

**Architecture of a Community Infrastructure
for Coastal Inundation Prediction & Response**

Philip Bogden and Tom Gale

Southeastern Universities Research Association, Washington, D.C. &

Gulf of Maine Ocean Observing System, Portland, ME

Gabrielle Allen and Jon MacLaren

Louisiana State University, Baton Rouge

Guy Almes and Gerald Creager

Texas A&M University, College Station

Joanne Bintz and L. Donelson Wright

Southeastern Universities Research Association, Washington, D.C.

Hans Graber and Neil Williams

University of Miami, Miami, FL

Sara Graves and Helen Conover

University of Alabama, Huntsville

Ken Galluppi¹ and Rick Luettich

University of North Carolina, Chapel Hill and ¹Renaissance Computing Institute, Chapel Hill, NC

William Perrie and Bechara Toulany

Bedford Institute of Oceanography, Nova Scotia, Canada

Y. Peter Sheng and Justin R. Davis

University of Florida, Gainesville

Harry Wang and David Forrest

Virginia Institute of Marine Science, Gloucester Point

The Southeastern Universities Research Association (SURA) has advanced the SURA Coastal Ocean Observing and Prediction (SCOOP) program as a multi-institution collaboration to design and prototype a modular, distributed system for real-time prediction and visualization of the coastal impacts from extreme atmospheric events, including hurricane inundation and waves. The SCOOP program vision is a community “cyberinfrastructure” that enables advances in the science of environmental prediction and coastal hazard planning. The system architecture is a coordinated and distributed network of interoperable, modularized components that include numerical models, information catalogs, distributed archives, computing resources, and network infrastructure. The components are linked over the Internet by standardized web-service interfaces in a service-oriented architecture (SOA). The design philosophy allows geographically disparate partnering institutions to provide complementary data-provider and integration services. The overall system enables coordinated sharing of resources, tools, and ideas among a virtual community of coastal and computer scientists. The distributed design builds on the notion that standards enable innovation, and seeks to leverage successes of the World Wide Web by creating an environment that nurtures interaction between the research community, the private sector and government agencies working together on behalf of the nation.

1 Introduction

The 2005 Atlantic hurricane season was remarkable for its early beginning, the number of named storms (27), and for having the most intense hurricane on record in the Atlantic (NCDC, 2006). Undoubtedly the season will be most remembered for Hurricane Katrina, and its devastation of southern Mississippi and Louisiana. Hurricane Katrina alone is estimated to have caused more than 1,800 deaths and record damages of over \$81 billion (Knabb et al, 2006). Even more sobering is the conclusion that the destructive potential of these storms is increasing as an effect of climate change [Emanuel, 2005]. With hurricane impacts so devastating and costly, the potential benefits of timely hazard predictions are correspondingly high. One obstacle to realizing these benefits involves getting relevant information to emergency responders well in advance of hurricane landfall (Malone and Hemsley, 2007). Responders need time to weigh the costs of unnecessary evacuation - often estimated at one million dollars per mile of coastline evacuated (Smith, 1999; Whitehead, 2000) - against the potential loss of life and property due to inadequate preparation. This manuscript describes the design of an information architecture for a community science initiative aimed at enhancing our nation's capacity for coastal hazard prediction and disaster mitigation.

With funding provided by the National Atmospheric and Oceanic Administration (NOAA) and the Office of Naval Research (ONR), the Coastal Research Committee of the Southeastern Universities Research Association (SURA) has advanced the SURA Coastal Ocean Observing and Prediction (SCOOP) program (<http://scoop.sura.org/>). SCOOP is a multi-institution (<http://scoop.sura.org/partners.html>) collaboration to prototype a distributed information network of shared resources focused on the challenges of predicting coastal hazards. The prototype

network will broaden access to the requisite data, models, computational resources, and other key components of a comprehensive real-time environmental prediction system. Many of the primary data sources and modeling capabilities are already used by operational agencies, private enterprise, and research institutions. Data from these systems are increasingly available due to the adoption of community standards from organizations such as the World Wide Web Consortium (W3C; <http://www.w3.org>) and the Open Geospatial Consortium (OGC; <http://www.opengeospatial.org>). The SCOOP vision is to leverage and integrate the disparate data sources in an information architecture that enables transformational science and provides innovative science products.

The SCOOP system design uses a service-oriented architecture (SOA). SOA is increasingly common in science (Foster, 2005), and comprises a design philosophy as opposed to an information technology. The architecture has been variously referred to as “cyberinfrastructure” (Atkins, et al, 2003) and “e-Infrastructure” (Foster, 2005). The “Grid” is a closely related term coined in the mid 1990s to describe the application-neutral concept of a distributed computing infrastructure (Foster et al, 2001). Examples include the NSF TeraGrid (<http://www.teragrid.org>), the Open Science Grid (<http://www.opensciencegrid.org>), and the relatively new SURAGrid (http://www.sura.org/programs/sura_grid.html). The SCOOP program focuses on coastal science applications that use distributed computational and storage resources.

The premise is that a shared, community cyberinfrastructure can engage the intellect and resources of some of the nation’s top research universities, and invite involvement of the private sector and operational agencies, as appropriate. The approach to creating a coordinated system of networked components builds from lessons learned in creating the World Wide Web. The underlying technologies are common in computer science but relatively new in the coastal

sciences. As a design philosophy, SOA underlies the “system of systems” concept for observing the earth (<http://www.epa.gov/geoss>), and is found in two recently proposed conceptual architectures (http://www.ocean.us/IOOS_Arch_Proposals) for the Integrated Ocean Observing System (IOOS). SOA is arguably the only cost-effective approach to integration of legacy systems because the design involves standards-based service interfaces to connect heterogeneous components, but does not require standardizing the underlying components themselves.

