8.3 - A FOURIER-BASED POST-PROCESSING SCHEME FOR IMPROVING NWS EXTRA-TROPICAL STORM SURGE GUIDANCE

Alexander Hrabski[1*], Arthur Taylor[2†] and Huiqing Liu[2]
1. University of Michigan, Department of Naval Architecture and Marine Engineering, Ann Arbor, MI
2. NOAA / OSTI / Meteorological Development Laboratory, Silver Spring, MD
(*) Presented at the 103rd AMS Annual Meeting, Denver CO, January 10, 2023

1. INTRODUCTION

The Extra-Tropical Storm Surge Model (ETSS) was developed by the National Weather Service's (NWS) Meteorological Development Lab (MDL) to provide storm surge guidance for coastal regions (Kim et al. 1996). In its present day form, ETSS offers 102 hours of both gridded and station-based Total Water Level (TWL) predictions due to storm surge and tide. Many of the included stations are equipped with observations, and the incorporation of station-based observations for the purpose of TWL bias-correction has led to significant error reduction over previous versions of the model (Rios-Berrios et al. 2010; Schuster et al. 2015). These gains are in part realized by incorporating these station-based observations into a post-processing step for station-based guidance.

The operational station-based post-processor assesses the error of the model TWL against observations over a 5-day time history. The average 5-day error is then used to correct TWL predictions. While this procedure does an excellent job at correcting for slow fluctuations in mean water level, it was not designed to mitigate time-varying errors. However, a closer look at ~50 stations (primarily along the east coast) reveals a systematic, time varying error in which oscillations at tidal frequencies dominate the model error. These erroneous oscillations in model guidance can be in excess of 1 foot peak-to-trough amplitude, and they can lead to errors in both the magnitude and arrival time of peak TWL during surge events.

In this work, we develop a Fourier-based post-processing scheme that detects and mitigates these erroneous oscillations in ETSS TWL guidance. Our novel post-processor first uses a Fourier-based signal-to-noise indicator to determine the presence of erroneous oscillations at tidal frequencies. When present, the post-processor generates a filter that connects the tidal signal to oscillations in model error from the recent past. The filter is then used to mitigate the erroneous oscillations in TWL guidance. The skill of this new post-processing methodology is demonstrated via a statistical analysis of historical data, showing significant improvements over the operational post-processor at affected stations.

The remainder of this document is organized as follows. First, we will detail the relevant aspects of ETSS and the operational post-processing scheme. Next, we will demonstrate the erroneous oscillations and present the Fourier-based post-processing methodology. The subsequent section will present the statistical analysis of post-processor performance, and we will conclude with comments on the source of the erroneous oscillations and recommendations for future work.

2. STATION-BASED GUIDANCE AND THE OPERATIONAL ETSS POST-PROCESSOR

The Extra-Tropical Storm Surge model (ETSS) incorporates a number of quantities that determine Total Water Level (TWL). For station-based guidance, these critical quantities are: tides and surge, and (where available) observations. Figure 1 depicts station-based guidance for tides and surge at Pilot's Station East, Louisiana. We will refer to both the hindcast and forecast regions of this station-based guidance, corresponding to times before and after (respectively) the time and date for which the guidance was prepared. In the later sections, the "forecast" region will be more loosely interpreted.

For station-based guidance, the tides are modeled via 37 tidal constituents that are specific to each station (Schuster et al. 2015). These 37 cosine functions are summed together to create a model for the way TWL changes at the station due to the tides. Figure 1 depicts an example of a tidal signal $T(t)$ constructed via the 37 constituents for Pilot Station East.

(†) Corresponding author: Arthur Taylor, 1325 East-West Hwy, Silver Spring, MD 20910-3280; e-mail: Arthur.Taylor@noaa.gov

