Building the ADCIRC Collaborative

A federated approach to storm surge modeling and model output distribution to enable better decision support

A RENCI WHITE PAPER

renci

CONTACT INFORMATION

Brian Blanton
Telephone: 919.445.9620
Email: brian_blanton@renci.org

Summary

Strong coastal storms, including both tropical cyclones (e.g., hurricanes) and extra-tropical storms (e.g., nor'easters) routinely threaten human lives and property. Over the past decade, 16 storms striking the Atlantic coast have caused more than $1 billion in damage each, with other storms costing many hundreds of millions more. Hurricane Katrina in 2005, one of the most devastating in U.S. history, caused 1,833 deaths and $125 billion in damage (National Climatic Data Center 2013). Much of the damage from such storms comes from flooding due to storm surge (an abnormal rise in ocean water caused by a coastal storm), and destruction caused by large wind-generated waves.

Accurate, detailed, and rapid storm predictions can help save lives and prevent property damage by informing evacuations and other emergency preparedness and response activities. Storm impact models are also critical to long-term planning for coastal development, infrastructure, and insurance. Current models used to predict storm impacts require tradeoffs between detail, accuracy, and speed. ADCIRC is a powerful model that specializes in predicting the response of the coastal ocean, including storm surge and wave action, caused by coastal storms. Although ADCIRC is being deployed at several major research centers and has proved highly valuable for research and long-term planning, computing resource challenges limit its general utility for making timely predictions during actual storm events. Establishing an ADCIRC collaborative—a working relationship among research centers with ADCIRC capabilities—would optimize the deployment of this modeling system and enable real-time analysis for the benefit of communities affected by coastal storms.

At a Glance

- Strong coastal storms, in the form of tropical cyclones and extra-tropical storms, cause enormous damage. Accurate prediction of a storm's impacts can help save lives and reduce property damage.
- ADCIRC is a powerful model that predicts the impacts of storm surge and wave action associated with coastal storms at a level of detail and accuracy not available from other modeling systems.
- An ADCIRC collaborative—a working relationship among research centers with ADCIRC capabilities—will allow effective and efficient deployment of ADCIRC to support real-time storm prediction and on-the-ground preparedness and response.

The Challenge

Many factors combine to make the consequences of coastal storms difficult to predict. They include a confluence of multiple atmospheric and oceanographic factors—such as air and ocean temperature, wind speed, wind direction, tides, storm surge, and wave action—which then interact with an array of built and natural structures. The impact of storm surge, which is often responsible for the extensive flooding associated with these storms, depends on the size, intensity, and trajectory of the storm, as well as the topography of the land and the presence of structures such as dunes, levies, or floodwalls. Similarly, the impact of wind-generated waves, responsible for much of the damage to built structures, depends on wind speed and oceanographic factors, as well as the design of coastal communities, building architecture, and other factors. The ability to accurately predict flooding and structural damage is critical for planning road closures and evacuation routing, shelter and equipment placement, and recovery and restoration efforts—all of which impact a storm's human and material toll.

Two models are most frequently used for predicting storm surge associated with coastal storms. The model developed by the National Oceanic and Atmospheric Administration (NOAA) is the Sea, Lake and Overland Surges from Hurricanes (SLOSH) model. SLOSH can be run in a short enough time interval to allow the completion of hundreds of model runs based on

The Team

Written by **Brian Blanton**, RENCI Senior Research Scientist and Oceanographer, with contributions from **Rick Luettich**, Director, UNC Hazards Center and UNC Institute of Marine Sciences, **Jessica Losego**, UNC Institute for the Environment, and **Jason Fleming**, Principal Engineer at Seahorse Coastal Consulting. Additional writing support from **Anne Frances Johnson**.

