Designing a Dynamic Data Driven Application System for Coastal and Environmental Modeling

Gabrielle Allen[1], Philip Bogden[2], Richard A. Luettich, Jr.[3], Edward Seidel[4], and Robert Twilley[5]

[1] Center for Computation & Technology and Department of Computer Science, Louisiana State University gallen@cct.lsu.edu
[2] Gulf of Maine Ocean Observing System bogden@gomoos.org
[3] University of North Carolina at Chapel Hill, Institute of Marine Sciences rick_luettich@unc.edu
[4] Center for Computation & Technology and Department of Physics, Louisiana State University eseidel@cct.lsu.edu
[5] Department of Oceanography and Coastal Science, Louisiana State University rtwilley@lsu.edu

1 Abstract

The economically important Louisiana Coastal Area (LCA) is susceptible to hurricane activity which is increasingly aggravated by the continuing erosion of wetlands. Various programs are aimed at building sophisticated models of meteorological, coastal, and ecological processes. The emerging paradigm of Dynamic Data Driven Application Systems (DDDAS) can be applied to these models leading to new scenarios for integrated, real-time simulations that include feedback control with sensors and simulations. This paper describes the motivation and components for a comprehensive DDDAS for coastal and environmental modeling and the implications this has for scientific libraries and high performance computing.

2 Introduction

The Louisiana Coastal Area (LCA) is one of the world's most environmentally damaged ecosystems. In the past century nearly one-third of the wetlands in the LCA, approximately 4500 km^2, has been lost. In its current state of decay, it is expected that by 2050 only one-third of the wetlands in this region will remain. Further erosion of the wetlands will be not only a massive environmental catastrophe, but also adversely affect the livelihoods of millions of people.

The economic importance of the LCA is substantial, containing 25% of U.S. coastal wetlands; 40% of contiguous U.S. salt marshes; 30% of U.S. total fish catch; 17% of U.S. oil. Further, 25% of U.S. national natural gas supplies come from Gulf coast waters. LCA ports handle more tonnage than any port worldwide, and LCA is home to the only "superport" in the contiguous U.S.

Beyond economic loss, LCA erosion has devastating effects on inhabitants of the coastal region, especially in New Orleans. This culture-rich city and its surroundings includes over one million people, and lies in a basin 6 meters under sea level, making it extremely vulnerable to hurricanes and tropical storms. On 29th August 2005 Hurricane Katrina (Figure 1) hit New Orleans, with storm surge and flooding resulting in a tragic loss of life and destruction of property and infrastructure. Soon after, Hurricane Rita caused similar devastation in the much less populated area of southwest Louisiana, and once again parts of New Orleans were under water. In both cases entire communities were destroyed.
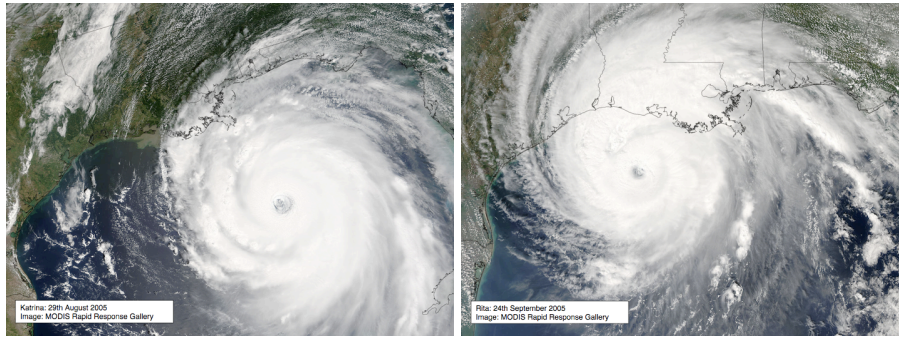

Fig. 1. From left to right, satellite images of Hurricanes Katrina and Rita which made dramatic landfall on the southeast US coast in 2005. Katrina resulted in the loss of nearly 2000 lives and caused some $120 billion of property damage. The storm size at landfall was 460 miles, with 145mph winds (Category 3), and storm surges of up to 22 feet. [Image credits: MODIS Rapid Response Gallery]

The entire LCA is very amenable to comprehensive computer modeling. Today there exist many sophisticated models for this area with processes ranging from ocean circulation and wave propagation to sediment transport and salinity in wetlands. Typically these are isolated models, and currently there is no established standard for tools that forecast an ecosystem response to restoration projects. A generic approach that encompasses the entirety of a restoration project's goals, capable of simulating all relevant interacting processes from erosion to storm surge to biodiversity in the ecosystem, is needed. This framework would need the capability to integrate data from observation and sensor systems and a complete set of detailed computer models, With such a system, more sophisticated decisions for short and long term planning and policy, as well as for immediate emergency response can be made. Decision making would be improved by using algorithms based on both data and detailed, realistic models that can dynamically adapt to real-life scenarios.

An emerging paradigm in computational science that seeks to integrate data streams and computational models, adapting them as needed to control or provide critical information about the behavior of complex systems is called Dy-

namic, Data-Driven, Application Systems, or DDDAS [1]. DDDAS is finding use in a vast range of complex problems, such as manufacturing supply chains in business, forest fire control [2], or combustion in engines. All these problems have complex systems that can in principle be simulated, and whose response to inputs can be predicted. These responses can then be used to make decisions that may affect or control the behavior of the system itself, or of other systems that depend on it.

The LCA makes an ideal case-study for DDDAS capabilities, as coastal properties change on relatively rapid time-scales. Unlike most ecosystems, due to extensive levee, dam, channel, and controllable water diversion projects over the last century, it is possible to regard the LCA as an *experimental* system. The ecosystem can be, and is regularly, changed and regulated on time scales of hours and days, providing unique opportunities for real-time data acquisition, simulation, monitoring, and control.