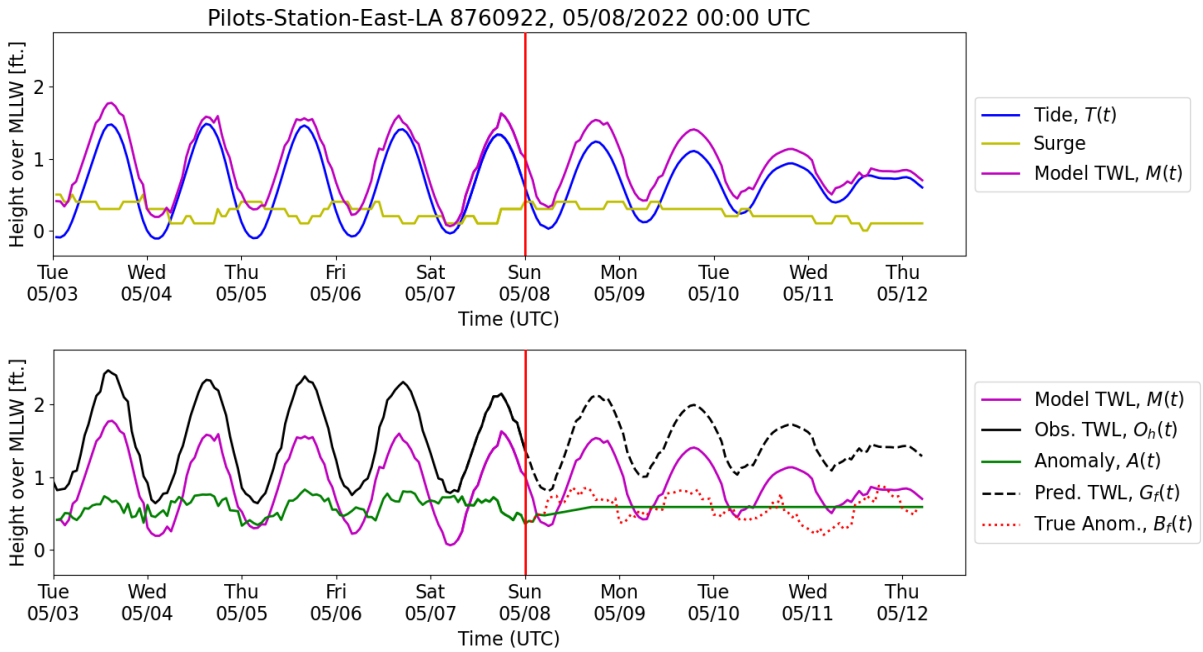

Figure 1. (top) An example of station-based guidance with no observations, including tide, surge, and their sum, the model Total Water Level (TWL). All quantities are with respect to the Mean Lower Low Water datum (MLLW). (bottom) An example of station-based guidance including observations, bias correction, and validation (true anomaly). The vertical red lines indicate the time for which the guidance was prepared. All data to the left of the red line will be considered hindcast (subscript h), and all data to the right of the red line will be considered forecast (subscript f). A quantity without a subscript refers to both the hindcast and forecast components of that quantity.

Figure 1 also contains predictions for surge. ETSS uses the Sea, Lake and Overland Surges from Hurricanes model (SLOSH) driven by winds from the NWS Global Forecasting System model (GFS) to generate the surge signal. The tide and surge signals are then summed together to create the Model TWL, which serves as guidance when observations are unavailable.

When a station is equipped with observations, additional guidance is available, as depicted in the second panel of Figure 1. This begins with the observations over the past 5 days, which may not agree with the model predictions over the past 5 days. Let the hindcast model TWL be $M_h(t)$ and the observations be $O_h(t)$. Here we use a subscript h to denote that a quantity is taken from the hindcast region of figure 2. The model error can be captured by the hindcast anomaly $A_h(t)$, which is simply the difference between observation and model predictions:

$$A_h(t) = O_h(t) - M_h(t). \quad (1)$$

In figure 2, the anomaly over the previous 5 days stays relatively constant at close to +6 inches. It is at this stage that the operational post-processor is used. An assumption is now made that the error over the next 102 hours will be the same as the mean error over the last 5 days. If we denote a quantity that includes the forecast region via a subscript f, then our assumption can be expressed as $A_f(t) = AVG(A_h(t))$. This amounts to projecting the average value of hindcast anomaly into the future as a constant value, as depicted in figure 2. The small variation in $A_f(t)$ at the beginning of the forecast has two causes: First, there is a short delay in ETSS dissemination while the GFS and ETSS models run, during which observations are available to compute the true anomaly. After this time period, a relaxation term is included that prevents jump-discontinuities in the TWL guidance. This forecasted anomaly is now added back onto the model TWL forecast to create the final TWL guidance, $G_f(t)$. Put mathematically,

$$G_f(t) = M_f(t) + A_f(t). \quad (2)$$

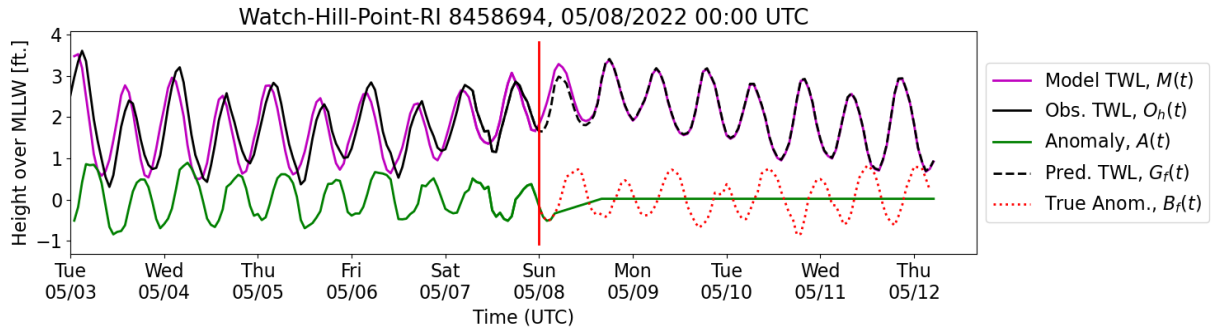

Figure 2. An example of station-based guidance with strong oscillations in the anomaly at tidal frequencies.

