

Featured Application:

Storm-Surge Modeling with ADCIRC

ABSTRACT: This application is a finite element method shallow water model for computing tidal and storm surge water level and depth-averaged currents associated with these phenomena. Grid environments such as SURAgird are ideal for the ensembles in applications like ADCIRC, which is one of a set included under SCOOP (SURA Coastal Ocean Observing and Prediction) program that is focused on improving predictions of coastal phenomena. Project Partners: University of North Carolina Marine Science, Renaissance Computing Institute (RENCI), MCNC, SAIC

Understanding the Science

Devastating weather events such as Hurricane Katrina have underscored and made the need for increased forecast accuracy all the more urgent. Scientists and information technologists have long been working together to meet this need in an area that includes numerous forecast models, a large range of variables to be measured and analyzed, and a diverse community of organizations working on particular pieces of the weather forecasting puzzle. Among these organizations is the North Carolina SCOOP¹ team, which is comprised of the RENAISSANCE Computing Institute (RENCI)², MCNC³ and the University of North Carolina (UNC) Institute of Marine Sciences⁴. The NC SCOOP team is working toward an integrated, distributed data management system that can handle real-time data feeds, schedule and execute a set of model runs, manage the model input and output data, and make large scale model results and status available to a larger audience through a web-based portal.

The NC SCOOP team has built a distributed software infrastructure to run a storm surge (a coastal phenomena driven by extra-tropical cyclones and hurricanes) model in a grid environment. The solution builds on existing standard grid and portal technologies including the Globus toolkit, Open Grid Computing Environment (OGCE) and lessons learned from grid computing efforts in other science domains. Specifically, the team has implemented specific techniques for resource management and increased fault tolerance due to the sensitivity of the application. This framework was developed as a component of Southeastern Universities Research Association's (SURA) Southeastern Coastal Ocean Observing and Prediction (SCOOP)¹ program, which is creating an open-access grid environment for the

southeastern coastal zone to help integrate regional coastal observing and modeling systems.

Application Characteristics

ADCIRC is at the core application of the infrastructure in NC SCOOP's Storm-Surge Ensemble Prediction System. The Advanced Circulation Model (ADCIRC) is a finite element method shallow water simulation model for computing tidal and storm surge water level and depth-averaged currents associated with these phenomena. To increase potential forecast accuracy, it is useful to run storm-surge models with different forcing conditions (i.e., wind fields), all of which constitute an ensemble run. The NC SCOOP Storm-Surge modeling application consists of a group of ADCIRC simulations driven by an ensemble of wind fields that represent a forecast of a current North Atlantic basin tropical storm or hurricane. The wind fields are generated by other SCOOP partners and distributed as files that contain the hurricane pressure and winds. Other wind sources include the National Centers for Environmental Prediction North American Mesoscale forecast run and their hurricane-specific forecast run. When the wind fields arrive, the ADCIRC workflow is triggered

Recently, large-scale ocean and meteorological modeling has resulted in the use of grid resources and high performance environments for running these models. Ensemble simulation runs, such as those used in ADCIRC, typically require a large amount of compute resources and are thus ideal for deployment in grid environments like SURAgird. By grid-enabling ADCIRC, not only is the NC SCOOP team is able to take advantage of "compute on demand" cycles, their using grid technologies to help them find compute resources to run the jobs on based on factors such as the compute load on available resources, the number of CPUs available, the resolution of the model's FEM and the number of ensemble members. This granularity is particularly important since the ADCIRC workflow is not manually initiated by human input but rather by incoming wind fields for a current North Atlantic storm or hurricane (as supplied by other SCOOP partners), so effective automated execution is key.

¹ SCOOP is a distributed project that includes: Gulf of Maine Ocean Observing System, Bedford Institute of Oceanography, Louisiana State University, Texas A&M, University of Miami, University of Alabama in Huntsville, University of North Carolina, University of Florida and Virginia Institute of Marine Science. <http://scoop.sura.org/>

² <http://www.renci.org/>

³ <http://www.mcnc.org/>

⁴ <http://www.marine.unc.edu/>

SURAgriD Deployment

The NC SCOOP team wanted to use more resources than they had available to them through MCNC and UNC. More resources mean they can retrieve data from ADCIRC model runs more quickly and can also run more models and at higher resolutions. SURAgriD was selected as one of NC SCOOP's additional grid resource environments and the NC SCOOP team has found that the additional resources they're now using result in ensemble runs with a higher predictive value.