To effectively model the LCA region, a much more comprehensive and dynamic approach than currently available is needed. This includes the ability to couple models, invoke dynamic algorithms based on streams of sensor and satellite data, locate appropriate data and computational resources, and create necessary workflows on demand, all in real-time. Such an environment could help better plan restoration strategies, improve ecological forecasting, placement of future sensors, control of water diversion for salinity control, or predict/control harmful algal blooms, and support sea rescue, oil spill response, and shipping forecasts. In extreme situations, such as approaching hurricanes, results from multiple coupled ensemble models, compared with observations, can be used for greatly improved emergency warnings. Input from sensors and control of dams can be optimized to both improve the forecasts and actually reduce flooding. The comprehensive modeling system must be able to handle multiple time and length scales, from hours (storms) to years (restoration) and from meters (estuaries) to kilometers (Gulf of Mexico).

This paper describes the motivation and components for a comprehensive DDDAS for coastal and environmental modeling. Section 3 provides an overview of DDDAS including two detailed application scenarios, Section 4 describes different application communities that provide the scientific and operational background for DDDAS systems, and Section 5 details requirements and needs for the framework.

3 Dynamic Data Driven Application Systems

Simulation codes used today, across the physical and engineering sciences, typically allow only static workflows. Input data and parameter files must be created in advance, and are read by the simulation code at start up, and after this point the user can no longer interact with the running code, except to terminate a run. Integration with observing systems, data archives, and experiments is usu-

ally done manually if at all, through static, simplified input files, derived from data archives.

DDDAS describes new complex, and inherently multidisciplinary, application scenarios where simulations can dynamically ingest and respond to real-time data from measuring devices, experimental equipment, or other simulations. In these scenarios, simulation codes are in turn also able to control these varied inputs, providing for advanced control loops integrated with simulation codes. Implementing these scenarios requires advances in simulation codes, algorithms, computer systems and measuring devices.

In 2000 the National Science Foundation held a DDDAS Workshop [3], which included numerous application scenarios which could advance both science and society by incorporating these ideas. Application areas described at the workshop included control of forest fires, predicting the spread of contaminants, improving transportation systems and supply chains, and enabling oil exploration. Recent developments in cyberinfrastructure, including data archives, information systems, metaschedulers, etc, have helped provide surrounding infrastructure needed to implement some of these ideas in practice, which were further developed in another NSF DDDAS workshop in 2006 [4].

In the remainder of this section we describe how DDDAS ideas can be developed and applied to specific problems in storm surge prediction for emergency response and controlled water diversion for ecological purposes. From these application scenarios, we derive a set of requirements on the computational infrastructure, apply them to these and related problems, and suggest new algorithms that need to be developed to advance DDDAS applications generally. While motivated by these problems, many of the basic ideas and requirements are common to other disciplines requiring DDDAS techniques, motivating the need to develop a general DDDAS toolkit.

3.1 Event Driven Hurricane Predictions

When advisories from the National Hurricane Center indicate that a storm in the Atlantic or Gulf of Mexico may make landfall in a region impacting Louisiana, government officials, based on information provided by model predictions (Figure 2) and balancing a number of economic and social factors, must decide whether to evacuate New Orleans and surrounding towns and areas. Such advisories are provided every six hours, starting from some five days before the storm is predicted to make landfall. Evacuation notices for large cities like New Orleans need to be given 72 hours in advance. Here we outline a high level description of a complex DDDAS scenario, illustrated in Figure 3, which provides hurricane predictions using ensemble modeling.

A suddenly strengthening tropical depression tracked by satellite changes direction, worrying officials. The Louisiana Hurricane Center issues an alert to state researchers and an advanced autonomic modeling system begins the complex process of predicting and validating the hurri-
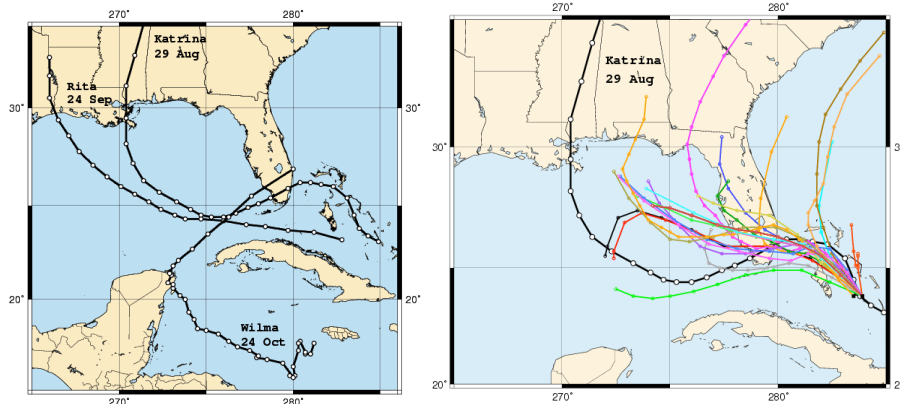
Fig. 2. Forecasting the effects of hurricanes and tropical storms in a timely manner is imperative for emergency planning. The paths and intensity of the devastating hurricanes Katrina, Rita and Wilma [left] during 2005, as with other storms, are forecast from five days before expected landfall using a number of different numerical and statistical models [right]. The validity of model predictions depend on many factors such as the storm properties, location and environment.