The error in the predicted TWL can be quantified by measuring the true anomaly of the model TWL with respect to future observations. This allows us to define $B_f(t)$, the true anomaly, which is

$$B_f(t) = O_f(t) - M_f(t), \quad (3)$$

Where $O_f(t)$ is taken from future observation data (not to be confused with a forecast for observations). It is straightforward to show that if $A_f(t) = B_f(t)$, then the error in the predicted TWL is zero (i.e., the guidance is perfect). Also, the difference between the forecasted anomaly and the true anomaly gives the error in guidance. In figure 1, it can be seen that the forecasted anomaly and the true anomaly differ at most by only a few inches. This indicates that the error in the predicted TWL is small throughout the 102 hours of forecast. We are now poised to discuss the problem of oscillating error.

3. OSCILLATING ANOMALIES AND FOURIER-BASED POST-PROCESSING

We have just shown that the operational post-processor is capable of mitigating minimally fluctuating errors by forecasting anomaly as the 5-day mean of $A_h(t)$. Not all anomalies are constant, however. Out of the 277 stations in the system with observations, we have identified approximately 50 stations where oscillations dominate the anomaly. One such station is Watch Hill Point, Rhode Island, which is shown in figure 2.

At this station, we see that there is a clear, large-amplitude oscillation in the hindcast anomaly that cannot be extrapolated into the forecast anomaly by the operational post-processor. The oscillations have a peak-to-trough amplitude of over 1 foot and have clear tidal frequency content. An equivalent interpretation of this error is that there is a near constant phase-shift between $O_h(t)$ and $M_h(t)$. This type of error can manifest as particularly problematic whenever the peak of a surge event coincides with high tide: a phase shift in the model could lead to error in the time (and magnitude) of highest water.

In order to forecast oscillations in the anomaly, it will be useful to connect the amplitude, frequency, and phase of the oscillations to quantities that we can reliably forecast. Because these oscillations occur at tidal frequencies, we choose the hindcast tide signal $T_h(t)$ and forecast tide signal $T_f(t)$ as references. The general strategy is as follows:

1. Determine the amplitude of oscillations in $A_h(t)$ at each tidal frequency.
2. Find the relative phase between $A_h(t)$ and $T_h(t)$ at each tidal frequency.
3. Use the $T_f(t)$ (which can be computed via tidal constituents) to construct a forecast anomaly $A_f(t)$ with the amplitudes from step (1) and relative phases from step (2). Perform this step for each tidal frequency.

We implement this strategy via the Discrete Fourier Transform (DFT). This requires that we take the DFT of the hindcast tide signal, which itself is constructed from 37 sinusoids. However, the 37 tidal constituents have very precise (and sometimes closely spaced) frequencies. To produce an accurate reconstruction of the tide signal via DFT, we require a longer time history than 5 days. We find that a 14 day window offers excellent frequency resolution, in large part because many tidal constituents are close-to-periodic over 14 day intervals. On the other hand, a 14 day window is still short enough as not to include "old" oscillation data that might not be representative of the present.

