Integrated Modeling of the Battlespace Environment

The goal of the Battlespace Environments Institute (BEI) is to integrate Earth and space modeling capabilities into a seamless, whole-Earth common modeling infrastructure that facilitates interservice development of multiple, mission-specific environmental simulations and supports battlefield decisions, improves interoperability, and reduces operating costs.

Characterizing the natural environment is crucial to US Department of Defense (DoD) mission planning because understanding battlespace conditions enhances both safety and warfighting effectiveness. Historically, DoD production centers have used stand-alone models—such as those for weather and ocean conditions—which have associated maintenance costs. Although such models continue to improve, alone they can provide only an incomplete representation of the environmental conditions that might impact a DoD mission.

1521-9615/10/$26.00 © 2010 IEEE
Copublished by the IEEE CS and the AIP

Tim Campbell, Richard Allard, Ruth Preller, Lucy Smedstad, and Alan Wallcraft
Naval Research Laboratory, Stennis Space Center
Sue Chen and Hao Jin
Naval Research Laboratory, Monterey
Saša Gaberšek
University Corporation for Atmospheric Research
Richard Hodur
Science Applications International Corporation
Joseph Reich
Air Force Weather Agency, Offutt Air Force Base
Craig D. "Ghee" Fry
Exploration Physics International
Vince Eccles
Space Environments Corporation
Hwai-Ping Cheng, Jing-Ru C. Cheng, and Robert Hunter
US Army Engineer Research & Development Center, Vicksburg
Cecelia DeLuca
NOAA Cooperative Institute for Research in Environmental Sciences
Gerhard Theurich
Science Applications International Corporation

Environmental processes interact on multiple time scales, and many such processes interact on time scales that are short enough to be significant to the DoD. Our environmental subsystems therefore must be coupled into a larger interacting system. The problem, however, is that any single service lacks adequate resources to develop a complete, coupled prediction capability for the battlespace environment.

Creating coupled modeling systems using a standard DoD modeling framework will foster collaborative efforts throughout the DoD and facilitate partnerships with outside organizations. We established the Battlespace Environments Institute (BEI) with a vision of multi-agency and multiservice collaboration for the rapid development and transition of new models to support mission planning. BEI stakeholders include the US Navy, Air Force, and Army; National Aeronautics and Space Administration (NASA); Department of Energy; Department of Commerce; and National Science Foundation.

BEI's goal is to integrate Earth and space modeling capabilities into a seamless, whole-Earth common modeling infrastructure to allow interservice development of multiple, mission-specific environmental simulations to support battlefield decisions, improve interoperability, and reduce operating costs. To develop a whole-Earth infrastructure that excluded components originating outside the DoD would be cost prohibitive. Given that the environmental community had already invested heavily in the Earth System Modeling Framework (www.earthsystemmodeling.org), the DoD decided to not only use ESMF as the basis for a common modeling infrastructure, but also to invest in ESMF to address DoD-specific needs.

ESMF provides the basic software layer for implementing a whole-Earth environment. However, it doesn't mandate how components interact. A major BEI task has thus been to design and implement rules that a component must follow to be part of the DoD whole-Earth system. To encourage progress in this area, BEI has focused resources on a few projects aimed at developing prototypes of restricted domains (such as littoral and air-ocean environments). It will then expand the projects' experiences to a whole-Earth environment modeling capability.

Here, we offer an overview of the ESMF software architecture and design strategies, and then highlight specific BEI projects focused on various whole-Earth system subdomains.

Earth System Modeling Framework

The ESMF is open-source software for building climate- and weather-related modeling components and coupling them together to form applications. ESMF was motivated by the desire to exchange modeling components among centers and to reduce costs and efforts by sharing codes. Existing software-framework efforts heavily influenced ESMF's design.[1] The project is distinguished by its strong emphasis on community governance and distributed development, and by a diverse customer base that includes modeling groups from universities, major US research centers, the US National Weather Service, the DoD, and NASA. The ESMF development team is centered at the US National Oceanic and Atmospheric Administration (NOAA) Cooperative Institute for Research in Environmental Sciences.

Architectural Overview

The ESMF architecture is based on the components concept. At its simplest, a software component is a code with a well-defined calling interface and a coherent function.[2] Component-based design is a natural fit for Earth system modeling because components are ideally suited to represent a system with substantial, distinct, and interacting domains—such as atmosphere, land, sea ice, and ocean. Further, because Earth system domains are often studied and modeled as collections of subprocesses (radiation and chemistry in an atmosphere, for example), it's convenient to model Earth system applications as a nested-components hierarchy.

Component-based software is also well suited to the way Earth system models are developed and used. Individual specialists typically develop a model's multiple domains and processes as separate codes. Creating a viable environment application requires the integration, testing, and tuning of the pieces—a scientifically and technically formidable task. When we can represent each piece as a component with a standard interface and behavior, then integration—at least at the technical level—is more straightforward. Similarly, standard interfaces help foster component interoperability and component use in different contexts. This is a primary concern for modelers because they're motivated to explore and maintain alternative algorithm versions (such as different implementations of the atmosphere's governing fluid equations), whole physical domains (such as oceans), parameterizations (such as convection schemes), and configurations (such as stand-alone versions of physical domains).

