

Unidata Community Equipment Awards Cover Sheet

Proposal Title:

Computer Hardware Support for Incorporating Unidata Technologies into a Storm Surge Prediction System for Guiding Decision-Making under Present and Future Climate Scenarios

Date: March 22, 2010

Principal Investigator Name: Craig Mattocks

Title: Dr./Research Associate/Atmospheric Research Scientist

Institution: University of NC at Chapel Hill, Institute for the Environment

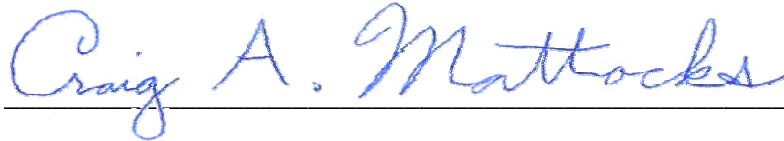
Telephone number: 919-966-9922

FAX number: 919-966-9920

Street Address: 337 W. Rosemary Street, Chapel Hill, NC 27599-1105

Email address: cmattock@email.unc.edu

Signature of PI:



Craig A. Mattocks

Name of Institution Official: Tony Waldrop, Ph.D.

Title: Vice Chancellor for Research and Economic Development

Telephone number: 919-966-3411

FAX number: 919-962-3352/5011

Email address: resadminosr@unc.edu

Signature of University Official:

B. Project Summary

This proposal seeks computer hardware to support the integration of key Unidata technologies (RAMADDA, LDM, IDV) into two numerical modeling projects at UNC's Institute for the Environment: (1) a real-time, event-triggered storm surge prediction system that was developed for the state of North Carolina to assist emergency managers, policy makers and other government officials with evacuation planning, decision-making and resource deployment during tropical storm landfall, and (2) a climate-driven coastal flood modeling system (CFMS) that will help populations adapt to the threat of sea level rise (SLR) during tropical cyclone events under current and future global climate warming scenarios by planning and building sustainable, hazard-resilient coastal communities.

These numerical modeling systems generate unique, ultra-high resolution output datasets that not only contribute to the advancement of meteorology and oceanography but would be useful to a diverse array of scientific disciplines (hydrology, geography, ecosystem studies, estuarine biology, habitat restoration, societal impact studies) and government agencies (emergency management, land use planning, wildlife management, park services, departments of transportation, water, sewer and communication systems). Unidata's technologies will enable this data to be shared to a broad user community through Internet Data Distribution (IDD).

A single hurricane simulation from UNC-IE's storm surge prediction system, which has been running quasi-operationally during the last four hurricane seasons, produces large volumes of output, in both numerical (ASCII text and NetCDF output files) and graphical formats (Google Earth kmz files, graphics images, web-based animations). Examples can be viewed on the OLAS (Ocean Land Atmosphere System) testbed web site (<http://www.unc.edu/~cmattock/olas/web/olas.html>). Multiple simulations will be run to construct forecast ensembles from idealized tropical cyclones that approach the coastline from a range of directions at different speeds and intensities in order to construct storm catalogs for present and future climate scenarios based on projected sea level rise in the decades 2010, ..., 2100. Therefore, a graphical user interface (GUI) for human browsing, searching, and querying of data stored in SQL databases and managing the associated metadata will be essential to enable users to quickly and easily retrieve information crucial to guiding decision-making processes. Unidata provides all of these capabilities in the RAMADDA (Repository for Archiving, Managing and Accessing Diverse DATA) package, a tool designed to provide an open and extensible data repository framework. RAMADDA will be deployed on the Unidata-funded computer system at UNC-IE and customized to serve OLAS and CFMS simulation results through a web browser-based interface.