cane path. Realtime data from sensor networks on buoys, drilling platforms, and aircraft, across the Gulf of Mexico, together with satellite imagery, provide varied resolution data on ocean temperature, current, wave height, wind direction and temperature. This data is fed continuously into a ensemble modeling tool which, using various optimization techniques from a standard toolkit and taking into account resource information, automatically and dynamically task farms dozens of simulations, monitored in real-time. Each simulation represents a complex workflow, with closely coupled models for atmospheric winds, ocean currents, surface waves and storm surges. The different models and algorithms within them, are dynamically chosen depending on physical conditions (e.g., mud or sand bottoms) and required output sensitivity. Data assimilation methods are applied to observational data for boundary conditions and improved input data. Validation methods are used to compare data between different ensemble runs and live monitoring data, with tracking of data as it flows through the different solvers providing more information for dynamic decisions. Studying ensemble data from remotely monitored simulations, researchers help steer computations to ignore faulty or missing input data. Known sensitivity to uncertain sensor data is propagated through the coupled ensemble models quantifying uncertainty. Sophisticated comparison with current satellite data is made with synthesized data from ensemble models to determine in real-time which models/components are most reliable, and a final high resolution model is run to predict 72 hours in advance the detailed loca-
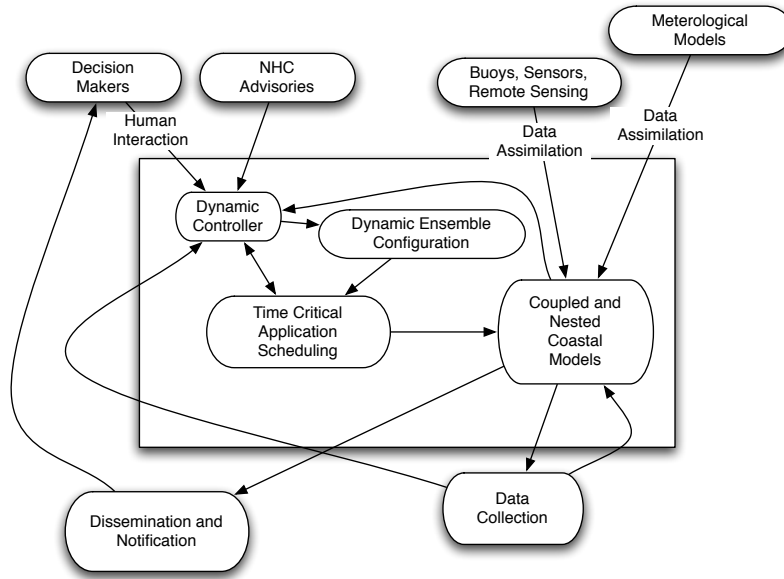
Fig. 3. Predictions of the coastal response to hurricanes, including storm surge and wave height, motivate a control system involving various data sources, human interaction, and time limited scheduling.

> *tion and severity of the storm surge. Louisiana's Office of Emergency Preparedness disseminates interactive maps of the projected storm surge and initiates contingency plans including impending evacuations and road closures.*

3.2 Ecological Scenario

A different DDDAS scenario, involving ecological processes, builds on unique features of the Louisiana ecosystem, which is dynamically controllable via gates in the *Breton Sound diversion*. When opened, this control diversion structure allows the Mississippi River to flow into a wide area of threatened coastal wetland. Currently human intuition is the primary control factor for the gates, targeted at controlling salinity, but this $150M diversion could be scientifically controlled, through automated infrastructure linked to sophisticated models and the increasing array of sensors and satellite images (Figure 4, ultimately treating the whole wetland environment in Louisiana as a steerable DDDAS optimization system.

> *High-resolution models, coupling hydrodynamic, geomorphic, ecological and water quality components, are dynamically configured through uniform interfaces in a cascade of hydrodynamic models. Runtime toolkits*

enable dynamic injection of basin sensor data, at runtime, as it becomes available. Results from ensembles of different model cascade combinations, run across regional and national grids, are compared with actual data, allowing optimal validity calibration, choosing the best model cascade combination. Improved results from the hydrodynamic models initiate coupled high-resolution ecological models. These models generate 3-day forecasts for ecological conditions pertaining to chlorophyll content, nutrient concentration, salinity, and turbidity in the basin. These biogeochemical attributes can be used to provide an index of potential eutrophication under specific control diversion structure operations; given the set of real-time boundary conditions of the landscape (especially wind fields). Measurements obtained from the sensors in the basin and complementary satellite data in combination with model-model comparisons could be used to assess and improve the efficiency of the models.

Forecasts from such a framework will help engineers provide ecologically compatible flow conditions to enhance restoration while minimizing eutrophication problems. This is a step towards an automated control scenario, which taken further, determines the optimal flow control variables to achieve desired effects on different ecological parameters. This requires close interaction between the sensors, operational structures and the dynamic modeling framework. More complex scenarios such as effects of opening of multiple structures and their combined effect on the ecosystem could be determined beforehand, thus preventing unintended damages. An intelligent decision making system could then mine results obtained by various permutations and combinations at all possible levels and construct the optimal set of structure control parameters. Decisions could be made based on results generated by such DDDAS frameworks to efficiently control the diversion structures.