ESMF Components

ESMF has two types of components: *gridded components* (`ESMF_GridComp`) represent a model's scientific and computational functions, while *coupler components* (`ESMF_CplComp`) contain the operations necessary to transform and transfer data between them. Both gridded and coupler components are implemented in the Fortran interface as derived types with associated modules. Because ESMF doesn't currently contain prefabricated gridded or coupler components, users must write them. The ESMF documentation and source distribution provide tools and examples to guide users through this task.

Each major physical domain in an Earth system model is represented as an ESMF gridded component with a standardized calling interface and arguments. Physical processes or computational elements, such as radiative processes or I/O, also can be represented as gridded components. ESMF components can be nested, so that parent components can contain child components with progressively more specialized processes or refined grids.

As a model steps forward in time, the physical domains represented by gridded components must periodically transfer interfacial fluxes. The operations necessary to couple gridded components together might involve data redistribution, spectral or grid transformations, time averaging, and unit conversions. In ESMF, a coupler component encapsulates these interactions. Coupler and gridded components share the same standard interfaces and arguments. The interfaces' key data structure is the `ESMF_State` object, which holds the data to be transferred between components.

Each gridded component is associated with an import state containing the data required for it to run, and an export state containing the data it produces. Coupler components arrange and execute data transfer from producer-gridded components' export states into consumer-gridded components' import states. The same gridded component can be a producer or consumer at different times during model execution.

There's no single, generic coupler component for all ESMF applications. Modelers write coupler component internals using ESMF classes bundled with the framework. These classes include methods for time advancement, data redistribution, interpolation weight calculation, interpolation weight application through a sparse matrix multiply, and other common functions.

Users can write coupler components to transform data between a pair of gridded components or use a single coupler component to couple more than two gridded components. Multiple couplers can be included in a single modeling application. This is a natural strategy when the application is structured as a component hierarchy. Each level in the hierarchy usually has its own set of coupler components.

Design Goals and Strategies

Design goals for ESMF applications include the ability to

- use the same gridded component in multiple contexts,
- swap different gridded component implementations into an application, and
- assemble and extend coupled systems easily.

In short, the goal is software reuse and interoperability.

One design pattern that addresses these goals is the mediator pattern, in which one object encapsulates how a set of other objects interact.[3] The mediator serves as an intermediary and prevents objects from referring to each other explicitly. ESMF coupler components follow this pattern. It's an important aspect of the ESMF technical strategy because it lets users deploy an application's gridded components in multiple contexts—that is, it lets them be used in different coupled configurations without changing the source code. For example, a user might couple the same atmosphere to an ocean in a hurricane prediction model and to a data assimilation system in a numerical weather prediction model.

Another mediator pattern advantage is that it promotes a simplified view of intercomponent interactions. The mediator encapsulates all the complexities of data transformation between components. However, this can lead to excessive complexity within the mediator itself.[3] One approach to addressing this is to create multiple, simpler coupler components and predictably embed them in a hierarchical architecture. The degree to which we can apply a hierarchical approach depends on the intercomponent interactions' nature.

Computational environment and throughput requirements motivate a different set of design strategies. ESMF component wrappers must not impose significant overhead and must operate efficiently on a wide range of computer architectures, including desktop and petascale supercomputers. To satisfy these requirements, the ESMF software relies on memory-efficient and highly scalable algorithms, such as that by Karen Devine and her colleagues.[4] ESMF runs efficiently on tens of thousands of processors.

How users map a modeling application's components to computing resources can significantly impact performance. Strategies vary for different computer architectures, and ESMF is flexible enough to support multiple approaches. ESMF components can run sequentially (one following the other, on the same computing resources), concurrently (at the same time, on different computing resources), or by combining these execution modes. Most ESMF applications run as a single executable—that is, all components are combined into one program. Starting at a top-level driver, each level of an ESMF application controls the partitioning of its resources and the next lower level's component sequencing.

As we now describe, these goals and strategies have been implemented in various Earth system subdomains, including space and marine weather and the coastal watershed.

Space Weather Modeling

Giving military commanders actionable weather information is at the heart of the Air Force

Weather Agency's mission. Space weather is no different. To give commanders ample time to prioritize missions, a major goal of space weather is to forecast for the 120-hour air-tasking-order cycle. To accomplish this in the ionosphere, upstream information is required about both solar irradiance and the solar wind, which drives high-latitude currents that affect the plasma density at high, middle, and low latitudes. The end result is longer, more accurate forecasts of ionospheric conditions.

The Space Weather Modeling System (SWMS) is a BEI project that couples two space environment models under ESMF: the Hakamada-Akasofu-Fry version 2 (HAFv2) solar wind model and the global assimilation of ionospheric measurements (GAIM1) forecast component. The HAFv2 model ingests solar observations and provides the outputs to GAIM1 to forecast the time-dependent energy input into the high-latitude ionosphere. The resulting output must be consistent with the predevelopment code, scalable, and portable. The coupled HAFv2-GAIM1 will show payoff to AFWA operations by providing the first quantitative forecasts of ionospheric conditions that extend days into the future.

