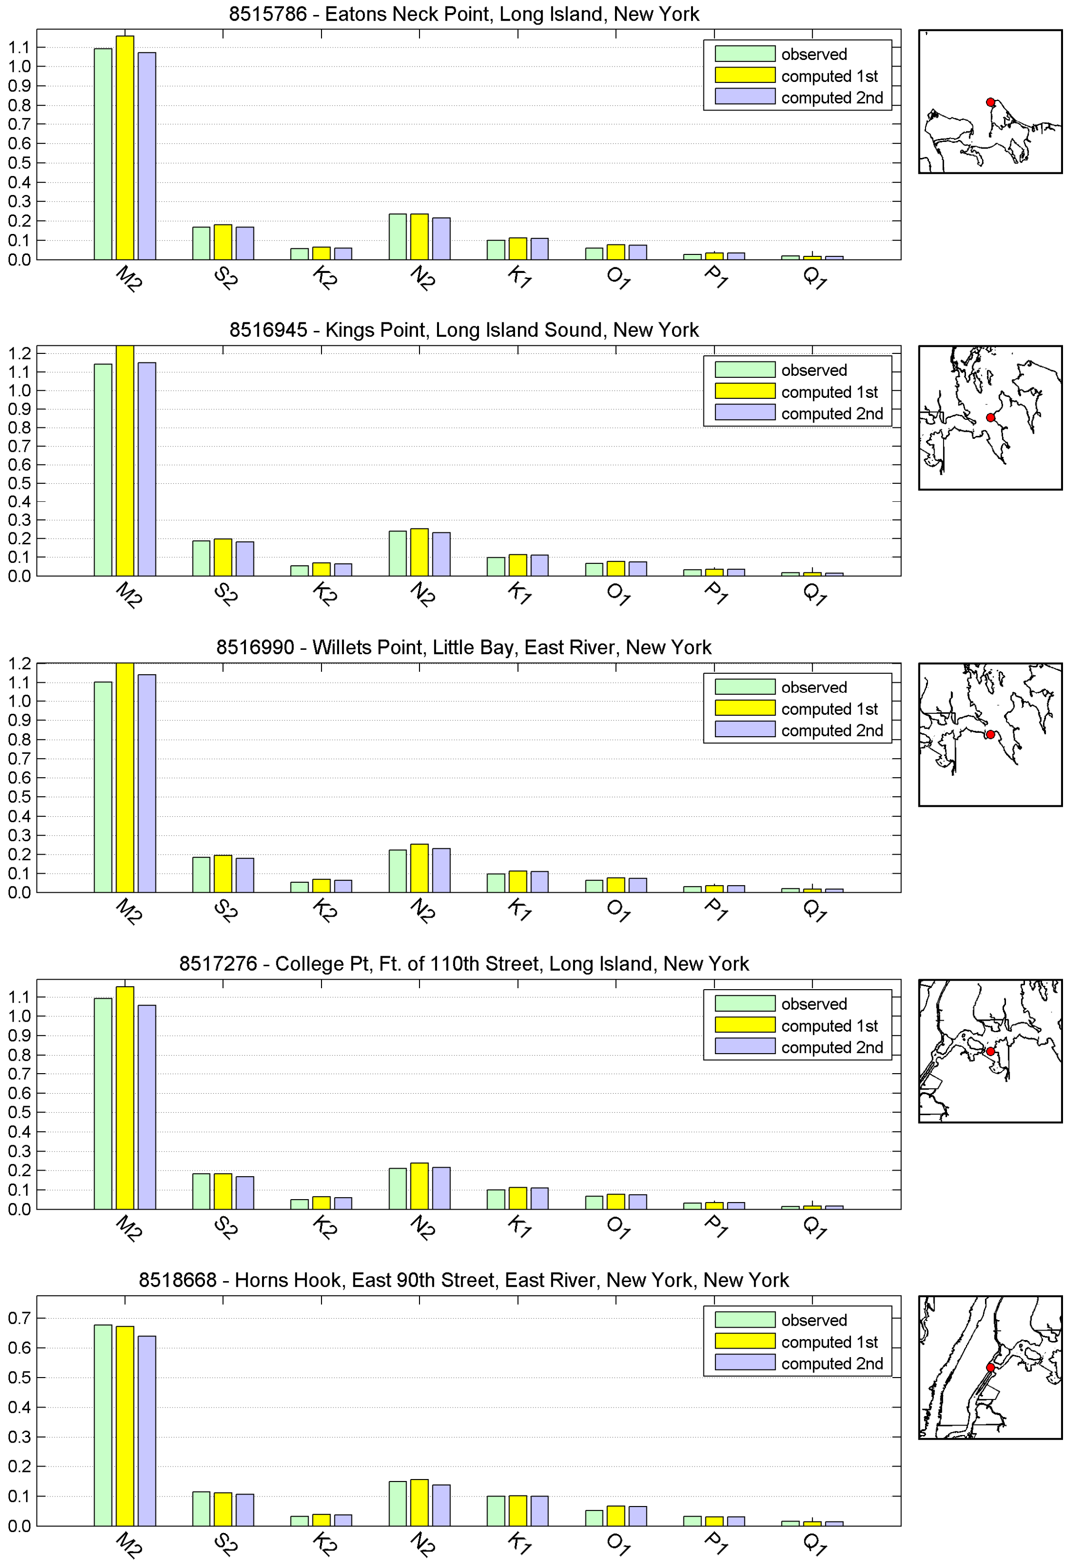

Figure A.6 – Amplitudes for main eight constituents of observed and simulated data (first and final simulation) per tidal station

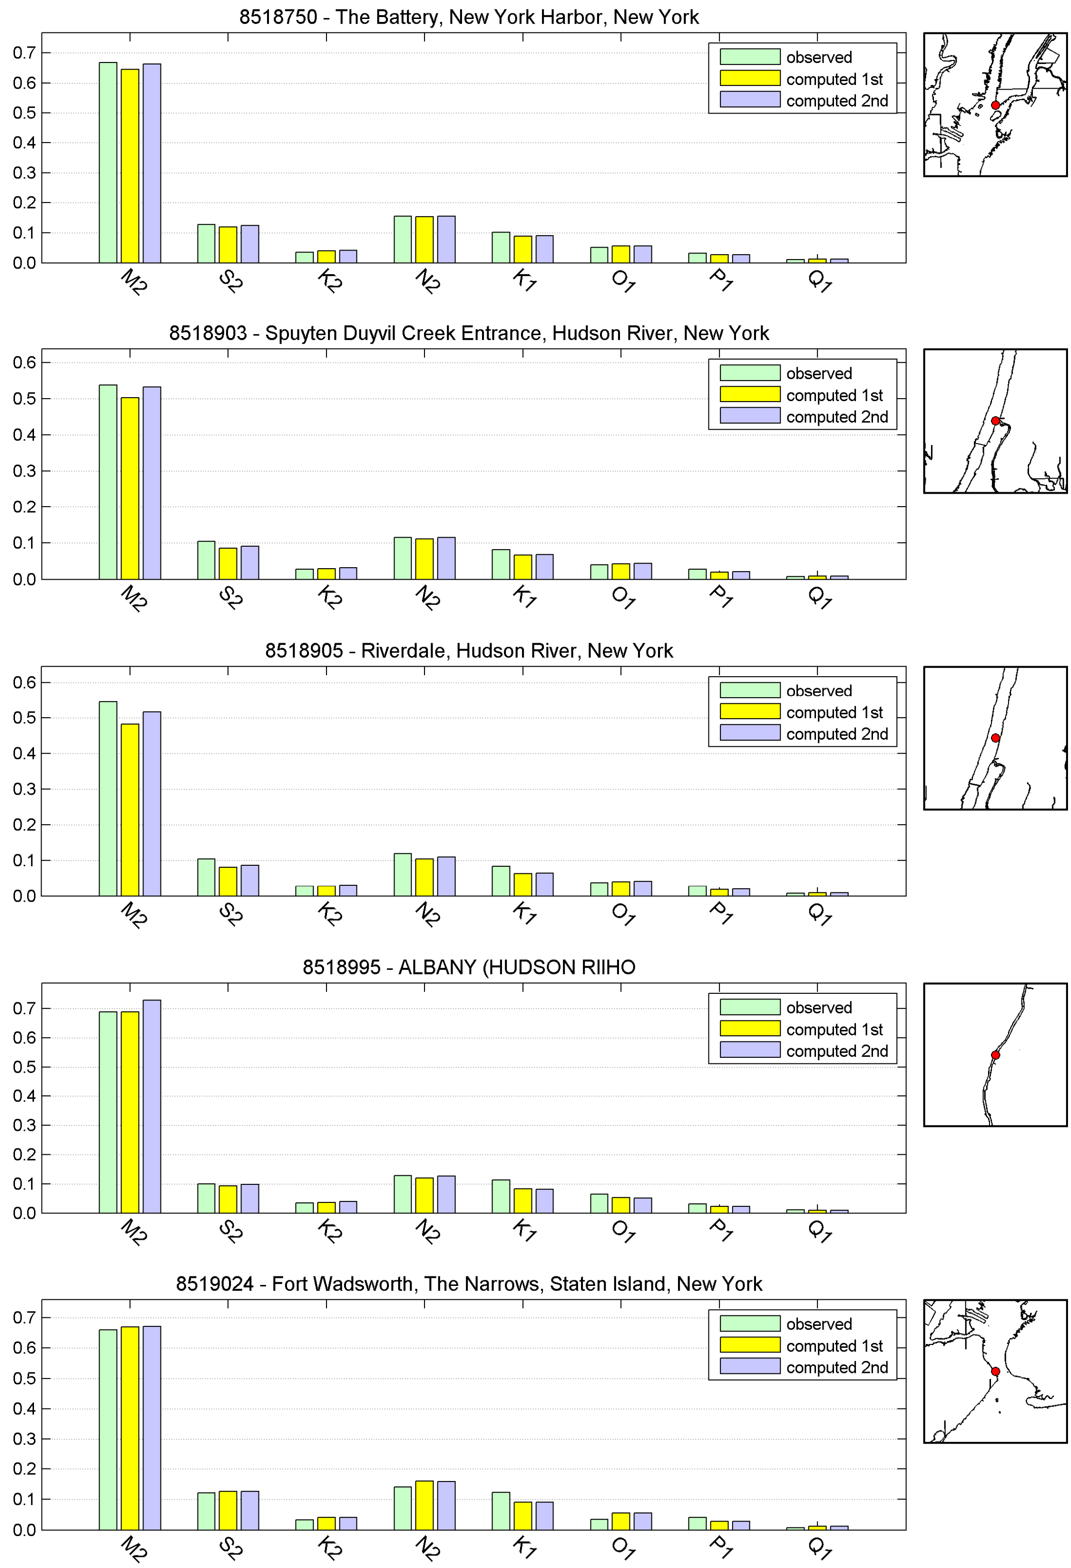

Figure A.7 – Amplitudes for main eight constituents of observed and simulated data (first and final simulation) per tidal station

