

## 4.1 LATEST DEVELOPMENTS IN THE NWS' SEA LAKE AND OVERLAND SURGES FROM HURRICANES MODEL

Arthur Taylor<sup>1\*</sup> and Huiqing Liu<sup>2#</sup>

1. NOAA / NWS / STI / Meteorological Development Laboratory, Silver Spring, MD

2. Ace Info Solutions, Inc. Reston, VA

(#) Presented at the 100<sup>th</sup> AMS Annual Meeting, Boston MA, January 14, 2020

### 1. INTRODUCTION

The Sea, Lake, and Overland Surges from Hurricanes (SLOSH) model is used by the National Weather Service (NWS) to produce storm surge guidance in several ways. SLOSH-based simulation studies form the basis of the "hazards analysis" portion of hurricane evacuation planning. The Probabilistic tropical cyclone storm Surge (P-Surge) model runs SLOSH to analyze the storm surge inundation from each of its approximately 630 ensemble wind inputs. Similarly, the Probabilistic Extra Tropical Storm Surge (P-ETSS) model runs SLOSH to analyze the storm surge inundation from the North American Ensemble Forecast System's 42 ensemble members.

One key aspect of SLOSH, for this ensemble modeling, has been its fast-computational speed. This was due to design choices necessitated by the technology available when it was developed in the 1980's. However, SLOSH computational domains have recently become significantly broader and finer. While the results are more accurate, they take longer to run on a single processor. Additionally, NWS is working to couple a wave model to SLOSH, which will slow down the computation by at least a factor of two. To allow real-time ensemble storm surge modeling (e.g., P-Surge and P-ETSS) to continue to produce results within one hour, SLOSH needs to move from a single processor to a multi-processor program, enabling it to scale based on the number of available processors.

This paper will discuss the design changes that the Meteorological Development Laboratory has made to transition SLOSH into a multi-processor program and will demonstrate the impact via simulations of Hurricanes Katrina in 2005 and Irene in 2011. Section 2 expands upon why real-time ensemble modeling is important and why SLOSH is used. Section 3 describes the current run-time challenge, while section 4 describes the solution. Section 5 provides run-time results for the two cases. The paper concludes in section 6 with a summary.

### 2. MOTIVATION


#### 2.1. Large Wind Uncertainty

The accuracy of the National Hurricane Center's storm surge forecasts is directly related to both the accuracy of the storm surge model (SLOSH) used and the accuracy of the input parameters provided to it. Jelesnianski et al. (1992) found that when the hurricane's track, intensity, and size are estimated as well as possible after the event, SLOSH was correct to within approximately 20% of high-water marks. This is particularly good, considering high water marks often vary by 20% for locations that are less than a mile apart (Jelesnianski et al. 1984).

Unfortunately, as seen in Taylor and Glahn (2008), the errors in the wind input provided to SLOSH cause storm surge errors which are much larger than 20%. For example, as shown in Figs. 1 and 2, the storm surge forecast for Advisory 54 of Hurricane Ivan in 2004, made 12 hours before landfall, predicted a surge of 10 to 12 feet for Mobile, Alabama. Due to errors in the predicted position and size of the hurricane, only

---

\* Corresponding author address: Arthur Taylor, MDL, Silver Spring, MD 20910-3280; e-mail: [arthur.taylor@noaa.gov](mailto:arthur.taylor@noaa.gov)

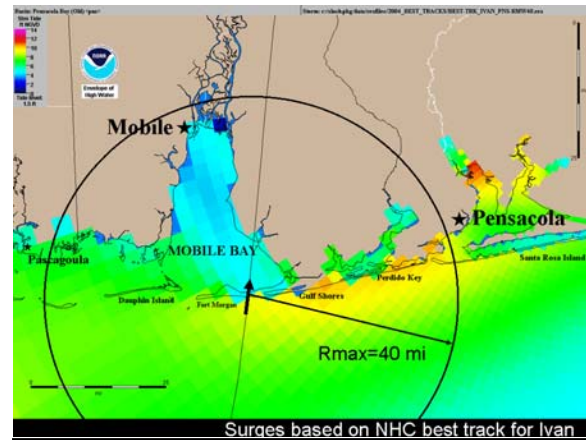


**Figure 1.** SLOSH storm surge forecast for Advisory 54 of Hurricane Ivan in 2004, made 12-hr before landfall.

3 to 5 feet actually occurred. Similarly, Pensacola Bay, Florida was forecast to have only 2 to 5 feet, but actually experienced 7 to 11 feet. Note, this is not a one-time phenomenon, as similar results can be seen for Hurricane Joaquin in 2015 in Liu and Taylor (2018). Thus, the storm surge information in the advisory needs to rely on an ensemble of perturbations of the wind inputs to account for the wind uncertainty.