Hakamada-Akasofu-Fry Version 2

Exploration Physics International (EXPI) developed HAFv2 to predict solar wind conditions at the Earth and elsewhere in the solar system days in advance.[5–7] The HAFv2 model is designed to track interplanetary disturbances' progress after solar events. Predicted solar wind parameters of speed, density, dynamic pressure, and interplanetary magnetic field are key inputs to numerical prediction models that forecast near-Earth space weather disturbances such as geomagnetic storms, enhanced energetic particle fluxes, and ionospheric disturbances.

The HAFv2 model is driven by solar event reports and by synoptic solar observations in the form of source surface maps of radial magnetic fields and speeds at 2.5 solar radii.[8,9] The source surface maps are introduced at the inner boundary to initialize the HAFv2 model run. Event reports are converted into time-dependent perturbations that modulate the model's inner boundary. Given these observation-based inputs, HAFv2 predicts a time series of solar wind values at Earth. Spacecraft such as the Advanced Composition Explorer located in a stable orbit upstream (sunward) of the Earth provide the ground-truth measurements for comparison with the solar wind predictions.

Global Assimilation of Ionospheric Measurements

The GAIMv2.3 effort[10,11] merges ionospheric observations with state-of-the-art ionospheric model results using a Gauss-Markov Kalman Filter to produce a real-time weather description of the ionosphere. The physics-based ionosphere model in GAIMv2.3—that is, its GAIM1 component—is the Space Environment Corporation's ionosphere forecast model. The IFM is a global ionosphere model from 90 to 1,400 kilometers. It solves for electron and ion densities as well as electron and ion temperatures.[12] The solar wind and solar spectrum are the primary inputs into IFM.

Coupling the Models

GAIMv2.3 provides a 24-hour forecast of ionospheric conditions assuming persistence of the current day's geophysical conditions. Without solar inputs, ionospheric forecasts relax to ionospheric climatology. One-way coupling of HAFv2 to GAIM1 links the solar storm drivers to the ionospheric response. HAFv2 provides the solar wind speed, density, and magnetic field to GAIM1, enabling multiday forecasts of ionospheric electron density, currents, and upper-atmosphere dynamics.

Figure 1 shows the HAFv2 forecasts from 1–29 January 2000, a highly active space weather period at the 11-year solar cycle's peak. HAFv2 simulates the time-dependent interplanetary magnetic field (Figure 1a) that results from the input solar conditions. The predicted solar wind quantities at Earth (Figure 1b) are fed as inputs to GAIM1. Figure 2 shows the GAIM1's subsequent prediction of the high-latitude response of Joule heating before and after the 27–28 January storm's onset.

SWMS development is a structured project with well-defined milestones, moving from partial to full adoption of ESMF. The SWMS's coupled HAFv2-GAIM1 components incorporate the necessary infrastructure and superstructure to achieve full ESMF adoption. The SWMS oversees initialization, running, coupling, and finalization of the HAFv2 and GAIM1 components via ESMF component calls. In addition to coupling, the project also imposes requirements for portability, scalability, and accuracy. We achieved HAFv2 and GAIM1 model scalability through extensive restructuring and implementing parallel algorithms.

Bringing the HAF and GAIM models together in SWMS has enabled significant improvements in space weather data processing and throughput.