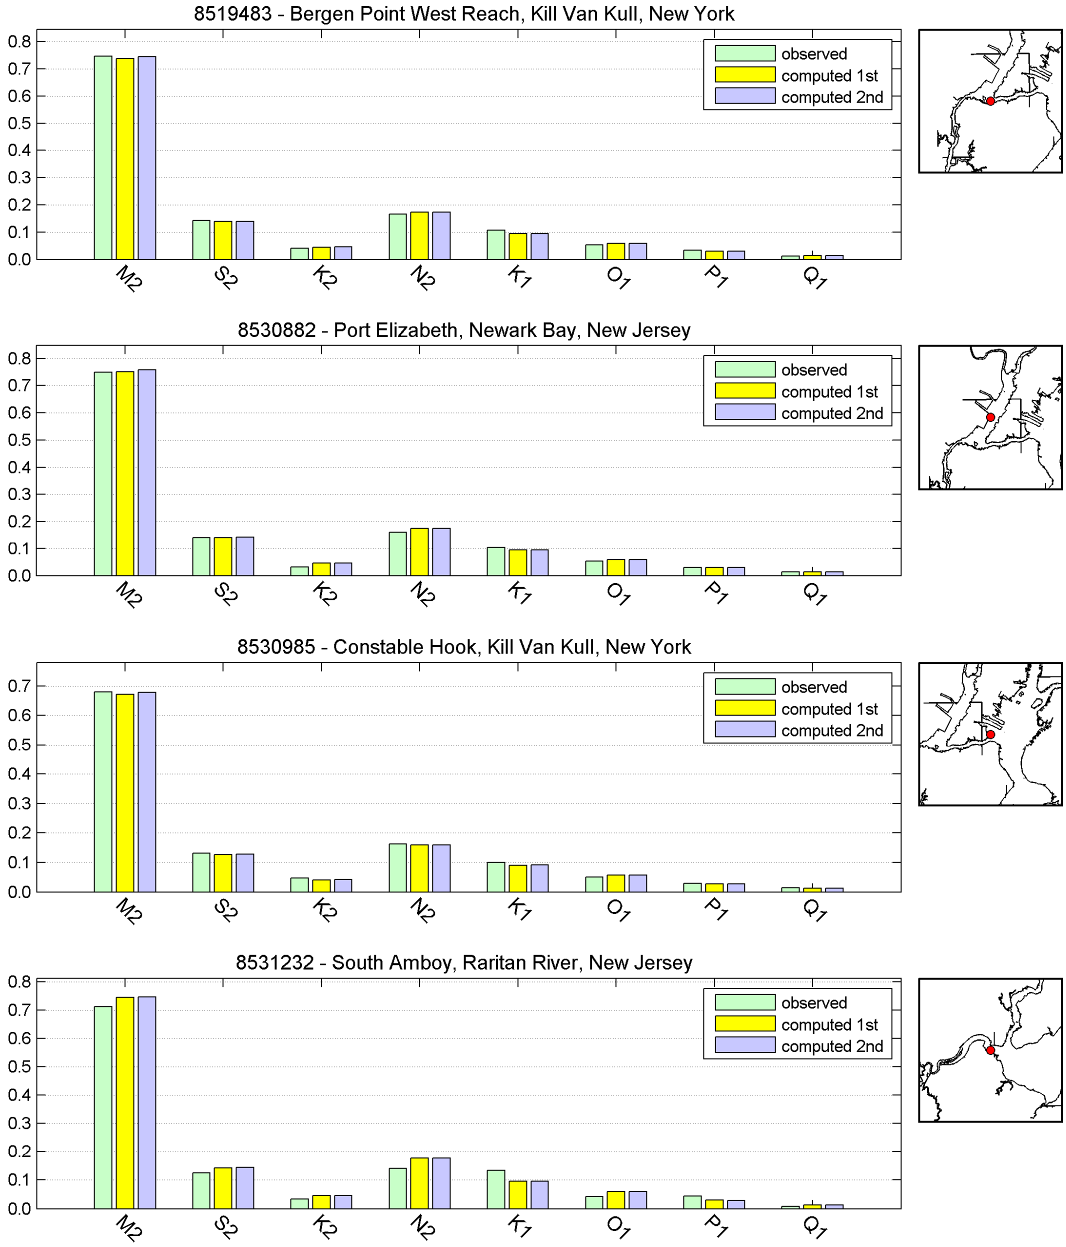

Figure A.8 – Amplitudes for main eight constituents of observed and simulated data (first and final simulation) per tidal station

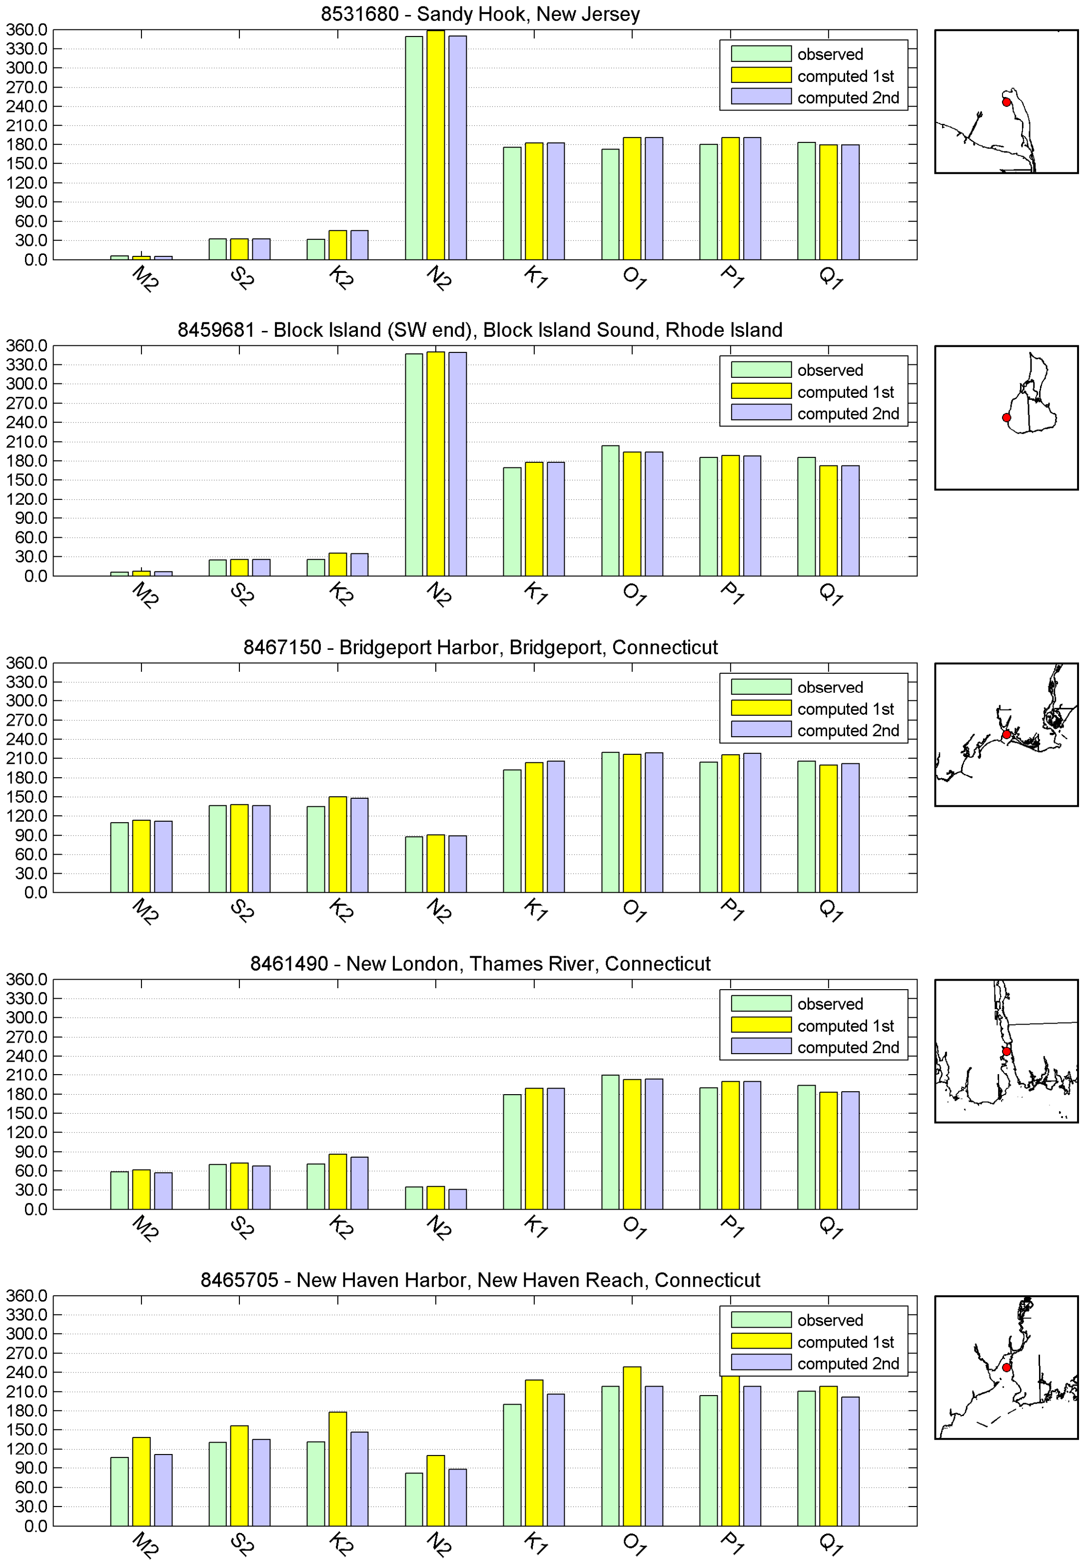

Figure A.9 – Phases for main eight constituents of observed and simulated data (first and final simulation) per tidal station