## 2.2. Ahead of Time Ensembles

SLOSH is currently used in ahead-of-time ensembles in the form of Maximum Envelope of Water (MEOWs) (Shaffer 1989). The idea is to calculate the potential storm surge for an area. To do so, a set of hypothetical storms are provided as input to SLOSH. The storms all have the same forward speed, forward direction, size, and Saffir-Simpson category. The difference between them is where the storm track crosses land (i.e., landfall location). The storms in a MEOW each make landfall within at most a few miles of the other storms (see Fig. 3). The results of each ensemble are combined together by taking the maximum value in each grid cell at any time from any of the hypothetical storms. The MEOWs are made well before the storm exists and are provided to the US Army Corps of Engineers. They in turn combine the MEOWs with population and transportation studies to



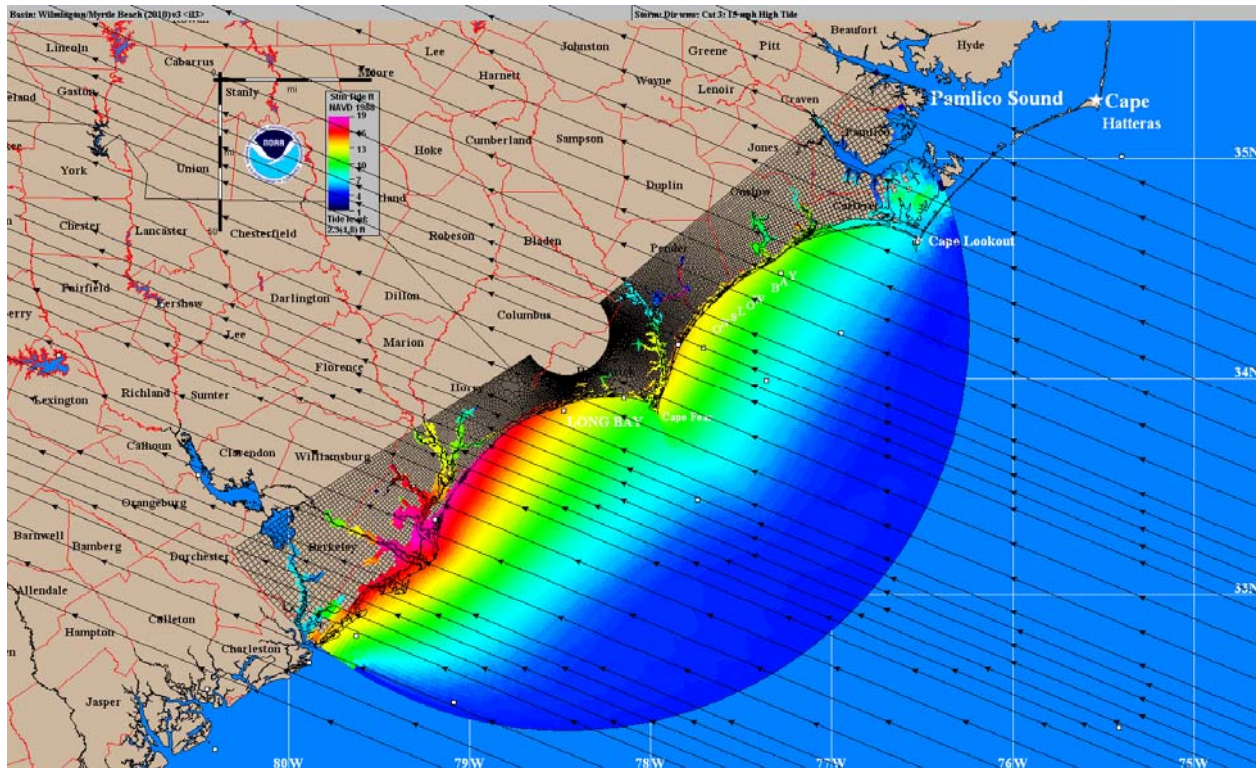
**Figure 2.** SLOSH storm surge hindcast created with the best available information for Hurricane Ivan in 2004.

develop evacuation plans which are provided to the Federal Emergency Management Agency.