Contact

Brian Blanton, 919.445.9620
brian_blanton@renci.org

Vol. 1, No. 1 in the RENCI White Paper Series, August 2013
Renaissance Computing Institute
University of North Carolina at Chapel Hill
919-445-9640
www.renci.org
@renci

different predicted storm trajectories and properties, which can then be used to create maps of the probability of flooding in storm-impacted areas (Taylor & Glahn 2008; Glahn, et al. 2009). However, two issues limit SLOSH's accuracy. First, SLOSH is not written to harness multiple computer CPUs simultaneously, so spatial resolution must be kept relatively coarse for simulations to fit onto one CPU and still complete quickly. Second, SLOSH represents spatial features using an orthogonal rectangular grid, which makes it impossible to achieve high resolution in the coastal region without over-resolving deeper water. This typically results in the use of model domains that are too small to fully capture the coastal response.

A second model, ADCIRC, can compute extremely high-resolution, detailed, and accurate predictions. It is written to use as many computer CPUs as are available, and because ADCIRC uses triangles as its descriptive basis, the model's spatial resolution can be very high in the coastal zone without over-resolving deeper, offshore water (Westerink, et al. 2008). Using this approach, ADCIRC can generate much more detailed flood forecast maps than SLOSH (see Figure 1). However, when configured to run at high spatial resolution, ADCIRC requires significant computing resources, making each run of the model more resource intensive than SLOSH.

While SLOSH and ADCIRC generally solve the same equations for predicting storm surge, ADCIRC incorporates two additional mechanisms directly into its simulations: astronomical tides and breaking wind-generated waves. Astronomical tides, which are caused by the gravitational attractions between the earth-moon-sun system, cause the periodic fluctuations of oceanic water level seen at the coast. The most dangerous storm surges and waves are those that occur at the same time as high tide; however, the timing of high tide is location-specific, so the time of peak tide in an inlet or sound may be quite different from the tide timing at a nearby beach.

ADCIRC dynamically incorporates the tide in each model run, generating a location-specific tidal response and therefore a more accurate prediction of storm surge response. In addition, ADCIRC accounts for the impact of wind-generated waves on storm surge. As wind-generated waves break, they can transfer momentum to the storm surge, increasing the water level of the storm surge by as much as 5 to 15 percent on wide, shallow continental shelves (Dean & Dalrymple 1991; Niedoroda, et al. 2008) and substantially more on narrow, steep shelves (Resio & Westerink 2008). ADCIRC is coupled to the SWAN wind-wave model and therefore directly models the wind-wave field and incorporates the momentum transfer from the waves to the storm surge. Accounting for astronomical tides and breaking wind-generated waves improves ADCIRC's accuracy; however, it also adds to the computing resources required to run the model.