Numerous control diversion systems are planned, each costing of order $150M. DDDAS capabilities could play an important role in optimizing the
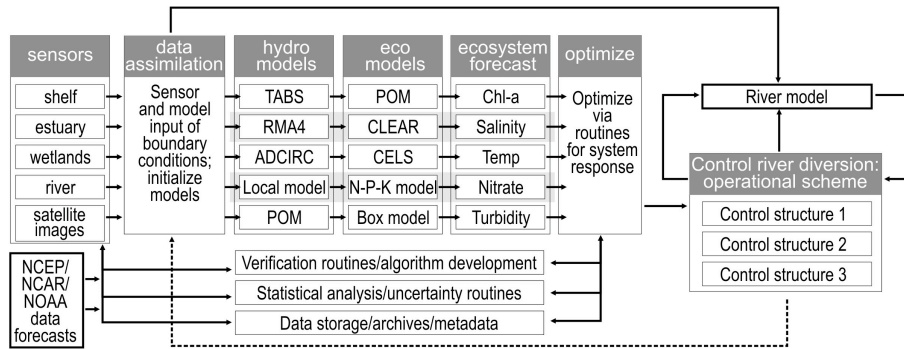

Fig. 4. *A complex ecological modeling DDDAS scenario designed for controlling river diversions.*

operation, design and placement of future structures. An extensive monitoring system is already in place, but is not yet linked to a modeling system. When the DDDAS systems are in place, they will be able to optimize dynamic placement of mobile sensors.

4 Synergistic Realtime Programs

We next describe three ongoing research and development programs (LPFS, CLEAR, and SCOOP) which are active in developing coastal modeling, assessment, and restoration projects in the Louisiana region. These programs form the foundation for the DDDAS research we describe in Section 5 below, which is aimed at advancing these and other programs that can benefit from these techniques.

4.1 Lake Pontchartrain Forecast System

During Hurricane Katrina storm surge water from Lake Pontchartrain flooded the city of New Orleans via breaches in outfall canals whose usual function is to drain rain water out of the city. To ensure this does not happen during future storms, the Army Corp of Engineers has developed a plan to close Interim Gated Structures at the mouths of three canals (17th Street, Orleans, and London, see Figure 5). However, closing these interim gates requires several hours and cannot be undertaken when storm force winds are present. Further, closing the gates should be delayed as long as possible, to allow the storm rain water to be drained from the city.

The Lake Pontchartrain Forecast System (LPFS) [5] has been developed by UNC, LSU and collaborators to provide timely information to the Army Corp to enable their decision making for closing the canal gates. LPFS provides an automated modeling system which is activated if an advisory from the National Hurricane Center (which are disseminated every 6 hours during storm activity) places the track of a storm within 271 nautical miles of the canal mouths (Figure 6). The system then deploys an ensemble of ADCIRC [6] runs (currently five runs are used which are forced by different winds corresponding to the consensus storm and 4 perturbations to this storm) on the compute resources of the Louisiana Optical Network Initiative (LONI) [7] where mechanisms are in place to ensure they have sufficient priority to complete within two hours. The results from the ensemble are integrated together and disseminated via protected web pages.

The operational, distributed LPFS system is providing an application use case for developing new technologies and policies for priority driven and deadline based computing including a general notification mechanism, preemptive scheduling on LONI, and dynamic application set threat levels to prescribe the priority of model runs (Figure 9).
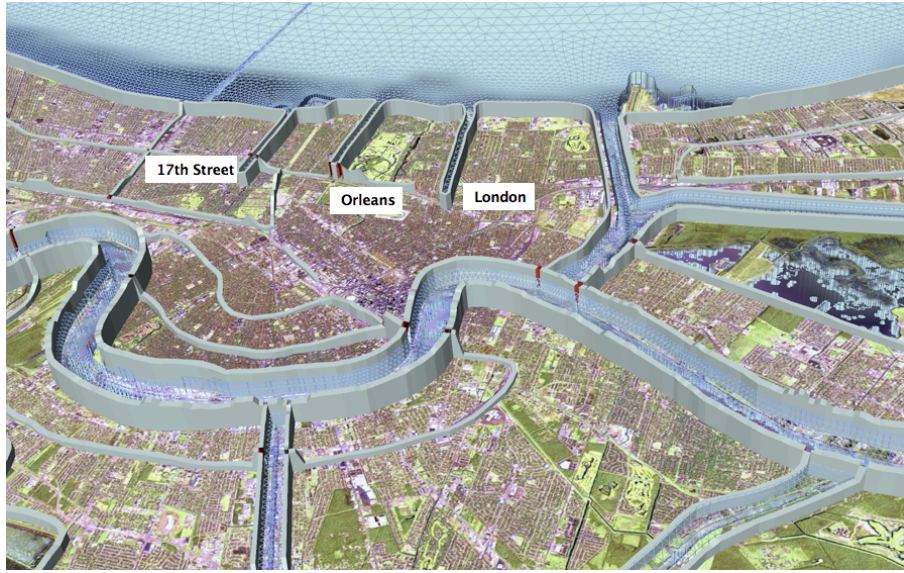
Fig. 5. Visualization of New Orleans superimposed with levees and computational mesh from the LPFS ADCIRC model. The three targeted canals (17th Street, Orleans, and London) are shown leading into Lake Pontchartrain on the north side of New Orleans. [Image credits: CCT Scientific Visualization group, Amanda Long, Werner Benger, Ana Buleu, Shalini Venkataraman, Steve Beck].

4.2 Louisiana CLEAR Program

The Coastal Louisiana Ecosystem Assessment and Restoration (CLEAR) program [8, 9] is developing ecological and predictive models to connect ecosystem needs and opportunity with engineering design. CLEAR is concentrating on a linked set of physical and ecological models. The physical models forecast endpoints of salinity, hydroperiod, and sediment distribution. From this geophysical footprint, predictions are then made of geomorphic features and shifts in land to water distribution. Ecological forecasts then simulate changes in habitat type, habitat use and conditions of water quality. This set of linked models provides the basis for an ecosystem forecasting system (geophysical processes, geomorphic features, water quality conditions, and ecological succession) which is used to evaluate alternative designs of engineering projects based on the projected response of the ecosystem.