Here, we provide a scientific rationale and requirements for the SCOOP architecture (including a description of the specific SCOOP applications for coastal inundation), describe design requirements and a generalized work process used to design the architecture, present two different views of the architecture itself, discuss the present status of the prototype implementation, and relate the SCOOP SOA to other types of modeling frameworks and parallel-computing paradigms. We conclude with some thoughts relevant to turning the SCOOP program prototype into an operational system.

2 Scientific Rationale and Requirements

2.1 Ensemble Modeling

Researchers continue to make substantial progress in predicting storm tracks and intensities, but significant challenges remain. Consider, for example, Figure 1; the sequence of images shows predicted trajectories for Hurricane Katrina from eleven different numerical weather-prediction models over the five days before landfall in New Orleans. The actual trajectory appears in black with open circles. There was large disagreement among the models when Katrina was five days from landfall. Four days before landfall, most models agreed, but incorrectly put the storm on track for the Florida panhandle. Not until two days before landfall did all models agree that Katrina would hit near New Orleans.

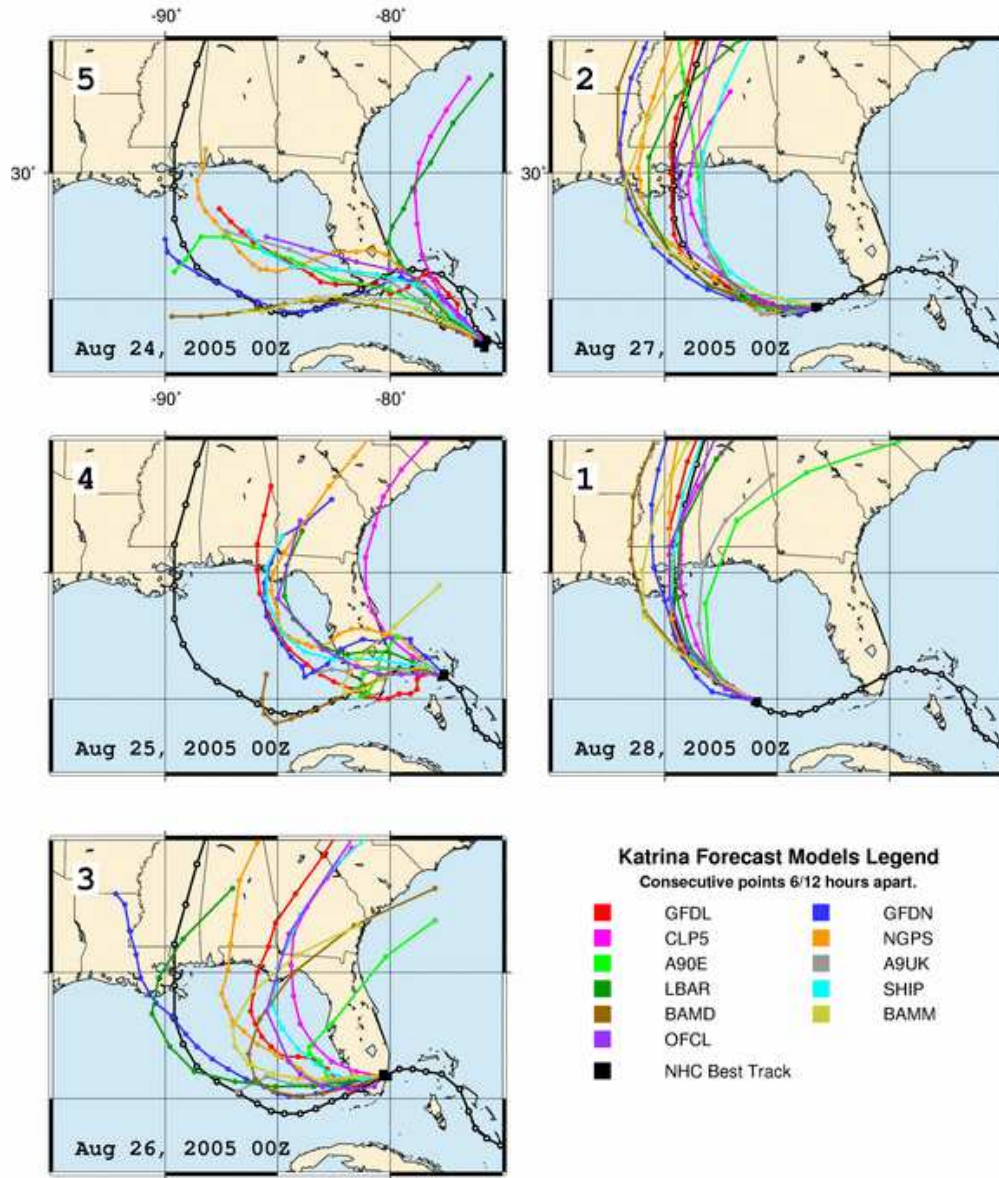


Figure 1. Hurricane Katrina forecast tracks and observed track (black with open circles) for the five days before landfall in New Orleans. The actual track is appears in black and is marked with open circles. The Source data used in the figure come from the NHC FTP service at: <ftp://ftp.nhc.noaa.gov/pub/atcf/>. Information about the forecast models and acronyms is available at: <http://www.nhc.noaa.gov/aboutmodels.shtml>. Figure courtesy of Shree Balasubramanian, LSU Center for Computation and Technology. The color legend can be viewed in the PDF version of this document available at: ***ADD URL FOR PDF FILE HERE***

The sequence illustrates that environmental prediction is a problem in probability and statistics. Emergency planners and responders need estimates of forecast probability, such as the likelihood that a specific region will flood and for how long. Predicting the single best storm track and its coastal impacts is a challenge in itself, but predicting probability distributions is much more difficult. The associated ensemble-modeling techniques require that computer simulations of an event be run many tens or hundreds of times—once for each plausible realization—thereby creating huge computational demands on any near-real-time, event-driven system.