C. Project Description

The storm surge prediction system's suite of coupled models consists of the ADCIRC (ADvanced CIRCulation) coastal ocean circulation model, the Advanced Hurricane version of the WRF (Weather Research and Forecasting) numerical weather prediction model, and the GWAVA (gradient wind asymmetric vortex analysis) wind model. ADCIRC (Luettich et al. 1996) is a finite-element hydrodynamic model used to simulate wind-driven storm surge, tides, riverine flow, inundation and sea level rise. It employs a wetting and drying methodology that allows wave propagation and retreat over land. The unstructured, triangular grid includes all

waters in the western Atlantic, Caribbean and Gulf of Mexico. Several high-resolution (less than 30 m) grid meshes (Fig. 1) are available for ADCIRC, with computational points draped across inlets and waterways, aligned with shoreline and elevation contours. The unstructured-mesh version of the SWAN (Simulating WAVes Nearshore) spectral wave model has recently been integrated with ADCIRC (Dietrich et al. 2010), so simulation of wave-circulation interactions is now possible. An example of the maximum water surface elevations predicted by ADCIRC during Hurricane Ophelia (2005) is shown in Fig. 2.

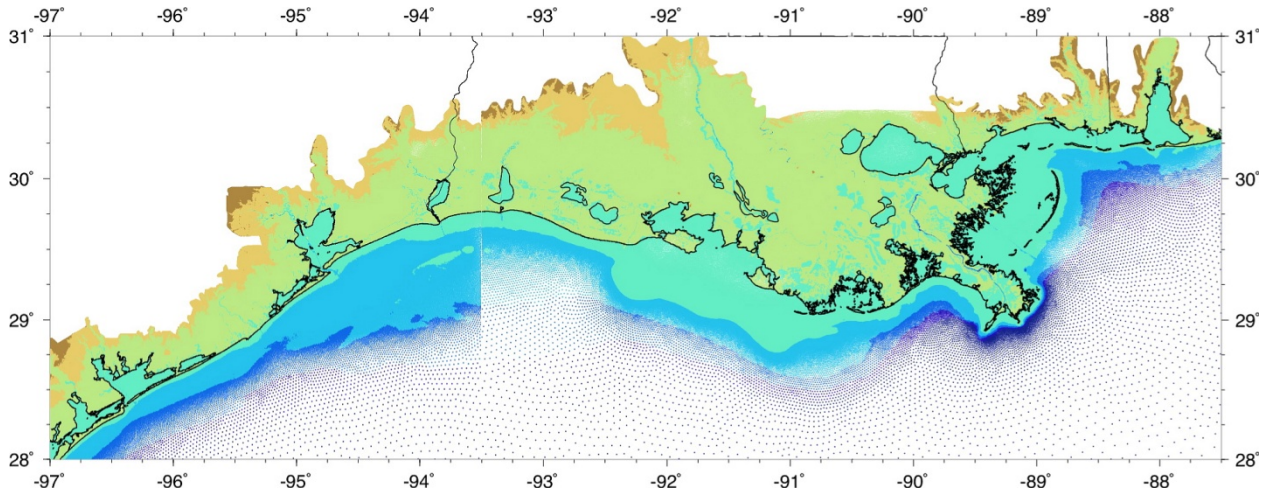


FIG 1. High-resolution Texas (left) and Louisiana SL15 (right) grid domains used in ADCIRC simulations of storm surge over the Gulf of Mexico. Each dot represents a grid node in the finite-element mesh, with colors indicative of bathymetry/topography.

GWAVA (Mattocks and Forbes 2008) is based on the Holland (1980) gradient wind model, with the added feature that the radius of maximum winds varies azimuthally around the cyclone to capture the asymmetry in the shape of the storm. A cross-isobar inflow angle that varies as a function of distance from the center of the storm is applied to represent surface friction. Since these parametric winds are generated on-the-fly from National Hurricane Center (NHC) forecast advisory or official/best track information in an extremely computationally efficient manner and are available at exact analytical resolution and true intensity, they can be directly coupled to an atmosphere, ocean or climate model at every time step and grid point while the model is running. This model is currently deployed in the OLAS testbed at UNC-IE to generate wind forcing for real-time storm surge predictions of tropical cyclones (Fig. 3).

WRF (Skamarock et al. 2005) is a next-generation mesoscale numerical weather prediction (NWP) system. The Advanced Hurricane WRF-ARW (Davis et al. 2008), a moving-nest vortex-tracking version of WRF, which includes drag saturation at high wind speeds and a one-dimensional columnar, mixed layer ocean model to more accurately simulate vertical momentum and heat exchange, will be used in the CFMS. A snapshot of a wind field from an Advanced Hurricane WRF simulation of Hurricane Ike (2008) is shown in Fig. 4.