Grid Workflow

There are two modes of operation to run ADCIRC jobs on SURAgriD resources (as well as other compute resources). The ADCIRC storm surge model⁵ is run in hindcast mode, which is one of the two modes⁶ the model can be run in. In hindcast mode, the modeler can either use an OGCE-based portal or a command line shell interface to launch the jobs to investigate prior hurricane data sets. While many SURAgriD resource and application owners participate in the Bridge Certificate Authority-based SURAgriD authentication and authorization trust fabric, such participation is not a requirement. Instead SURAgriD users, like those at NC SCOOP, can choose to manually exchange their user and site certificates with each SURAgriD resource site where they want to run jobs. SURAgriD resource administrators at each of site supporting ADCIRC follow a configuration process⁷ to ready their resource to support ADCIRC runs. The configuration includes the following Globus components: Globus gatekeeper for job submission, GridFTP for file transfer, PBS or LSF as the local scheduler, MDS for queue information sharing, and NWS for network information.

Once the NC SCOOP user has authenticated with a properly configured SURAgriD resource, the ADCIRC workflow proceeds as follows: all the data for a given ensemble member is received (e.g., wind files) and available grid resources are then discovered using a resource selection algorithm. Once the appropriate resources are located, the files for each ADCIRC simulation (e.g., wind, grid definition, initial conditions) are moved to the appropriate resource, the model run is executed and the output data is staged back to the user's site.

An important component of the ADCIRC model's architecture is the simple Java API algorithm the team developed using standard grid tools (such as Globus 3.2.1 and Java Cog). This algorithm locates the best resources for simulation runs based on the availability and current load of the grid resources. Additional components of the architecture include a central Application Coordinator that

⁵ For full model details and more complete grid component descriptions, see SCOOP Storm Surge Model

<http://www.videnet.gatech.edu/gtcb/draft/related-links/StormSurge.php>

⁶ The model has also been configured to automatically run on its own in "forecast" mode when it is triggered by real-time data arrival.

⁷ <https://www.ccs.uky.edu/scoop/>

orchestrates the data and job management actions and interacts with the Globus services and an application preparation component that readies the application bundle (tarball) that needs to be used on a remote resource (e.g., binary, input and other initialization files).

Lessons Learned

There were several the issues the NC SCOOP team initially encountered when working with SURAgriD resource sites to configure resources for ADCIRC, and most of these issues have since been documented and published⁷. Some remaining issues are being considered for inclusion in future changes to the ADCIRC model, including ways to ease configuration of pre-WS MDS and fuller model integration with Globus 4.x and above.

For more information about Storm-Surge Modeling with ADCIRC, contact:

Lavanya Ramakrishnan (RENCI) - lavanya@renci.org
Kenneth Galluppi (RENCI) – galluppi@renc.org

Collaborators


Lavanya Ramakrishnan, Renaissance Computing Institute
Brian O. Blanton, Science Applications International Corporation
Kenneth Galluppi, Renaissance Computing Institute
Gopi Kandaswamy, Renaissance Computing Institute
Howard M. Lander, Renaissance Computing Institute
Richard A. Luettich, Jr, UNC Chapel Hill Institute of Marine Sciences
Daniel A. Reed, Renaissance Computing Institute
Steven R. Thorpe, MCNC

Acknowledgements

This framework was developed as a component of Southeastern Universities Research Association's (SURA) Southeastern Coastal Ocean Observing and Prediction (SCOOP) program. The SCOOP program is a distributed project that includes Gulf of Maine Ocean Observing System, Bedford Institute of Oceanography, Louisiana State University, Texas A&M, University of Miami, University of Alabama in Huntsville, University of North Carolina, University of Florida and Virginia Institute of Marine Science. SCOOP is creating an open-access grid environment for the southeastern coastal zone to help integrate regional coastal observing and modeling systems.

Funding Support

SCOOP funding support has been provided by the Office of Naval Research, Award N00014-04-1-0721 and by the National Oceanic and Atmospheric Administration's NOAA Ocean Service, Award NA04NOS4730254. The NC SCOOP team thanks SCOOP team members: Philip Bogden (SURA, GoMOOS); Will Perrie, Bash Toulany (BIO); Eric Bridger (GoMOOS); Gabrielle Allen and Greg Stone, Jon McClaren, Bret Estrada, Chirag Dekate (LSU, Center for Computation and Technology); Gerald Creager and Larry Flournoy, Wei Zhao, Donna Cote and Matt Howard (TAMU); Sara Graves, Helen Conover, Ken Keiser, Matt Smith, and Marilyn Drewry (UAH); Peter Sheng, Justin Davis, Renato Figueiredo, and Vladimir Paramygin (UFL); Harry Wang, Jian Shen and David Forrest (VIMS); Hans Graber, Neil Williams and Geoff Samuels (UMiami); Charlton Purvis; and Mary Fran Yafchak, Don Riley, Don Wright and Joanne Bintz (SURA). Thanks are also extended to the various SURAgriD members including Steven Johnson (TAMU), Ian Chang-Yen (ULL), and Michael McEniry (UAH) for making resources available for the SCOOP grid and providing valuable system administrator support.

 is developing an infrastructure for access to heterogeneous resources and support for a dynamic and diverse application set.

For more information, or to join SURAgriD:

- <http://www.sura.org/SURAgriD>
- maryfran@sura.org