Additional challenges arise from observational requirements for real-time prediction. The requisite observational data collectively determine the reliability of the information products they are used to produce. However, a wide variety of organizations collect and (sometimes) disseminate relevant data in a correspondingly wide variety of non-standard formats. The lack of standards makes data flow into any prediction system extremely complex.

From environmental sensors to supercomputers to decision-makers, observations and predictions must be turned around and disseminated in a matter of hours to have their most critical societal benefits. The U.S. Commission on Ocean Policy has articulated the rationale for a national Integrated Ocean Observing System (IOOS) that would enable the scientific community to take on such challenges:

Coastal and ocean observations provide critical information for protecting human lives and property from marine hazards, enhancing national and homeland security, predicting global climate change, improving ocean health, and providing for the protection, sustainable use, and enjoyment of ocean resources. While the technology currently exists to integrate data gathered from a variety of sensors deployed on buoys, gliders, ships, and satellites,

implementation of a sustained, national Integrated Ocean Observation System (IOOS) is overdue and should begin immediately. (U.S. Commission on Ocean Policy, 2004)

The first step in building the system involves integrating diverse data flows from a variety of ocean observing initiatives around the country (Ocean.US, 2006). Utilizing some of the first Web-service specifications from the OGC, SURA supported a community effort that made substantial progress toward demonstrating interoperability of data and data products using web-based Geographical Information Systems (GIS). The effort resulted in a technical proof of concept for service-oriented aggregation with a portal at <http://www.openioos.org>. This paper describes the relatively complex next step in a system design where integrated data flows support an open-access, scalable, modular, and distributed real-time environmental prediction system.

2.2 The SCOOP Use-Case Scenarios

The SCOOP architecture is designed to support an application set that includes three distinct use-case scenarios. An “event-driven, ensemble-prediction” scenario uses warnings of extreme events (e.g., as issued by the National Hurricane Center (NHC) during an active hurricane) to trigger automated, on-demand, ensemble-model calculations of inundation and waves. These calculations are designed to run quickly so they can support the urgent and immediate needs of real-time hazard planning and response. A “24/7” scenario runs constantly, and provides continually updated day-to-day forecasts of waves and water level. This scenario could serve a variety of practical purposes. In SCOOP applications, the 24/7 scenario provides initialization and boundary conditions for other models in the infrastructure. A third “retrospective” scenario supports research and analysis of past events.

The SCOOP architecture is designed to support all three scenarios without reconfiguration of the underlying components. The extent to which a single architecture can

support all three scenarios determines the extent to which a production version of the system could blur the distinction between research and operations, to the benefit of both.

The process chain for all scenarios starts with data sources that “drive” the system. These include NHC warnings, wind-field predictions from operational service providers, and *in situ* sensor observations of variables such as water level and wave height. The end of the process chain includes predictions of coastal impacts that yield value-added and quality-controlled information products. The products are visualized on a Web browser and disseminated in data formats that support GIS-compatible decision-support tools.

Workflow for any specific use-case scenario will involve a unique set of data sources selected from a broad range of possible inputs. For example, the SCOOP system ingests hurricane wind predictions and daily wind forecasts from meteorological models coming from a variety of NOAA and Navy programs. These winds drive the 24/7 coastal surge and waves models. Observational data sources used for verification and quality control come from the NOAA National Water Level Observation Network (NWLON), and wave measurements come from the National Data Buoy Center (NDBC), as well as various university-based regional ocean-observing systems along the U.S. coastal zone. Since the objective is to build and maintain a community resource, SCOOP places strong emphasis on system requirements that enable reusability of any of these data sources for multiple purposes. Consequently, results must be archived and registered to enable data discovery. This includes a catalog that documents provenance and quality assurance and supports discovery and access, especially by those who were not responsible for generating the data.

2.3 Community modeling & community models

The SCOOP program goal is to create a collaborative environment for community modeling that engages scientists from different disciplines and organizations. The associated design requirements recognize that the coastal modeling community has not settled on any single community model, nor should it. Different models often have different strengths, and the best models will likely evolve from healthy competition and peer review by a community that has a variety of model-development efforts. Rather than advocate one model over another, the design goal is to provide a collaborative environment that facilitates controlled inter-comparisons that might ultimately give rise to better models and model products. Therefore, numerical models are considered interchangeable components in the SCOOP system infrastructure.




Figure 2. The numerical modeling grids employed by the SCOOP Program and the regions over which they are deployed.

Figure 2 shows a variety of numerical modeling grids and the acronyms used by the various models that are being integrated into the SCOOP infrastructure. The collection includes

regular grids associated with surface-wave models: Simulating Waves Nearshore (SWAN), WAVEWATCH III (WW3), and Wave Analysis Model (WAM), plus finite-element and curvilinear grids associated with storm-surge models: Advanced Circulation Model (ADCIRC), Eulerian-Lagrangian CIRCulation (ELCIRC), and Curvilinear-grid Hydrodynamics 3D (CH3D; <http://ch3d.coastal.ufl.edu>). Model “integration” into the SCOOP infrastructure involves a range of activities including translation and interpolation interfaces that allow these grids, the associated model inputs and outputs, and the models themselves to be combined and run in a variety of ways.