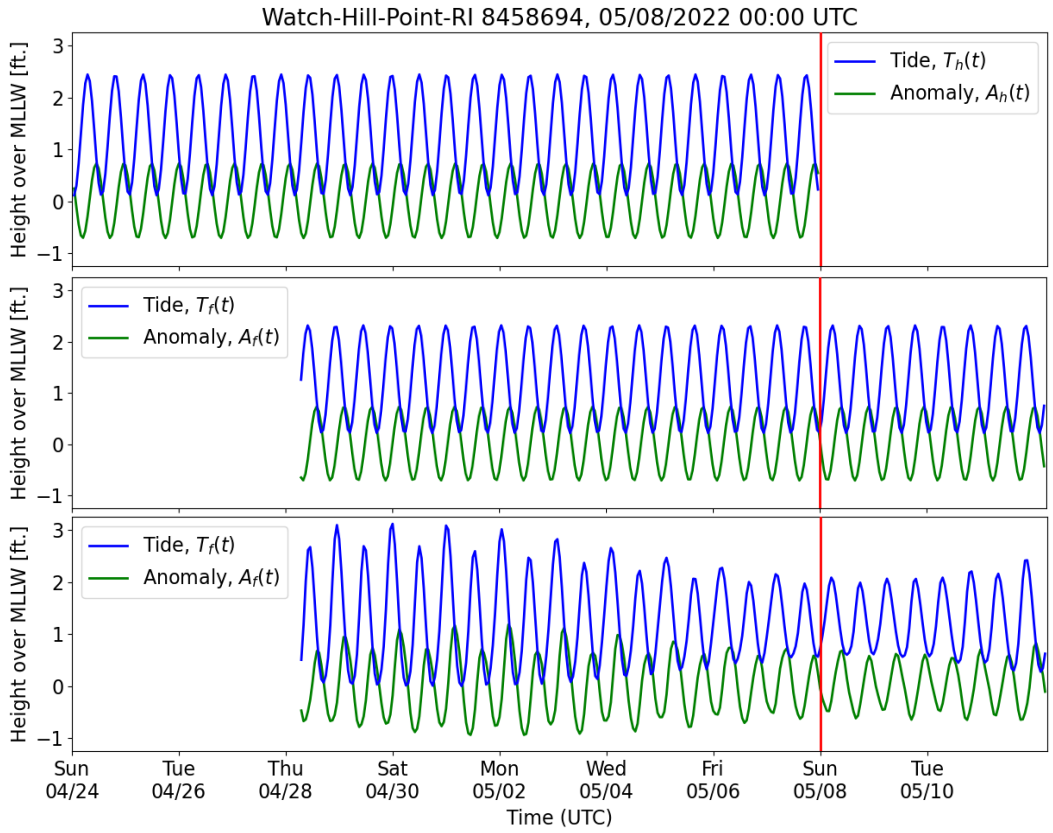

Figure 3. (top) The primary frequency component of the hindcast tides $T_h(t)$ and the corresponding frequency component of the hindcast anomaly $A_h(t)$. (center) The primary frequency component of the forecasted tide signal, propagated via the tidal constituents. The forecasted anomaly is constructed via the amplitude of $A_h(t)$ and relative phase of $A_h(t)$ with respect to $T_h(t)$, applied to $T_f(t)$. (bottom) The procedure described by the top two panels, applied to all tidal frequency components and summed together.

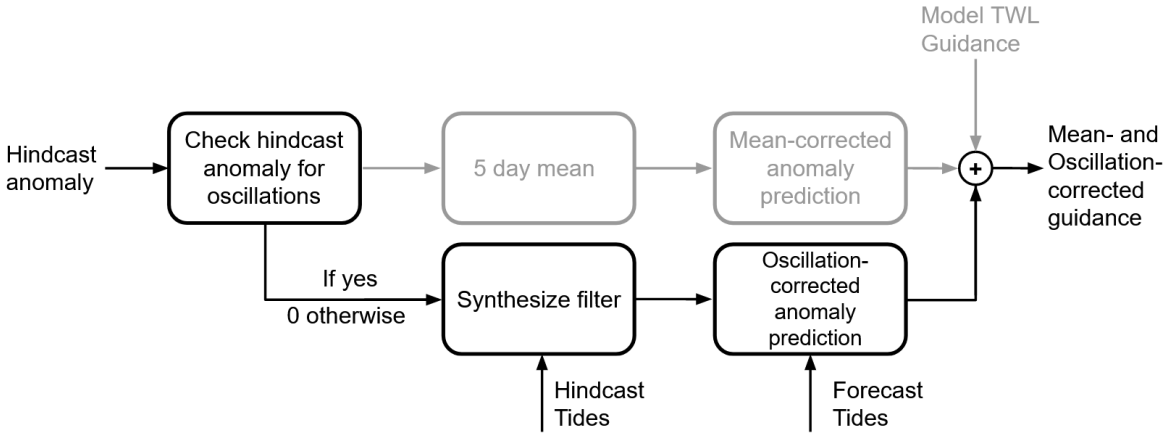

Figure 4. A diagram of the post-processing scheme. Grey indicates the existing operational post-processor, whereas black indicates the post-processing scheme developed in this work.

