

Performance of the Unstructured-Mesh, SWAN+ADCIRC Model in Computing Hurricane Waves and Surge

J.C. Dietrich · S. Tanaka · J.J. Westerink · C.N. Dawson · R.A. Luetlich Jr. · M. Zijlema · L.H. Holthuijsen · J.M. Smith · L.G. Westerink · H.J. Westerink

Received: 29 April 2011 / Revised: 28 August 2011 / Accepted: 25 October 2011
© Springer Science+Business Media, LLC 2011

Abstract Coupling wave and circulation models is vital in order to define shelf, nearshore and inland hydrodynamics during a hurricane. The intricacies of the inland floodplain domain, level of required mesh resolution and physics make these complex computations very cycle-intensive. Nonetheless, fast wall-clock times are important, especially when forecasting an incoming hurricane.

We examine the performance of the unstructured-mesh, SWAN+ADCIRC wave and circulation model applied to a high-resolution, 5M-vertex, finite-element SL16 mesh of the Gulf of Mexico and Louisiana. This multi-process, multi-scale modeling system has been integrated by utilizing inter-model communication that is intra-core. The modeling system is validated through hindcasts of Hurricanes Katrina and Rita (2005), Gustav and Ike (2008) and comprehensive comparisons to wave and water level measurements throughout the region. The performance is tested on a variety of platforms, via the examination of output file

J.C. Dietrich (✉) · S. Tanaka · J.J. Westerink · L.G. Westerink · H.J. Westerink
Department of Civil Engineering and Geological Sciences, University of Notre Dame, Notre Dame, IN, USA
e-mail: dietrich@ices.utexas.edu

Present address:

J.C. Dietrich
Institute for Computational Engineering and Sciences, University of Texas at Austin, 1 University Station C0200, Austin, TX 78712, USA

C.N. Dawson
Institute for Computational Engineering and Sciences, University of Texas at Austin, 1 University Station C0200, Austin, TX 78712, USA

R.A. Luetlich Jr.
Institute of Marine Sciences, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA

M. Zijlema · L.H. Holthuijsen
Faculty of Civil Engineering and Geosciences, Technical University at Delft, Delft, The Netherlands

J.M. Smith
Coastal Hydraulics Laboratory, U.S. Army Corps of Engineers, Vicksburg, MS, USA

requirements and management, and the establishment of wall-clock times and scalability using up to 9,216 cores. Hindcasts of waves and storm surge can be computed efficiently, by solving for as many as $2.3 \cdot 10^{12}$ unknowns per day of simulation, in as little as 10 minutes of wall-clock time.

Keywords Hurricane waves · Storm surge · Unstructured meshes

1 Introduction

Several recent hurricanes have damaged critical infrastructure in New Orleans and Southern Louisiana. In 2005, Katrina caused devastating flooding within the city itself and created storm surge along the Mississippi-Alabama coastline that was the largest ever measured in the continental United States [14], while Rita flooded large portions of the marshes and bayous in southwestern Louisiana [4, 11]. In 2008, Gustav made landfall in southeastern Louisiana and threatened New Orleans with wave overtopping of its levee protection system [13], and Ike made landfall in Galveston but created currents and extensive flooding along the coastlines of all of Louisiana and eastern Texas [19].

These hurricanes created complex environments of waves, currents and storm surge throughout the region. In the deep water of the Gulf of Mexico, large, long waves were developed that propagated as swell in all directions. These waves have been measured at buoys to have peak periods up to 25 s and significant heights up to 17.5 m, and their significant heights have been estimated to be as large as 20–25 m closer to the hurricane track. In regions where the continental shelf is narrow, such as at the bird's foot of the Mississippi River, these large waves approach closely to the shoreline before breaking due to rapid changes in bathymetry. Behind the breaking zones and inside the marshes and bayous of Southern Louisiana, the wave environment is completely different, with wind sea generated locally but limited by depth and bottom friction to periods less than 5 s and significant heights less than 1.5 m.

The storm surge also varies widely from its generation on the continental shelf to its interaction with the nearshore estuaries, floodplains and channels. Currents of 2 m s^{-1} or greater can exist on the shelf, around the delta and the barrier islands, and within the natural and man-made passes and channels that connect New Orleans to the Gulf. Water levels reached 4–5 m relative to NAVD88 (2004.65) along the coastline of southwest Louisiana during Rita, 3–4 m along the Mississippi River levees during several storms, and up to 8.8 m along the Mississippi-Alabama coastline during Katrina.

Waves and storm surge interact strongly, despite being separated in frequency space. Surface gravity waves (or short waves), such as wind sea and swell, have periods ranging from 0.5–30 s, whereas longer waves, such as storm surge and tides, can have periods ranging from minutes to months. However, short waves are impacted by circulation; currents can shift wave energy due to the Doppler effect, and water levels affect propagation and dissipation due to depth-limited breaking. The transformation of short waves exerts radiation stress gradients, which drive currents and surge. Water levels can be increased by as much as 35 percent due to local wave-driven set-up [11, 33].