3 The SCOOP Work Process

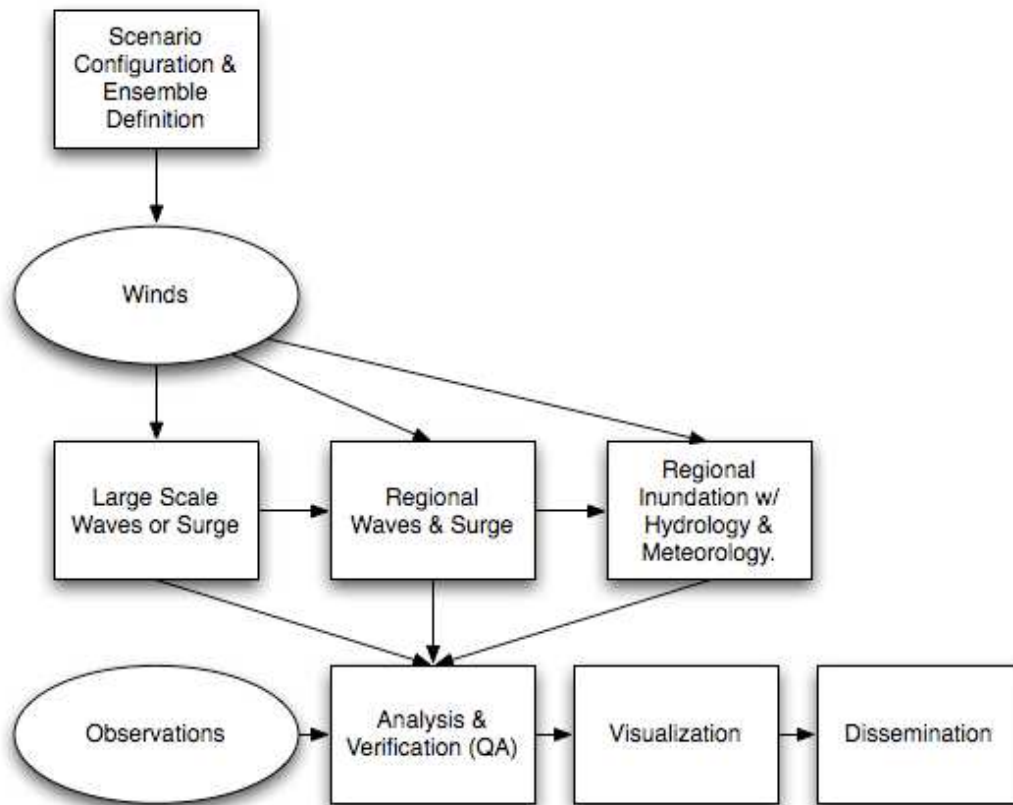


Figure 3. Generalized work process in SCOOP.

3.1 Generalized Work Process

Figure 3 provides a high-level view of the general, end-to-end work process that describes the entire SCOOP application set. The process begins at the top of the diagram with a definition of the specific scenario configuration. All three use-case scenarios involve some type of wind ensemble (represented by the oval labeled “winds”). The 24/7 scenario includes the degenerate case where the wind ensemble has only one member. A working prototype of the event-driven ensemble-prediction scenario – computationally the most challenging of the supported scenarios – uses a 40-member wind ensemble derived from NHC warnings. The parameters for hurricane-trajectory uncertainties (the probability cone) are used to create an ensemble of five-day wind field predictions from a well-known hurricane-wind formulation (Holland, 1980).

The next step in the work process involves coupling individual wind-ensemble members with one or more numerical models for storm surge or wind waves. The SCOOP team is integrating a hierarchy of increasingly complex models and model combinations, as indicated by the three rectangles in the third row of Figure 3. The simplest members of the hierarchy include two large-scale models for the entire North Atlantic and Gulf of Mexico: ADCIRC (Luettich, et al, 1992) provides tidal water levels and wind-driven storm surge, and WW3 (Tolman 2002, Tolman, et al., 2002) provides ocean surface waves. The large-scale models are uncoupled, that is, they run independently of one another.

The next level of complexity in the model hierarchy involves regional, high-resolution models of coastal regions and estuaries along the Gulf of Mexico and the east coast of North America – the box labeled “Regional Waves and Surge” in Figure 3. In some cases, these regional models use open-ocean boundary conditions derived from the output of the large-scale models, as indicated by the arrows connecting the adjacent boxes in row 3 of Figure 3. For

example, ELCIRC (Zhang et al., 2004a; Wang et al., 2004; Shen et al., 2006) supports a high-resolution forecast of street-level products in the Chesapeake Bay region, with boundary conditions coming from large-scale ADCIRC. Other regional formulations include CH3D (Sheng, 1987; Sheng, 1990), which is configured for Charlotte Harbor, Tampa Bay and eastern Florida, as well as the entire Northern Gulf. For waves in regions of shoaling bathymetry, SWAN (Booij, et al., 1999) is embedded at a high-resolution with boundary conditions coming from the large-scale wave model.

The most sophisticated model suites involve greater coupling between the wave and surge models. A model suite in development for the North Carolina coast will use a high-resolution, regional wave model (SWAN) to provide momentum flux from breaking waves to a high-resolution storm-surge model (ADCIRC) along the North Carolina coast, and the surge calculation will provide water levels for SWAN.

In general, a numerical model or model suite is combined with each member of the wind-ensemble. For an N-element ensemble, there are N numerical calculations run as N independent jobs on any of a variety of supercomputers available in a “Grid-enabled” network such as SURAGrid. This part of the work process is implicit in Figure 3.

SCOOP architecture requirements extend well beyond the need to run multiple jobs on distributed computers. Once the jobs are run, output from each job is transported to a distributed archive for analysis and visualization. The lowest row in the work-process diagram indicates that observations from environmental sensors for example, coastal water levels from NOAA/NWLON, wave measurements from NOAA/NDBC and coastal observing systems including the Gulf Coast Ocean Observing System (GCOOS), the SouthEast U.S. Atlantic Coastal Ocean Observing System (SEACOOS) and the Gulf of Maine Ocean Observing System

(GoMOOS), support various analysis and verification steps. Results are then made available for visualization and dissemination by a variety of web-accessible service and portal interfaces.