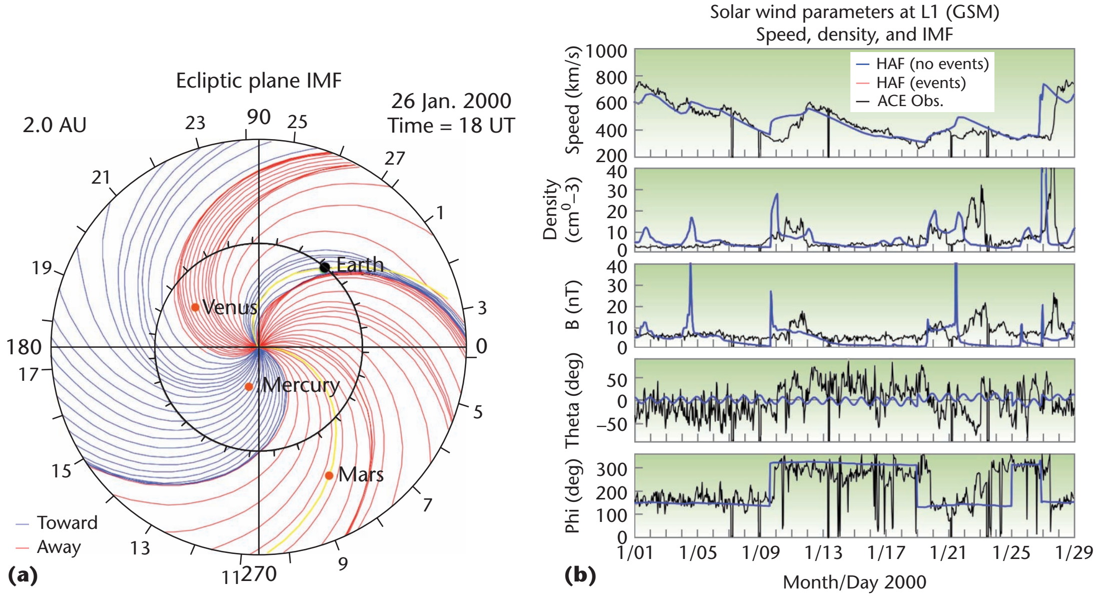

Figure 1. Hakamada-Akasofu-Fry version 2 (HAFv2) solar wind forecasts from 1–29 January 2000. (a) HAFv2 simulation of the interplanetary magnetic field (IMF) in the ecliptic plane to 2 astronomical units (AU) for 26 January 2000. The Earth's orbit is at 1 AU; its location is shown by the black dot. The red lines represent outward-directed regions of IMF, while the blue represent inward-directed regions. (b) HAFv2 simulation of solar wind speed, density, and IMF magnitude (*B*), north-south angle (θ), and azimuthal angle (φ) during 1–29 January 2000.

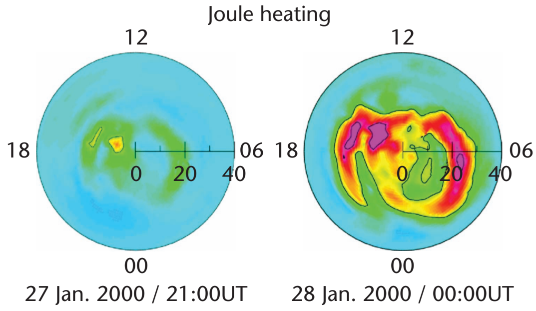

Figure 2. The global assimilation of ionospheric measurements' (GAIM1's) HAFv2-driven predictions of the high-latitude response of Joule heating before and after the 27–28 January 2000 storm's onset. The prediction was calculated in the northern polar region using the HAFv2-GAIM1 coupled models. The Joule heating magnitude is a key component of the ionospheric forecast.

Modifying the codes for the high-performance computing (HPC) environment brings other new capabilities, including the ability to

- ingest diverse data sets at a higher resolution and cadence,
- use denser computational grids, and
- perform ensemble forecasts.

The result is not just more accurate ionospheric forecasts for DoD missions, but also improved solar wind, geomagnetic, and thermospheric forecasts for the DoD and other government and commercial users.

Weather and Marine Prediction

Given the continuing global warming trend, there's a pressing need to better understand the interactions between the atmosphere and ocean. How these two systems respond together to the rising temperature dictates our ability to project future climate change and its impact on shorter time-scale prediction. Coupling different, sophisticated atmosphere, ocean, and wave models to form one superior system is therefore an important approach for capturing many physical and dynamical processes that govern the air-sea interaction.

The Naval Research Laboratory (NRL) Coupled Ocean and Atmosphere Mesoscale Prediction System (COAMPS) is a high-resolution, fully coupled air-ocean-wave system.[13] As Figure 3 shows, COAMPS' ocean circulation model is the limited-area version of the Navy Coastal Ocean Model (NCOM)[14,15] and it can incorporate the Simulating Waves Nearshore (SWAN) and Wavewatch III wave models. Efforts are underway to also include the Hybrid Coordinate Ocean Model (HYCOM; www.hycom.org).

At the heart of COAMPS is a driver/coupler that controls the time-stepping and coordinates field exchange between components. For each component pair, the coupler computes a sparse matrix that combines the weights for interpolating between grids and the extrapolation weights for treating land-sea boundary mismatches. The ESMF parallel sparse matrix multiply efficiently handles the grid transformations at each coupling interval. In contrast to loosely coupled systems, the coupling's scalability and efficiency allows for tight model integration. The coupling interval can be as small as the least common multiple of the coupled components' time steps. In typical applications, with coupling intervals of about two to three atmospheric time steps, the coupling overhead is less than 1 percent of the overall computation time. Background ocean and wave components are included to support flexible model setup and improved relocation capability.

Development of the fully coupled air-ocean system was completed in mid-2008 and represents the first limited-area weather and ocean prediction system that uses ESMF to couple air and ocean models. Sue Chen and her colleagues offer a detailed description of the system along with results of two test cases.[16] The hurricane Katrina test case is of particular interest because hurricane and tropical cyclones have a tremendous impact on the safety of coastal communities and DoD operations.

One highlight of the researchers' study was the coupled system's ability to simulate realistic hurricane intensity reduction and structural change because of the hurricane-induced trailing ocean cold wake. The atmospheric response to the ocean cold wake was significantly reduced heat and moisture fluxes from the ocean and an increase in the storm's flow asymmetry. Analysis of the ocean temperature budget suggested that within the ocean's mixed layer, vertical advection (upwelling) and wind mixing contribute equally to the cold wake's generation near the storm's center. Additional cooling found to arise in the cold wake along its right-front quadrant was because of horizontal advection of colder water forced upward along the storm track. These results clearly indicate the inadequacy of applying 1D mixed-layer models—which ignore horizontal advection—to capture the full impact of the air-ocean interaction, which impacts the sea-surface-temperature cooling's development and structure in a mature hurricane's wake.