A conceptually similar method of utilizing pre-made results was done by Smith et al. (2012). Their idea was to determine which of a set of pre-run hypothetical storms best fit the current forecast storm. Given a sub-set of likely storms, they could look up the results and statistically combine them together. The hypothetical storms could model both surge and waves.

The advantage of both the MEOW and Smith's data-mining concepts is that they require minimal computation during the storm as they are made ahead-of-time. The challenge is they don't handle time very well, because the hypothetical storms are run before the timing of the actual storm is known. The lack of timing information makes it challenging to model tidal water levels (highly dependent on time), abnormal water levels (e.g., sea level rise, disruptions of currents, etc.), and external wind-fields.

The MEOW's handle tidal and abnormal water levels by initializing the water to high-tide. This results in a conservative over-estimate, which is sufficient given the planning purpose of the MEOWs. Smith's data-mining effort could linearly superimpose the tidal and abnormal



**Figure 3.** A sample Maximum Envelope of Water (MEOW) for a Category 3 hurricane in the Wilmington basin at high tide moving west-north-west at 15 miles per hour.

water levels, but that wouldn't model non-linear tide surge interactions, nor calculate the correct inundation. Thus, both examples of ahead-of-time ensembles can account for tidal and abnormal water levels; however, they can't do so as well as a real-time ensemble. Additionally, neither method can handle external wind fields.

### 2.3. Real-Time Ensembles

In order to do a real-time storm surge ensemble for forecasting purposes, the results need to be available within an hour. Delays beyond an hour would result in the data being overcome by the next forecast cycle. Additionally, the forecast needs to contain at least 4 days of information as evacuation decisions are made 3 to 5 days before landfall. Furthermore, it needs to provide coverage for all US storm surge vulnerable areas and run on the National Centers for Environmental Prediction's (NCEP) operational super-computer.

In terms of modeling, the storm surge ensemble system should have a reasonable selection of wind cases, preferably based on error statistics. If a parametric wind model is used, it should be integrated within an external wind field. Ideally the surge-model component would explicitly model surge, tide, waves, and rainfall/river flooding. Additionally, the ensemble system should have some form of observational analysis for initial water levels. Finally, there would be a readily available process of maintaining the computational domains.

This is a tall order. To date, the best solution for this are SLOSH-based ensemble systems, as design choices for SLOSH, necessitated by the technology available when it was developed in the 1980's, make it very fast with a small computational footprint. That allows the ensemble system to diagnose numerous wind permutations with the requisite 4-day forecasts while running on NCEP's machines. SLOSH has coverage for most of the US, only lacking coverage for some US islands in the Pacific

Ocean, and that coverage is maintained via the National Hurricane Program.

That said, real-time ensembles based on SLOSH don't meet all the requirements. For instance, NWS is currently working to enable SLOSH to explicitly model waves caused by each ensemble member (Yang 2020). Additionally, NWS is working on adding an initial water condition based on observations to the extra-tropical storm surge ensemble (Liu 2020) which eventually will be added to the tropical storm surge ensemble. Finally, long term plans include enabling the SLOSH-based ensemble systems to handle the rainfall/river flooding requirement as well as nesting the parametric wind field within an external wind grid. Still, despite its flaws, no other storm surge model can meet as many of the run-time and ensemble modeling requirements as SLOSH can.

### 3. PROBLEM

One recent trend with SLOSH is to create broader basins with a fine grid resolution. The logic for this stems from Hurricane Ike which made landfall in Galveston Texas in 2008. The high-resolution basin available at the time was too narrow to capture Hurricane Ike, so it significantly under-forecasted the event. The alternative was to use the broader Gulf of Mexico basin, which captured Hurricane Ike, but was too coarse to resolve the details of the flooding. The desire to add waves has accelerated the trend.

This trend is exemplified by the fact that in 2006, P-Surge used 35 basins which, if they were all run, required 2,369,815 cell Calculations per model Minute (C/M). In 2020, P-Surge requires 28 basins which, if they are all run, require 14,862,102 C/M, or 6 times more calculations. So, the trend is for fewer, but broader and finer, grids. This will continue as we implement the next generation of grids with basins like: Super Texas (1,920,000 C/M), New Orleans (99,225 -> 2,232,240 C/M), and New York (233,887 -> 4,760,550 C/M). Thus in 2022, P-Surge may require 23,441,780 C/M, or be 10 times more computationally expensive than the

P-Surge of 2006. Unfortunately, it is likely to be even more computationally expensive than that, as NWS is planning to add more physics (such as waves) and more domains (such as Puerto Rico and Hawaii).