3.2 A Specific Workflow Example

Figure 4 provides a relatively detailed view of workflow for the event-triggered, ensemble-modeling scenario. The figure shows various system components organized so that time proceeds from left to right. Data flow proceeds vertically in the figure from networked resources along the bottom to application-specific components higher in the diagram. Arrows in the conceptual diagram reflect a combination of service messaging (dashed), data transport (solid) and process flow (thick solid).

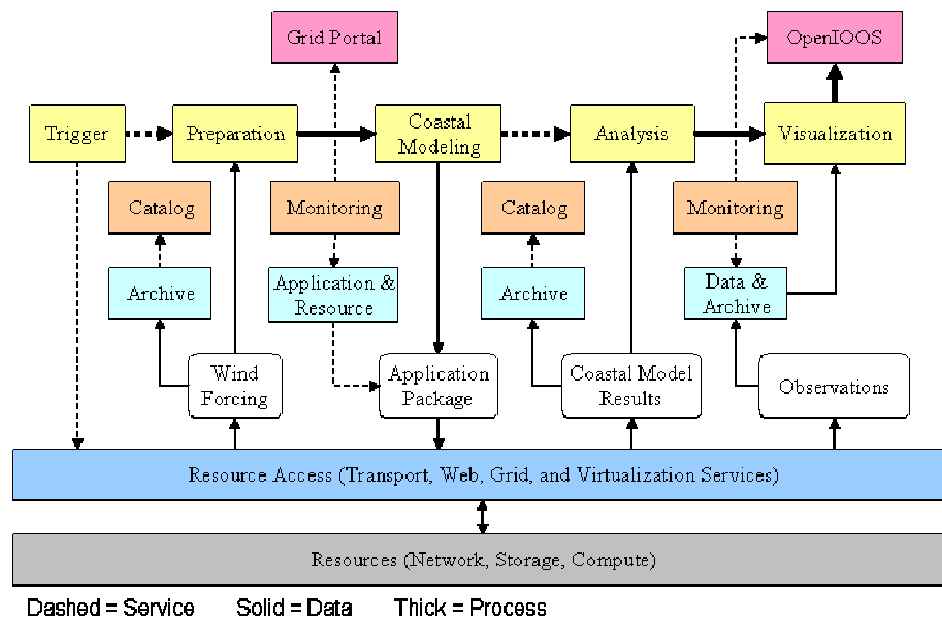


Figure 4. Conceptual workflow diagram for the event-triggered ensemble-modeling scenario.

The sequence is triggered by an automated warning from the NHC. Wind fields are created at one site, and the data enter the workflow through the resource-access layer that connects to the internet. The winds are then archived and registered in the catalog.

The arrival of the forecast-wind ensemble also triggers a sequence of events that spawn the collection of model runs. The sequence includes translation filters that assure the models can ingest the data. The resource selection service (i.e., resource management) identifies and selects the available computational resources, and launches the model runs across the network of distributed computing resources. Meanwhile, observation data (including water levels, waves, streamflow and winds) are obtained from the appropriate source locations and/or archive.

As model results are generated, they are pushed to the archives, cataloged and processed with various verification and analysis tools. Finally, visualization services are used to process the model results so that they can be displayed on web portals (e.g., <http://www.OpenIOOS.org>) or integrated into an OGC-compliant GIS.

Monitoring and workflow services (not explicit Figure 4) facilitate automation of the complete sequence. Other scenarios involve workflows that combine the same components in appropriate configurations.

4 The SCOOP Architecture

This section provides two high-level views of the target SCOOP architecture – not all components are supported in the present implementation of the working prototype.

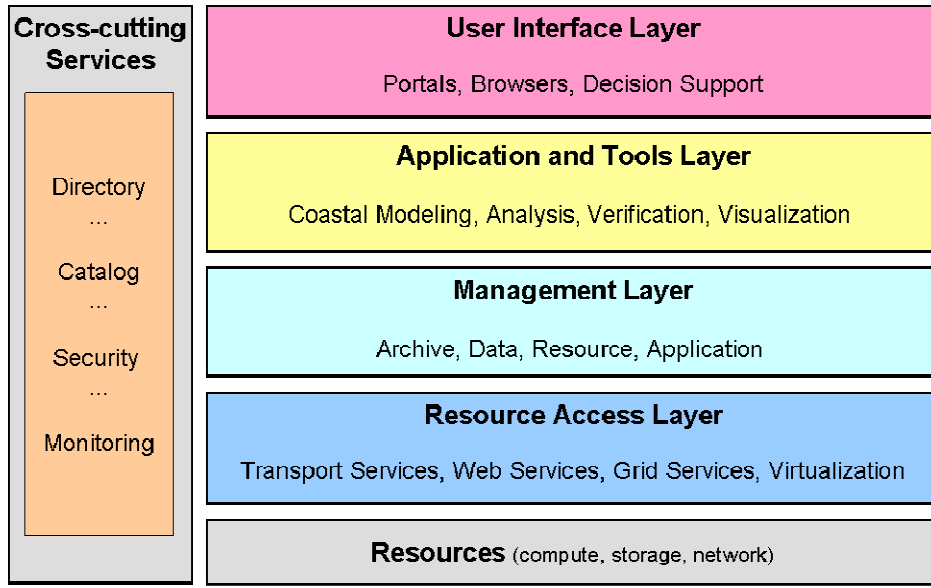


Figure 5. Layer-stack diagram for the SCOOP SOA. The color scheme reflects the colors used for individual components in Figure 4.

4.1 Layer-Stack Diagram

In general, layer-stack diagrams provide the hierarchical framework for functionality of an SOA. The layer diagram in Figure 5 is a high-level overview of the SCOOP architecture. Each layer in the stack represents a collection of components with closely related functions. The various components within each layer provide functionality that supports information flow across adjacent layers. Connections between layers can be made in a variety of ways in order to support any number of different use scenarios. The color-scheme for the layers in the stack corresponds to the colors used for individual components in Figure 4.