FIG. 2. Maximum water elevation (m) over Pamlico Sound, NC during Hurricane Ophelia (2005), as simulated by the ADCIRC storm surge prediction model driven by winds from the GWAVA wind model.

The OLAS operational workflow is depicted in Fig. 5. One long tidal spin-up simulation of at least 60 days (150 days preferred) in length is required prior to each hurricane season to provide “hot start” capability. Thereafter, the system operates in three modes: 1) daily background (tides only) mode, in which the tides are updated every 24 hours, 2) event nowcast (current state) mode, in which a wind field is applied and the system is gradually ramped up to full strength to allow spurious transient oscillations generated by initialization shock to dissipate and a balanced flow to be attained, and 3) event forecast (future state) mode, which is switched on when a tropical cyclone threatens the U.S. coastline and 5-day predictions are launched four times per day. In addition, one pure 10-day tidal simulation is run at the onset when a storm first enters the ADCIRC grid domain to allow the calculation of storm surge from the total water surface elevation (storm surge = total water elevation – astronomical tide).

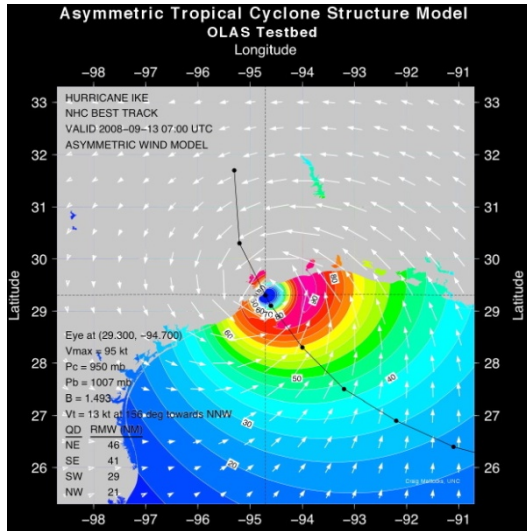


FIG 3. Wind speed (kt) at landfall generated by the GWAVA asymmetric wind/tropical cyclone structure model from NHC best track/forecast advisory information during Hurricane Ike (2008).




FIG 4. Snapshot of winds (ms-1) from a simulation of Hurricane Ike (2008) using the moving-nest, vortex-tracking Advanced Hurricane version of the WRF numerical weather prediction model. This simulation incorporates temporally varying 1 km SPoRT MODIS SST analyses at its lower boundary.

C.1 Key Unidata Technologies

The following key Unidata technologies will be employed to accomplish these projects' goals.

C.1.1 NetCDF

NetCDF (network Common Data Form) is a set of software libraries and machine-independent data formats that support the creation, access, and sharing of array-oriented scientific data. Both the WRF numerical weather prediction model and the ADCIRC coastal ocean model generate output in NetCDF format. NetCDF's cross-platform portability is ideal for operational applications because it allows check-pointing (restarting or hot-starting) of numerical simulations across multiple supercomputing clusters in the case of a computer failure or when job submittal queues are saturated. It also allows the pre-processing of model initial/boundary conditions, model execution, and post-processing of model output to be conducted on separate computer systems.

C.1.2 Local Data Manager (LDM)

Designed for event-driven data distribution, the Unidata Local Data Manager (LDM) is a collection of cooperating software programs that select, capture, manage, and distribute data products. The system, which is used in the Unidata Internet Data Distribution (IDD) project, acquires data and shares it with other networked computers. How specific types of data passed to the LDM server are processed is determined by data identifiers and a configuration file. Processing actions include capturing the data in files and running programs on the data.