Hurricanes also act over a wide range of spatial scales. Waves and storm surge are generated in the deep waters of the Gulf and on the continental shelf, respectively, and then they propagate and transform in the complex nearshore environment due to rapid changes in bathymetry and bottom friction. Wave dissipation can be spread over large, smoothly-varying shelves, or it can be focused near the barrier islands or other breaking zones.

Storm surge is pushed over the sounds and marshes and then interacts with the levees, lakes and channels in the region. In hurricane modeling applications, these spatial scales necessitate the use of unstructured meshes, so resolution can be varied from kilometers in the deeper Gulf, to hundreds of meters on the continental shelf and behind the breaking zones, to tens of meters near the small-scale features. These meshes are highly efficient, because resolution is increased only in regions where the solution gradients are large.

Wave and circulation models have been developed to apply unstructured meshes during hurricane simulations. The unstructured-mesh version of the Simulating WAVes Nearshore (SWAN) model employs an analog of the Gauss-Seidel sweeping method used by the structured-mesh version, in order to propagate efficiently the wave action density [44]. SWAN incorporates the source/sink terms for nearshore wave physics, such as the triad nonlinear interactions, bottom friction and depth-limited breaking, in addition to the physics of quadruplet nonlinear interactions and whitewater. The ADvanced CIRCulation (ADCIRC) model for tidal-, wind- and density-driven circulation employs a continuous-Galerkin solution of the Generalized Wave Continuity Equation (GWCE), and the two-dimensional version employs the depth-integrated momentum equations on an unstructured, finite-element mesh. ADCIRC has been validated for several hurricanes in Southern Louisiana [4, 11, 43], and it has been used extensively by the US Army Corps of Engineers (USACE), the Federal Emergency Management Agency (FEMA) and local agencies to design levee and flood mitigation systems, and to evaluate hurricane flooding risk.

SWAN and ADCIRC have been integrated and coupled closely so that they run on the same global unstructured mesh [12]. The resulting SWAN+ADCIRC model employs an identical, unstructured sub-mesh on each computational core, so that information can be passed between these two models without interpolation. Inter-model communication is performed intra-core, via local cache or memory, and thus is highly efficient. In addition, both models make use of the same parallel infrastructure and pass information at boundaries between the local sub-meshes, so that the intra-model communication is inter-core. Thus the sharing of the same unstructured mesh allows the coupled model to be both accurate (by increasing mesh resolution in regions with large spatial gradients) and efficient (by eliminating the need for costly interpolation and extensive global communication as information is passed between models). SWAN+ADCIRC is well-positioned to simulate waves and surge from their generation in deep and shallow water to their dissipation nearshore.

Hurricane forecasting applications demand both accuracy and efficiency. Model results must be reliable for a wide range of storm characteristics, and thus a high-resolution mesh should be employed to resolve the complex geometry throughout the region. But model results must also be timely, often on the order of 1 h, so that they can be useful to emergency management officials to aid with decision-making. In this work, we validate the accuracy and efficiency of SWAN+ADCIRC on the Southern Louisiana (SL16) unstructured mesh, which employs 5M vertices and 10M finite elements to provide a high-resolution description of Southern Louisiana. At every vertex in this mesh, SWAN computes wave action densities in 36 directions and 40 frequency bins at intervals of 600 s, while ADCIRC computes water levels and currents at intervals of 1 s, and thus the coupled model is solving for $2.3 \cdot 10^{12}$ unknowns per day of simulation. The model is validated against measured waves and storm surge during the four recent hurricanes to impact the region, namely Katrina and Rita (2005), and Gustav and Ike (2008). Validation results show SWAN+ADCIRC simulates accurately the evolution of waves and surge in this region.

Benchmarking results show SWAN+ADCIRC is efficient to thousands of computational cores. As meshes continue to grow in size and complexity, it is imperative that models are scalable on massively-parallel computational resources. The coupling paradigm employed by SWAN+ADCIRC complements the high scalability of the component models, and the coupled model also manages well its file output through the use of dedicated writer cores. When employed on the SL16 mesh, the SWAN+ADCIRC model maintains its scalability to 8,196 computational cores, simulating accurately and efficiently the coupled hurricane waves and surge in as little as 10 wall-clock minutes per day of simulation.