The uppermost “user-interface layer” interacts directly with end users and end-user tools, isolating them from the complexities of the underlying system architecture. Components in this layer may be tailored to the workflows of one or more specific scenarios. Examples include web portals that spawn user-customized workflows and decision-support tools that automatically ingest information products through an OGC-compliant, GIS-compatible interface.

The “application and tools layer” contains the various numerical models for predicting and analyzing environmental phenomena. This layer includes data translation and visualization service components, and workflow-configuration tools. Components within this layer generally interact with one another in a fashion that is coordinated by the workflow tools.

The “management layer” enables workflows by coordinating components in the applications and tools layer and assuring that they have the necessary resources. In this role, the management layer coordinates access to physical resources in the system. Management services generally vary from one scenario to another. Thus, management components need to be flexible and readily configurable to accommodate a range of possible workflows.

The “resource-access layer” provides standard protocols that connect the management layer to physical computing, storage, and network resources. These protocols facilitate management and access to the distributed resources. Components include a combination of middleware protocols (e.g., Globus), web-service and data-transport protocols (e.g., Simple Object Access Protocol (SOAP) and Local Data Manager (LDM), and lower-level networking protocols (e.g., Hyper-Text Transfer Protocol (HTTP)).

The “resources layer” includes all the physical resources for computing, storage, and network connectivity. For SCOOP, physical resources and system components are distributed at different institutions over the Internet. Distribution has advantages, but it also introduces special challenges that are overcome with cross-cutting components. For example, resource-management challenges arise in heterogeneous environments (e.g., different hardware platforms may require different operating systems). Virtualization interfaces within the resource-access layer provide one increasingly popular technique for dealing with such heterogeneities, but there is no panacea. Benefits to this distributed approach include an ability to access a large pool of

resources, thereby significantly increasing the probability of having sufficient resources with very short notice. Also, geographically distributed resources are much less susceptible to localized severe storm events and associated power outages, ensuring access to required resources when most needed.

4.2 Node Connectivity Diagram

The node connectivity diagram (Figure 6) provides yet another view of the SCOOP target architecture. Nodes in the figure indicate some of the main service components. Arrows connecting nodes show control and messaging interactions. This is a high-level conceptual diagram – actual software-engineering diagrams would describe specific and detailed interactions and messaging sequences that are represented only conceptually in Figure 6.

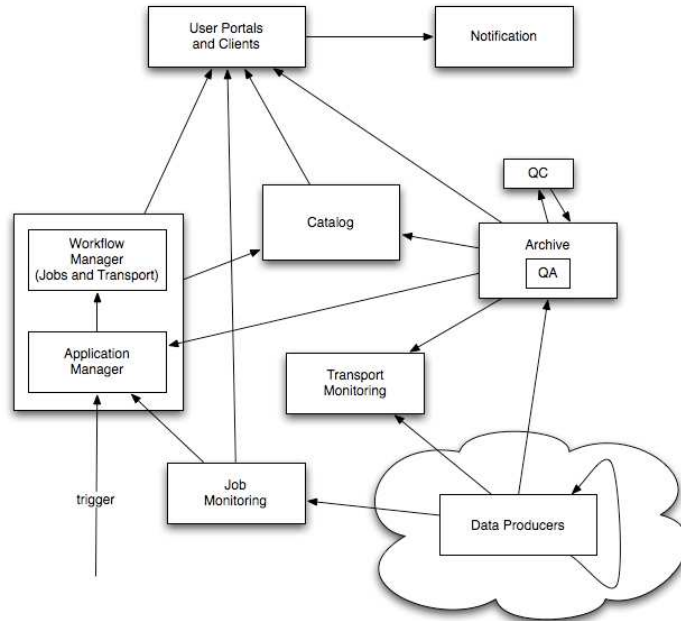


Figure 6. Node connectivity diagram.

The application manager provides a high-level job-management service that knows about specific SCOOP applications. Such domain-specific knowledge makes this component different from a general workflow or resource manager (e.g., Condor, Pegasus). The application manager tells the workflow manager what needs to be done for the particular SCOOP application.

The quality assurance (QA) component reflects the fact that contributing to the archive must undergo quality assurance procedures to insure integrity of the data and transfers (checksums, etc.) and to document data provenance. The quality control (QC) component represents automated procedures and tools for quality control, including comparisons between model output and actual observations, which might trigger quality control flags or warnings when predicted values exceed threshold values.

The “user portals and clients” node represents a variety of interfaces into the system. These include limited-access portals for overall system administration and control, portals tailored to scientific researchers who want ready access to search and access all kinds of results in the archive, and end-user portals that provide customized and relatively limited access to end products of specific value to users.