The development and application of the models supports a strong adaptive management approach of existing and future projects, defining a body of knowledge by which hypotheses and assumptions can be continuously evaluated to incrementally reduce uncertainty in the model codes and thus improve the accuracy of the ecosystem response.
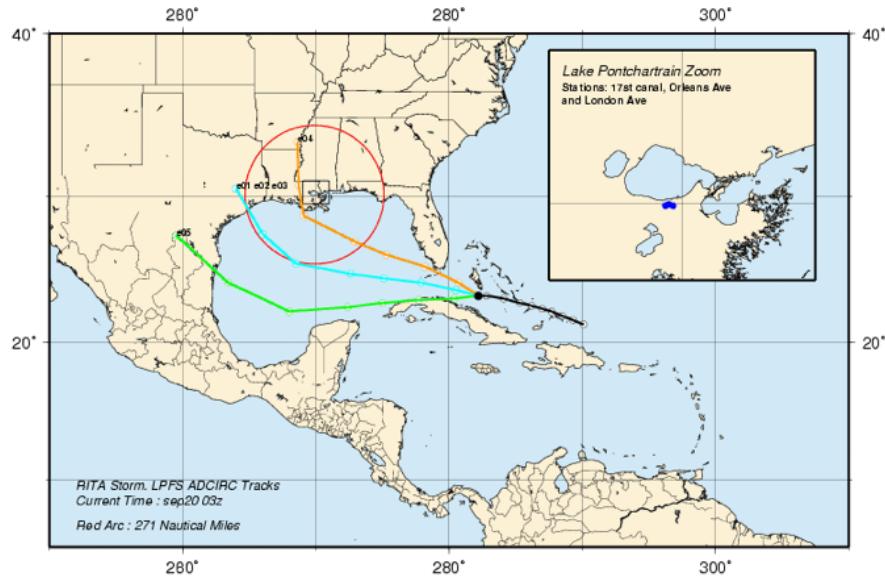
Fig. 6. The Lake Pontchartrain Forecast System deploys an ensemble of models to provide predictions for water level and wind speed at the mouths of three New Orleans Canals. The system is activated automatically if any of the ensemble tracks generated from an advisory from the National Hurricane Center places the hurricane track within 271 nautical miles from the canal mouths (inside the red circle in this image), and provides forecasts within two hours.

CLEAR has developed a modeling tool to evaluate restoration alternatives along with environmental benefits using a combination of modules that predict physical processes, geomorphic features, and ecological succession. The CLEAR program is now continuing to develop conceptual and simulation models to further develop an ecosystem forecasting system that integrates the elements of a comprehensive monitoring and adaptive management program within the LCA to evaluate environmental benefits in the coastal ecosystems of the Mississippi Delta. This system will help answer such questions as what will happen to the Mississippi River Deltaic Plain under different scenarios of restoration alternatives, and what will be the benefits to society?

4.3 SURA Coastal Ocean Observing and Prediction

SURA is a consortium of 62 universities whose mission is to "nurture scientific discovery and grow the scientific capacity of our region and the nation." The SCOOP Program [10, 11] is a SURA initiative in Coastal Research that involves a diverse collaboration of coastal modelers and computer scientists working with government agencies to create an open integrated network of distributed

sensors, data and computer models. The SCOOP architecture will serve the needs of the nation for a broad array of services oriented toward applications and research involving coastal environmental prediction. The SCOOP program implements the SURA mission to foster collaboration among its member institutions and to encourage new ideas for collaboration. Thus, by developing this broad community-oriented cyberinfrastructure, the SCOOP program will facilitate the implementation and coordination of projects such as CLEAR and LPFS, and also provide a rich environment for research projects in computer science such as DDDAS. The SCOOP program aims to go one step farther. When aligned with the mission goals of operational agencies focussed on practical applications, the SCOOP cyberinfrastructure should enable transformational collaborations with the research community for use-inspired research that has a fast track to provide tangible benefits for the nation.

SCOOP activities are driven by the need for improved forecasts and real-time information for severe storm events, such as tropical storms and hurricanes. The recent catastrophes in the southeast US following the triad of hurricanes Katrina, Rita and Wilma have highlighted the pressing need for timely and accurate forecasts as well as improved coordination and information transfer between domain experts, policy makers and emergency responders.

SCOOP covers a wide range of activities with the central aim of providing a service-oriented cyberinfrastructure for the community, to be achieved by modularizing critical components, providing standard interfaces and data descriptions, and leveraging new Grid technologies. This cyberinfrastructure will include components for data archiving, integration, translation and transport, model coupling and workflow, event notification and resource brokering. Rather than developing a single community model or toolkit, using framework approaches such as the Earth System Modeling Framework (ESMF) [12] or the Cactus Code [13, 14], SCOOP is building interfaces to allow existing models to communicate with each other with coarse-grained connectivity.

The SCOOP community currently engages in distributed coastal modeling across the southeastern US, including both the Atlantic and Gulf of Mexico coasts. Various coastal hydrodynamic models are run on an operational (24/7/365) basis to study physical phenomena such as wave dynamics, storm surge and current flow. The computational models, include[1] Wave Watch 3 (WW3), Wave Model (WAM), Simulating Waves Nearshore (SWAN), ADvanced CIRCulation (ADCIRC) model, ElCIRC, and CH3D. Atmospheric model results from models such as NAM, NOGAPS, COAMPS, and analytical models provide the wind forcing which feed into the coastal hydrodynamical models. Most SCOOP models are run on an operational basis at least once

---

[1] Wave Watch 3 (WW3): `http://polar.ncep.noaa.gov/waves/wavewatch/wavewatch.html`, Wave Model (WAM): `https://www.fnmoc.navy.mil/PUBLIC/WAM/wam det.html`, Simulating Waves Nearshore (SWAN): `http://www.wldelft.nl/soft/swan/`, ADvanced CIRCulation (ADCIRC) model: `http://www.adcirc.org`, EL-CIRC: `http://www.ccalmr.ogi.edu/CORIE/modeling/elcirc/`, CH3D: `http://users.coastal.ufl.edu/ pete/CH3D/ch3d.html`.