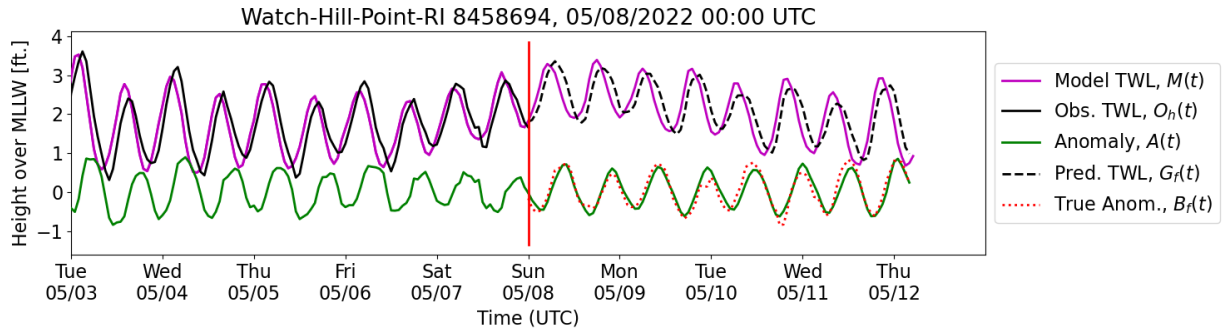

Figure 5. The Watch Hill Point example with the Fourier-based post-processing methodology. The signal-to-noise ratio (S2N) is 15.5, well exceeding the threshold value of 1.5.

Additionally, to keep the algorithm relatively simple, we use the same time window size for hindcast and "forecast" data. Therefore, the forecast window will be the hindcast window shifted by 102 hours. Note that this means there will be a slight abuse of our notation: $T_f(t)$ and $A_f(t)$ include hindcast data when they refer to data taken from a 14 day window.

These concepts, along with the rest of the forecasting strategy, is illustrated in figure 3. A technical description of the filtering process is provided in Appendix A, alongside the supporting mathematics. Critically, this procedure assumes that tidal oscillations are present in the anomaly. If this procedure is applied to a station where these erroneous oscillations are not present, it has the potential to do more harm than good. Therefore, it is important to develop a method to detect when this post-processor should be used. One such method is discussed in the next section.

4. DETECTING OSCILLATIONS IN THE HINDCAST ANOMALY

Because we have already set up a Fourier-based filtering strategy, we rely on the Fourier transform to aid us in detecting the erroneous oscillations. In particular, we use a Fourier-based signal-to-noise ratio (S2N) in the hindcast anomaly as an indicator. We define "signal" to be any frequency within 2 degrees per hour (deg/hr) of tidal constituent frequency, and any other frequency as "noise". When S2N exceeds 1.5, we activate the new post-processing step. A technical description of the S2N indicator (with suggestions) is provided in Appendix B. We will now discuss how this new post-processing strategy can be applied to the existing ETSS model, and then we will demonstrate its use.

5. APPLICATION

This method is designed to be run alongside the operational post-processing scheme, as depicted in figure 4. First, the S2N-based indicator is used to assess whether clear tidal oscillations are present in the $A_h(t)$ (over the 14 day window). If present, a filter is synthesized from the $T_h(t)$ (again, over a 14 day window) and $A_h(t)$, and that filter is used to generate $A_f(t)$. By nature of the DFT, the oscillating anomaly forecast has zero-mean over the 14 day window. Therefore, the 5-day mean anomaly correction is still necessary. The output of the filter is added to the 5-day mean anomaly to generate combined oscillation- and mean-corrected anomaly forecast. In the event that no oscillations are detected in the hindcast anomaly, only the 5 day mean anomaly correction is applied.

The resulting anomaly forecast and predicted TWL for the Watch Hill Point, Rhode Island example are depicted in figure 5. When compared to figure 2, an improvement in the agreement between the forecasted anomaly and true anomaly is evident. Equivalently, the phase shift between the model TWL and observations in the hindcast is preserved in the forecast. Now that we have demonstrated that this Fourier-based method has the potential to forecast the oscillations in the anomaly, we perform a statistical analysis using historical data to quantify the performance of the method across the ETSS system.