2 Computational Wave and Circulation Models

2.1 SWAN

Individual wind waves exist on length and time scales that are too small to be resolved when computational models are applied to large domains. Thus, instead of resolving the phases of the waves and their associated processes, SWAN represents the wave field as an phase-averaged spectrum [3]. The wave action density $N(t, \lambda, \varphi, \sigma, \theta)$ is allowed to evolve in time (t), geographic space (λ, φ) and spectral space (with relative frequencies σ and directions θ), as governed by the action balance equation:

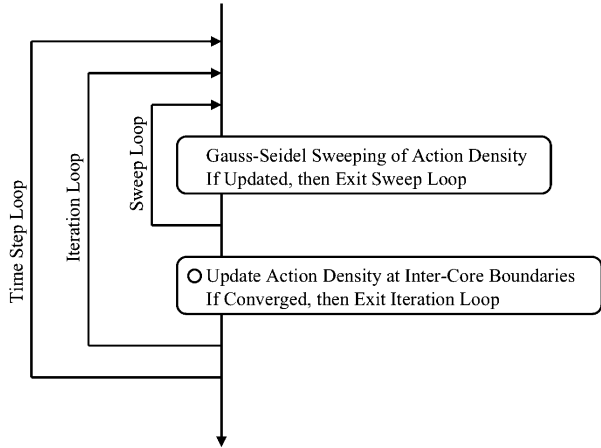
$$\begin{aligned} \frac{\partial N}{\partial t} + \frac{\partial}{\partial \lambda} [(c_\lambda + U)N] + \cos^{-1} \varphi \frac{\partial}{\partial \varphi} [(c_\varphi + V)N \cos \varphi] \\ + \frac{\partial}{\partial \theta} [c_\theta N] + \frac{\partial}{\partial \sigma} [c_\sigma N] = \frac{S_{tot}}{\sigma} \end{aligned}$$

where (c_λ, c_φ) is the group velocity, (U, V) is the ambient current, and c_θ and c_σ are the propagation velocities in the θ - and σ -spaces. The source terms S_{tot} represent wave growth by wind; action lost due to whitecapping, surf breaking and bottom friction; and action exchanged between spectral components due to nonlinear effects in deep and shallow water.

The structured-mesh version of SWAN employs a Gauss-Seidel sweeping algorithm to propagate the wave action density in geographic space, and an analog of that algorithm was used recently to extend SWAN to run on unstructured meshes [44]. The mesh vertices are ordered so that SWAN can sweep through them in alternating directions and update the action density from neighboring vertices. When the action density has been updated everywhere in the mesh, SWAN proceeds to the next iteration on that time step, and continues to iterate until the solution converges. A schematic representation of SWAN's solution algorithm is shown in Fig. 1.

In a parallel computing environment, after the action density has been updated at all vertices within a computational sub-mesh, it is updated at the shared boundaries between neighboring sub-meshes. This method does not require the global communication associated with the solution of matrix systems, and thus it is highly scalable. The local communication occurs at the end of each iteration, and SWAN will iterate until the action density has converged on its time step. In the present hindcasts, SWAN iterates until 95 percent of the vertices in the global domain have converged, to a maximum of 20 iterations. The solution is considered to have converged if the absolute change in the significant wave height (H_s) from one iteration to the next is less than 0.005 or the relative change in H_s is less than 0.01 and the curvature of the iteration curve of H_s normalized with H_s is less than 0.005. At

Fig. 1 Schematic of SWAN’s solution algorithm. Local communication to neighboring computational cores is marked with a *circle*



every geographic mesh vertex, SWAN computes the action density in 36 directions and 40 frequency bins at an interval of 600 s, solving for 10^{12} unknowns per day of simulation on the SL16 mesh.

2.2 ADCIRC

ADCIRC solves forms of the shallow-water equations (SWE) for water levels ζ and the vertically-integrated momentum equations for currents \vec{U} [8, 20, 23, 43]. The model applies the continuous-Galerkin finite-element method with linear C_0 triangular elements to discretize and solve the SWE on unstructured meshes, and thus it allows localized refinement in regions where the solution gradients are largest. The temporal discretization is different for the continuity and momentum equations [37]. For the continuity equation, the time derivatives are discretized over three levels, so that the solution for the future water level requires knowledge of the present and past water levels. For the momentum equation, the temporal discretization is explicit for all terms except Coriolis, which uses an average of the present and future velocities. At every geographic mesh vertex, ADCIRC solves for the water level and the two components of the current at an interval of 1 s, solving for $1.3 \cdot 10^{12}$ unknowns per day of simulation on the SL16 mesh.

ADCIRC computes water levels via the solution of the Generalized Wave Continuity Equation (GWCE), which is a combined and differentiated form of the continuity and momentum equations:

$$\frac{\partial^2 \zeta}{\partial t^2} + \tau_0 \frac{\partial \zeta}{\partial t} + S_p \frac{\partial \tilde{J}_\lambda}{\partial \lambda} + \frac{\partial \tilde{J}_\phi}{\partial \phi} - S_p U H \frac{\partial \tau_0}{\partial \lambda} - V H \frac{\partial \tau_0}{\partial \phi} = 0$$

where:

$$\begin{aligned} \tilde{J}_\lambda = & -S_p Q_\lambda \frac{\partial U}{\partial \lambda} - Q_\phi \frac{\partial U}{\partial \phi} + f Q_\phi - \frac{g}{2} S_p \frac{\partial \zeta^2}{\partial \lambda} - g S_p H \frac{\partial}{\partial \lambda} \left[\frac{P_s}{g \rho_0} - \alpha \eta \right] \\ & + \frac{\tau_{s\lambda, winds} + \tau_{s\lambda, waves} - \tau_{b\lambda}}{\rho_0} + (M_\lambda - D_\lambda) \\ & + U \frac{\partial \zeta}{\partial t} + \tau_0 Q_\lambda - g S_p H \frac{\partial \zeta}{\partial \lambda} \end{aligned}$$