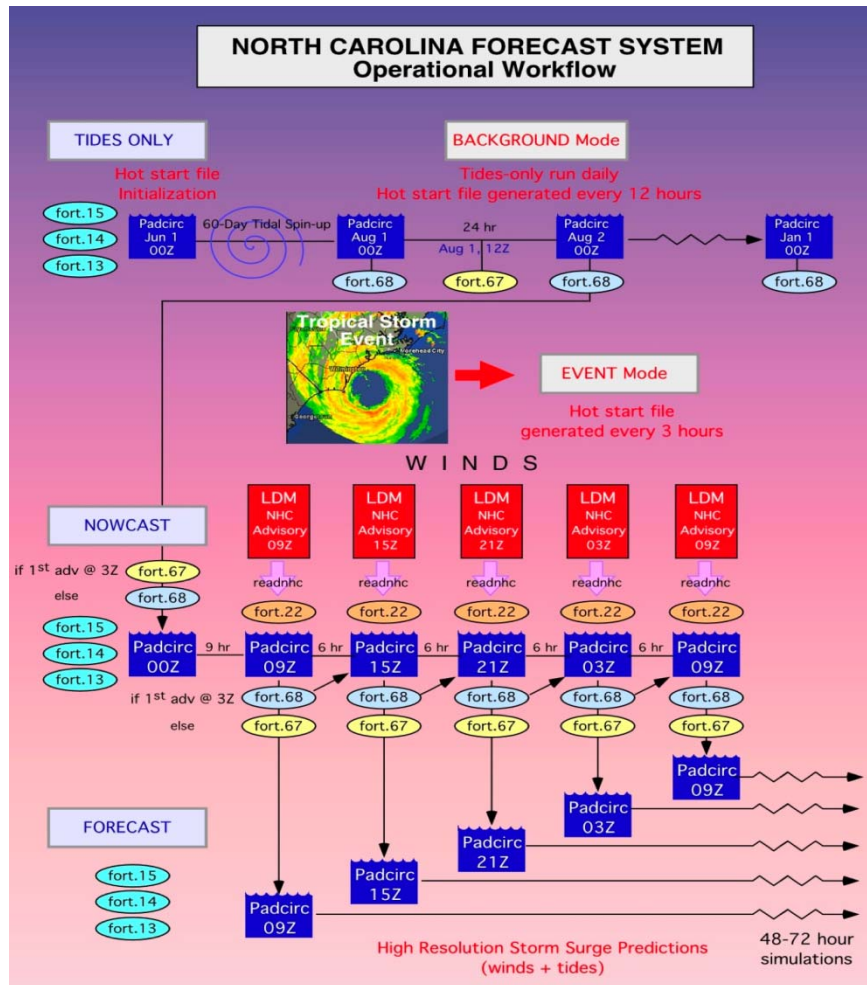


FIG. 5. Schematic depiction of the NCFS operational workflow, showing its three different modes of operation.

The NHC forecast advisory data is transmitted directly through the LDM in real-time. These forecast advisories are fed to the LDM within 1-2 minutes after the Hurricane Specialists at NHC “hit the button” that submits the forecast to the National Weather Service (NWS) telecommunications gateway. Sent through AWIPS, then placed on the NOAAPORT, they are transmitted through the high-bandwidth Family of Services (FOS) Domestic Data Plus stream to subscribed LDMs. In order to capture this data, a line containing a regular expression will be inserted into the LDM's pqact configuration-file to detect the AWIPS and World Meteorological Organization (WMO) headers that identify these products as soon as they enter the data stream, then the LDM PIPE command will be used to automatically start up a unix shell script that preprocesses this data and initiates the storm surge workflow.

C.1.2 IDV and the VisAD OpenGL Visualization Engine

The Integrated Data Viewer (IDV) is a geoscience display and analysis software system. It brings together the ability to display and work with satellite imagery, gridded data (for example, numerical weather prediction model output), surface observations, balloon soundings, NWS WSR-88D Level II and Level III RADAR data, and NOAA National Profiler Network data, all

within a unified interface. It also provides 3-D views of the earth system and allows users to interactively slice, dice, and probe the data, creating cross-sections, profiles, animations and value read-outs of multi-dimensional data sets. The IDV can display any Earth-located data if it is provided in a known format. It uses the VisAD (Visualization for Algorithm Development) library, a Java component library for interactive and collaborative visualization and analysis of numerical data. Users will be able to launch IDV through a hyperlink in the storm surge portal to a Java Web Start JNLP file with a single mouse click, then display, analyze and create animations from datasets in the vast catalog of storm scenarios.