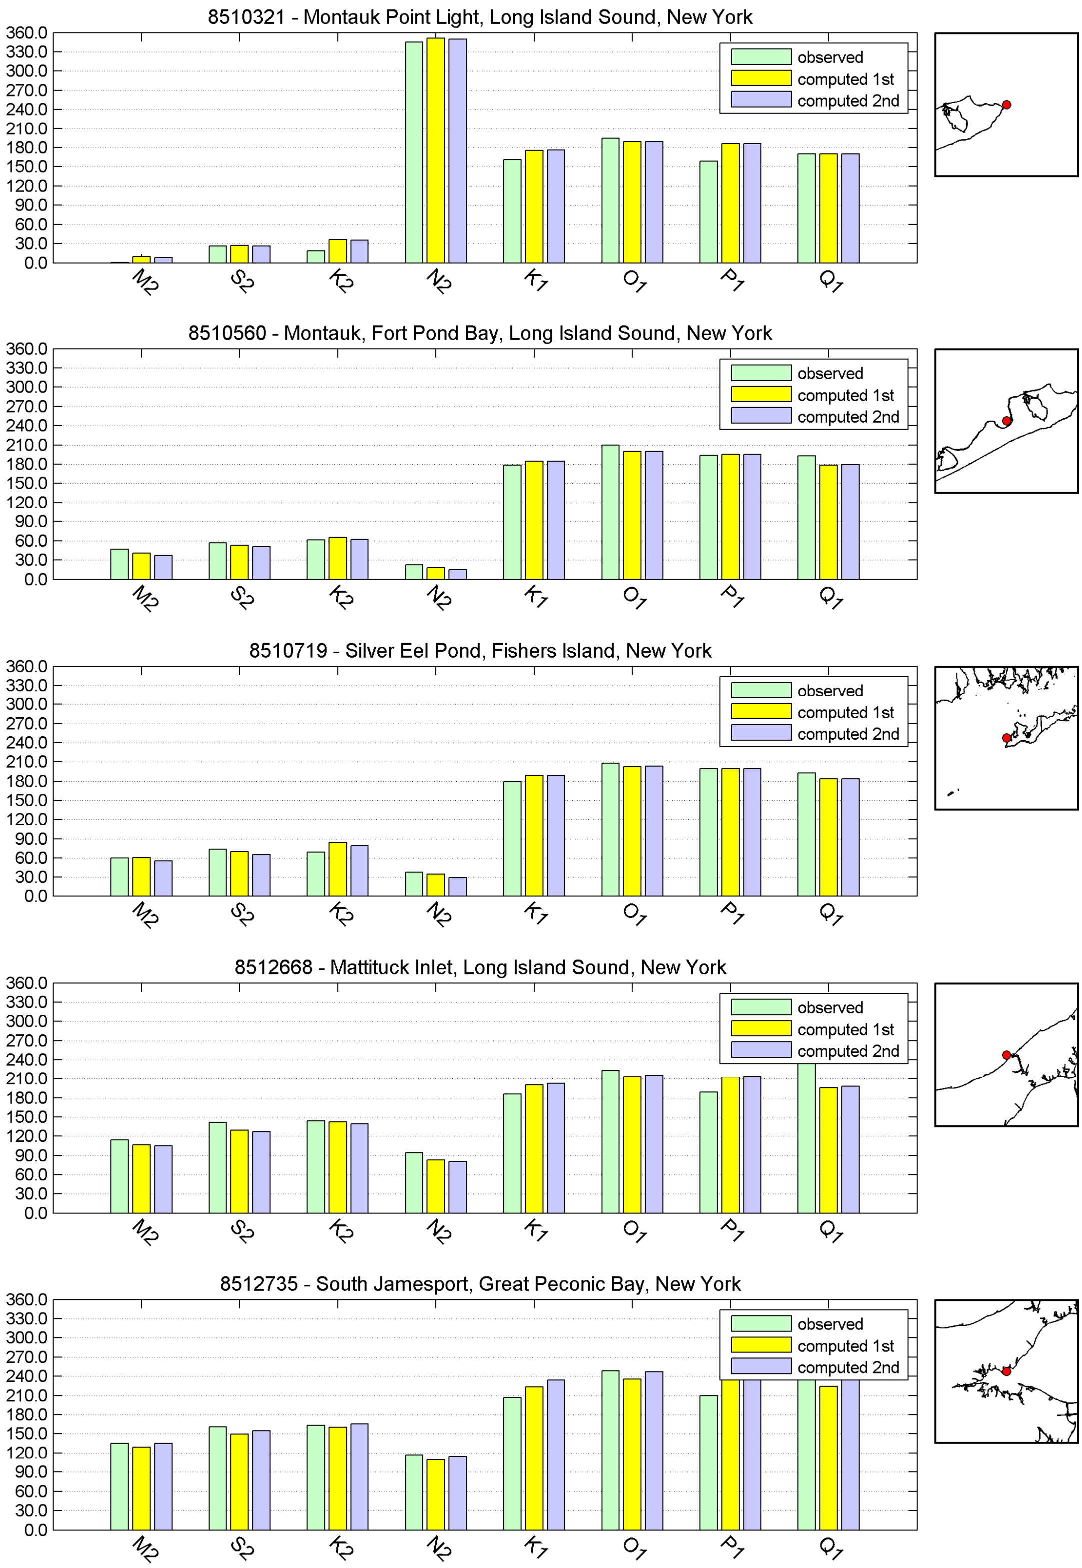

Figure A.10 – Phases for main eight constituents of observed and simulated data (first and final simulation) per tidal station

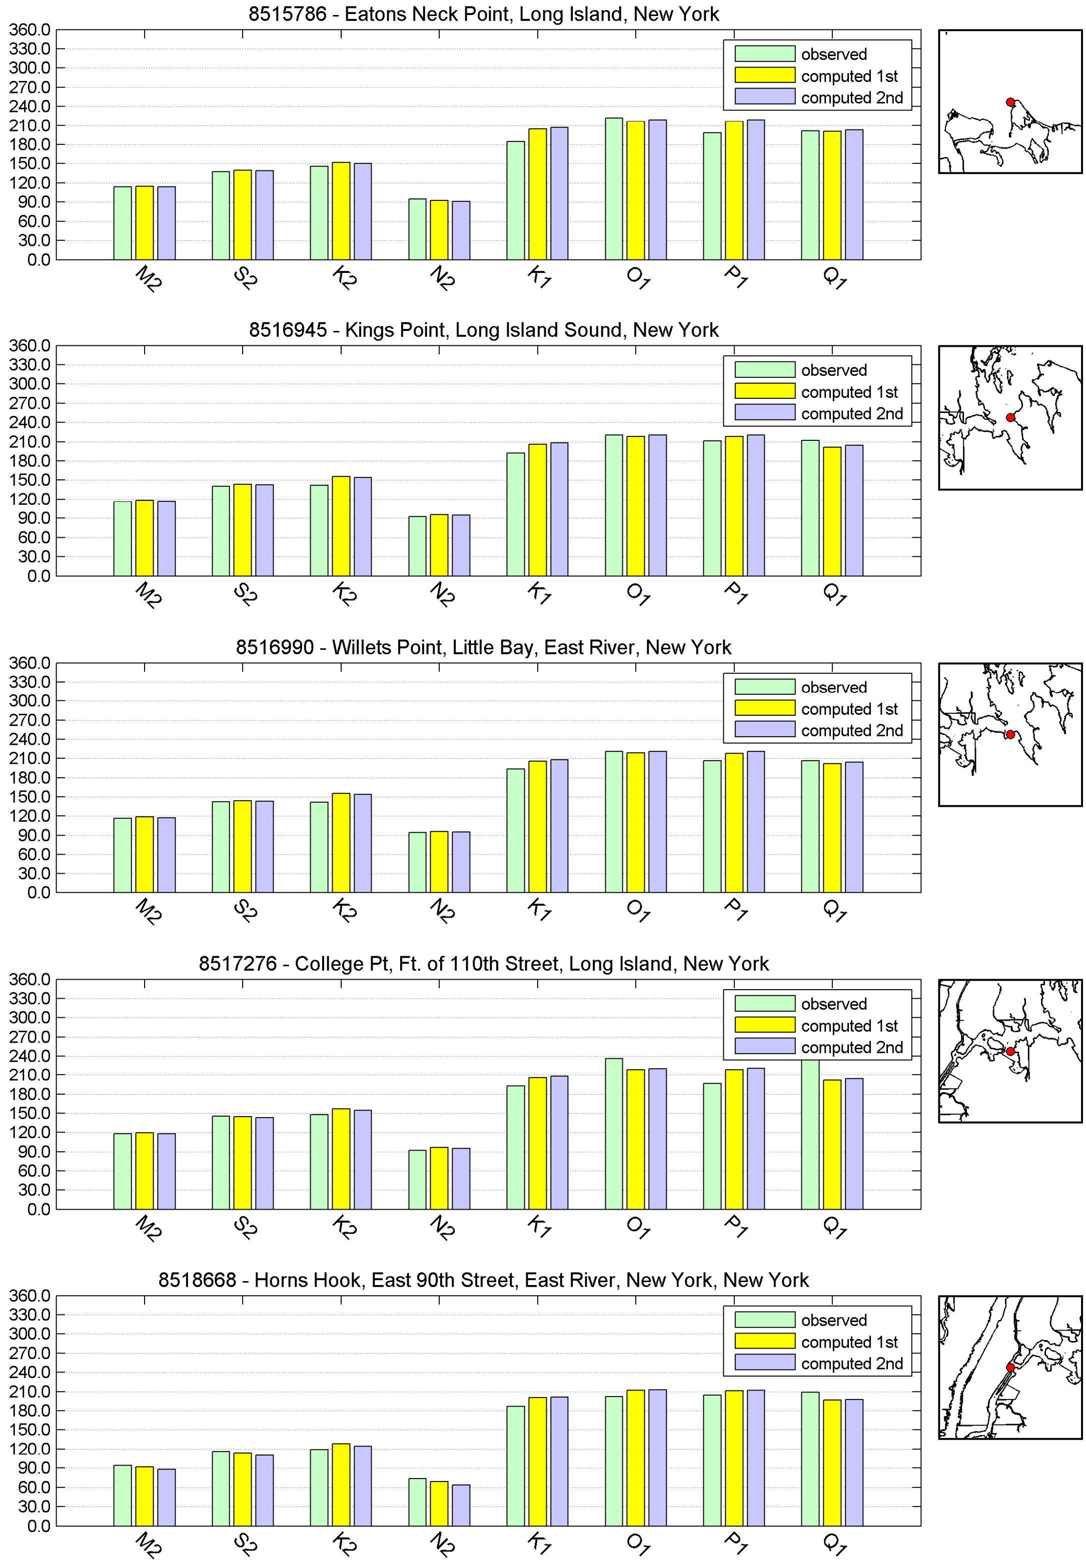

Figure A.11 – Phases for main eight constituents of observed and simulated data (first and final simulation) per tidal station

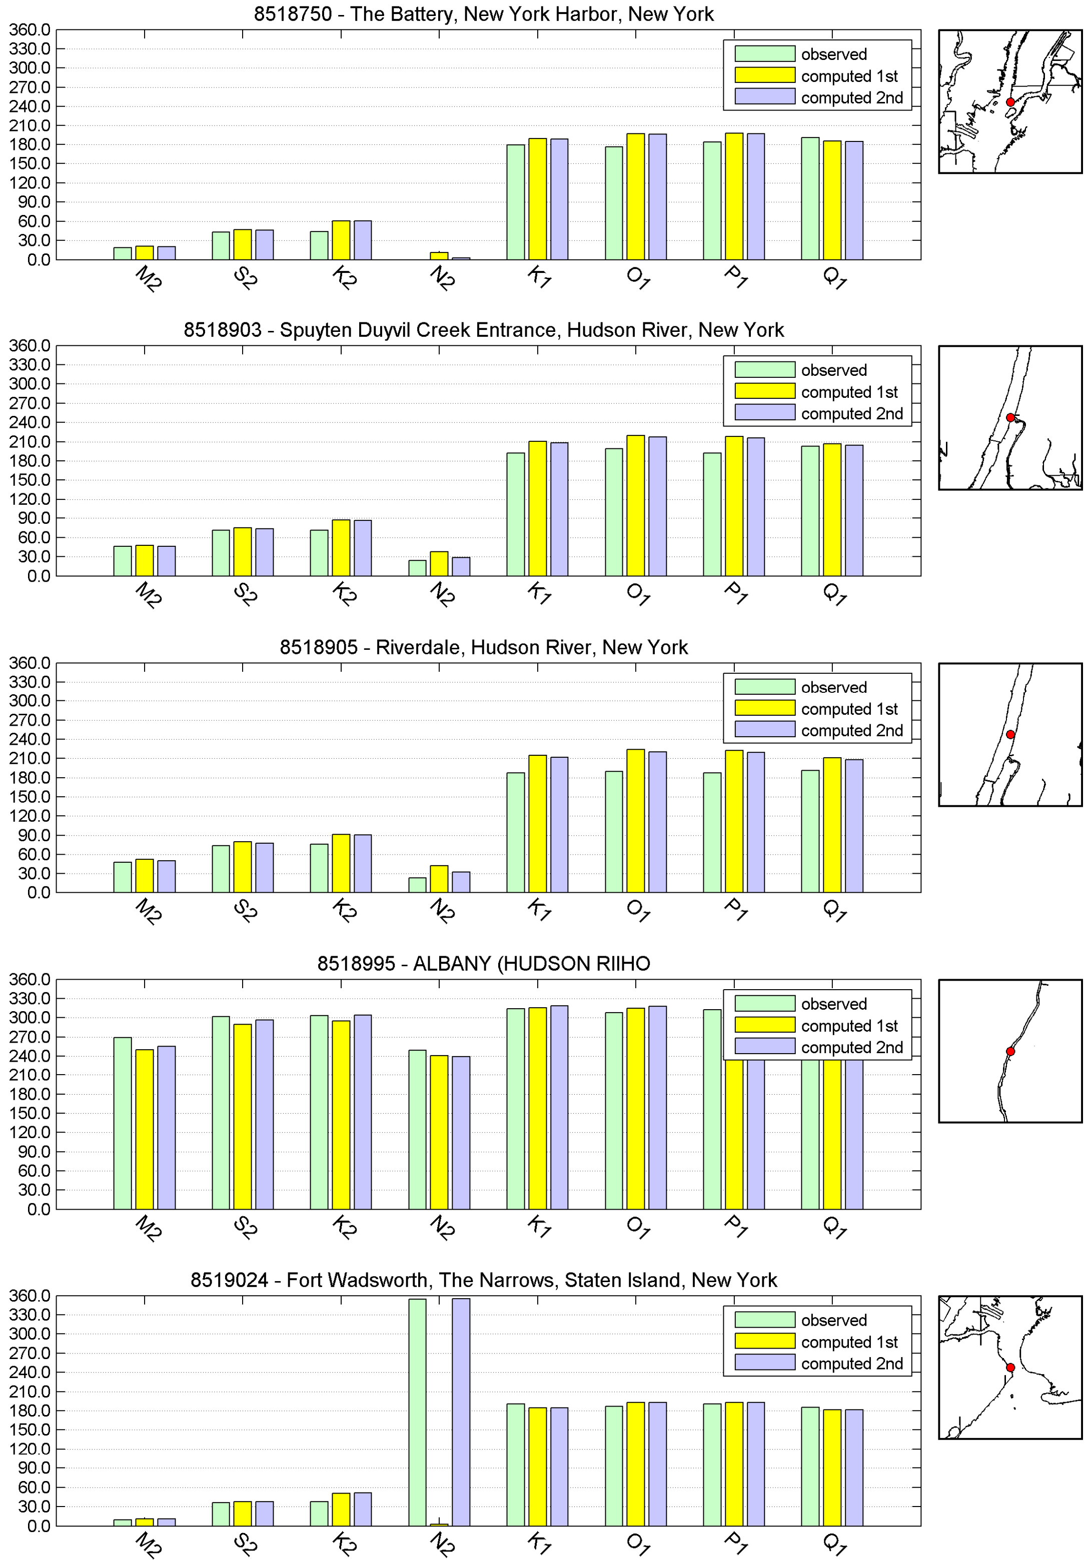

Figure A.12 – Phases for main eight constituents of observed and simulated data (first and final simulation) per tidal station