and:

$$\begin{aligned} \tilde{J}_\varphi = & -S_p Q_\lambda \frac{\partial V}{\partial \lambda} - Q_\varphi \frac{\partial V}{\partial \varphi} - f Q_\lambda - \frac{g}{2} \frac{\partial \zeta^2}{\partial \varphi} - gH \frac{\partial}{\partial \varphi} \left[\frac{P_s}{g\rho_0} - \alpha\eta \right] \\ & + \frac{\tau_{s\varphi,winds} + \tau_{s\varphi,waves} - \tau_{b\varphi}}{\rho_0} + (M_\varphi - D_\varphi) \\ & + V \frac{\partial \zeta}{\partial t} + \tau_0 Q_\varphi - gH \frac{\partial \zeta}{\partial \varphi} \end{aligned}$$

and the currents are obtained from the vertically-integrated momentum equations:

$$\begin{aligned} \frac{\partial U}{\partial t} + S_p U \frac{\partial U}{\partial \lambda} + V \frac{\partial U}{\partial \varphi} - fV \\ = -gS_p \frac{\partial}{\partial \lambda} \left[\zeta + \frac{P_s}{g\rho_0} - \alpha\eta \right] + \frac{\tau_{s\lambda,winds} + \tau_{s\lambda,waves} - \tau_{b\lambda}}{\rho_0 H} + \frac{M_\lambda - D_\lambda}{H} \end{aligned}$$

and:

$$\begin{aligned} \frac{\partial V}{\partial t} + S_p U \frac{\partial V}{\partial \lambda} + V \frac{\partial V}{\partial \varphi} + fU \\ = -g \frac{\partial}{\partial \varphi} \left[\zeta + \frac{P_s}{g\rho_0} - \alpha\eta \right] + \frac{\tau_{s\varphi,winds} + \tau_{s\varphi,waves} - \tau_{b\varphi}}{\rho_0 H} + \frac{M_\varphi - D_\varphi}{H} \end{aligned}$$

where $H = \zeta + h$ is total water depth; ζ is the deviation of the water surface from the mean; h is bathymetric depth; $S_p = \cos \varphi_0 / \cos \varphi$ is a spherical coordinate conversion factor and φ_0 is a reference latitude; U and V are depth-integrated currents in the x - and y -directions, respectively; $Q_\lambda = UH$ and $Q_\varphi = VH$ are fluxes per unit width; f is the Coriolis parameter; g is gravitational acceleration; P_s is atmospheric pressure at the surface; ρ_0 is the reference density of water; η is the Newtonian equilibrium tidal potential and α is the effective earth elasticity factor; $\tau_{s,winds}$ and $\tau_{s,waves}$ are surface stresses due to winds and waves, respectively; τ_b is bottom stress; M are lateral stress gradients; D are momentum dispersion terms; and τ_0 is a numerical parameter that optimizes the phase propagation properties [1, 20].

ADCIRC utilizes a wetting and drying algorithm that activates and deactivates entire elements during the inundation and recession of coastal topography [9, 10, 21, 22]. Dry vertices are wetted when a balance is satisfied between the water level gradients and bottom friction relative to the adjacent wet vertices on a triangular element. Wet vertices are dried when the total water depth decreases below a minimum wetness height. If all three vertices on an element are wet, then the element is also considered wet; otherwise, the element is dry. The wetting and drying algorithm is employed once per time step, and thus the coastline location can change by only one layer of elements at a time.

The GWCE can be solved implicitly or explicitly, as represented schematically in Figs. 2 and 3. The implicit solution requires the assembly of a matrix system and the application of the Jacobi Conjugate Gradient (JCG) method, which iterates to convergence. The explicit solution utilizes a lumped diagonal in the mass matrix to solve directly the system; it is faster per time step than the implicit solution, but it may require a smaller time step for stability. (However, the time step required for stability by the explicit solution is often similar to the time step required for accuracy by the implicit solution, because the accuracy is limited by the very fast flows of 2–3 m s⁻¹ and the propagation speeds of any