C.1.3 RAMADDA

The Repository for Archiving, Managing and Accessing Diverse Data (**RAMADDA**) is a Unidata package that provides a suite of comprehensive data management, archiving and repository services. It is based on a relational database and supports browsing and searching over metadata. It also supports data file ingest and organization, metadata creation, data access control, search and browse capabilities. A high-level software architecture diagram for RAMADDA is presented in Fig. 6.

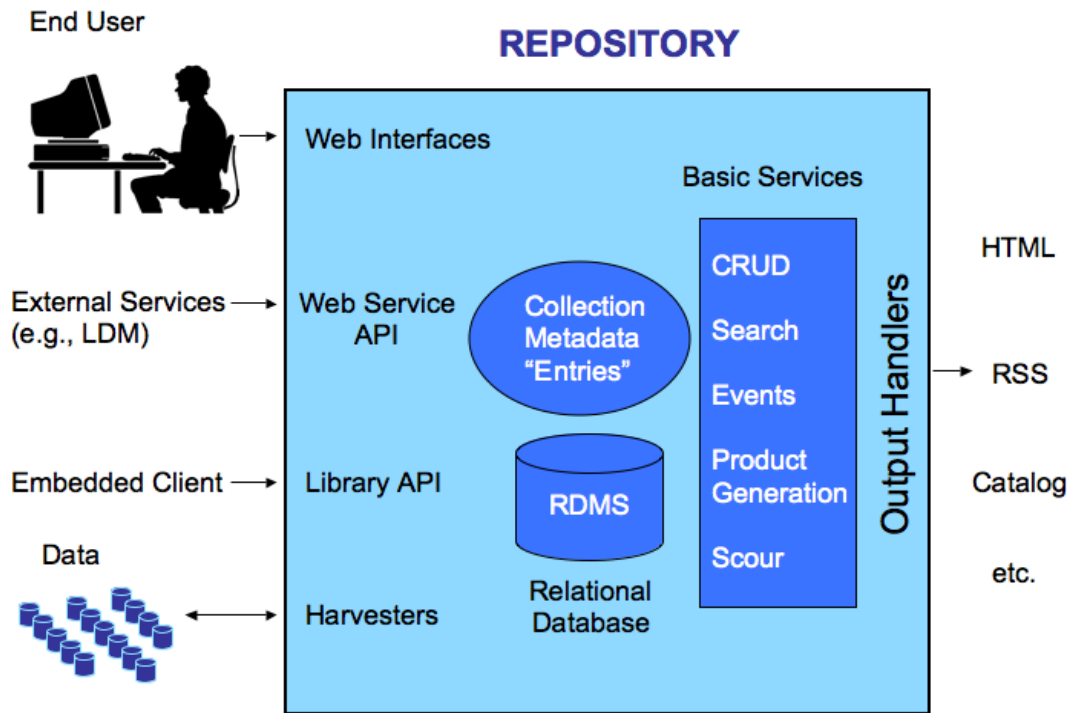


FIG. 6. High-level software architecture diagram for RAMADDA (Repository for Archiving, Managing and Accessing Diverse DATA). The goal of this Unidata project is to provide an open and extensible data repository framework/server. RAMADDA is Java-based and uses a relational database as its underlying metadata storage facility. It has been designed to run both stand-alone as well as running under a Tomcat server.

This technology will enable the generation of sophisticated search tools that can also generate new data products. For example, one could create a configurable search GUI that constructs an ensemble of idealized hurricane tracks from weighted combinations of input parameter perturbations (maximum sustained winds, radius of maximum winds, translation speed, Holland B shape parameter, SST, etc.) for one specific storm or derived from cumulative density

functions (CDFs) computed from the complete historical record of tropical cyclones. These tracks could then be used to drive the storm surge forecasting system.