Model	Description	Domains	Mesh
ADCIRC: Advanced Circulation Model	2D surge and current	Entire SE	
CH3D: Curvilinear grid Hydrodynamics model in 3D	3D velocities, temperature, salinity, surge	Western Florida	
ElCIRC: Eulerian-Lagrangian Circulation Model	3D baroclinic circulation across rivers to oceans	Chesapeake Bay	
WW3: Wave Watch III	Water depth & wave propagation	Entire SE, parts of Gulf of Mexico	
WAM: Wave Model	Water depth & wave propagation	Gulf of Mexico	
SWAN: Simulating Waves Nearshore	Water depth & wave propagation	Coastal Louisiana	

Table 1. Primary coastal models used in the SCOOP project and the large scale and regional domains on which they are deployed. These models use a variety of meshes and numerical algorithms to simulate different physical properties.

a day and generate 72 hour forecasts. The models are currently run at specific sites with existing local expertise, although the project is implementing a Grid-based infrastructure for future coordinated distributed deployment.

In addition to operational modeling, extreme events such as hurricanes or tropical storms initiate additional automated model workflows. Advisories from the National Hurricane Center (NHC) about impending tropical storms or hurricanes are used to trigger automated workflows that start with the generation of high resolution wind fields around the center of the energetic event. These wind fields then initiate hydrodynamic models at different sites. The current scenario involves the running of ADCIRC (at University of North Carolina), CH3D (at University of Florida), and ELCIRC (at Virginia Institute of Marine Sciences), WW3 (at Bedford Institute of Oceanography and Louisiana State University). Figure 7 shows the different regions for which the SCOOP models provide forecasts.

The resulting data fields obtained from both the operational and storm event scenarios are distributed to the SCOOP partners for local visualization and further analysis, and are also archived for further use in a high available archive [15].
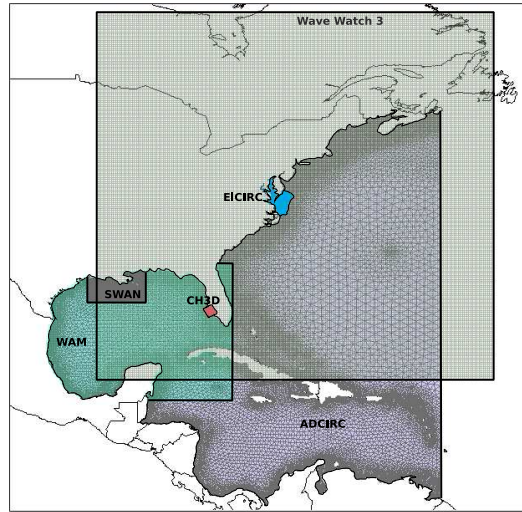
Fig. 7. Illustration of the different geographical regions over which the SCOOP ensemble of coastal models are currently deployed.

5 Components

Implementing the dynamic scenarios described in this paper will require technical advances across the areas of simulation codes, algorithms, computer systems and measuring devices. In this section, we focus on technical issues related to the various components that are needed for these scenarios. From these components, one can build many DDDAS applications, such as those described here, and others.

5.1 Varied Data Sources

Data from many sources needs to be integrated with the various models for accurate simulations of the complex physical systems under study. For example, wind fields are crucial for these applications, and can be provided both from different observational sources and computer models. Each wind field will have different uncertainties; improving the quality/resolution of input data, on demand, can lower uncertainty in forecasts. Implementing these scenarios requires the ability to dynamically create a customized ensemble of both analytic and model wind fields, validated and improved with available sensor data, for specified regions, complete with uncertainty functions that can be propagated through models for sensitivity analysis.

Data from sensors, typically collected by regional observing systems such as WAVCIS in Louisiana or GoMOOS across the Gulf of Maine, needs to be quickly available for real-time verification and data assimilation.

5.2 Data Management

Services for finding, transporting, and potentially translating data are needed to enable current ensemble scenarios and then to scale to complex workflows of coupled interacting models. To support emergency and real-time computing scenrios, data sources need to be highly available, while data transport needs to be fault tolerant with guaranteed quality of service. With new optical networks in place, which can transport 10 Gigabits per second on a single "lambda", we need mechanisms for dynamically reserving and provisioning networks as part of workflows, as well as data scheduling capabilities which ensure that data is in the right place at the right time.

Metadata describing the huge amounts of data, distributed across multiple sources, is also crucial. This should fully describe the data and its history, including information about the models which created it, or the systems which observed it, to provide complete provenance information which can be used to verify and understand results.

5.3 Model-Model Coupling

In both scenarios, cascades of coupled models are used, at different levels; e.g. circulation models, wave models, transport models, etc. Beyond simply defining interfaces between these, one must develop techniques to track uncertainties throughout entire cascades of models, create and optimize workflows on Grids as storms approach or as dams are manipulated, and invoke models or model components preferentially, based on algorithm performance and features for conditions indicated by input data.

5.4 Ensembles

Cascades of models, coupled with multiple components at each stage of the cascade, give rise to potentially hundreds of combinations. One may not know *a priori* which combinations give the best results. Automated and configurable ensemble modeling across grid resources, with continuous validation of results against observations and model-model comparisons, will be critical to dynamically refining predictions on the fly. It will also be necessary to develop algorithms for dynamic configuration and creation of ensembles to provide predictions with a specifiable, required accuracy. In designing these ensembles, the system should also take into account the availability and "cost" of computational resources, noting that the available resources may also depend on the seriousness and urgency of the situation being modeled. For the hurricane scenario, the system should also react to the potential threat of the storm, for example a Category 5 Hurricane could require a highly level of accuracy and quality of services than a Category 3 Hurricane.