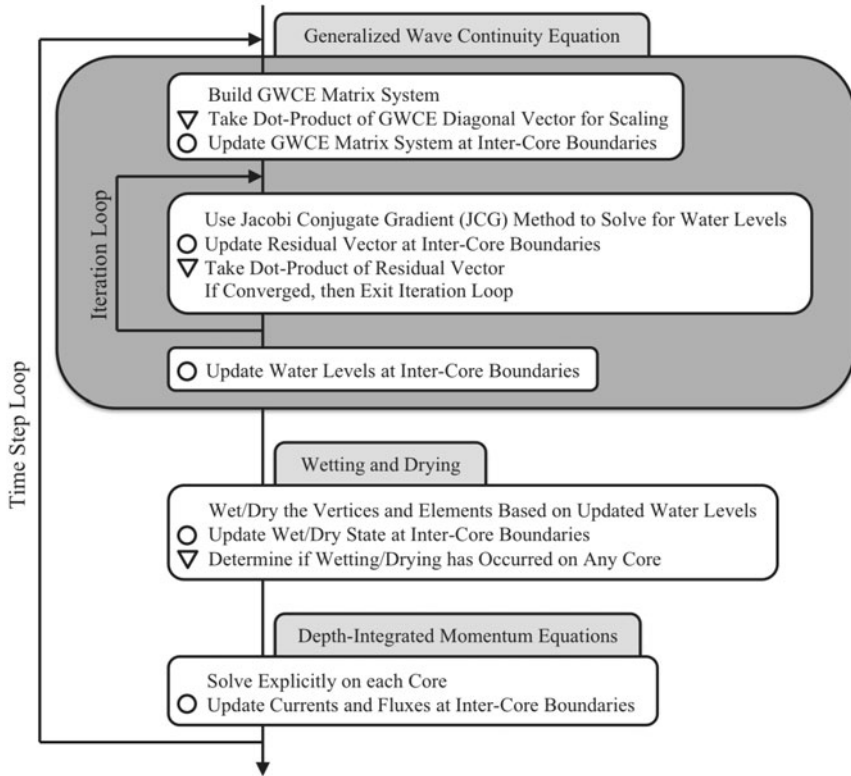


Fig. 2 Schematic of ADCIRC’s algorithm with implicit solution of the GWCE. Local communication to neighboring computational cores is marked with *circles*, while global communication over all computational cores is marked with *triangles*

wetting/drying fronts.) After the new water levels are computed, wetting and drying is allowed, and then the vertically-integrated momentum equations are solved explicitly for the currents.

In a parallel computing environment, ADCIRC’s solution algorithm can require local and global communication between computational cores. Several quantities, including the water levels and currents, must be updated along the overlapping boundaries of neighboring sub-meshes, but this communication is local and highly scalable. However, the implicit solution of the GWCE requires global communication to: take the dot product of the diagonal vector for scaling of the GWCE matrix system; take the dot product of the residual vector after each JCG iteration; and determine if the global GWCE matrix system needs to be re-sized due to any occurrence of wetting or drying within the global domain. These instances of global communication utilize the *MPI_ALLREDUCE* command to collect information and broadcast it to all of the computational cores, and thus they hinder the scalability of ADCIRC. As the number of cores increases, this global communication becomes more costly [37]. The explicit solution of the GWCE does not require any global communication, because it is matrix-free, and thus it is highly scalable.

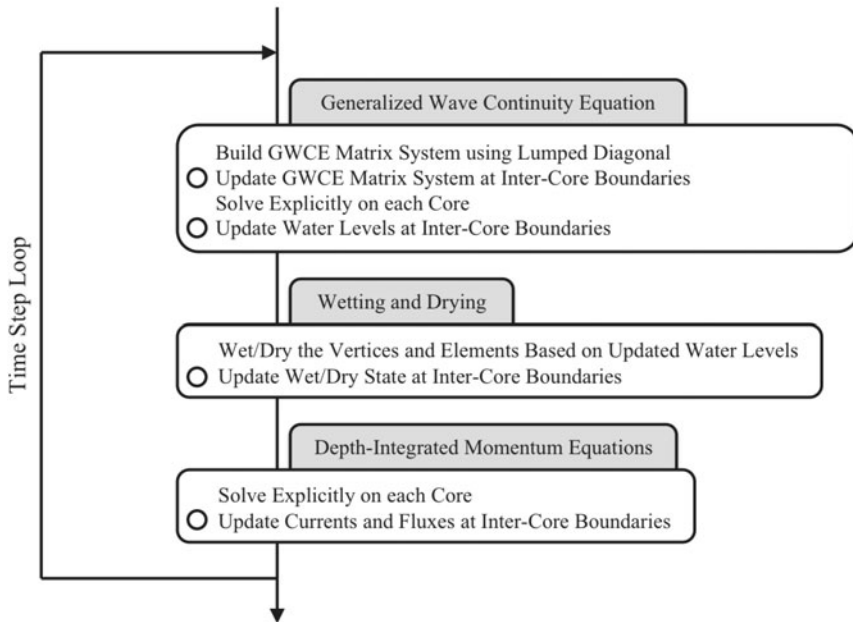


Fig. 3 Schematic of ADCIRC's algorithm with explicit solution of the GWCE. Local communication to neighboring computational cores is marked with *circles*

2.3 Model Coupling

SWAN and ADCIRC are coupled so that they run on the same computational core and on the same unstructured sub-mesh [12]. ADCIRC passes wind velocities, water levels and currents through local cache or memory to SWAN, which utilizes those quantities to force its computations. At the end of each of its time steps, SWAN computes the wave radiation stresses and their gradients, and then it passes those gradients as a forcing function to ADCIRC.