For strong wind conditions, the air-sea energy exchange from ocean surface waves is non-negligible. Therefore, accurately depicting air-ocean interaction in a coupled system must include a wave model. We've extended the air-ocean coupled system to include the SWAN wave model to study the air-sea interaction under hurricane conditions.

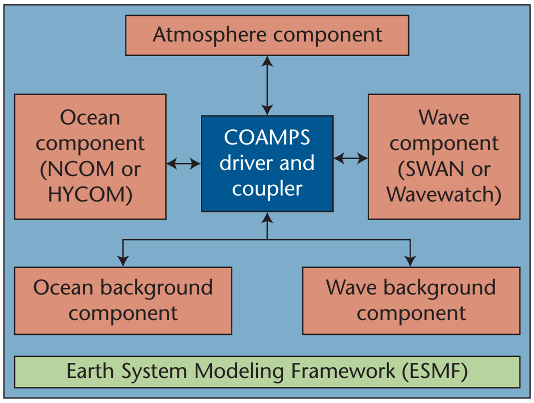

Figure 3. The Coupled Ocean and Atmosphere Mesoscale Prediction System (COAMPS) integrates air, ocean, and wave models for weather and marine predictions. The system includes a limited area version of the Navy Coastal Ocean Model (NCOM) and can incorporate the Simulating Waves Nearshore (SWAN) and Wavewatch III wave models.

We ran a series of sensitivity tests on the Katrina test case to investigate the impact of the atmospheric wind, sea-surface height, and ocean current influences on wave growth. Comparing the significant wave height with several buoys from the National Buoy Data Center suggests reasonable agreement with the observed wave state. As Figure 4a shows, the coupled model has a longer and higher wave developed along the storm front quadrant. Compared to the run without the ocean current effect on waves, using the ocean current model reduced the hurricane-induced wave growth in the storm front quadrant (Figure 4b). Although the precise wave feedback to the atmospheric and ocean models is still in basic research and has yet to be implemented, our results do show some wave-growth sensitivity to the air-ocean-wave coupling.

The NRL in-house research and development projects have already experienced a significant impact since they successfully transitioned to the ESMF-based coupled air-ocean system. NRL scientists have applied the coupled system to study and validate many different types of weather and ocean scenarios that are strongly influenced by air-sea interaction, including

- the cold ocean upwelling along the US west coast and South Chile;

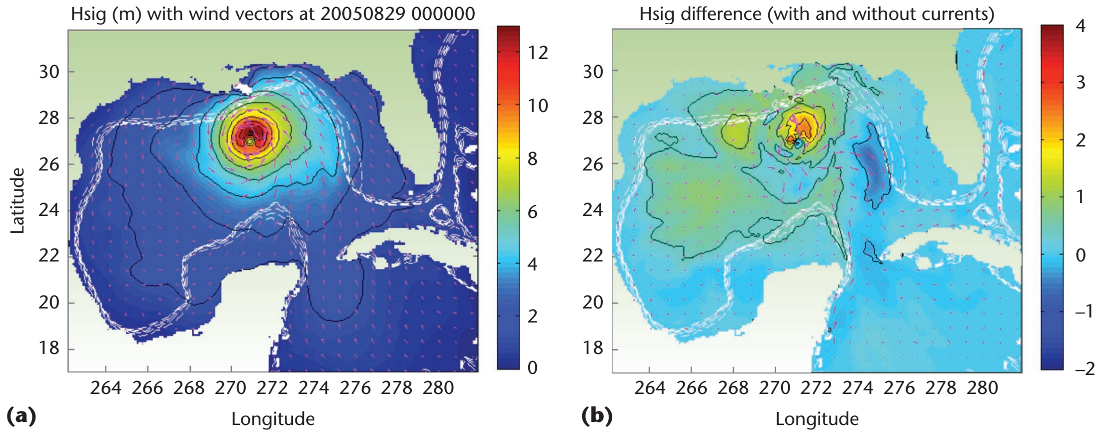

Figure 4. The impact of wind, sea-surface height, and ocean currents on wave growth. (a) There were significant wave heights (m) and 10-meter atmospheric winds (red vector) before the fully coupled model-simulated hurricane Katrina made landfall. (b) The significant difference in wave height between model runs with and without current input to the wave model.

- a study of the Kuroshio extension current in the Western Pacific;
- an analysis of the Mistral wind jets in the Adriatic;
- Madden-Julian Oscillation studies in the Indian Ocean; and
- tropical cyclones in the Pacific, Indian Ocean, and Atlantic basins.

NRL scientists can also seek new funding opportunities not previously possible before the fully coupled mode was developed. BEI's air-ocean-wave coupling framework is a foundation for further development and research of next-generation limited-area high-resolution weather and marine prediction models for the US Navy. The project's new technology will also be transferred to non-DoD partners. The COAMPS coupler, for example, will be used as a prototype in a joint university, NOAA, and NRL initiative to build a unified air-sea interface module for future-generation hurricane research and operational models in the US.