5.5 Steering

The scenarios described above both involve automated steering (at both the component and workflow levels) to adjust models to physical properties and the system being modeled. In the hurricane scenario, one could steer sensor inputs for improved accuracy. For the ecological scenario, it would allow one to steer the "experiment"– the active river diversions–which can be controlled by the models to improve water quality, reduce salinity, control algal blooms, etc. The remote steering of model codes, for example to change output parameters to provide verification data, or to initiating the reading of new improved data, will require advances to the model software. Beyond introducing the technical capabilities for steering parameters (which often requires the involvement of domain experts), the mechanism for steering should require authentication for security, and changes need to be properly logged to ensure reproducibility.

5.6 Visualization and Notification

Visualization will be used for different purposes. Detailed visualizations, integrating multiple data and simulation sources showing the predicted effect of a storm (see Figure 8) will be important both for scientific understanding and analysis and for public awareness. For example, the visualization in Figure 8, which is scientifically accurate, would be very effective in impressing on the public the urgency of evacuation on the general population. Interactive and collaborative 3-D visualization for scientific insight will stress high speed networks, real-time algorithms and advanced clients. Visualizations including verification analysis and real-time sensor information will be crucially important.

Notification mechanisms which can automatically inform scientist, system administrators and emergency responders need to be robust and configurable sending the relevant information for a particular device e.g. cell phone, instant message client, email. The automated systems we envisage will require human intervention and confirmation at different points, and here it is important that the system include mechanisms which require authenticated response and have intelligent fallback mechanisms.

5.7 Priority and Deadline Based Scheduling:

Events such as hurricanes cannot be planned, and must be responded to at a time of their choosing. This type of scenario, and the need to run multiple models, concurrently with data streams and analysis tools, leads to new types of requirements on scheduling and reservation systems: priority, deadline-based, and co-scheduling. Not only must computational resources be made available on demand, with a specific deadline and guarantee for results, multiple resources must be scheduled simultaneously and/or in sequence. Furthermore, these resources go beyond traditional computing resources, and now need to include archival data, file systems, networks, visualization systems, and so on.
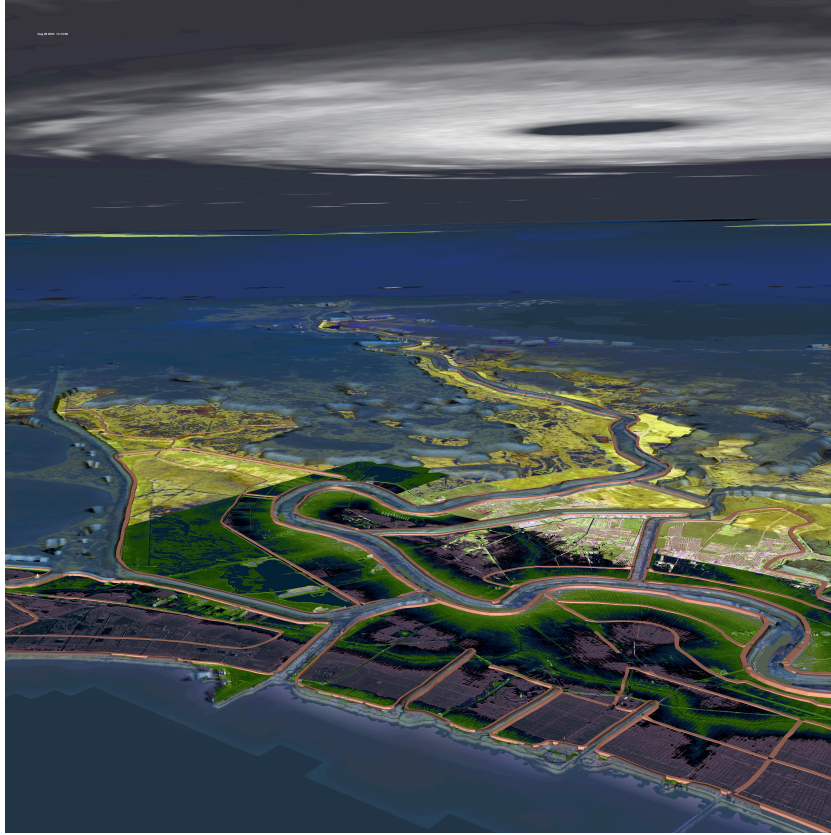
Fig. 8. Combined visualization of Hurricane Katrina approaching New Orleans, showing the water surge as predicted by the ADCIRC simulation model, cloud density as observed in infrared by the GOES-12 satellite, terrain height information obtained from high-resolution LIDAR measurements (green: above sea level, magenta: below sea level, blue: sea level) with the levee information of the ADCIRC model, overlaid on aereal photography and GIS terrain information. [Image credits: CCT Scientific Visualization group, Amanda Long, Werner Benger, Ana Buleu, Shalini Venkataraman, Steve Beck].