ADCIRC is Courant-limited algorithmically due to its explicit features, and it is also limited by the propagation speed of any wetting fronts (because only one layer of triangular elements is activated at the wetting front during a time step), so its time step must be relatively small. However, SWAN is unconditionally stable, and thus its time step can be relatively large. The coupling interval is taken to be the same as the SWAN time step. The present hindcasts utilize a coupling interval and SWAN time step of 600 s, and the ADCIRC time step is 1 s. The two models are run sequentially, so that each computational core is alternating running either SWAN or ADCIRC.

ADCIRC runs first on the coupling interval, and it uses the SWAN radiation stress gradients from the previous interval to extrapolate forward its wave forcing in time. After its time steps on the coupling interval, ADCIRC passes wind velocities, water levels, currents and roughness lengths to SWAN. Then SWAN is run on the same interval, using the average of the ADCIRC variables from the interval in its computations. After its time step, SWAN computes the radiation stress gradients and passes them to ADCIRC, which then begins the process anew on the next interval.

This coupling paradigm maximizes efficiency in a parallel computing environment. Figure 4 shows a schematic of the communication between models and cores. SWAN and




Fig. 4 Schematic of parallel communication between models and cores [12]. *Dashed lines* indicate communication for all vertices within a sub-mesh, and are inter-model and intra-core. *Solid lines* indicate communication for the edge-layer-based vertices between sub-meshes, and are intra-model and inter-core

ADCIRC utilize the same local sub-mesh, and thus there is a one-to-one correspondence between the geographic locations of the mesh vertices. No interpolation is required; water levels, currents, wind velocities and radiation stress gradients can be passed directly through local cache or memory. Inter-model communication is intra-core, and thus unaffected by the parallel computing environment. Intra-model communication is handled as described above, with SWAN utilizing only local, neighbor-to-neighbor communication along overlapping sub-mesh boundaries, and with ADCIRC utilizing local (and global, for the implicit solution,) communication. It should be noted that this coupling paradigm avoids the costs associated with the interpolation of variables or forcing information that would be necessary for heterogeneous meshes; in that case, the inter-model communication would be inter-core and increasingly expensive. At a minimum, it would involve many sub-meshes and cores, and this would put significantly more data traffic on the network, thus limiting the scalability.

2.4 File Output

The coupled SWAN+ADCIRC model utilizes ADCIRC's ability to dedicate cores for file output [37]. Computational cores run SWAN+ADCIRC, using the coupling paradigm described above, but a small subset of so-called "writer" cores is set aside for writing the large, global output files to the disk subsystem. When the computational cores are ready to output a time snap of forcing data and/or computed solutions, they send these data to a set of specialized writer cores, which collect the data into large, global arrays. Then each array is written to the appropriate output file. Once the data has been collected by the writer cores, the computational cores can proceed with the next stage of their calculations, while the writer cores write independently. When the computational cores are ready to write the next time snap, the previously-assigned writer cores may not have finished their task and be ready to receive the next set of global arrays. For this reason, options for sequential batches of writer cores have been implemented so that the writer cores work in a rotation.

This output paradigm is most efficient when there is at least one dedicated writer core for each global output file. For example, the present hindcasts require SWAN+ADCIRC to generate 10 global output files containing: water levels, currents, wind pressures, wind velocities, radiation stress gradients, significant wave heights, mean wave directions, two definitions of mean wave periods, and peak wave periods. Thus, when file output is enabled in the timing results that follow, SWAN+ADCIRC utilizes at least 10 dedicated writer cores. On large numbers of cores, it is also important to employ the sequential approach to the writer cores, because the computational cores may produce data for the next time snap before the writer cores finish writing data from the previous time snap. The sequential approach expands the number of writer cores, so that more than one set of 10 writer cores is available to receive data [37].

3 Unstructured Mesh

Unstructured meshes allow for resolution to be increased in regions where gradients are large, such as in regions with rapid changes in bathymetry and/or topography, variability in bottom roughness or some other parameter, wave dissipation over short distances, or the build-up of surge against a raised feature. Resolution is varied from kilometers in the deep ocean, to hundreds of meters on the continental shelf, to tens of meters in the marshes, floodplains and channels onshore. The SL16 mesh provides this variability in resolution, as shown in Figs. 5, 6, 7 and 8. The mesh extends outward to the Atlantic Ocean, so that storms can be started inside the computational domain but far from the area of interest. The mesh sizes are 4–6 km in the Gulf of Mexico, to capture the generation and propagation of hurricane waves in SWAN. On the continental shelf and in the wave breaking zones, the resolution varies down to 200 m, to capture the wave dissipation in SWAN and the accurate transfer of wave radiation stress gradients to ADCIRC. The mesh sizes are no larger than 200 m in the marshes and floodplains of Southern Louisiana, and they vary down to 20–50 m in the small-scale channels, such as the distributaries of the Mississippi River delta. The SL16 mesh contains 5,036,960 vertices and 9,949,317 triangular elements.