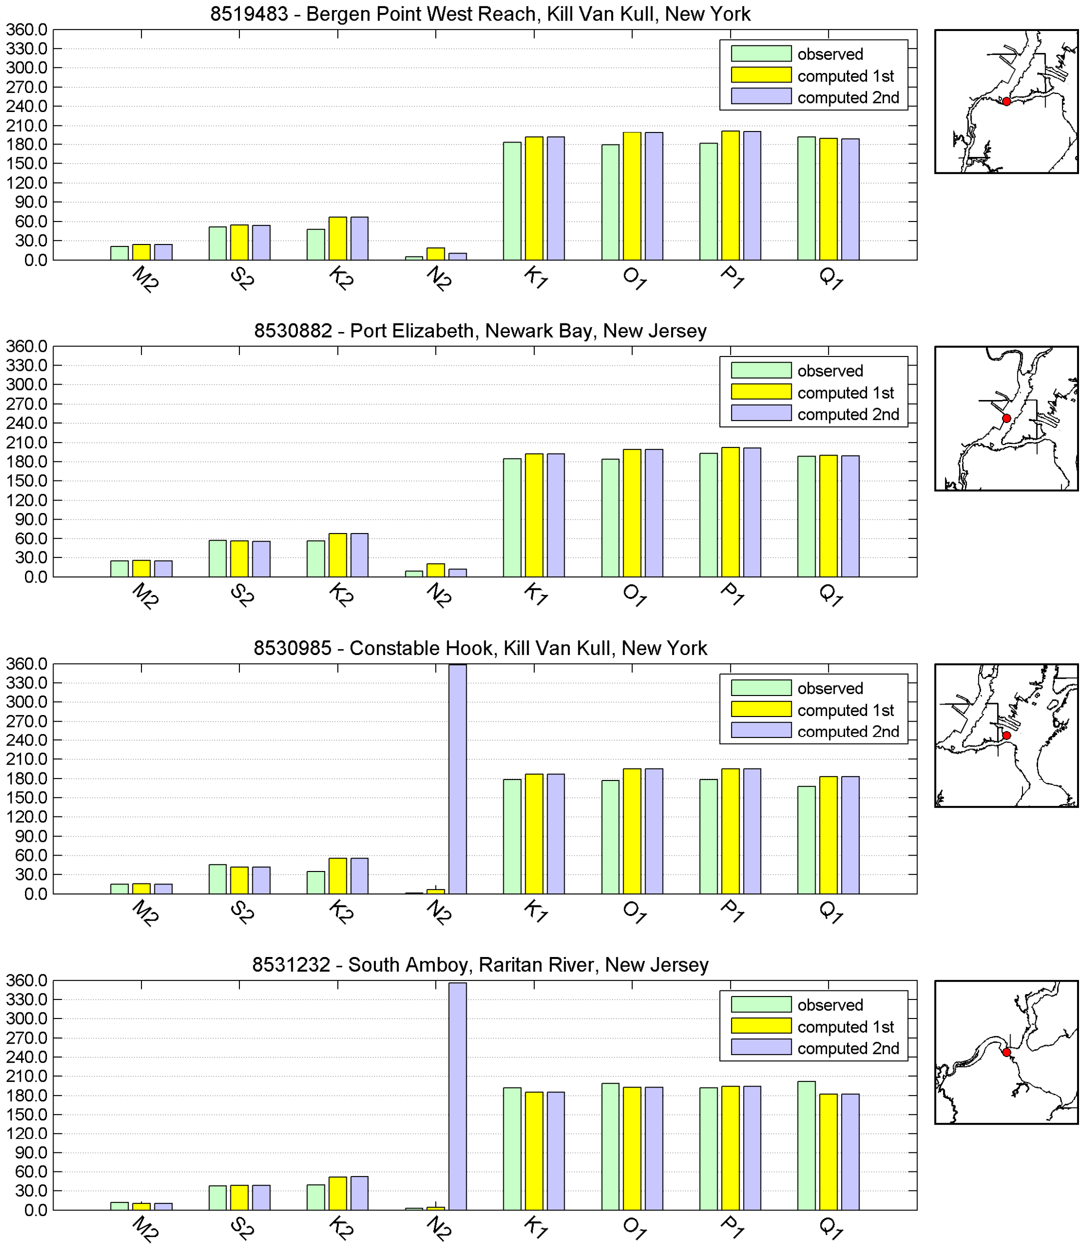

Figure A.13 – Phases for main eight constituents of observed and simulated data (first and final simulation) per tidal station

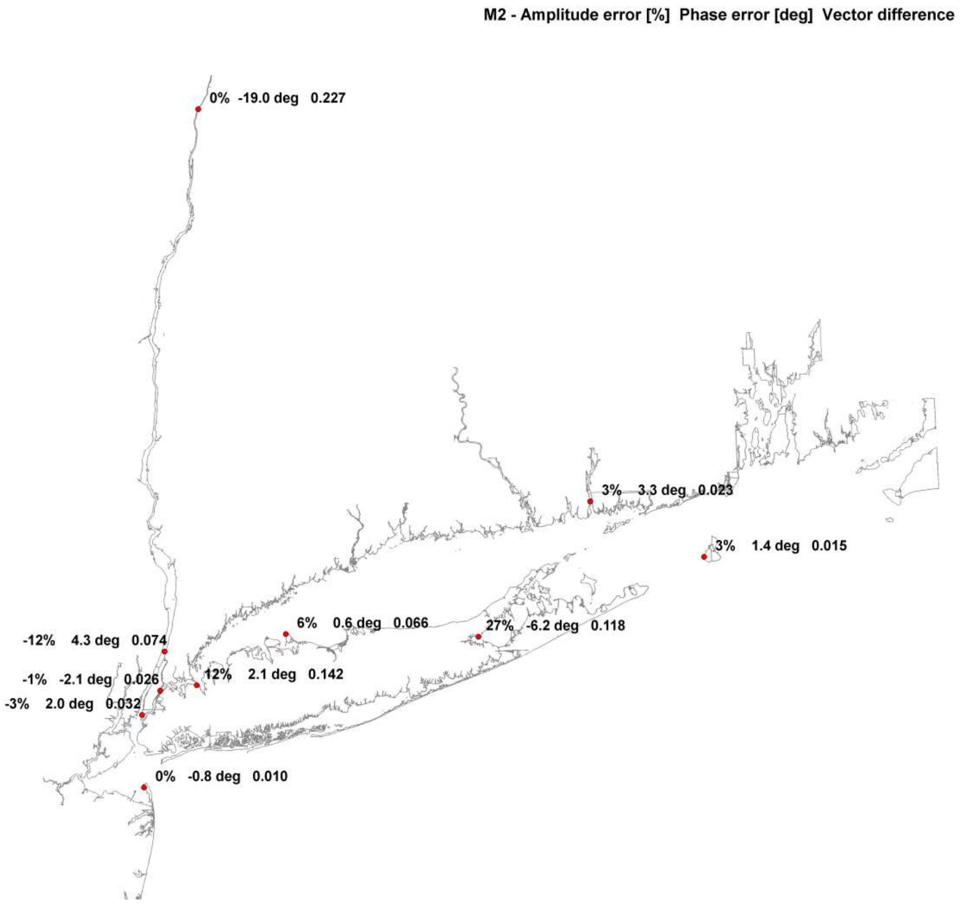

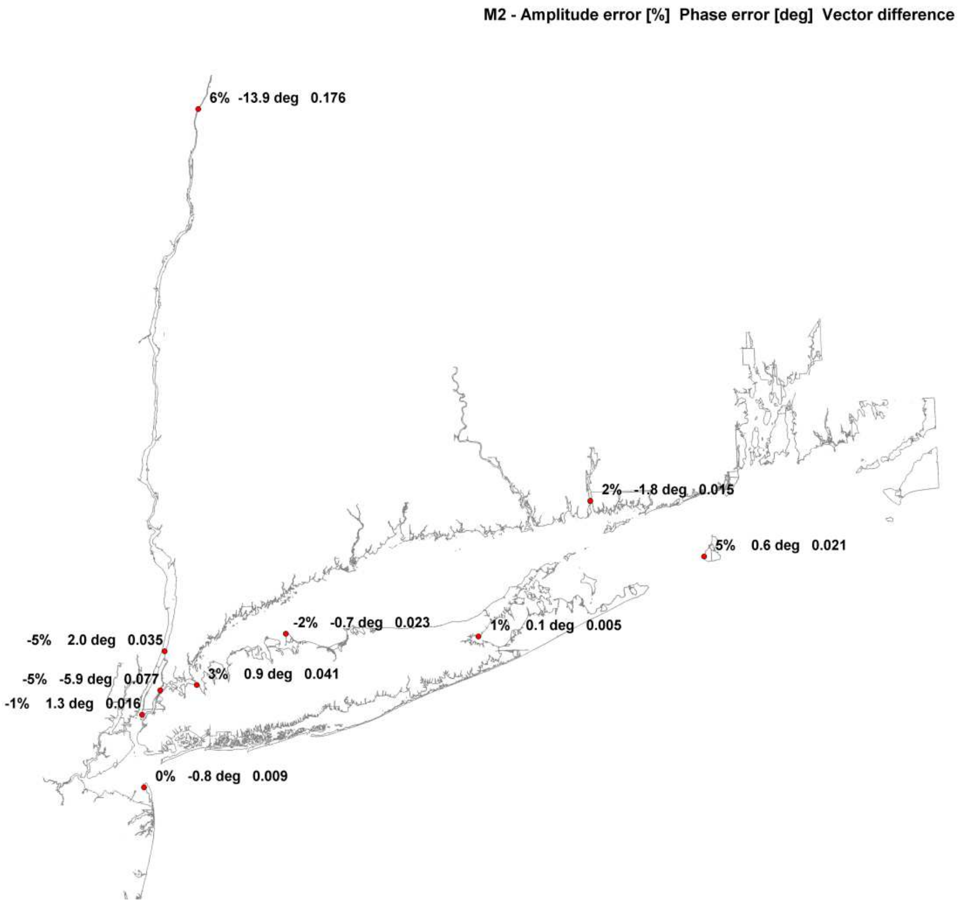

Figure A.14– Differences between observed and computed data of M2 constituent for first simulation [above] and final simulation [beneath]. A positive phase or amplitude error means higher computed than simulated data.