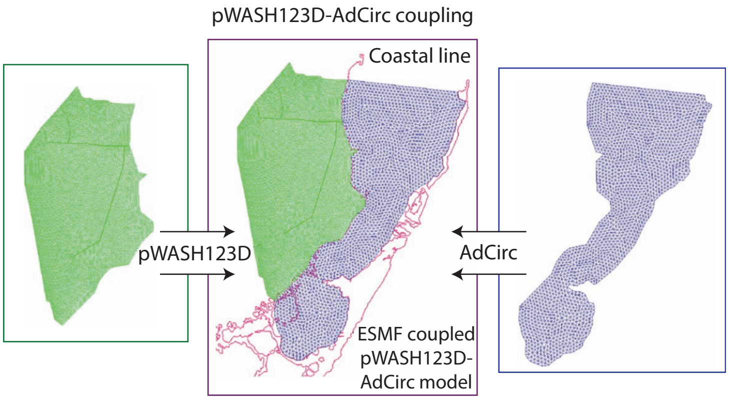

Figure 5. Coupled watershed and coastal ocean model grids for the Biscayne Bay region. Coupling the watershed model (pWASH123D) with the coastal ocean model (AdCirc) improves modeling and prediction of hydrodynamics in these areas.

Coupled Watershed Nearshore Modeling

Realistic modeling of flow and transport in hydrosystems is important to infrastructure planning, environmental remediation, and ecosystem restoration. In the coastal areas and in estuaries, where fresh and salt waters meet and interact, the ability to simulate coupled salinity transport and density-dependent flow is essential for accurately modeling water quantity and quality. Although there are many channel, overland, groundwater, watershed, nearshore, and ocean models that can compute both water flow and salinity transport, most of them are stand-alone models and few were designed or later parallelized for HPC. As a result, their applicability to real-world problems is restricted.

To improve modeling of the watershed and coastal regimes, the Army Engineer Research and Development Center in Vicksburg, Mississippi, collaborated with the NRL at the Stennis Space Center to couple a watershed model (pWASH123D) with a coastal ocean model (AdCirc). Figure 5 shows an example application of the coupled watershed nearshore model for Florida's Biscayne Bay region.

The pWASH123D[17,18] model is a physically based finite-element numerical model that computes watershed systems' water flow and simulates them as combinations of 1D channel networks, 2D overland regimes, and 3D subsurface media. The interactions between different media (between 1- and 2D, 2- and 3D, and 1- and 3D) impose flux continuity and state variable continuity on the medium interfaces. The pWASH123D model aims to efficiently simulate the regional scale of real-world problems on HPC machines. We implement different parallel algorithms and partitioning strategies in different components to maintain load balance and reduce communication overhead. The model implemented is for parallel, distributed memory platforms and consists of a mix of C/C++ and Fortran code.

AdCirc (www.adcirc.org) is a coastal circulation and storm surge model that uses finite-element computer programs for solving time-dependent, free surface circulation and transport problems in 2D and 3D. The model implements the continuous Galerkin finite-element method based on the generalized wave-continuity equation. AdCirc is written in Fortran and developed for parallel, distributed memory platforms; typical applications include

- modeling tides and wind-driven circulation,
- analyzing hurricane storm surges and flooding,
- dredging feasibility and material disposal studies,
- larval transport studies, and
- nearshore marine operations.

One challenge this project faced involved data exchange in memory between two models using unstructured meshes written in different programming languages. When this project was initiated, ESMF didn't have functionality for unstructured meshes. To allow the watershed modeling project to evolve with ESMF, we decided to use the framework to construct the components and manage data flow to and from the coupler. In the coupler, however, we used DBuilder[19] to interpolate between the two model grids. The DBuilder toolkit is a parallel data management library for scientific applications that's part of the pWASH123D infrastructure. Developers use DBuilder to implement parallel versions of their codes; the toolkit provides a simple and consistent interface that hides many of the programming details associated with domain partitioning, parallel data management, domain coupling, and invoking parallel linear solvers. Because DBuilder supports coupling independent domains in a single model, adopting it to exchange data among multiple models was a straightforward task.

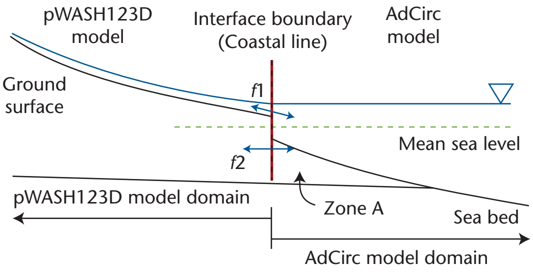

Figure 6. A side view of the interface between the pWASH123D and AdCirc model domains. Here, f1 represents the boundary flux from surface flow, and f2 represents the subsurface flow.

Coupling uses the coastal (shore) line as the coupling interface. Figure 6 shows a side view perpendicular to the coupling interface. In Figure 5, the elevation drop on the AdCirc's side of the interface boundary represents a possible scenario in which AdCirc is run without wetting and drying and, hence, imposes a minimum water depth. We use the water surface elevation simulated by AdCirc at the coupling interface as a Dirichlet-type boundary condition for computing channel flow and overland flow in the surface water system. With the hydrostatic assumption, this computed surface elevation is distributed to the subsurface nodes beneath the coastal line for computing subsurface flow.