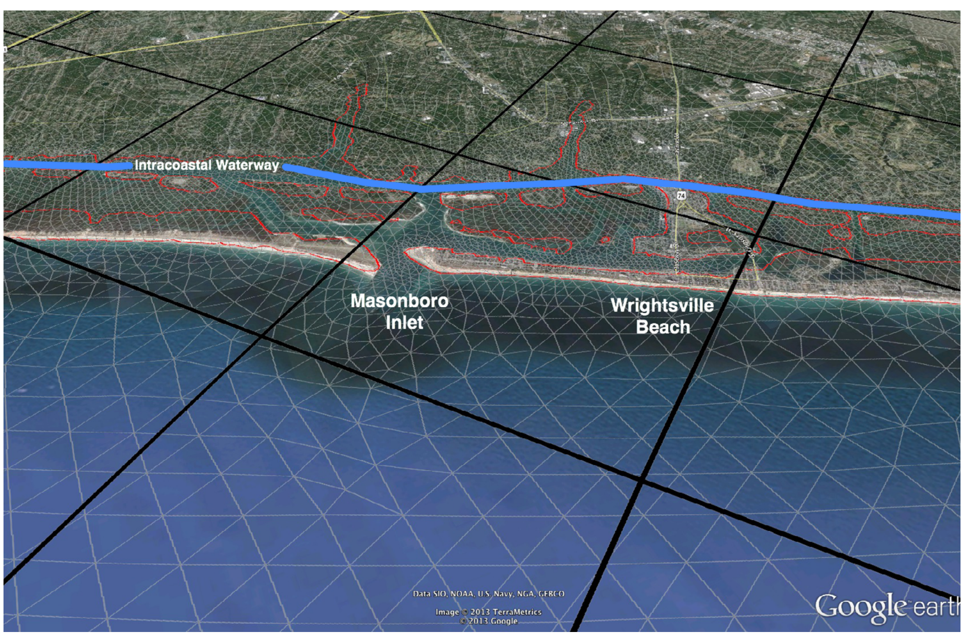

Figure 1: ADCIRC (gray triangles) and SLOSH (black squares) model grids in the Wrightsville Beach, NC area. Because ADCIRC uses triangles as its descriptive basis, it allows flood prediction at a much greater level of detail than the coarser SLOSH model, which uses squares. The red line is the coastline in ADCIRC's grid.

The ability to compute a sizable ensemble of high-resolution ADCIRC runs would provide information on prediction uncertainty and, given enough runs (e.g., 100 or more), could allow the creation of extremely detailed probabilistic flood and wave-impact predictions to guide decision making during coastal storms. However, several key challenges limit the utility of ADCIRC for real-time ensemble simulations:

1. *Horsepower:* Depending on the size of the area being modeled, a typical five-day long forecast run can take an hour or more on several hundred to several thousand computer

processors. Running multiple simultaneous simulations requires more computing power than computing centers have available at any one time.

2. *Cost:* Due to its high demands on computing resources, ADCIRC can impose substantial costs in personnel and computing time.
3. *Time:* To be useful for decision making during an actual storm, predictions must be generated in a matter of hours or less. Few, if any, computing centers have the available capacity to perform a sizable ensemble of high-resolution ADCIRC runs in this time frame. In addition, the nature of tropical storm events means computing resources would need to be available on demand with a short warning period—a significant obstacle for heavily-used computing facilities.
4. *Output Translation: Forecasters,* emergency managers and other key decision makers need predictive data in formats that fit into decision making tools that are already in place. Thus, ADCIRC output data needs to be translated into forms that are practical for use on the ground.

Federal computing facilities (such as the Department of Defense's Cray systems in Vicksburg, MS, and NOAA's Jet systems in Boulder, CO) have the necessary hardware to rapidly create ADCIRC ensemble predictions, but these clusters and supercomputers are heavily used and on-demand access is generally limited. At the federal level, there has historically been much more emphasis placed on computer simulations of the meteorological components of a hurricane itself, with less emphasis on computing the impacts of these events in terms of storm surge, waves, and even damage.

ADCIRC Key Features

ADCIRC was developed in the early 1990s by Rick Luettich of the University of North Carolina at Chapel Hill and Joannes Westerink of the University of Notre Dame (Luettich, et al. 1992; Westerink, et al. 2008). While they have continued to lead model development (e.g., Dietrich, et al. 2011; Tanaka, et al. 2011), in recent years the broader ADCIRC community, comprised of federal government, academic, and industry users and researchers, has contributed to ADCIRC's continued development and improvement, as well. The Renaissance Computing Institute (RENCI) helps to maintain the ADCIRC code and serves as the main ADCIRC code repository.

To accurately predict storm surge at high resolution, ADCIRC incorporates a variety of factors affecting the physical forces behind storm surge (forcing mechanisms), which areas will be inundated (space), and when the greatest impacts will occur (time). These factors include:

Forcing mechanisms

- Surface wind and pressure from meteorological models or hurricane forecast tracks
- Wave breaking forces
- Astronomical tides
- River discharge

Space

- Linear triangular finite elements
- High-resolution topography and bathymetry
- Hydraulic structures and raised features (levees, roadways, channels)
- Local land roughness and tree canopy
- Parallel implementation via MPI, Domain Decomposition (METIS)