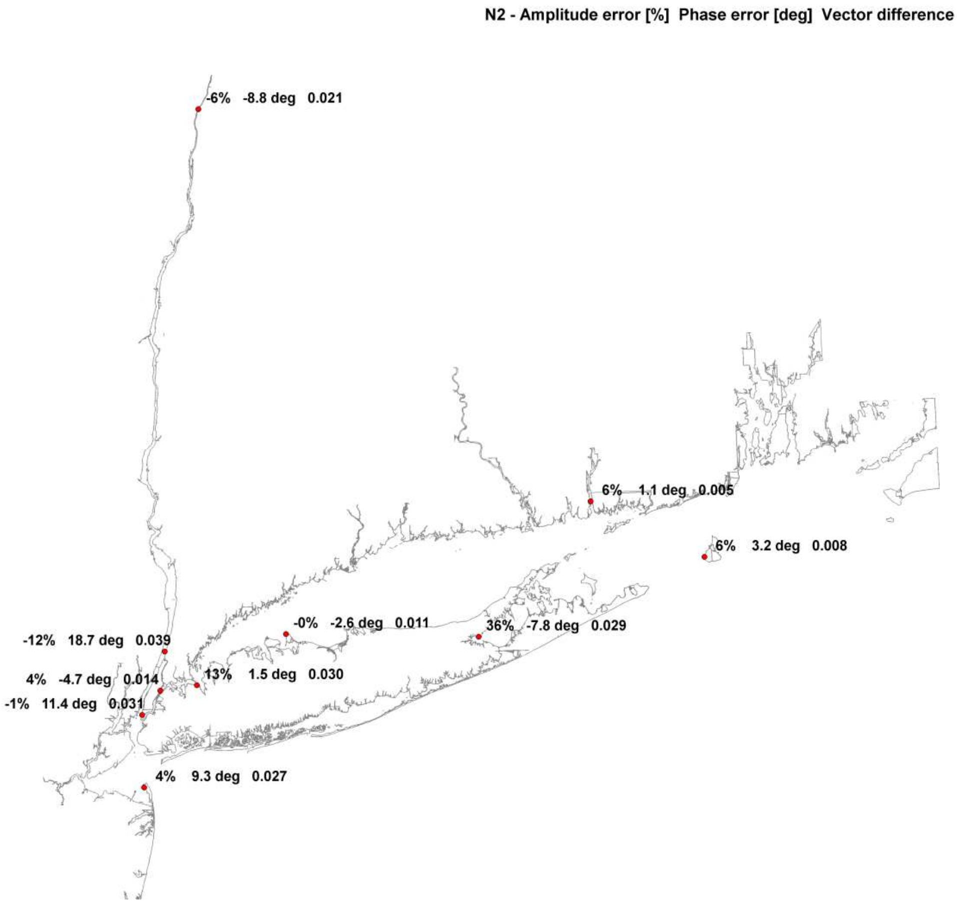

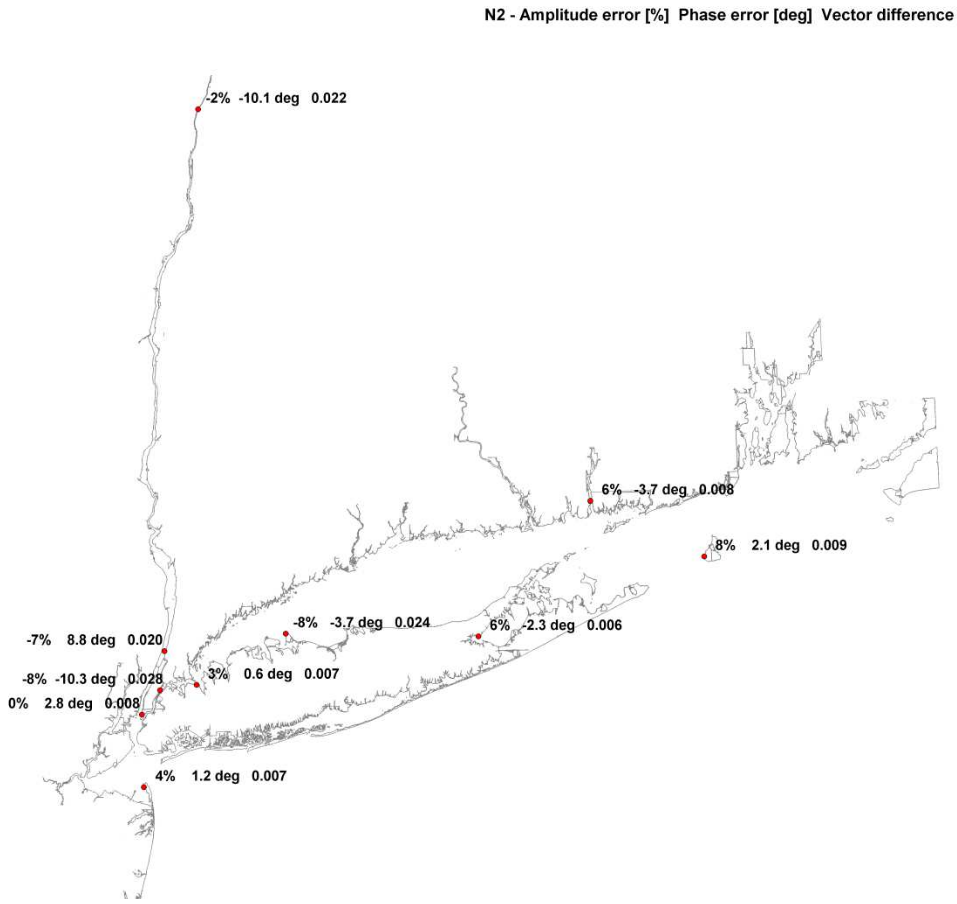

Figure A.15– Differences between observed and computed data of N2 constituent for first simulation [above] and final simulation [beneath]. A positive phase or amplitude error means higher computed than simulated data.

S2 - Amplitude error [%] Phase error [deg] Vector difference

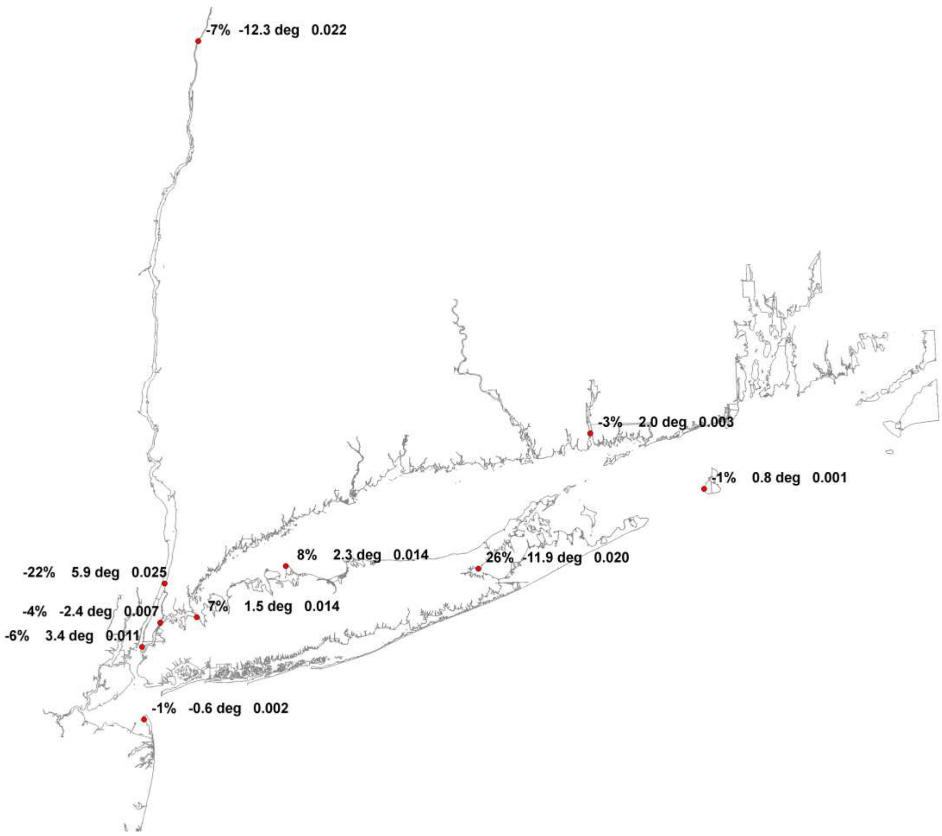

S2 - Amplitude error [%] Phase error [deg] Vector difference

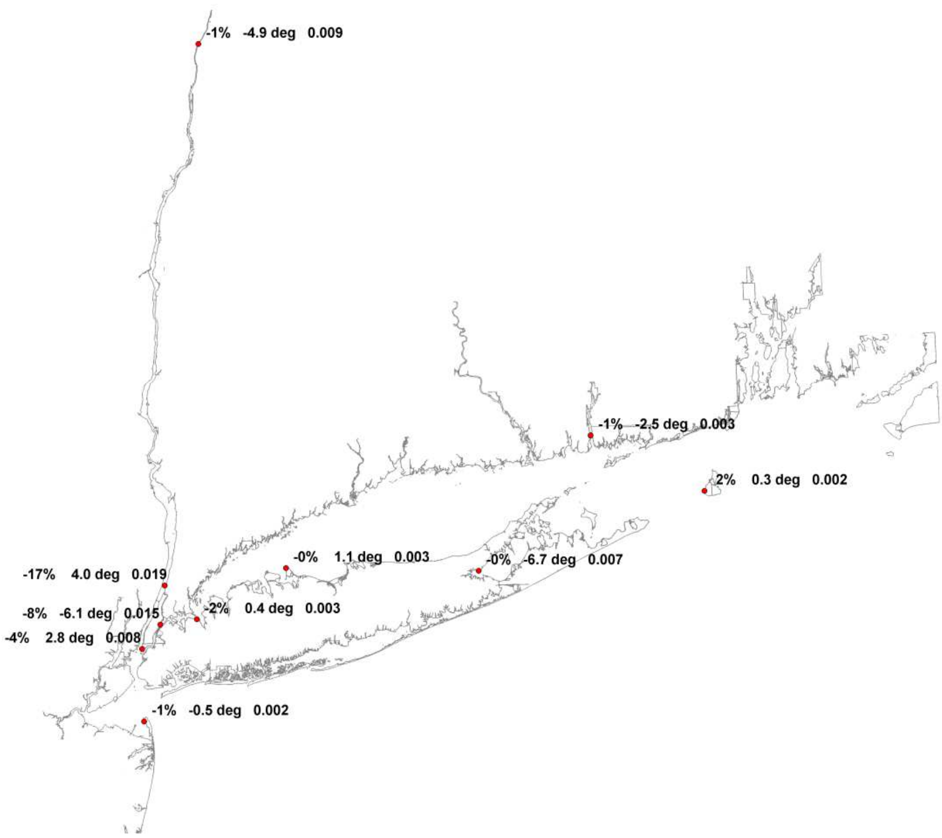

Figure A.16– Differences between observed and computed data of S2 constituent for first simulation [above] and final simulation [beneath]. A positive phase or amplitude error means higher computed than simulated data.