Figures 5–8 show the SL16 mesh bathymetry and topography, which were applied as described in [13]. Hydraulic friction is computed in ADCIRC using a Manning's n formulation, and SWAN converts the same Manning's n values to roughness lengths that vary in space and time, for use in a friction formulation from [24]. The details of this bottom friction formulation are described in [13].

4 Validation

SWAN+ADCIRC was validated comprehensively on the SL16 mesh for the four recent historical storms: Katrina and Rita (2005), and Gustav and Ike (2008). As shown in Figs. 5 and 6, the tracks of these storms provide good coverage of Southern Louisiana, with Katrina and Gustav impacting the Mississippi River and New Orleans in the southeast, and Rita and Ike impacting the low-lying marshes and bayous throughout Southern Louisiana. Also, and just as importantly, the wind forcing data and measurement data associated with these storms is extensive and a good test of the models' skill at simulating waves and surge.

Wind fields for three of the four storms were data-assimilated by using NOAA's Hurricane Research Division Wind Analysis (H*Wind) system [28, 29, 32], and then blended with larger scale winds using the Interactive Objective Kinematic Analysis (IOKA) system




Fig. 5 Bathymetry/topography (m) of the SL16 mesh. The tracks of the four historical storms are shown

[5, 6]. The Rita winds are based solely on the IOKA system. ADCIRC interpolates these wind fields to its unstructured mesh, and then it shares the wind velocities with SWAN. Both models apply a wind drag coefficient that depends on storm sector [13, 30, 31].

In the present hindcasts, the SWAN time step (and thus the coupling interval) is taken to be 600 s. SWAN utilizes 36 directional bins of constant width 10° , and 40 frequency bins that increase logarithmically over the range 0.031–1.42 Hz. Wind input is based on the formulation from [36], whitecapping is applied via the modified expression of [34], and nonlinear quadruplet interactions are computed using the Discrete Interaction Approximation of [16]. In shallow water, bottom friction is parameterized as described above, while depth-induced breaking is computed with a spectral version of the model due to [2] with the breaking index $\gamma = 0.73$. Wave refraction is enabled only in the northern Gulf, where the mesh resolution is sufficient.

ADCIRC utilizes an implicit solution of the GWCE, with a time step of 1 s. Its water levels are adjusted for the regional difference between local mean sea level (LMSL) and NAVD88 (2004.65) and the seasonal fluctuation in sea level in the Gulf of Mexico. Bottom friction is parameterized using a Manning's n formulation, as described above. The Mississippi and Atchafalaya Rivers are forced with flow rates that are average conditions during the storms. In addition, seven tidal constituents are forced on the open boundary in the Atlantic Ocean. The rivers and tides are allowed to ramp to a dynamic equilibrium for 18 days prior to the start of the wind and wave forcing.




Fig. 6 Mesh resolution (m) of the SL16 mesh




Fig. 7 Bathymetry/topography (m) of the SL16 mesh in the northern Gulf of Mexico. The tracks of the four historical storms are shown

4.1 Historical Storms

The northern Gulf of Mexico has been impacted by several intense hurricanes during the past few years. These storms are excellent for validation of SWAN+ADCIRC, because they




Fig. 8 Mesh resolution (m) of the SL16 mesh in the northern Gulf of Mexico

are characterized by a wealth of measurement data for waves and water levels. Furthermore, their tracks and landfall locations varied along the coastline from Mississippi to Texas (Fig. 7), and thus the system responded differently to each storm.

Katrina devastated large portions of Southern Louisiana, Mississippi and Alabama during the 2005 hurricane season [4, 11]. It strengthened to a Category 5 storm while in the Gulf, and it generated large waves with measured significant heights of 17.5 m offshore of the Mississippi River delta (Fig. 9a). Its easterly winds pushed storm surge across the Louisiana-Mississippi continental shelf and against the protruding delta, where it collected against the levees of the Mississippi River. The storm made landfall in Plaquemines Parish and tracked northward, over the levees and sounds to the east of New Orleans. As the storm moved through the system, it pushed northward the surge that had collected against the river levees; this surge moved over the marshes and built along the coastlines of Mississippi and Alabama (Fig. 9b). The recorded peak surge of 8.8 m is the largest ever measured in the United States [14].

Rita was a strong Category 5 storm, and it generated large waves as it moved across the Gulf. It was a threat to impact Galveston and Houston, but it turned northward and made landfall near the border between Texas and Louisiana. As Rita moved onto the continental shelf near southwestern Louisiana, it created storm surge along the coastline that peaked at 4.7 m [4, 11]. As the storm made landfall, the winds shifted to blow northward, pushing the surge over the low-lying marshes and bayou in the southwestern part of the state (Fig. 10). Rita is a good validation test because it allows for an examination of the model's skill in representing the friction-dominated flooding and recession processes.