Time

- Finite difference, Courant–Friedrichs–Lewy time-step limited
- User-specified model output rate; hourly snapshots of water level and waves are typical

Ideas into Action: An ADCIRC Collaborative

ADCIRC is poised to transition from a valuable research/design model into a powerful tool for forecast guidance and decision making. ADCIRC already contains all of the technical elements needed to generate on demand either single deterministic results or an ensemble of results at high resolution and translate them for real-world, real-time forecasting. But no single computing center can easily achieve this goal alone. What is needed is a collaborative partnership among computing centers equipped with ADCIRC. The University of North Carolina at Chapel Hill, University of Texas, Louisiana State University, College of Staten Island at the City University of New York, and NOAA, are all experienced with ADCIRC and comprise a logical group to form the basis of such a partnership.

This partnership will pave the way for ADCIRC to be more efficiently and effectively deployed—and for prediction results to be distributed to forecast users—whenever a strong coastal storm poses a threat to the U.S. coast. The establishment of the ADCIRC collaborative is envisioned in two phases.

Phase I

Phase 1 would connect ADCIRC outputs from all regional centers to create a unified view of data from all ADCIRC runs. To achieve this, all sites would contribute data to a shared data management system using standard data formats, metadata descriptions, and data access protocols. One approach would be to manage the data using iRODS (integrated Rules-Oriented Data System, an advanced data management middleware, http://irods.org), and use OPeNDAP (www.opendap.org) and Unidata's THREDDS (www.unidata.ucar.edu/projects/THREDDS) for data cataloging and access.

Ultimately, the shared system would translate the output data into Web-based visualizations and shapefiles for use by forecasters and emergency managers. Connecting and standardizing the outputs from all sites running ADCIRC would allow critical information to be accessed through a single point-of-entry, regardless of which center generated the output data or which areas were impacted by the storm.

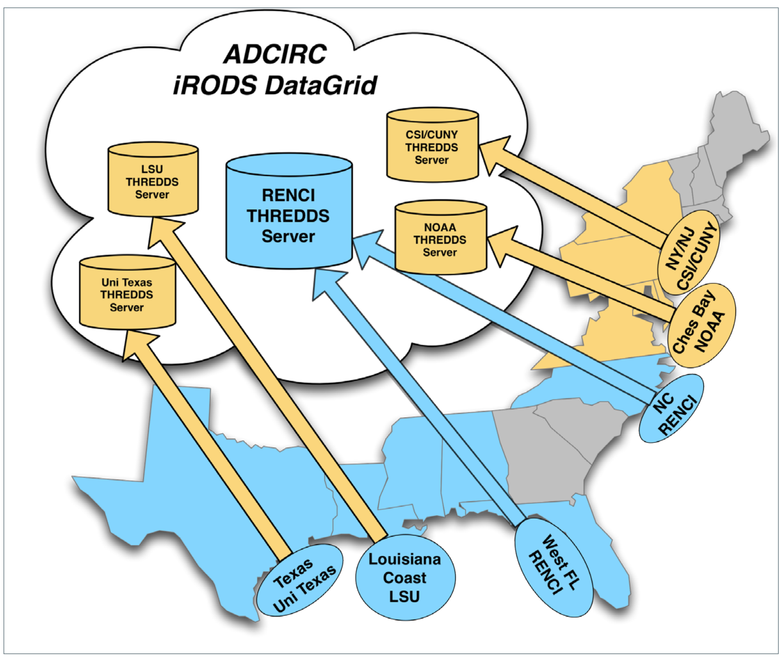

Figure 2: Conceptual diagram of Phase 1 of the ADCIRC Collaborative. Existing ADCIRC systems are shown in blue and cover the coastal regions covered by the ovals. Planned systems are shown in orange. Each system publishes ADCIRC model results to a THREDDS server. Currently, only the RENCI THREDDS server is operating. The participating groups are the University of Texas at Austin, Louisiana State University, RENCI, NOAA's Coast Survey Development Laboratory, and the College of Staten Island at the City University of New York. The collection of ADCIRC results is represented by the cloud and is maintained by the iRODS (integrated Rule-Oriented Data System) data management middleware.