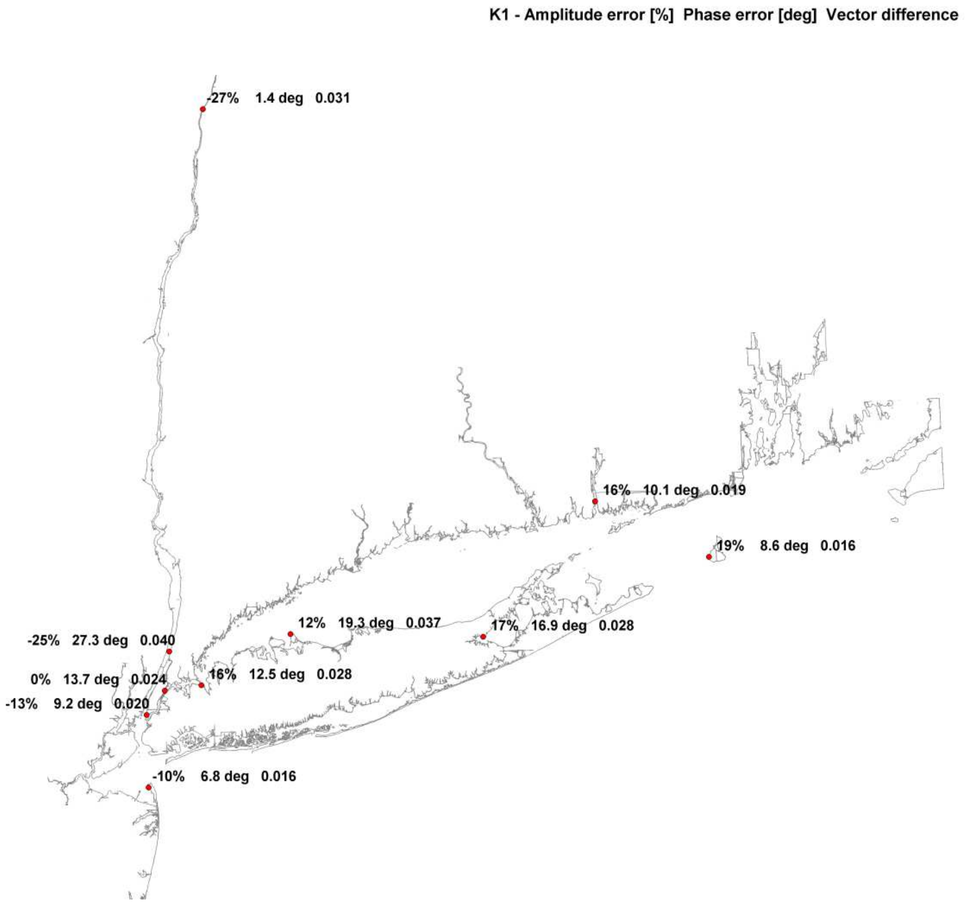

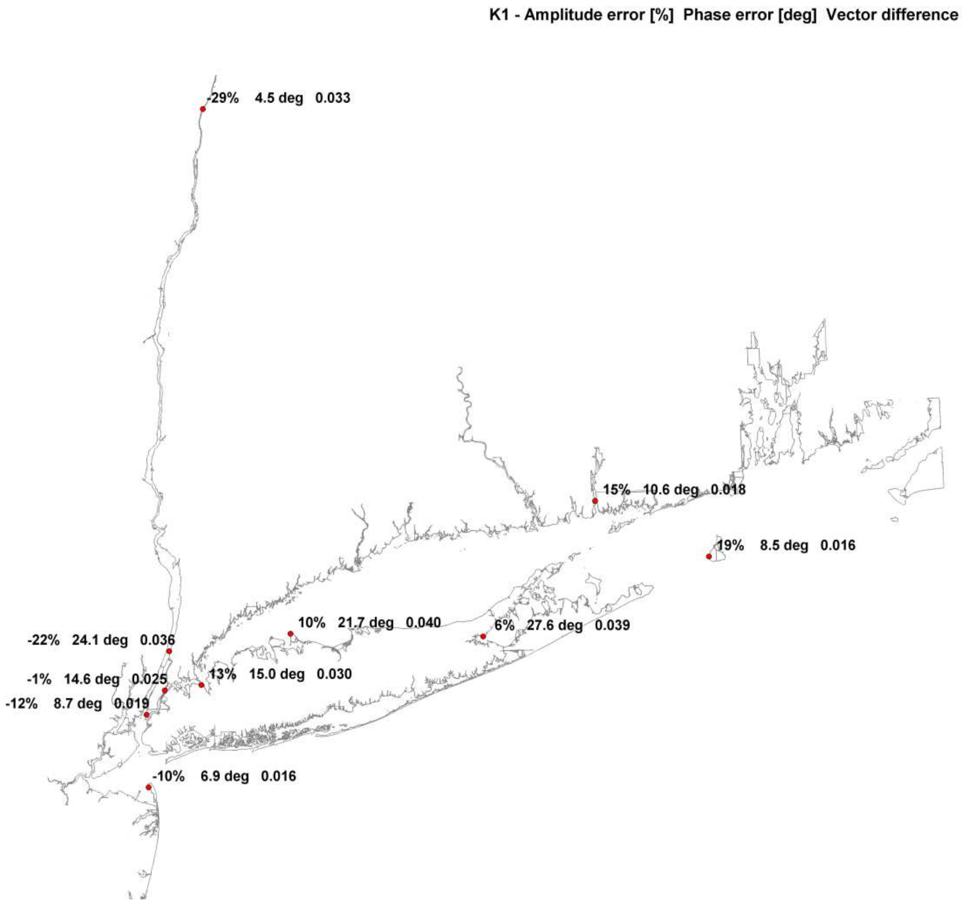

Figure A.17– Differences between observed and computed data of K1 constituent for first simulation [above] and final simulation [beneath]. A positive phase or amplitude error means higher computed than simulated data.

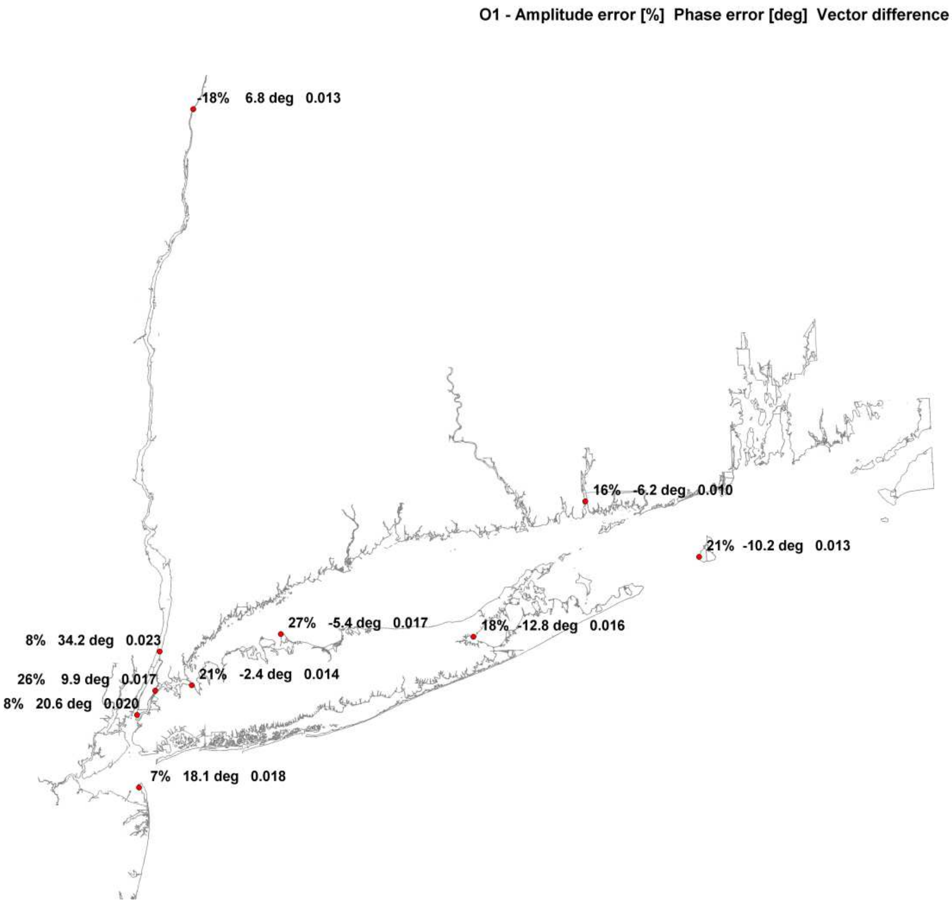

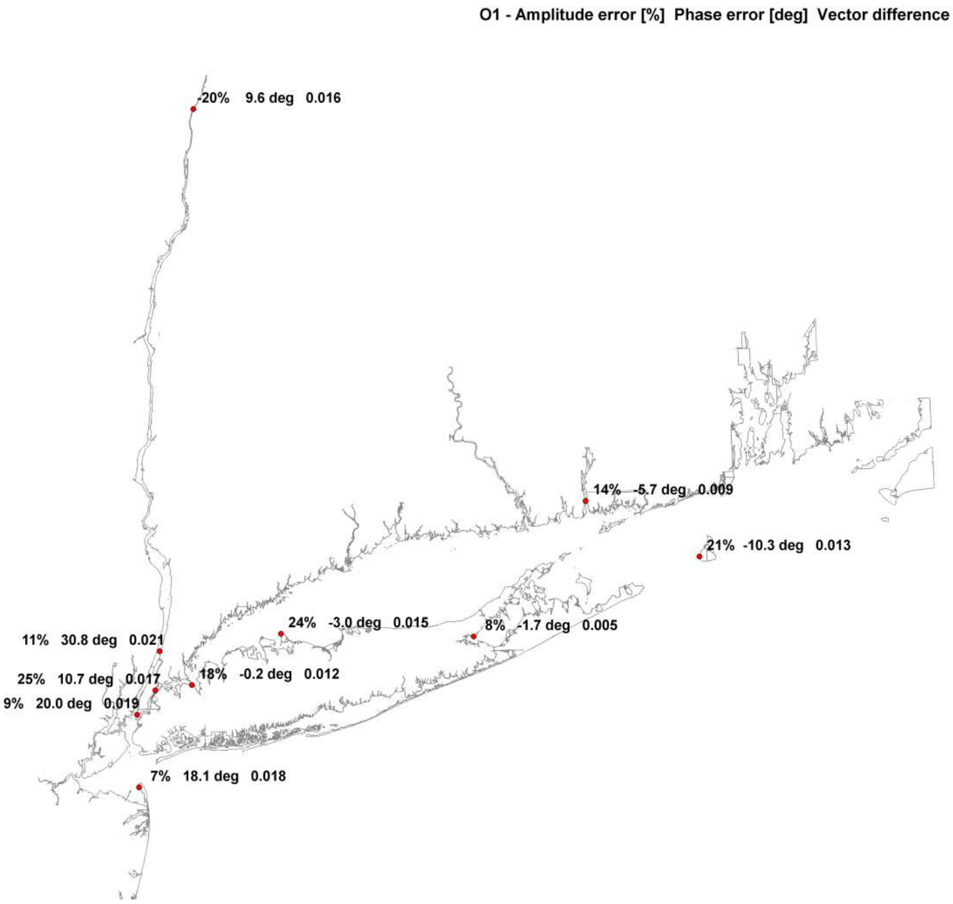

Figure A.18– Differences between observed and computed data of O1 constituent for first simulation [above] and final simulation [beneath]. A positive phase or amplitude error means higher computed than simulated data.