D. Budget

The computer system on which the RAMADDA storm surge portal and LDM software (Table 1) will be installed was configured on the University of North Carolina’s RAM Shop web site (http://store.apple.com/us_educ_24729/), which offers Apple equipment at an educational discount of approximately 10%. Since this purchase will exceed \$5,000, no indirect costs will be charged. Substantial cost savings (\$2,175) will also be achieved by purchasing the 32 GB RAM memory upgrade from a 3rd party vendor (Other World Computing). It is likely that, by the time this proposal is funded, Apple will have upgraded this to a dual hexacore (12-core, 24-thread) Xeon 5600 Gulftown (i7-980X) system running at 3.33 GHz. Ample disk space (two 2TB drives and a RAID controller) is requested to support rapid access to large model output datasets through the RAMADDA interface and to insert these unique datasets into downstream LDM feeds for interested users.

TABLE 1. Computer system that will host the RAMADDA storm surge portal and the LDM real-time weather data stream.

Item	Cost
Apple Mac Pro Dual 2.93GHz Quad-Core Intel Xeon 5500 Nehalem	\$5339
Mac Pro RAID Card	\$630
Two 2TB 7200-rpm Serial ATA 3Gb/s	\$810
4x NVIDIA GeForce GT 120 512MB	\$405
One 18x SuperDrive	–
Apple Cinema HD Display (30" flat panel)	\$1,599
Apple Keyboard with Numeric Keypad (English) and User's Guide	–
Apple Mouse	–
AirPort Extreme Wi-Fi Card with 802.11n	\$45
AppleCare 3-Year Protection Plan for Mac Pro	\$199
32 GB DDR3 1066 RAM upgrade (3rd party, ramseeker.com)	\$1,155
Total	\$10,182

The project PI, who served as Chief of Computer Operations at the National Hurricane Center, has extensive training and experience as a unix systems administrator. He will donate his time to maintain the computer system, install and configure the Unidata software, and monitor the LDM real-time weather data stream. In addition, UNC’s Information Technology Services (ITS) provides expert staff to help troubleshoot computer system problems, install software, and automatically backup the contents of hard disks to a mass storage silo with a capacity of 175 TB.

Substantial UNC-IE computing resources will be leveraged for use in this project. The PI recently received a grant of 80,000 Service Units (CPU hours) from the National Science Foundation's TeraGrid for running simulations on NCAR Frost (<http://tinyurl.com/yflazap>), TACC Ranger (<http://tinyurl.com/yl4qcco>), NCSA Abe (<http://tinyurl.com/ygljcm9>) and NICS Kraken (<http://tinyurl.com/yg443qe>), plus 10 TB of data storage Indiana University's High Performance Storage System (HPSS). The numerical weather prediction model (WRF) and the coastal ocean model (ADCIRC) will be run on these systems and the Topsail Linux supercomputing cluster (4168 processors) at UNC-CH. Complete details of UNC-ITS resources can be found at <http://its.unc.edu/Research/hardware/index.htm>.

One crucial role that the Unidata-granted computer system will play is scanning the LDM real-time weather data stream for NHC forecast advisories and spawning event-triggered storm surge forecasts. This capability has been tested and employed in collaborative projects with other institutions, but it is not currently available at UNC-IE.

Another capability that is essential to the storm surge prediction and coastal flood model system workflows, but is not allowed on most supercomputing clusters due to security restrictions, is running scheduled, time-critical cron-driven unix scripts. For example, daily tidal simulations must be run to keep the initial and boundary conditions up-to-date, in anticipation of a tropical cyclone developing or entering the model grid domain. Also, the data pre-processing, model execution, post-processing, graphics image generation, and web site update stages must be carefully synchronized to successfully run the workflows to completion. This will be coordinated on the Unidata-awarded computer system, which will serve as the control center for launching jobs on remote supercomputing clusters and monitoring the workflows.