MDL has already hit the run-time limits in P-Surge since a single track in the latest South Florida basin can take more than 1 hour to run for a single processor. To make it work, MDL had to split the basin into thirds. Unfortunately, that is a sub-optimal solution as computations on one sub-basin are unable to influence computations on the other sub-basins.


### 4. SOLUTION

A better solution is to enable multiple processors to work as a team to solve a single storm and basin problem. This is done via the Message Passing Interface (MPI) library and domain decomposition as seen in Fig 4.

To use MPI, we need to determine what information to pass between the processors and how often. With the parallelization of SLOSH via MPI (SLOSH-MPI), we chose to: (a) pass a halo or set of "shadow cells" along the boundary of each subdomain each time-step, and (b) pass the whole subdomain when it is time for the grid to be written to disk. Each time-step, neighboring subdomains calculate results for their neighboring halo cells, but the results are treated as "scratch space" and replaced by results from the neighbor's actual grid cells that overlap the halo ones. This works as long as the computation per communication time-step stays within the subdomain plus the surrounding halo.

SLOSH goes through three stages each time-step: Continuity, Momentum, and Smoothing. We choose to communicate the halo boundary at the end of all three stages rather than between some combinations of the various stages. The intent is to reduce the amount of communication thereby increasing efficiency. As the equations for Continuity and Momentum both require one extra grid cell (i.e. halo cell) on the boundary, the surrounding number of halo cells





**Figure 4.** A sample domain decomposition for an 11x11 grid into 9 4x4 grids with 1 grid cell overlap.

(i.e. halo-width), needs to be at least 2 grid cells. The challenge is the Smoothing stage. Originally, we thought the Smoothing stage would require increasing the halo-width by a single grid cell. Unfortunately, we couldn't get identical results for Hurricane Katrina in the New Orleans basin (HMS8) without using a total halo-width of 7. Our hypothesis is that Smoothing calculations are harder for some basins, particularly when dealing with large amounts of inundation.

Once we'd chosen the halo-width and were ready to modify the code, we had to choose how to update the indexes for all the loops. One option was to change the start of each loop from 1 to a variable based on the subdomain. This appeared dangerous as there was a lot of code to modify, and any index mistake would be subtle. Instead we chose to have each processor read in the entire basin and then shift the area of its focus to start at 1,1. This reduced the changes to the loops, but we had to be careful while updating the indexes of the vector data (e.g. barriers, cuts, flows, etc).

Another aspect to consider when dealing with domain decomposition is the topology. For instance, you can break up a rectangular domain amongst 4 processors by splitting the rows into 4

groups (4x1), splitting the columns into 4 groups (1x4), or splitting the rows and columns into 2 groups each (2x2). The SLOSH-MPI algorithm breaks up the domains based on a call to `MPI_Dims_create()`, which automatically creates a balanced distribution of processes. So, in the case of 4 processes it would split the rows and columns into 2 groups each (2x2).