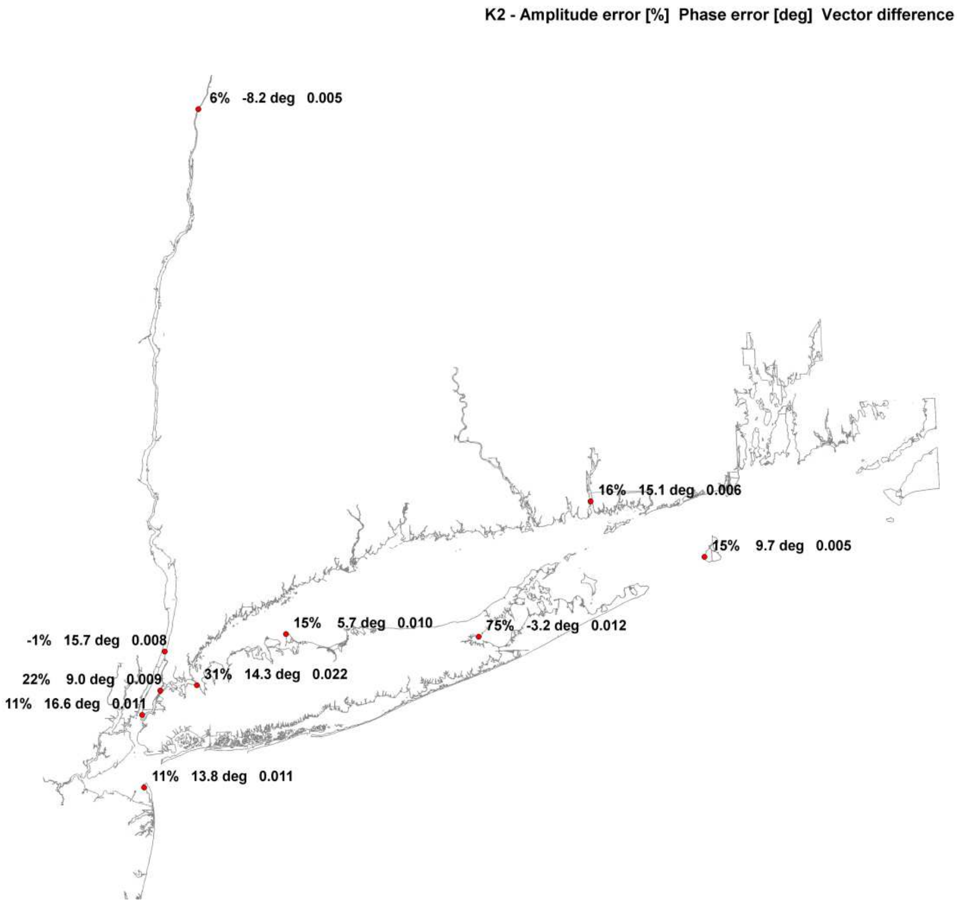

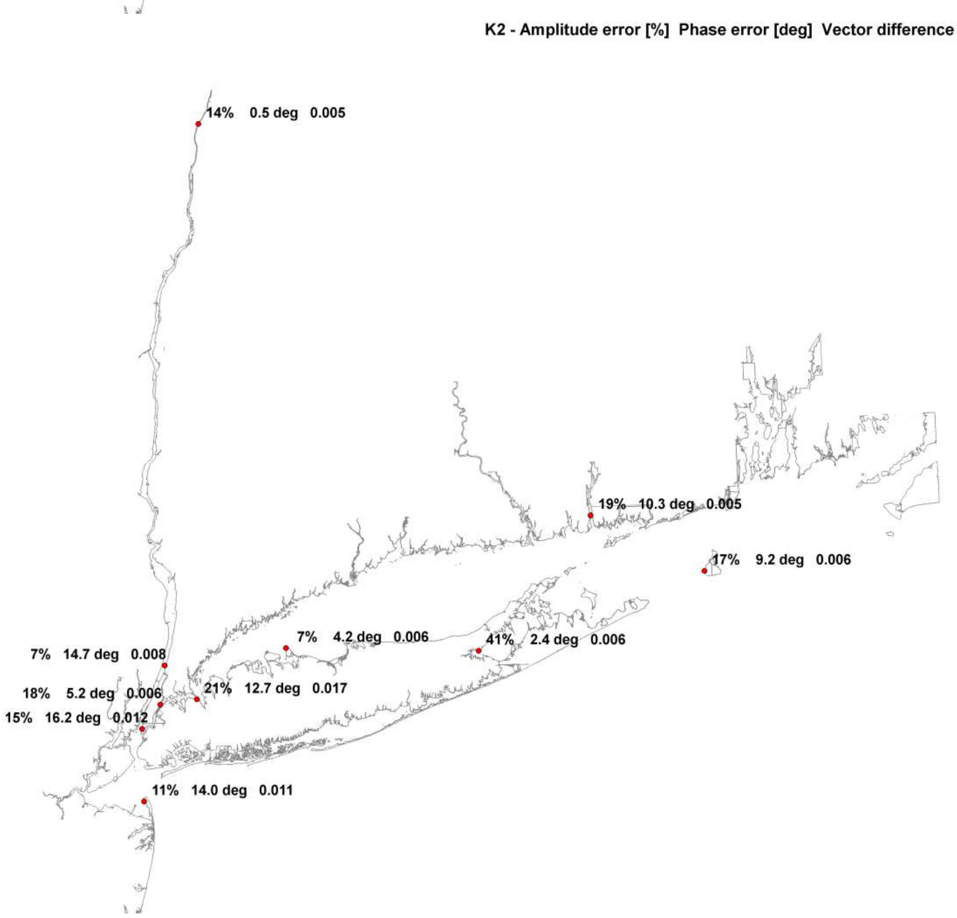

Figure A.19– Differences between observed and computed data of K2 constituent for first simulation [above] and final simulation [beneath]. A positive phase or amplitude error means higher computed than simulated data.

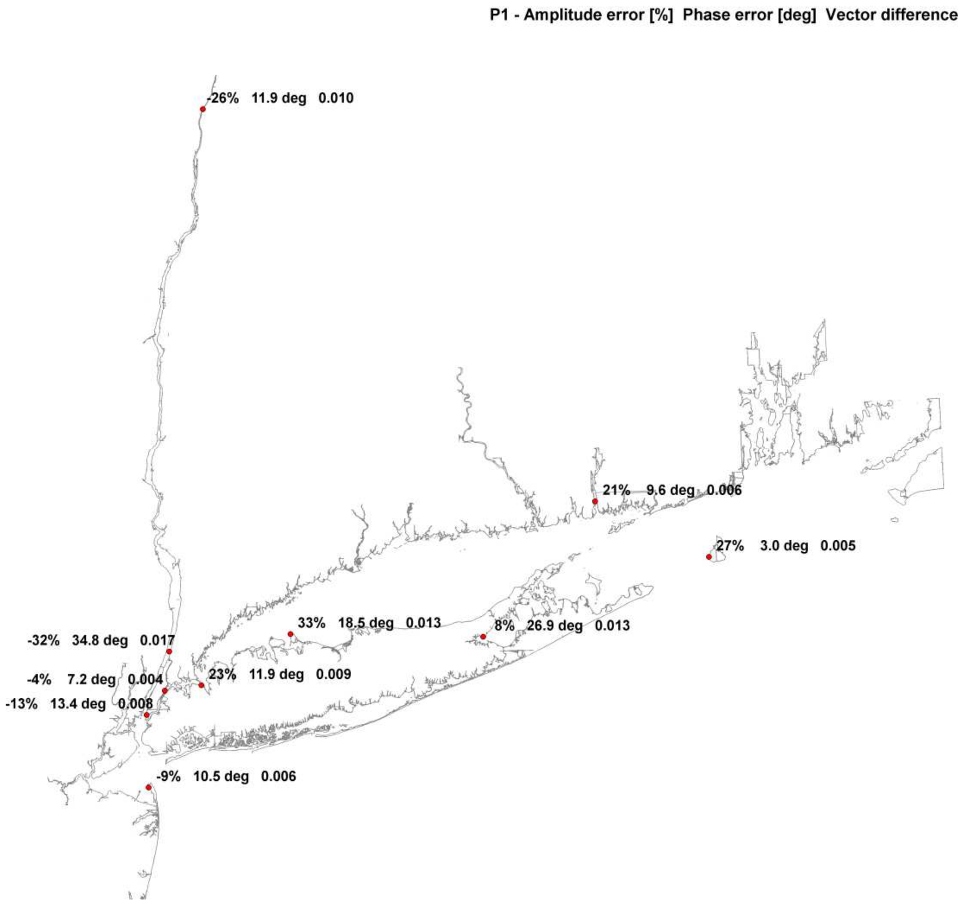

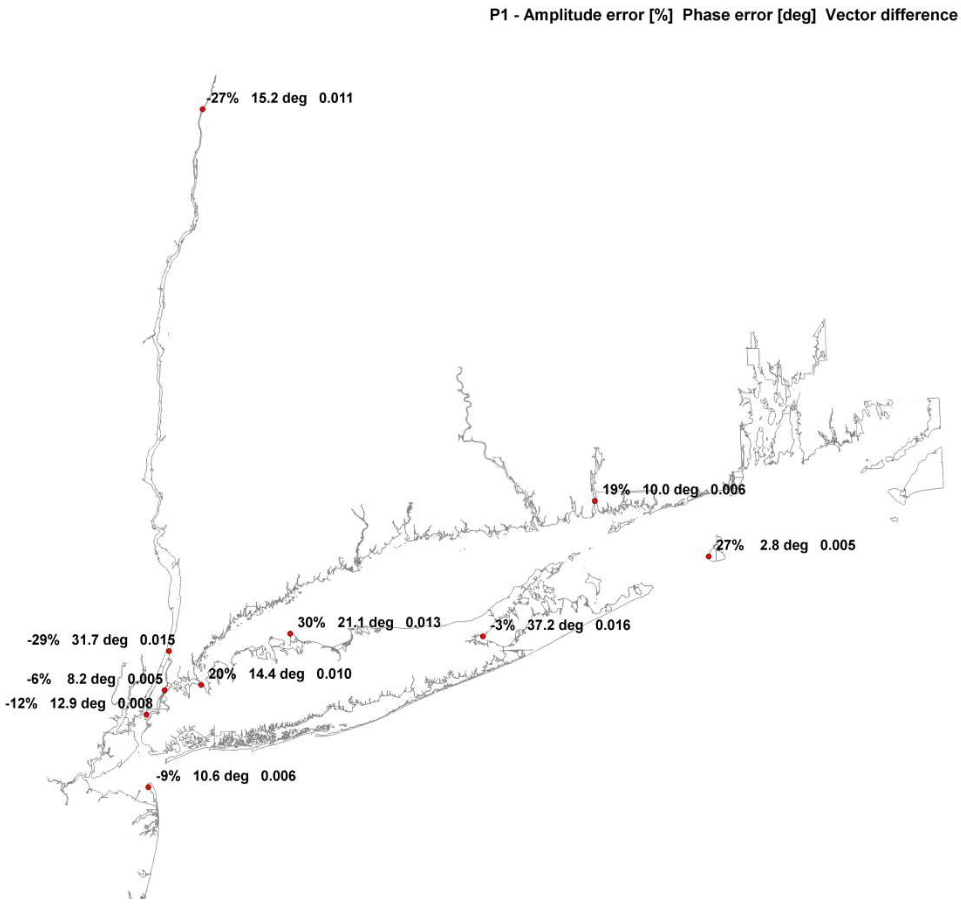

Figure A.20– Differences between observed and computed data of P1 constituent for first simulation [above] and final simulation [beneath]. A positive phase or amplitude error means higher computed than simulated data.