Phase 2

In Phase 1, only ADCIRC outputs would be dealt with collectively. Phase 2 would allow ADCIRC sites to gather inputs and run the model collectively, thus enabling the collaborative to realize the ultimate goal of producing ensemble calculations at high resolution on demand. Pooling the computing resources of multiple sites in this way would not only allow the timely generation of more simulations of a storm-threatened area, it would also provide an important structure for handling two unlikely (but not impossible) scenarios: that two or more coastal storms would threaten different parts of the coast at the same time, or that an ADCIRC site would go offline during a storm.

ADCIRC already contains all of the technical elements needed to generate on demand either single, deterministic results or an ensemble of results at high resolution and translate them for real-world, real-time forecasting. But no single computing center can easily achieve this goal.

At the conclusion of this phase, all participating ADCIRC systems would be able to respond to an impending storm threat anywhere along the coast by conducting multiple simultaneous model runs. The outputs from these model runs would then be integrated into high-resolution probabilistic predictions and translated into usable formats to inform decision making.

Achieving Phase 2 will require research on two fronts. First, research is needed to determine the best approach for the shared system to request remote resources, stage the needed model input data to the remote site, submit the parallel computation, and retrieve the model output. This research is already underway at RENCI, in projects such as ExoGENI (www.exogeni.net) and applications of the iRODS software technology. Second, a formal policy is needed through which each ADCIRC site would pre-designate computational power for on-demand use when a storm threatens the coast, and that would allow remote ADCIRC systems to use shared resources. This is a significant hurdle given high demands placed on university computer and data centers and competition with other time-sensitive computational needs.

The Upshot

An ADCIRC collaborative would address many of the challenges that have limited the real-world application of this powerful software thus far:

1. *Horsepower:* The collaborative would greatly increase the computational resources available for running ADCIRC during an actual storm event. By drawing upon the computing capacity of multiple computing centers simultaneously, the collaborative would allow coordinated, redundant, parallel model runs, increasing the likelihood of generating useful results in the required time window. It also would allow modeling of perturbations to account for different paths the storm could take, ultimately enabling the creation of ensemble-based flooding predictions. Finally, it would ensure ADCIRC could continue to run in multiple states even if one center drops offline due to storm impacts.
2. *Cost:* A collaborative structure allows all centers to reap benefits from shared infrastructure investments, reducing the overall cost to each center. For example, the shared output system developed in Phase 1 would eliminate the need for each center to develop its own method of translating data outputs for use by on-the-ground decision makers, reducing redundancy and thus lowering costs.
3. *Time:* The collaborative would allow multiple sites to run ADCIRC simultaneously, thus generating useful results much more quickly than would be feasible at any single site. In addition, the agreements put in place in Phase 2 would ensure computing resources can be immediately redirected for ADCIRC when a storm poses a threat, allowing model outputs to be generated in time to make a difference.
4. *Output Translation:* Rather than requiring data users to work directly with data producers, the collaborative output systems developed in Phase 1 would help data contributors offer practicable information for users without needing to create the interfaces themselves. This would result in a unified, streamlined product available to forecasters and emergency managers.

This partnership will pave the way for ADCRIC to be more efficiently and effectively deployed—and for prediction results to be distributed to forecast users—whenever a strong coastal storm poses a threat to the U.S. coast.

What Will It Take?

To fully unleash ADCIRC's potential, a new collaborative structure is proposed that channels the collective resources of all ADCIRC-equipped centers when a strong storm threatens the coast. Such a structure would largely overcome current computing resource limitations and ultimately allow the production of ensemble calculations at high resolution on demand. The extremely detailed, accurate storm impact predictions generated through this effort would directly benefit emergency managers and the public—saving lives and money.