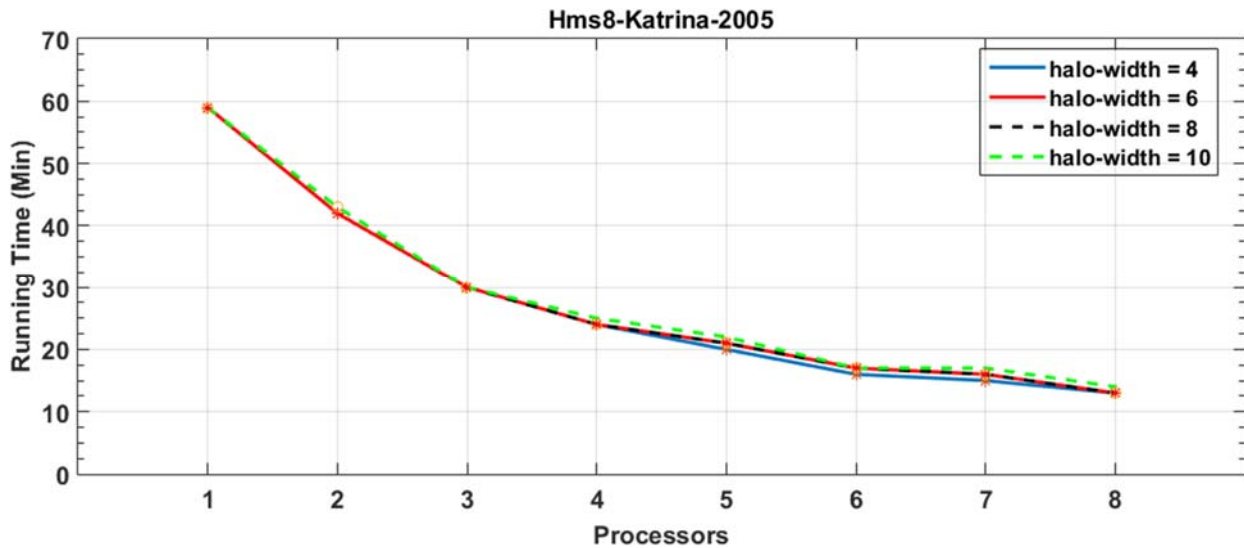
## 5. RESULTS

The results for Hurricane Katrina in 2005 in the New Orleans basin (HMS8) show a reduction of run-time as the number of processors increases (Fig 5). As seen in Table 1, the percent improvement of adding a single processor varies from 4% to 29%. Additionally, the influence of the subdomain topology can be seen in the sudden reduction in improvement when adding a 7<sup>th</sup> processor. While adding an 8<sup>th</sup> processor produced diminished returns, it is worth trying to add more for this basin and storm.

The results for Hurricane Irene in 2011 in the Chesapeake Bay basin (CP5) as seen in Table 2 also show a reduction of run-time as number of processors increases. Here, the percent improvement per additional processor ranges from -3% to 25%. The influence of the topology can be seen again when the 6<sup>th</sup> processor is

added. In this case we likely don't want to increase the number of processors any further than 8. Both tables also show that increasing the

halo-width results in a roughly linear increase in the run-time.



**Figure 5.** Run time for Hurricane Katrina in the New Orleans basin (HMS8) vs number of processors for halo-widths of 4, 6, 8, and 10 cells.

**Table 1.** Run time (minutes) and percent improvement (from 1 fewer processor) for Hurricane Katrina in the New Orleans basin (HMS8). Columns are the number of processors.

	1	2	3	4	5	6	7	8
<b>Halo-width = 4</b>	58.8	42.1 28%	29.7 29%	23.6 21%	20.0 15%	16.5 18%	15.0 9%	12.8 15%
<b>Halo-width = 6</b>	58.8	42.3 28%	30.1 29%	23.9 21%	20.5 14%	16.7 19%	15.6 7%	13.1 16%
<b>Halo-width = 8</b>	58.8	42.6 28%	30.3 29%	24.3 20%	21.1 13%	17.0 19%	16.2 5%	13.4 17%
<b>Halo-width = 10</b>	58.8	42.9 27%	30.5 29%	24.6 19%	21.7 12%	17.4 20%	16.7 4%	13.7 18%

**Table 2.** Same as in Table 1, but for Hurricane Irene in the Chesapeake Bay basin (CP5).

	1	2	3	4	5	6	7	8
<b>Halo-width = 4</b>	20.4	16.4 20%	13.1 20%	11.4 13%	8.9 22%	8.9 0%	6.7 25%	6.9 -3%
<b>Halo-width = 6</b>	20.4	16.5 19%	13.2 20%	11.5 13%	9.0 22%	9.0 0%	7.0 22%	7.1 -1%
<b>Halo-width = 8</b>	20.4	16.7 18%	13.3 20%	11.7 12%	9.2 21%	9.2 0%	7.4 20%	7.2 3%
<b>Halo-width = 10</b>	20.4	16.8 18%	13.4 20%	11.8 12%	9.4 20%	9.4 0%	7.7 18%	7.4 4%

## 6. SUMMARY

The parallelization of SLOSH via MPI was successful and will enable a number of enhancements to the real-time SLOSH based ensemble storm surge systems one of which is the use of larger, finer resolution basins. Another enhancement is enabling the ability to scale past the time it takes a single processor to handle a storm in a single basin. Finally, this parallelization will give the flexibility needed to explicitly calculate waves within the ensemble members.