6 Conclusion

We have described some of the new technology components and capabilities which will be needed to enable the DDDAS scenarios outlined in Section 3. However, it is important to recognize that along with these advances, several concurrent developments are needed for these applications to be developed by communities, and actually run on computational environments. For example, the relevant scientific communities need to embrace a model of sharing data, code components, as well as computational resources. As this sociological change
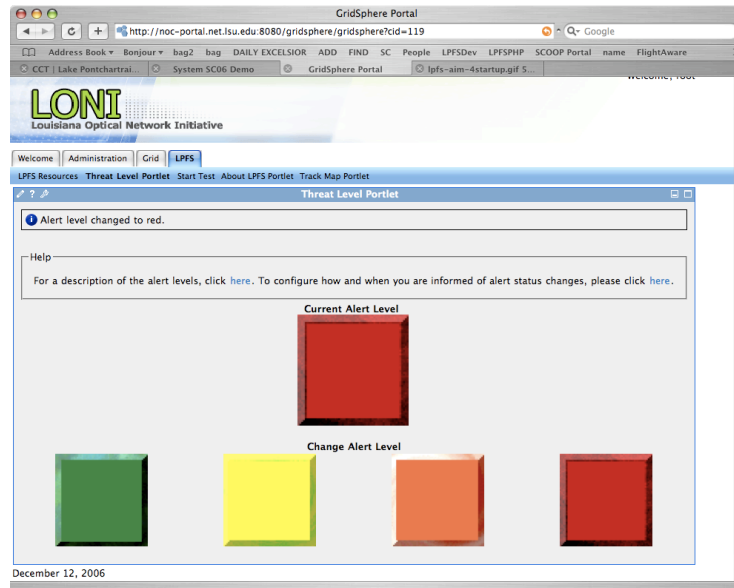
Fig. 9. The LPFS system has been augmented by a threat level system. The threat level can be changed by trusted applications or scientists, and is then used to both notify system administrators, customers and scientists of changing circumstances and to set policies on compute resources for running on demand jobs. This diagram shows a portal interface to the threat level system.

occurs, the communities will also need to adopt more modern software engineering practices and cyberinfrastructure in order ultimately to take advantage of the kinds of components we are developing. At the same time, it will be crucial for component and scientific library developers to adopt practices and interfaces that allow them to be used in this portable, dynamic environment. Finally, and just as importantly, it will be critical for policies to be adopted at computing centers that enable this kind of event-driven computing and data streaming, so that the real-time scenarios we describe can actually be carried out in practice; computational resources of various kinds will have to be available on demand, with policies that take into account the priority of certain jobs, such as hurricane storm surge computations.

Acknowledgements

We would like to acknowledge our colleagues in the SCOOP, CLEAR and LPFS projects who have carried out some of the work described here and contributed to ideas for a DDDAS system for coastal and environmental modeling.

References

1. F. Darema. Dynamic Data Driven Applications Systems: A New Paradigm for Application Simulations and Measurements. In *Lecture Notes in Computer Science, Computational Science — ICCS 2004*, pages 662–669. Springer, 2004.
2. C. C. Douglas, J. D. Beezley, J. Coen, Deng Li, Wei Li, A. K. Mandel, J. Mandel, G. Qin, and A. Vodacek. Demonstrating the Validity of a Wildfire DDDAS. In *Lecture Notes in Computer Science, Computational Science — ICCS 2006*, pages 522–529. Springer-Verlag, 2004.
3. Report from NSF DDDAS Workshop, March 2000, Washington. `http://www.nsf.gov/cise/cns//dddas/dd_das_work_shop_rprt.pdf`.
4. Report from NSF DDDAS Workshop, January 2006, Washington. `http://www.nsf.gov/cise/cns/dddas/2006Workshop/wkshpreport.pdf`.
5. Lake Pontchartrain Forecast System (LPFS). `http://www.cct.lsu.edu/projects/LPFS`.
6. The ADvanced CIRCulation model (ADCIRC): `http://www.adcirc.org/`.
7. Louisiana Optical Network Initiative (LONI) : . `http://www.loni.org`.
8. Coastal Louisiana Ecosystem Assessment and Restoration (CLEAR): Building a coastal ecosystem forecasting system. `http://www.clear.lsu.edu`.
9. Louisiana coastal area. louisiana — ecosystem restoration study — november 2004. volume 4. appendix c — hydrodynamic and ecological modeling. Technical report, 2004. `http://www.lca.gov/appc.aspx`.
10. SURA Coastal Ocean Observing Program (SCOOP) : `http://scoop.sura.org`.
11. Bogden, P., Allen, G., Stone, G., Bintz, J., Graber, H., Graves, S., Luettich, R., Reed, D., Sheng, P., Wang, H., Zhao, W. The Southeastern University Research Association Coastal Ocean Observing and Prediction Program: Integrating Marine Science and Information Technology. In *Proceedings of the OCEANS 2005 MTS/IEEE Conference, Sept 18-23, 2005, Washington, D.C.*, 2005. `http://scoop.sura.org/documents/050301-04.pdf`.
12. Earth System Modeling Framework. `http://www.esmf.ucar.edu`.
13. Tom Goodale, Gabrielle Allen, Gerd Lanfermann, Joan Massó, Thomas Radke, Edward Seidel, and John Shalf. The Cactus framework and toolkit: Design and applications. In *High Performance Computing for Computational Science - VECPAR 2002, 5th International Conference, Porto, Portugal, June 26-28, 2002*, pages 197–227, Berlin, 2003. Springer.
14. Fokke Dijkstra and Aad van der Steen. Integration of Two Ocean Models. In *Special Issue of Concurrency and Computation, Practice & Experience*, volume 18, pages 193–202. Wiley, 2005.

15. Jon MacLaren, Gabrielle Allen, Chirag Dekate, Dayong Huang, Andrei Hutanu, and Chongjie Zhang. Shelter from the Storm: Building a Safe Archive in a Hostile World. In *Proceedings of the The Second International Workshop on Grid Computing and its Application to Data Analysis (GADA'05)*, Agia Napa, Cyprus, 2005. Springer Verlag.