Implementation of Phase I will require two critical components and cyberinfrastructure experts. First, a capable web programmer must build a front-end portal to display the desired information coming from the participating ADCIRC centers. This would include a "system status" portal that gathers real-time information from participating centers using an AMQP implementation like RabbitMQ to track the status of simulations (for example, pending, running, complete, failed), collect ensemble member information, and link to the output files on data servers.

Second, a technician with knowledge of the methods used with ADCIRC must help maintain the automated software on participating computing resources and solve unanticipated issues that arise in federating multiple ADCIRC instances through the collaborative data grid using iRODS.

Phase II will require maintenance of Phase I activities and the expansion of the web-based results to include effective display of ensemble results and probabilistic assessments, and to organize and implement a coordinated set of model runs across distributed resources. The resource sharing aspects of Phase II will require policy agreements between participating centers and administrators to enable cross-center communication and cooperation, as well as adoption of compute resource queues that allow preemptive and on-demand simulations during storm events.

References

Dean, R. G. & Dalrymple, R. A. (1991). Water wave mechanics for engineers and scientists. World Scientific.

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

Glahn, B., Taylor, A., Kurkowski, N., & Shaffer, W. A. (2009). The role of the slosh model in National Weather Service storm surge forecasting. National Weather Digest, 33(1), 3-14.

Luettich, R. A., Westerink, J. J., & Scheffner, N. W. (1992). ADCIRC: an advanced three-dimensional circulation model for shelves coasts and estuaries, report 1: theory and methodology of ADCIRC-2DDI and ADCIRC-3DL. Dredging Research Program Technical Report DRP-92-6, U.S. Army Waterways Experiment Station, Vicksburg, MS, 137pp.

National Climatic Data Center. Billion-Dollar Weather/Climate Disasters. http://www.ncdc.noaa.gov/billions/overview [Accessed July 2, 2013]

Niedoroda, A., Das, H., Slinn, D., Dean, R., Weaver, R., Reed, C., & Smith, J. (2008). The role of wave set-up during extreme storms. Proceedings of the 31st International Conference on Coastal Engineering, 1, 250–261.

Resio, D.T. & Westerink, J.J. (2008). Modeling the physics of storm surge. Physics Today, 61, 33-38.

Tanaka, S., Bunya, S., Westerink, J.J., Dawson, C., Luettich, R.A. (2011). Scalability of an unstructured grid continuous Galerkin based hurricane storm surge model. Journal of Scientific Computing, 46, 329-358.

Taylor, A. A., & Glahn, B. (2008). Probabilistic guidance for hurricane storm surge. http://www.nws.noaa.gov/mdl/psurge/talks/PsurgeAMS_20080115.pdf [Accessed July 2, 2013]

Westerink, J., Luettich, R., Feyen, J., Atkinson, J., Dawson, C., Roberts, H., Powell, M., Dunion, J., Kubatko, E., & Pourtaheri, H. (2008). A basin- to channel-scale unstructured grid hurricane storm surge model applied to Southern Louisiana. Monthly Weather Review, 136, 833-864.

About RENCI

RENCI is an institute of the University of North Carolina at Chapel Hill that develops and deploys advanced technologies to enable research discoveries and practical innovations. RENCI partners with researchers, policy makers, and technology leaders to engage and solve the challenging problems that affect North Carolina, our nation and the world. The institute was launched in 2004 as a collaborative effort involving UNC Chapel Hill, Duke University and North Carolina State University. For more information, see www.renci.org.

How to reference this paper:

Blanton, B., Luettich, R., Losego, J., Fleming, J. & Johnson, A. F. (2013): Building the ADCIRC Collaborative. RENCI, University of North Carolina at Chapel Hill. Text. http://dx.doi.org/10.7921/G0RN35SG

renci
RESEARCH ENGAGEMENT INNOVATION

THE UNIVERSITY *of* NORTH CAROLINA *at* CHAPEL HILL