There were various design choices in the parallelization process including: (a) how frequently to communicate, (b) how wide a halo to establish, (c) whether to shift the grids or update all loop indexes, and (d) what topology to use for the domain decomposition. The halo-width decision was explored to determine its impact. Other decisions, such as topology or communication frequency, could be explored in the future. For instance, the best way to decompose a grid may be basin dependent as it may involve balancing the number of wet/dry cells. Similarly, it may be better to communicate a halo-width of 2 before the Continuity and Momentum stages and then a halo-width of 5 before the Smoothing stage. The results further show that the optimal number of processors may be basin dependent, so for ensemble systems, consideration should be given to running different basins with different numbers of processors.

In the near term, MDL needs to resolve how to handle basins with a periodic boundary condition (e.g.,  $UB(I,1) = UB(I,N)$  for  $I = 1 \dots M$ ). An example of this is the Puerto Rico basin, which allows water to flow from one edge to the opposite edge. Additionally, MDL will be working to add spatially varying bottom friction to SLOSH.

## 7. REFERENCES

Jelesnianski, C. P., J. Chen, W. A. Shaffer, and A. J. Gilad, 1984: SLOSH- a hurricane storm surge forecasting model. Preprints *Oceans*

84, Washington, D.C., Marine Technology Society and IEEE/Oceanic Engineering Society, 314-317, [https://slosh.nws.noaa.gov/docs/Jelesnianski\\_1984\\_SLOSH\\_HurcnFcstModel.pdf](https://slosh.nws.noaa.gov/docs/Jelesnianski_1984_SLOSH_HurcnFcstModel.pdf).

Jelesnianski, C. P., J. Chen, and W. A. Shaffer, 1992: SLOSH: Sea, Lake and Overland Surges from Hurricanes. NOAA Technical Report, NWS 48, National Oceanic and Atmospheric Administration, U. S. Department of Commerce, 71 pp., [https://slosh.nws.noaa.gov/docs/SLOSH\\_TR48.pdf](https://slosh.nws.noaa.gov/docs/SLOSH_TR48.pdf).

Liu, H., and A. Taylor, 2018: Development of the NWS' Probabilistic Extra-Tropical Storm Surge Model and Post Processing Methodology. *16th Symposium on the Coastal Environment*, Austin, TX, Amer. Meteor. Soc., 1.2, 8 pp., [https://ams.confex.com/ams/98Annual/mediatefile/Manuscript/Paper329410/P-ETSS1p0\\_AMS\\_abstract\\_extended.pdf](https://ams.confex.com/ams/98Annual/mediatefile/Manuscript/Paper329410/P-ETSS1p0_AMS_abstract_extended.pdf).

\_\_\_\_\_, and A. Taylor, 2020: Latest Development in the NWS' Probabilistic Extra-Tropical Storm Surge Model. *18th Symposium on the Coastal Environment*, Boston, MA, Amer. Meteor. Soc., 4.2, 10 pp., <https://ams.confex.com/ams/2020Annual/webprogram/Paper370559.html>.

Shaffer, W. A., C. P. Jelesnianski, and J. Chen, 1989: Hurricane Storm Surge Forecasting. *11th Conference on Probability and Statistics in Atmospheric Sciences*, Monterey, CA, Amer. Meteor. Soc., pp. J53-J58, <https://slosh.nws.noaa.gov/docs/Shaffer1989HurricaneStormSurgeForecasts.pdf>.

Smith, J. M., A. B. Kennedy, J. J. Westerink, A. A. Taflanidis, and K. F. Cheung, 2012: Hawaii Hurricane Wave and Surge Modeling and Fast Forecasting. *Int. Conf. Coastal Eng.*, vol. 1, no. 33, p 8, <https://doi.org/10.9753/icce.v33.management.8>.

Taylor, A., and Glahn, B., 2008: Probabilistic Guidance for Hurricane Storm Surge. *19th Conference on Probability and Statistics*, New Orleans, LA, Amer. Meteor. Soc., 7.4., <https://ams.confex.com/ams/pdfpapers/132793.pdf>.

Yang, D., Van der Westhuysen, A., Rhome, J., and Fritz, C., 2020: Efficient Wave-Surge Coupling with SLOSH-Wave for Hispaniola. *18th Symposium on the Coastal Environment*, Boston, MA, Amer. Meteor. Soc., 4.4., <https://ams.confex.com/ams/2020Annual/webprogram/Paper370526.html>.