E. Project Milestones

1. **(Month 1)** The proposed computer equipment will be purchased immediately after receipt of the award from Unidata.
2. **(Month 1)** At the outset of the project, the PI will attend Unidata software training courses offered annually in Boulder, CO (travel costs will be paid by a complementary NASA grant).
3. **(Months 2-3)** The LDM will be installed on the Unidata-funded computer at UNC-IE and configured to receive NHC forecast advisories and trigger real-time storm surge forecasts.
4. **(Months 2-3)** RAMADDA will be deployed on the Unidata-funded computer to serve numerical model results (netCDF format) and metadata to interested members of the scientific community.
5. **(Month 4)** Unidata's VisAD-based Integrated Data Viewer (IDV), will be integrated into the RAMADDA storm surge portal to enable visualization of WRF and ADCIRC numerical model output.
6. **(Months 5-10)** Products will be deployed in RAMADDA as the objectives of the U.S. Navy and NASA sponsored projects are achieved. These milestones include: 1) the implementation of new algorithms in the numerical models and workflows to more accurately simulate hurricanes and storm surge, 2) the generation of baseline hindcast simulations of present day hurricanes and their validation, 3) reconfiguration of WRF and ADCIRC, then running simulations of future climate scenarios for the decades 2010, ..., 2100, 4) conducting a thorough system evaluation, and 5) documentation of forecast system performance. These tasks will be updated each year with the incorporation of more scenarios

into the storm catalog.

7. **(Month 11)** The RAMADDA storm surge portal will be opened to the broader scientific, emergency management, and local/state government communities.
8. **(Month 12)** Major revisions of the RAMADDA graphical user interface will be conducted annually after receiving feedback from end users (e.g., emergency managers, habitat restoration specialists, navigation engineers, climate/resilience specialists, and nutrient/hypoxia experts from the Mississippi/Alabama Sea Grant Consortium) during scheduled demos, workshops, and training sessions held at the NASA-Marshall SPoRT facility.
9. **(Month 12)** A brief final report, in the form of a short article, will be submitted to the Unidata Program Center. It will describe: 1) how the computer hardware was employed, 2) the key capabilities that Unidata software provided, 3) the benefits of participating in the Unidata Equipment Awards program, 4) how working closely with Unidata accelerated the development and deployment of the storm surge portal, 5) the impact of delivering unique ultra-high resolution numerical model datasets to the educational, research, emergency management, and ecosystem restoration communities, and 6) assessments on approaches that worked/failed in order to improve the program.

F. References

- Davis, C., W. Wang, S.S. Chen, Y. Chen, K. Corbosiero, M. DeMaria, J. Dudhia, G. Holland, J. Klemp, J. Michalakes, H. Reeves, R. Rotunno, C. Snyder, and Q. Xiao, 2008: Prediction of landfalling hurricanes with the Advanced Hurricane WRF model. *Mon. Wea. Rev.*, **136**, 1990-2005.
- Dietrich, J.C., S. Bunya, J.J. Westerink, B.A. Ebersole, J.M. Smith, J.H. Atkinson, R. Jensen, D.T. Resio, R.A. Luetlich, C. Dawson, V.J. Cardone, A.T. Cox, M.D. Powell, H.J. Westerink, H.J. Roberts, 2010: A high-resolution coupled riverine flow, tide, wind, wind wave and storm surge model for Southern Louisiana and Mississippi: Part II - Synoptic description and analysis of hurricanes Katrina and Rita, *Mon. Wea. Rev.*, (in press).
- Holland, G.J., 1980: An analytical model of the wind and pressure profiles in hurricanes. *Mon. Wea. Rev.* **108**, 1212–1218.
- Luetlich, R. Jr., J. Hudgins and C. Goodall, 1996: Initial results from a combined tide and storm surge forecast model of the U.S. East Coast, Gulf of Mexico and Caribbean Sea. *Proceedings, 15th Conference on Weather Analysis and Forecasting*, American Meteorological Society, Norfolk, VA., August 1996, 547-550.
- Mattocks C. and C. Forbes, 2008: A real-time, event-triggered storm surge forecasting system for the state of North Carolina, *Ocean Modelling*, **25**, 95-119, doi:10.1016/j.ocemod.2008.06.008.
- Skamarock, W. C., J. B. Klemp, J. Dudhia, D. O. Gill, D. M. Barker, W. Wang, and J. G. Powers, 2005: A description of the Advanced Research WRF Version 2. NCAR Tech Notes-468+STR.