5 Discussion

5.1 Present Status

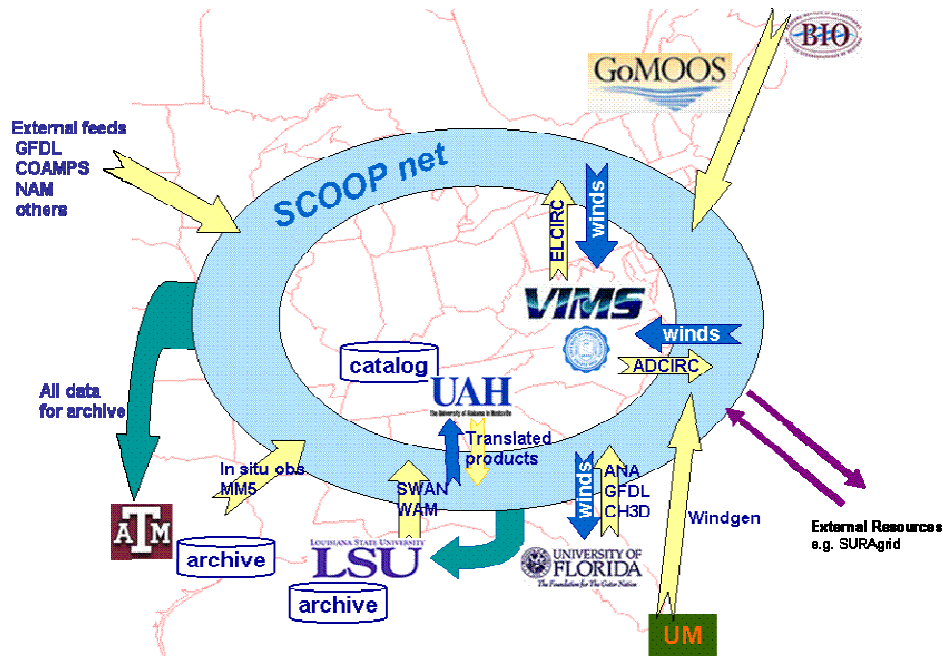


Figure 7. SCOOP architecture showing the institutional partners and their roles.

The SCOOP architecture has been deployed at over a half dozen SURA institutions (Figure 7), and there are working prototypes of the event-driven, ensemble prediction scenario and the 24/7 scenario. The retrospective scenario involves catalog discovery and access to the distributed archives, but as yet has no portal interface to automatically run models in a retrospective fashion. The intention is to leverage established Grid-enabled resources such as SURAGrid TeraGrid, the Open Science Grid, and others, as appropriate. As of this writing, several components from the set of SCOOP model suites were either running or in various stages of deployment on SURAGrid resources.

The SCOOP architecture design and implementation will allow these resource capabilities to be added incrementally. The system and its capabilities evolve with time, as do the

technologies that underlie the architecture design. But the scenarios and architecture evolve at a much slower rate and should therefore survive relatively rapid changes in technology.

To facilitate evolution of the system, the program is implementing a stable “production” environment, and companion “development” and “testing” environments that allow research and development to occur without affecting the production system. System innovation occurs by migrating new or updated components to production after testing.




Figure 8. SURAggrid participants as of March 2007.

SURAggrid (Figure 8) is represented by the resource (lowest) layer in the layer-stack diagram (Figure 5). SURAggrid includes a relatively large and rapidly growing number of SURA institutions. As one might expect in a service-oriented environment, new institutions can be added to SURAggrid on an ongoing basis that is ultimately transparent to the SCOOP application.

5.2 Community Modeling Frameworks – Loose and Tight Coupling

SOA is not a design requirement. Rather, it is a valuable feature that allows multiple institutions to contribute complementary components. SOA is a loosely coupled method to

implement a SCOOP modeling “framework,” meaning that it requires little effort to modularize into a collection of computational tasks for which there is no interdependency. In contrast, some frameworks are tightly coupled because they involve highly coupled models or computational tasks that require constant communication and data exchange, often for computational problems involving highly non-linear phenomena.

Loosely and tightly coupled frameworks should be complementary, however, loose coupling might be inappropriate for certain model suites that require a high degree of interaction between the various components in the suite. This type of application might best be implemented with a tightly coupled framework. Examples include Cactus (<http://www.cactuscode.org>) or the Earth Systems Modeling Framework [ESMF, <http://www.cisl.ucar.edu/research/2005/esmf.jsp>].

6 Conclusions

The SCOOP architecture is being implemented incrementally. The objective is to freeze overall system upgrades on a 12-month cycle, and then revisit system requirements in light of operational achievements and system capabilities. Subsequent development and implementation is based on the most recent production system. Consistent with this approach, Version 2.0 of the SCOOP architecture was in place as of July 1, 2005. In that version, a proof-of-concept SCOOP web portal based on OGC-compliant web services, <http://OpenIOOS.org>, showed real-time predictions of the storm surge in New Orleans and Mississippi before Hurricane Katrina made landfall.

The process of extracting functional requirements and using them to implement the system is iterative, and involves an ongoing dialog between system designers (computer

scientists) and system users (coastal scientists). Engaging the end-user community results from the application of adopted community standards, such as the OGC-compliant, GIS-compatible service interfaces. The SCOOP design team strives to develop a system architecture that provides functional value, yet withstands and outlives the underlying technologies used to implement the functions.

Another critical and overarching design principle involves adopting and implementing open community standards (Conover et al, 2006). This can facilitate effective partnerships between the academic, governmental, and private sectors. The distinction between open standards and open source software is also important. The goal within SCOOP is neither to force partners to adopt open source software, nor to make their software open source, but rather, to achieve interoperability among heterogeneous systems through the adoption of open standards for information exchange that can be implemented with either open source or proprietary solutions.

As the system matures, the small number of universities involved in the initial stages of deployment can readily grow to accommodate new participants and new science scenarios. The goal is to nurture broad participation in a range of increasingly diverse activities that leverage and contribute to the underlying infrastructure.

Acknowledgements

This work is part of the Southeastern Universities Research Association Coastal Ocean Observing and Prediction (SCOOP) Program funded by the Office of Naval Research, Award N00014-04-1-0721 and by NOAA Ocean Service, Award NA04NOS4730254. The SCOOP program is grateful to the SCOOP Implementation Team responsible for its development and operation: Yongcun Hu (BIO); Eric Bridger (GoMOOS); Bret Estrada and Chirag Dekate (LSU, Center for Computation and Technology); Steve Thorpe (MCNC); Tony Cook, Donna Cote, Matt Howard, Larry Flournoy, Charlton Purvis, and Srikanth Sastry (TAMU); Ken Keiser, Matt Smith, and Marilyn Drewry (UAH); Renato Figueiredo and Vladimir Paramygin (UFL);

Lavanya Ramakrishnan and Howard Lander (RENCI); Jian Shen (VIMS); Geoff Samuels (UMiami); Gary Crane, Don Riley and Mary Fran Yafchak (SURA).