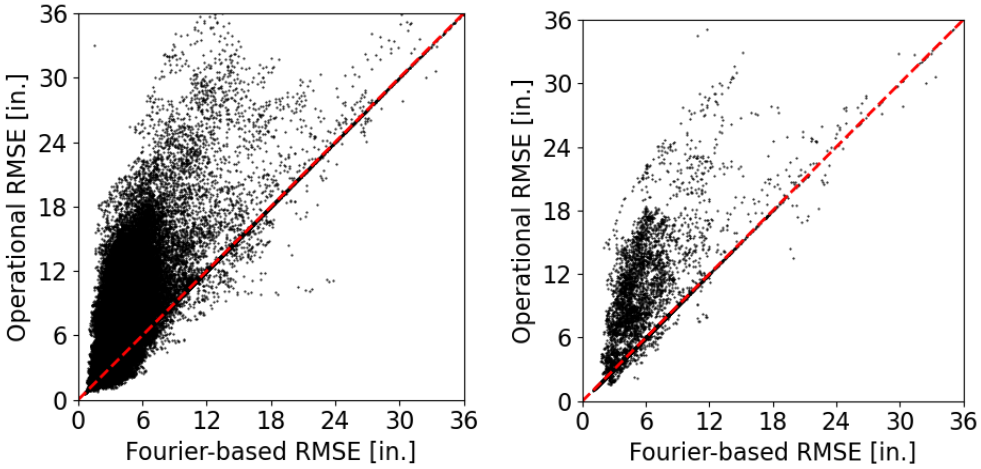

Figure 6. (left) The error in the operational post-processor vs. the Fourier-based post-processor over one year, in inches. Each point represents the error of a single forecast. The plot contains all forecasts whose data met quality standards from all stations. Points above the red dashed line indicate forecasts where the Fourier-based post-processor out-performed the operational post-processor. (right) the same scatter plot constructed only with forecasts where surge guidance exceeds 3 feet.

6. PERFORMANCE ANALYSIS

The performance of the post-processor is assessed over a one year period, beginning from 1 May 2021 and ending on 1 May 2022. The analysis includes all forecasts generated by ETSS in this time period at all stations with observations. We note, however, that no forecast was included that did not meet our standards for observation quality, which include criteria for sensor uptime over the 14 day window. A full description of the data cleaning process (alongside recommendations) is included in Appendix C.

We measure the root-mean-square error (RMSE) of guidance under both the operational post-processor as well as the Fourier-based post-processing scheme, where RMSE is defined as

$$\text{RMSE} = \sqrt{\frac{1}{T}\int_5^T (B_f(t) - A_f(t))^2 \, dt}. \qquad (4)$$

With $t = 0$ taken to be the start of the (true) forecast interval, we evaluate the RMSE beginning at $t = 5$ hours in order to omit the short time period where the forecasted anomaly is computed from observations by the ETSS system. We end the evaluation at $T = 102$ hours, the end of the forecast interval.

We begin the discussion of our statistical analysis with a summary of our findings. All averages in this section refer to the mean change of RMSE. Of the 277 stations in ETSS with observations, we find that 232 stations are unaffected by the Fourier-based post-processor. By “unaffected”, we mean that the change in RMSE at a given station was on average under 1 inch (over the course of the year). The remaining 45 stations all saw an average decrease in RMSE of an inch or greater. *Critically, this means that no stations saw a significant increase in RMSE on average.* The 45 stations that saw improvements are consistent with the ~50 stations that have demonstrated consistent tidal oscillations in their anomaly. Among the stations that saw improvements to RMSE on average, the average improvement was 4.22 inches, with a maximum average improvement of 12.5 inches.

The full results of the error analysis are aggregated as a scatter plot in figure 6, where the RMSE under both the operational and the Fourier-based post-processors are depicted on a per-forecast basis. We draw special attention to those forecasts that form a steep slope on the left-hand side of the plot. Along this slope, error under the Fourier-based post-processor remained small, despite being potentially quite large under the operational post-processor. On a per-forecast basis, the discussed performance improvements are illustrated in figure 7, which shows the change in RMSE for all

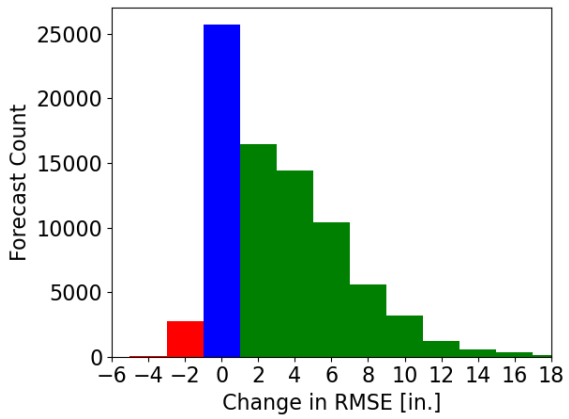

Figure 7. The change in RMSE under the Fourier-based post-processor, where positive (green) values represent a reduction of RMSE over the operational post-processor. Forecasts where guidance was unaffected (blue) and harmed (red) are also indicated.

forecasts where the Fourier-based post-processor modified guidance (i.e., when S2N > 1.5). In the majority of the instances where the Fourier-based post-processor modified guidance, RMSE was reduced, and in some cases the improvement exceeded a foot. On the other hand, we see very few instances where the RMSE of a forecast was increased by a significant amount (>1 inch). Thus, not only are there no stations that are meaningfully harmed by the Fourier-based post-processor on average, but in fact there are very few forecasts at all where RMSE is meaningfully increased by the Fourier-based post-processor.