The boundary flow associated with all the coupling interface's pWASH123D nodes are computed and then distributed to the interface's AdCirc boundary nodes as a flux-type boundary condition. The pWASH123D boundary nodal flow includes contributions from both surface and subsurface flow. In Figure 6, f1 represents the boundary flux from surface flow, and f2 represents the subsurface flow. In reality, part of f2 passes through Zone A and enters and exits the AdCirc model across the dashed line as sinks and sources, respectively. In this case, Zone A must be included in the modeled domain. The current-coupling approach, however, simplifies this situation by setting f2 the subsurface contribution to only the coupling interface's AdCirc boundary nodes.

We couple the models in a concurrent mode and time lag the boundary-forcing fields. At each coupling time step, we use the computed average boundary flux and water surface elevation from the previous coupling time step at the coupling interface's boundary conditions. The simulation proceeds to the next coupling time after both models have finished computing the current coupling time step.

The pWASH123D-AdCirc model is the first coupled unstructured-mesh model using ESMF. Developers can apply the model's features and capabilities to couple structured- and unstructured-mesh models. When compared to stand-alone models, the coupled model exhibits more accurate boundary conditions applied to the interface boundary between the watershed and the coastal models, especially when surface-subsurface and watershed coastal ocean interactions are significant during storm events. This better models the nearshore area's hydrosystem, improving our understanding of its complex interactions. This, in turn, benefits the construction and verification of water management plans and other water-related environmental issues. By adding salinity and reactive transport capabilities in the future, the coupled model will serve as a cutting-edge, design-level modeling tool for sustainable hydrosystem and ecosystem restoration.

In addition to improving ESMF, we've made significant steps toward the goal of integrating environmental modeling capabilities into ESMF as a common modeling infrastructure. Through BEI's focus projects, we've demonstrated payoff for DoD environmental research and operations, and made progress toward the overall goal of a whole-Earth environment modeling capability. However, many challenges remain. To achieve interoperability of Earth system components, it's necessary to define a common physical architecture—which physical processes each component contains and how those components are interconnected—and establish metadata and usage conventions. We're making progress in addressing these challenges through participation in larger community efforts, such as the Earth System Curator project[20] and The National Unified Operational Prediction Capability project (www.weather.gov/nuopc).

Acknowledgments

BEI is sponsored by the US DoD HPC Modernization Program (HPCMP) and we carried out our simulations at HPCMP supercomputing resource centers. We thank the NRL's Paul Martin, Paul May, James Doyle, and Eric Rogers for contributing both insight and code to the COAMPS project; and the ERDC's Jerry Lin for model construction, data analyses, and mesh generation. Support and funding for ESMF is through the DoD HPCMP, the NASA Modeling Analysis and Prediction Program, the NOAA National Weather Service and Climate Program Office, and the US National Science Foundation. COAMPS is a registered trademark of the US NRL.

References

1. C. Hill et al., "Architecture of the Earth System Modeling Framework," *Computing in Science & Eng.*, vol. 6, no. 1, 2004, pp. 18–28.
2. C. Szyperski, *Component Software*, Addison-Wesley, 2002.
3. E. Gamma et al., *Design Patterns: Elements of Reusable Object-Oriented Software*, Addison-Wesley, 1995.
4. K. Devine et al., "Zoltan: Data Management Services for Parallel Dynamic Applications," *Computing in Science & Eng.*, vol. 4, no. 2, 2002, pp. 90–97.
5. K. Hakamada and S.-I. Akasofu, "Simulation of Three-Dimensional Solar Wind Disturbances and Resulting Geomagnetic Storms," *Space Science Rev.*, vol. 31, no. 1, 1982, pp. 3–70.
6. C.D. Fry et al., "Improvements to the HAF Solar Wind Model for Space Weather Predictions," *J. Geophysical Research*, vol. 106, no. A10, 2001; doi:10.1029/2000JA000220.
7. C.D. Fry, "Forecasting Solar Wind Structures and Shock Arrival Times Using an Ensemble of Models," *J. Geophysical Research*, vol. 108, vol. A2, 2003; doi:10.1029/2002JA009474.
8. C.N. Arge and V. Pizzo, "Improvement in the Prediction of Solar Wind Conditions Using Near-Real-Time Solar Magnetic Field Updates," *J. Geophysical Research*, vol. 105, no. A5, 2000; doi:10.1029/1999JA000262.
9. C.N. Arge, "Stream Structure and Coronal Sources of the Solar Wind During the May 12th, 1997 CME," *J. Atmospheric and Solar-Terrestrial Physics*, vol. 66, nos. 15–16, 2004, pp. 1295–1309.
10. R.W. Schunk, L. Scherliess, and J.J. Sojka, "Recent Approaches to Modeling Ionospheric Weather," *Advances in Space Research*, vol. 31, no. 4, 2003, pp. 819–828.
11. R.W. Schunk et al., "Global Assimilation of Ionospheric Measurements (GAIM)," *Radio Science*, vol. 39, no. RS1SO2, 2004; doi:10.1029/2002RS002794.
12. R.W. Schunk, J.J. Sojka, and J.V. Eccles, *Expanded Capabilities for the Ionospheric Forecast Model*, tech. report AFRL-VS-HA-TR-98-0001, Space Environment Corp., 1997.
13. S. Chen et al., *COAMPS Version 3 Model Description*, tech. report NRL/PU/7500-04-448, Naval Research Laboratory, 2003.
14. P.J. Martin, *Description of the Navy Coastal Ocean Model Version 1.0*, tech. report NRL/FR/7322-00-9962, Naval Research Laboratory, 2000.
15. P.J. Martin et al., "User's Manual for the Navy Coastal Ocean Model (NCOM) Version 4.0," tech. report NRL/MR/7320-08-9151, Naval Research Laboratory, 2008.
16. S. Chen et al., "Effect of Two-Way Air-Sea Coupling in High and Low Wind Speed Regimes," *Monthly*

Weather Review, 2010; http://journals.ametsoc.org/doi/abs/10.1175/2009MWR3119.1.