References

Atkins, D.E., K.K. Droegemeier, S.I. Feldman, H. Garcia-Molina, M.L. Klein, D.G. Messerschmitt, P. Messina, J.P. Ostriker, and M.H. Wright. 2003. *Revolutionizing Science and Engineering Through Cyberinfrastructure*. Report of the National Science Foundation Blue Ribbon Advisory Panel on Cyberinfrastructure, Arlington, VA [Online] Available: <http://www.nsf.gov/od/oci/reports/toc.jsp> [March 22, 2007]

Booij, N, R.C. Ris, and L. H. Holthuijsen. 1999. A third-generation wave model for coastal regions. 1. Model description and validation. *J. Geophys. Res.*, 104 (C4): 7649-7666.

Conover, H., B. Beaumont, M. Drewry, S. Graves, K. Keiser, M. Maskey, M. Smith, P. Bogden, and J. Bintz. 2006. SCOOP Data Management: A Standards-Based Distributed System for Coastal Data and Modeling. In: *Proceedings of 2006 IEEE International Geoscience & Remote Sensing Symposium & 27th Canadian Symposium on Remote Sensing*, Marine Technology Society, Denver, CO.

Emanuel, K. 2005. Increasing destructiveness of tropical cyclones over the past 30 years. *Nature*, 436: 686-688.

Foster, I., C. Kesselman and S. Tuecke. 2001. The Anatomy of the Grid: Enabling Scalable Virtual Organizations. *International J. Supercomputer Applications*, 15(3)

Foster, I. 2005. Service-Oriented Science. *Science*, 308: 814-817.

Holland, G.J. 1980. An analytic model of the wind and pressure profiles in hurricanes. *Monthly Weather Review*. 8(108):1212-1218.

Knabb, R.D., J.R. Rhome, and D.P. Brown. 2006. *Tropical Cyclone Report: Hurricane Katrina, 23-30 August 2005*. National Hurricane Center, National Weather Service. [Online] Available: <http://www.nhc.noaa.gov/2005atlan.shtml> [March 22, 2007]

Luettich, R.A., J.J. Westerink, and N.W. Scheffner. 1992. *ADCIRC: An advanced three-dimensional circulation model for shelves, coasts and estuaries: Report 1: Theory and*

methodology of ADCIRC-2DDI and ADCIRC-3DL. U.S. Army Engineers Waterways Experiment Station Dredging Research Program Technical Report DRP-92-6. Vicksburg, MS

Malone, T. and M. Hemsley. 2007. Developing the IOOS for improved management and mitigation of coastal inundation. *Marine Technology Society Journal*, 40(4):45-55.

National Climatic Data Center. 2006. *Climate of 2005 Atlantic Hurricane Season*. U.S. Department of Commerce. [Online] Available: <http://www.ncdc.noaa.gov/oa/climate/research/2005/hurricanes05.html>. [March 22, 2007]

Ocean.US. 2006. *The First U.S. Integrated Ocean Observing System (IOOS) Development Plan: A Report of the National Ocean Research Leadership Council and the Interagency Committee on Ocean Science and Resource Management Integration*. The National Office for Integrated and Sustained Ocean Observations. Ocean.US Publication No. 9, Washington, D.C.

Shen, J, H. V. Wang, Sisson, M., Gong, W. (2006): "Storm tide simulation in the Chesapeake Bay using an unstructured grid model." *Estuarine, Coastal and Shelf Science*, 68 (1-2), 1-16

Sheng, Y.P. 1987. On Modeling Three-Dimensional Estuarine and Marine Hydrodynamics, in *Three-Dimensional Models of Marine and Estuarine Dynamics* (J.C.J. Nihoul and B.M. Jamart, Eds.), Elsevier Oceanography Series, Elsevier, pp. 35-54

Sheng, Y.P. 1990 Evolution of a Three-Dimensional Curvilinear-Grid Hydrodynamic Model for Estuaries, Lakes and Coastal Waters: CH3D. In: *Proceedings of the First International Conference on Estuarine and Coastal Modeling*, ASCE pp. 40-49 (1989).

Smith, K.T. 1999. Estimating the Costs of Hurricane Evacuation: A Study of Evacuation Behavior and Risk Interpretation using Combined Revealed and Stated Preference Household Data. Department of Economics, East Carolina University. [Online] Available: <http://www.ecu.edu/econ/ecer/kevinsmith.pdf> [March 22, 2007]

Tolman, H.L., 2002: *User Manual and system documentation of WAVEWATCH-III Version 2.22*. Technical Note. U.S. Department of Commerce, 133 pp.

Tolman, H. L., B. Balasubramanian, L. D. Burroughs, D. V. Chalikov, Y. Y. Chao, H. S. Chen, and V. M. Gerald. 2002: Development and implementation of wind generated ocean surface wave models at NCEP. *Weather and Forecasting*, 17: 311-333.

Wang, H. V., Joe Cho, Jian Shen, and Ya Ping Wang (2004): "What Has been learned about storm surge dynamics from Hurricane Isabel model simulations?" Hurricane Isabel in Perspective Conference, November 15-17, 2004, Baltimore, MD. p117-125

Whitehead, J.C. 2000. One Million Dollars a Mile? The Opportunity Costs of Hurricane Evacuation. Department of Economics, East Carolina University. [Online] <http://www.ecu.edu/econ/wp/00/ecu0005.pdf>. [March 22, 2007]

Zhang Y.-L., A.M. Baptista, and E.P. Myers, 2004. A cross-scale model for 3D baroclinic circulation in estuary-plume-shelf systems: I. Formulation and skill assessment. *Cont. Shelf Res.* 24:2187-2214