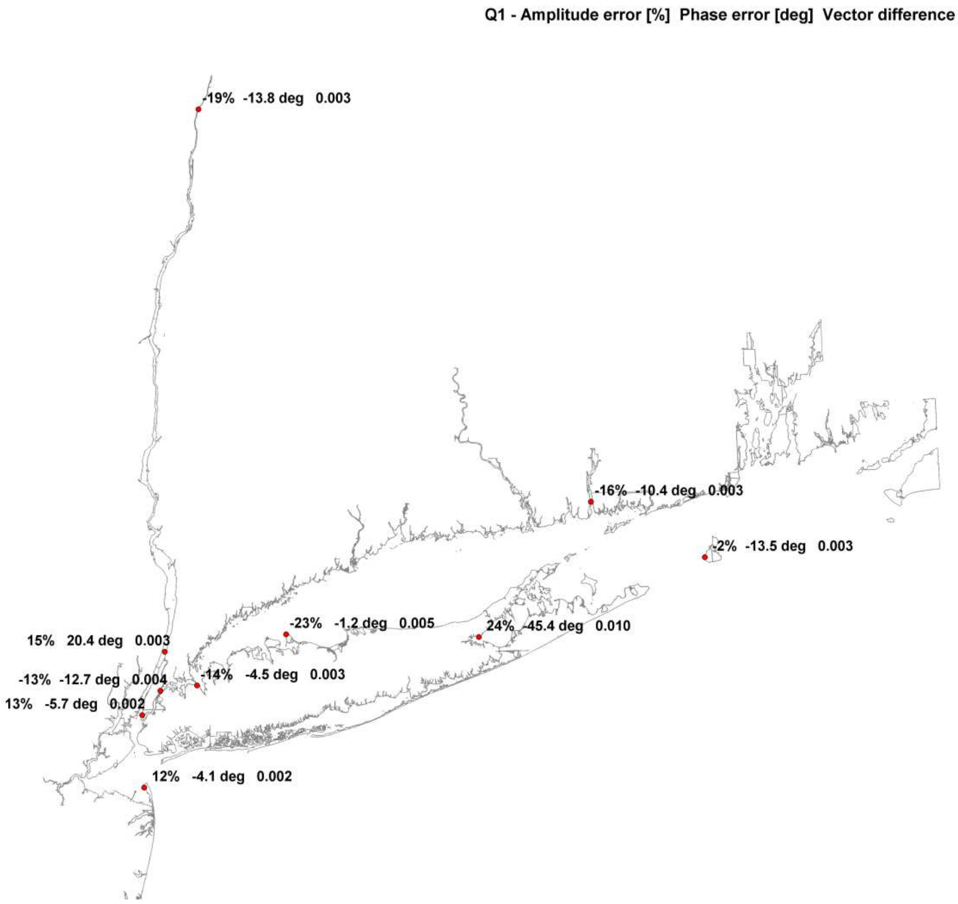

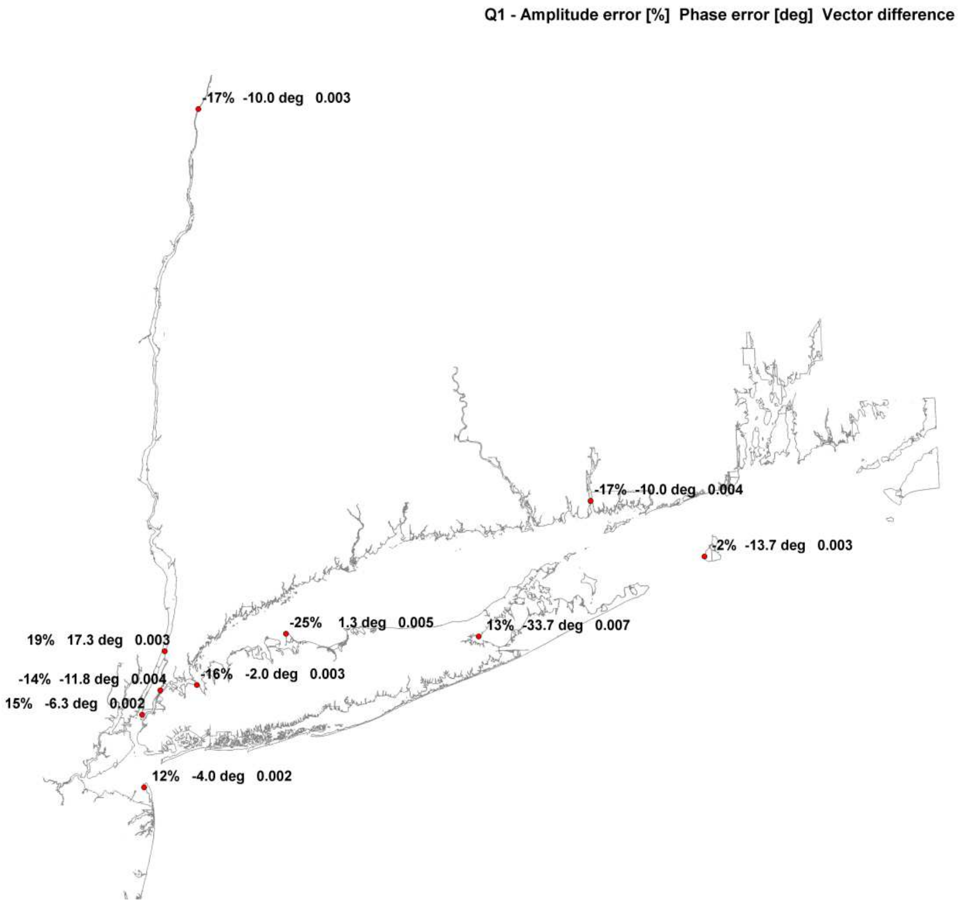

Figure A.21– Differences between observed and computed data of Q1 constituent for first simulation [above] and final simulation [beneath]. A positive phase or amplitude error means higher computed than simulated data.

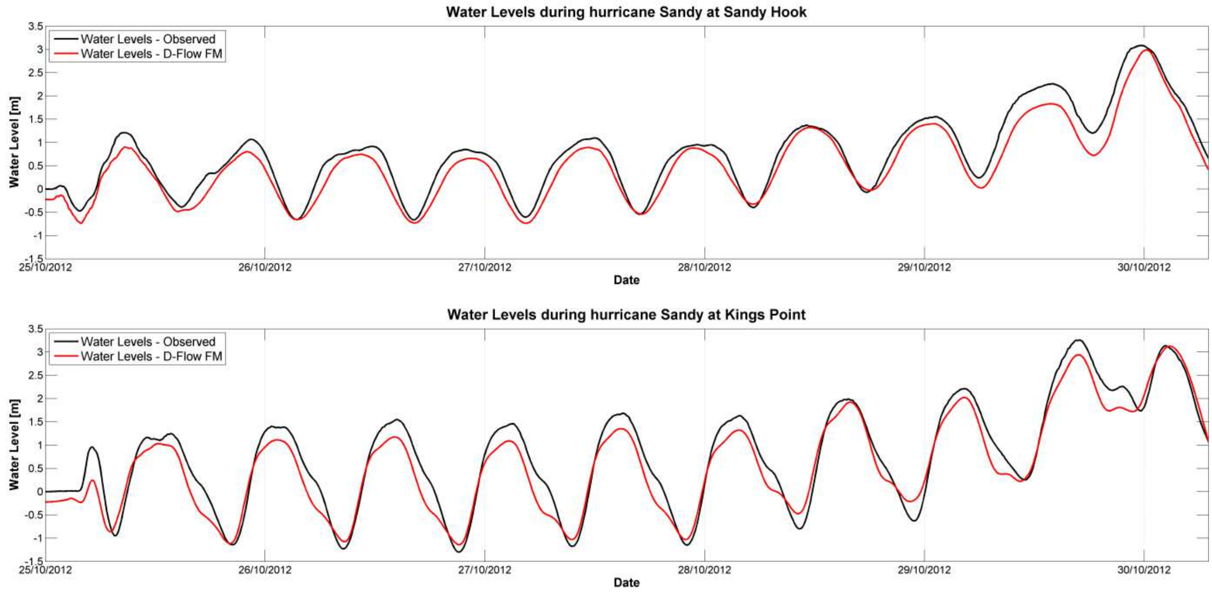

Figure A.22– Water levels during hurricane Sandy, both observed and computed, at Sandy Hook [above] and Kings Point [beneath].

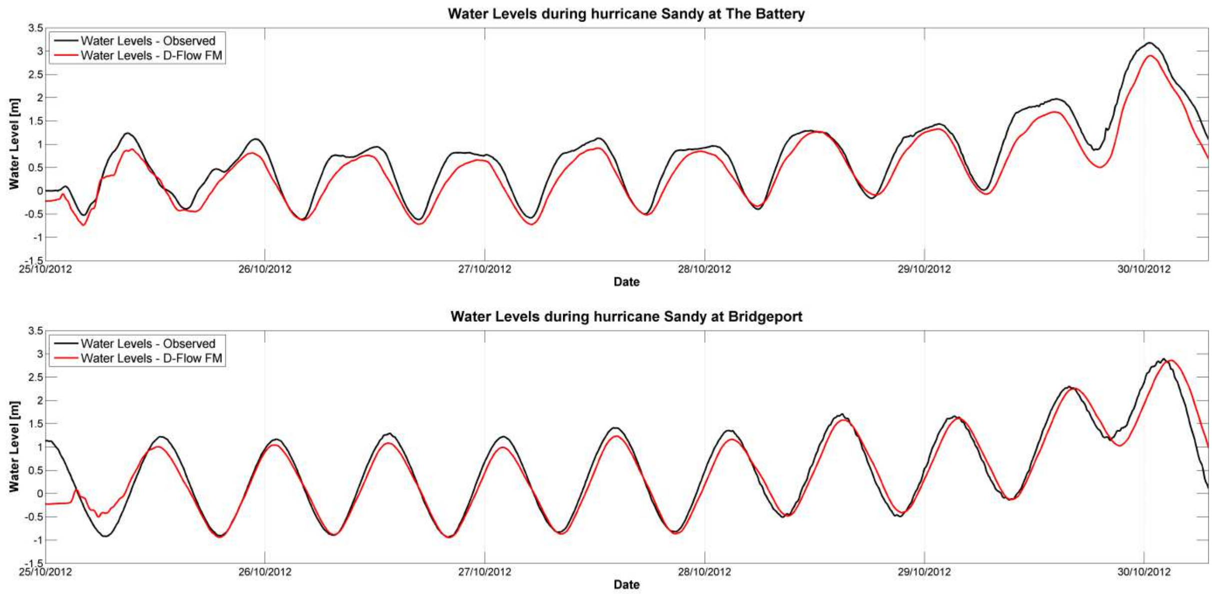

Figure A.23– Water levels during hurricane Sandy, both observed and computed, at The Battery [above] and Bridgeport [beneath].

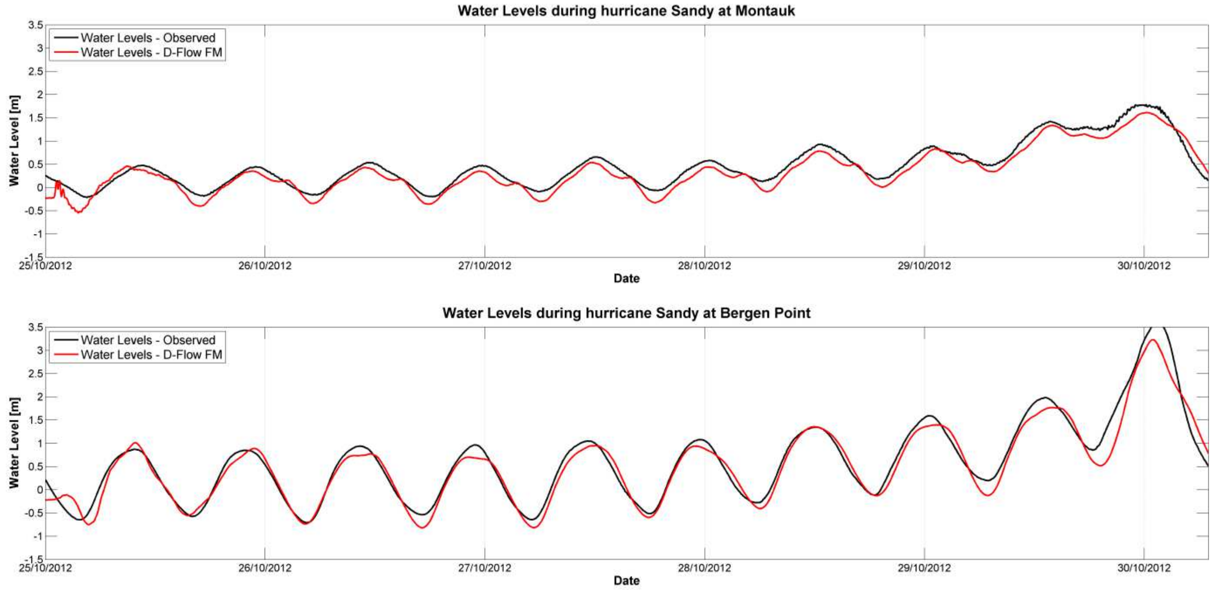

Figure A.24– Water levels during hurricane Sandy, both observed and computed, at Montauk [above] and Bergen Point [beneath].

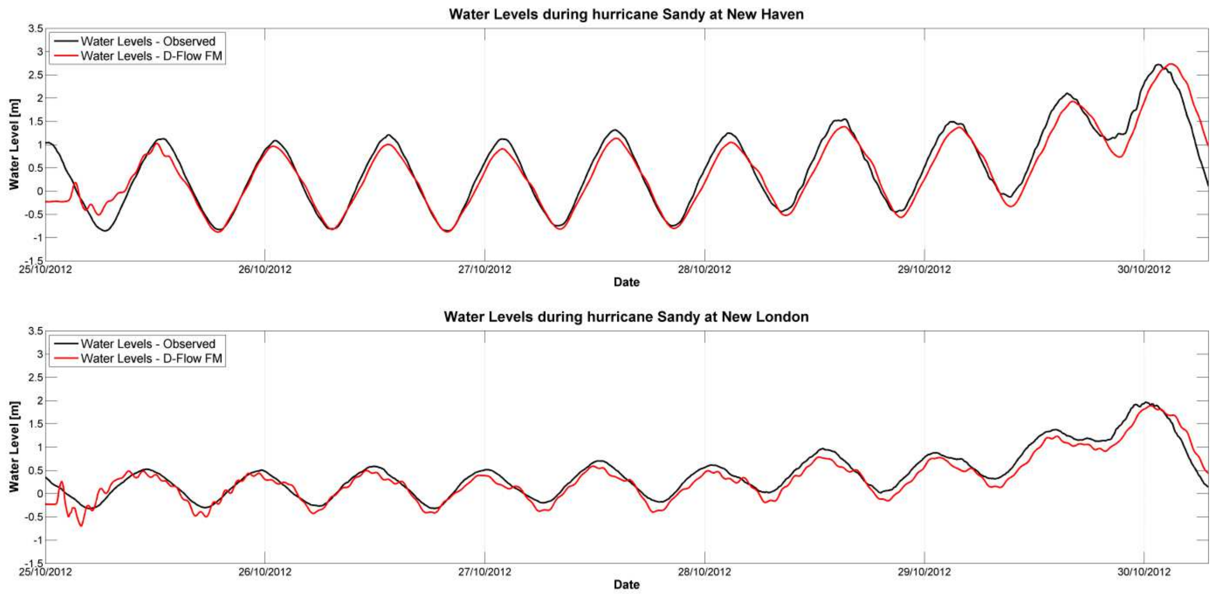

Figure A.25– Water levels during hurricane Sandy, both observed and computed, at New Haven [above] and New London [beneath].