7. CONCLUSIONS AND RECOMMENDATIONS

In this work, we have developed a novel, Fourier-based post-processing methodology for station-based guidance within the National Weather Service's (NWS) Extra-Tropical Storm Surge model (ETSS). The scheme successfully mitigates a systematic, time-varying error related to the tides at nearly all affected stations with minimal harm to forecast accuracy. This was demonstrated via a statistical error analysis over a year of ETSS station data. While we have found the post-processor to perform as intended in almost all tested cases, there are a few notable limitations.

First, the 14 day window used for filter synthesis does not provide sufficient frequency resolution to accurately capture complicated tidal patterns that unfold over month-long time-scales, such as in Washington state. Thus, in these locations, the filter can at times fail to correctly forecast oscillations in the anomaly. Fortunately, the RMSE due to erroneous, systematic oscillations in these locations is relatively small at the time of this writing. Second, the filter synthesis process is sensitive to missing data and errors in observations at a given station. Thus, a robust data cleaning process would need to be implemented alongside this post-processing methodology. We have found that linear interpolation works well for situations where there are missing observations over small time windows.

Finally, we conclude with a comment on the source of these erroneous oscillations within ETSS. Tidal frequency content in the anomaly would be expected to come from error within the tide model used in ETSS, and for station-based guidance, the 37 tidal constituents make up this model. While there are potential environmental factors that can drive error in the tidal constituent model (such as changing mean water levels), the consistent, systematic error we see at some stations suggests the tidal constituents used within ETSS may benefit from re-tuning. This is further supported by the method developed in this work, which acts very similarly to a re-tuning of the tidal constituents. Given that our post-processing methodology mitigates the erroneous oscillations effectively, there is good reason to believe that re-tuning the tidal constituents may achieve a similar level of error reduction. Nonetheless, the method developed in this work is adaptive in a way that a simple re-tuning of the tidal constituents cannot be.

8. APPENDICES

8A. Technical Specifications of the Linear Filter

The filter is developed using the Discrete Fourier Transform (DFT), which we will formally defined as

$$\hat{X}(k) = \sum_{n=0}^{N-1} X(n) \times exp(\frac{-2\pi j}{N} kn), \qquad (1A)$$

where $X(n)$ is the time-domain water height at time index n, N is the number of points in time, and j is the imaginary unit. The result is $\hat{X}(k)$, the Fourier coefficient of frequency mode $k \in [0, N-1]$. For the temporal spacing of data ($\Delta t = 1$ hour) and the number of data points included in the 14 day interval ($N = 336$), the k-th resolved frequency is given by $f(k) = 1.071\,k$ degrees per hour,

resolving up to maximum frequency of $f = 180$ degrees per hour.

A linear filter models the relationship between input $X(n)$ and output $Y(n)$ via the equation $\hat{Y}(k) = \hat{H}(k)\hat{X}(k)$. As a first step, we would like to find the transfer function $\hat{H}(k)$ that relates the hindcast tides to the hindcast anomaly. We then will use this transfer function (with modification) to relate the forecast tides to the forecast anomaly. Determining $\hat{H}(k)$ from the hindcast inputs and outputs of the filter will be referred to as synthesis, and can be achieved via the quotient

$$\hat{H}(k) = \widehat{A_h}(k)/\widehat{T_h}(k), \quad (2A)$$

Where $\widehat{A_h}$ and $\widehat{T_h}$ are the DFT of the hindcast anomaly and tide signals, respectively. We intend now to use this transfer function on the forecast tides to generate an anomaly forecast, i.e.,

$$\widehat{A_f}(k) = \hat{H}(k)\widehat{T_f}(k). \quad (3A)$$

The desired time-domain signal can be obtained via the inverse DFT. However, there are two key issues with applying the transfer function in its current form.

While the denominator of (2A) seldom takes a zero value for any k in practice, the synthesized transfer function, when applied to the forecast tides, does not guarantee that $|\widehat{A_h}(k)| = |\widehat{A_f}(k)|$. In other words, the filter is allowed to change the amplitude of the Fourier components of the anomaly signal. While this is desirable in some applications of linear filter theory, we would like to keep these amplitudes the same. Further, we also intend to predict only the oscillations at tidal frequencies in the anomaly. The current transfer function will include all resolved frequencies.