17. J.-R. Cheng et al., "A Parallel Software Development for Watershed Simulations," *Proc. Int'l Conf. Computational Science*, LNCS 3514, Springer Verlag, 2005, pp. 460–468.
18. R.M. Hunter and J.-R. Cheng, "DBuilder: A Parallel Data Management Toolkit for Scientific Applications," *Proc. Int'l Conf. Parallel and Distributed Processing Techniques and Applications* (PDATA'05), CSREA Press, 2005, pp. 825–831.
19. G.-T. Yeh et al., "A First-Principle, Physics-Based Watershed Model: WASH123D," *Watershed Models*, V.P. Singh and D.K. Frevert, eds., CRC Press, 2006, pp. 211–244.
20. R. Dunlap et al., "Earth System Curator: Metadata Infrastructure for Climate Modeling," *Earth Science Informatics*, vol. 1, nos. 3–4, 2008, pp. 131–149.

Tim Campbell *is a computer scientist in the Naval Research Laboratory's Oceanography Division at the Stennis Space Center in Mississippi. Contact him at tim.campbell@nrlssc.navy.mil.*

Richard Allard *is head of the Oceanography Division's Nearshore and Coupled Models section of the Naval Research Laboratory at the Stennis Space Center in Mississippi. Contact him at richard.allard@nrlssc.navy.mil.*

Ruth Preller *is superintendent of the Oceanography Division at the Naval Research Laboratory at the Stennis Space Center in Mississippi and head of the Battlespace Environments Institute. Contact her at ruth.preller@nrlssc.navy.mil.*

Lucy Smedstad *is a computer scientist in the Naval Research Laboratory's Oceanography Division at the Stennis Space Center in Mississippi. Contact her at lucy.smedstad@nrlssc.navy.mil.*

Alan Wallcraft *is a computer scientist in the Naval Research Laboratory's Oceanography Division at the Stennis Space Center in Mississippi. Contact him at alan.wallcraft@nrlssc.navy.mil.*

Sue Chen *is a meteorologist in the Naval Research Laboratory's Marine Meteorology Division in Monterey, California. Contact her at sue.chen@nrlmry.navy.mil.*

Hao Jin *is a meteorologist in the Naval Research Laboratory's Marine Meteorology Division in Monterey, California. Contact him at hao.jin@nrlmry.navy.mil.*

Saša Gaberšek *is UCAR project scientist in the Naval Research Laboratory's Marine Meteorology Division in Monterey, California. Contact him at sasa.gabersek.ctr@nrlmry.navy.mil.*

Richard Hodur *is a research meteorologist with Innovative Employee Solutions and an adjunct professor of meteorology at the University of Warsaw, Poland. Contact him at richard.hodur.ctr@nrlmry.navy.mil.*

Joseph Reich *is chief of the Space Weather Integration Team, 16th Weather Squadron, Offutt Air Force Base, Nebraska. Contact him at joseph.reich@offutt.af.mil.*

Craig D. "Ghee" Fry *is president and senior research physicist at Exploration Physics International in Huntsville, Alabama. Contact him at gfry@expi.com.*

Vince Eccles *is a research scientist at Space Environment Corporation in Providence, Utah. Contact him at vince@spacenv.com.*

Hwai-Ping Cheng *is a research hydraulic engineer at the US Army Engineer Research and Development Center in Vicksburg, Mississippi. Contact him at hwai-ping.cheng@usace.army.mil.*

Jing-Ru C. Cheng *is a computer scientist at the US Army Engineer Research and Development Center Information Technology Laboratory in Vicksburg, Mississippi. Contact her at Ruth.C.Cheng@usace.army.mil.*

Robert Hunter *is a computer engineer at the US Army Engineer Research and Development Center Information Technology Laboratory in Vicksburg, Mississippi. Contact him at Robert.M.Hunter@usace.army.mil.*

Cecelia DeLuca *leads the Earth System Modeling Framework (ESMF) project, Earth System Curator project, and the Global Interoperability Program at the US National Oceanic and Atmospheric Administration Cooperative Institute for Research in Environmental Sciences in Boulder, Colorado. Contact her at cecelia.deluca@noaa.gov.*

Gerhard Theurich *is the lead architect on the Earth System Modeling Framework core development team. Contact him at theurich@sourcespring.net.*

cn *Selected articles and columns from IEEE Computer Society publications are also available for free at http://ComputingNow.computer.org.*