Thus, we modify the transfer function as follows. First, we apply a bandpass filter to eliminate all frequency content that is non-tidal. We define a tidal frequency f_{mT} to be any frequency that is within 2 degrees per hour of $15m$ degrees per hour, where m is any positive integer. We also apply a low pass filter to eliminate non-physical, high frequency content from the forecast. Secondly, we apply a normalization factor to each mode of the transfer function to enforce that $|\widehat{A_h}(k)| = |\widehat{A_f}(k)|$. Formally, the new transfer function $\hat{H}'(k)$ is defined as

$$\hat{H}'(k) = \hat{H}(k) \times |T_h(k)/T_f(k)| \quad \text{if } |f(k) - f_{mT}| < 2, \text{ and } f(k) < 100$$

$$\hat{H}'(k) = 0 \quad \text{otherwise.} \quad (4A)$$

We note that the condition $|f(k) - f_{mT}| < 2$ is met if it holds for any m. Then, per equation (3A), $\widehat{A_f}(k) = \hat{H}'(k)\widehat{T_f}(k)$, and the desired forecast $A_f(n)$ is simply the inverse DFT of $\widehat{A_f}(k)$.

8B. *Signal-to-noise Indicator*

The formulation of our signal-to-noise ratio (S2N) is adapted from the definition of the linear filter. We measure S2N from the hindcast anomaly expressed in Fourier domain, $\widehat{A_h}(k)$. Our definition of "signal" is almost identical to the low- and band-pass filtering step (4A): The set of signal modes S, will be the set of modes k for which $|f(k) - f_{mT}| < 2$ and $f(k) < 100$. On the other hand, the set of all "noise" modes Q will be defined as the set of all k for which $|f(k) - f_{mT}| \geq 2$ and $f(k) < 100$. Then, S2N takes the definition

$$S2N = \left(\sum_{k\in S} |\widehat{A_h}(k)|^2\right)/\left(\sum_{k\in Q} |\widehat{A_h}(k)|^2\right). \quad (1B)$$

We remark that it may be beneficial to ignore very low frequency modes (say, 5 degrees per hour) in the definition of noise. This is because errors in surge predictions can lead to low frequency, high amplitude fluctuations in the anomaly that induce a small S2N despite clear tidal oscillations in the hindcast anomaly. As long as the tidal frequency content in the anomaly is well-separated (in frequency) from these low-frequency errors, we speculate that the filter may continue to mitigate erroneous tidal oscillations even in the presence of large-amplitude, low-frequency oscillations in the anomaly that are related to surge. We leave a detailed study of this hypothesis to future work.

8C. *Data Cleaning*

To accurately measure the anomaly, clean observation data is required. In practice, however, observations are not available with 100% reliability at all stations with the ETSS system. Our post-processing methodology is more sensitive to missing/erroneous observation data than the operational post-processor. Therefore, it is necessary to:

1. Check for missing observations. For the statistical analysis, this includes those future observations used assessing the true anomaly $B_f(t)$.
2. Use linear interpolation to estimate the Total Water Level (TWL) during short outages of observation.
3. Check for a few additional types of known issues with observations.

Here, we outline the data-cleaning strategy we employed for our statistical analysis. We recommend a similar strategy for future work.

We exclude data from our statistical study on a per-observation basis (rather than excluding entire stations). For a given set of guidance (both hindcast and forecast data), steps 1 and 2 are performed easily. We consider a small outage of observations (for which we perform interpolation) to be less than 7 hours in duration. We additionally exclude any guidance for which observations are down for more than 10% of either the 14-day hindcast or forecast windows.

In addition to outages, there are instances where the observations report non-physical behaviors, including large jumps and periods of perfectly constant TWL. To eliminate large jumps in the observations, we do not include any forecast for which the observed TWL changes more than 10 feet in 6 minutes. For our statistical analysis, this strategy also mitigates constant TWL errors, though a simple test for a constant level can also be constructed.

We note that the constant TWL issue is particularly disruptive to the discussed post-processing methodology. Constant observed TWL leads to the anomaly being dominated by the tide signal, activating the Fourier-based post-processor, which in turn attempts to make the predicted TWL constant to match the erroneously constant observations.

9. References

Kim, S. C., J. Chen, and W. A. Shaffer, 1996: An Operational Forecast Model for Extratropical Storm Surges along the U.S. East Coast. Preprints, Conference on Coastal Oceanic and Atmospheric Prediction, Atlanta, GA, Amer. Meteor. Soc., 281-286.

Rios-Berrios, R., 2010: Improving the Skill of the NWS's Extratropical Total Water Level Forecast System, 9th Annual Student Conference, Atlanta, GA, Amer. Meteor. Soc., S58. [https://ams.confex.com/ams/90annual/techprogram/paper_166793.htm].

Schuster R. and A. Taylor, 2015: Overhaul of MDL's Extratropical Storm Surge Post-Processing and Web Dissemination. Preprints, 13th Symposium on the Coastal Environment, *95th Amer. Meteor. Soc. Annual Meeting*, Phoenix, AZ, J12.5, 6 pp. [https://slosh.nws.noaa.gov/docs/data/2015_PostProcessing_ETSS.pdf]