Gustav reached Category 4 strength on the Saffir-Simpson scale while in the Caribbean Sea, but it weakened as it crossed western Cuba. It tracked northwestward across the Gulf, peaking at Category 3 strength before making landfall near Terrebonne Bay as a weak Category 2 storm. It generated large waves in the Gulf that propagated northward before breaking near the Mississippi River delta and the barrier islands offshore of Louisiana and Mississippi (Fig. 11a). Its eye never came closer than 75 km to New Orleans, but it created significant surge on all sides of the city. Surge was pushed into Lake Pontchartrain to the north, up the Mississippi River and into the canals that run into the heart of the city, and over the marshes to the southwest and east (Fig. 11b). The levee protection system was not breached, but there were reports of wave overtopping [13].

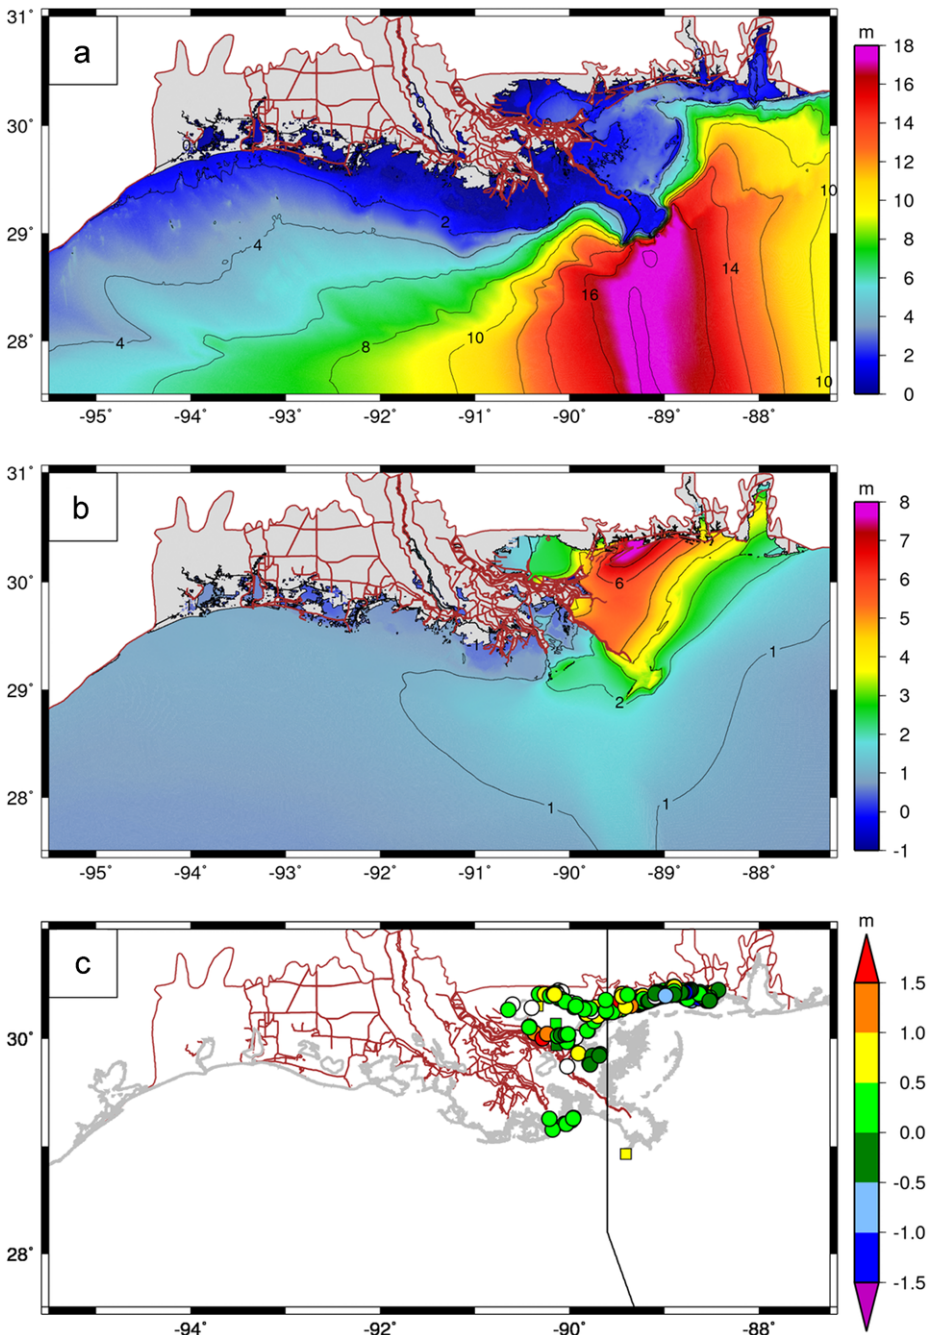


Fig. 9 Contour plots of maximum (a) wave heights (m) and (b) water levels relative to NAVD88 (2004.65) during Katrina in Southern Louisiana. Panel (c) indicates locations of the HWMs (circles) and hydrographs (squares) for Katrina in Southern Louisiana. The points are color-coded to show the errors between measured and modeled peak water levels; green points indicate matches within 0.5 m. White points indicate locations that were never wetted by ADCIRC

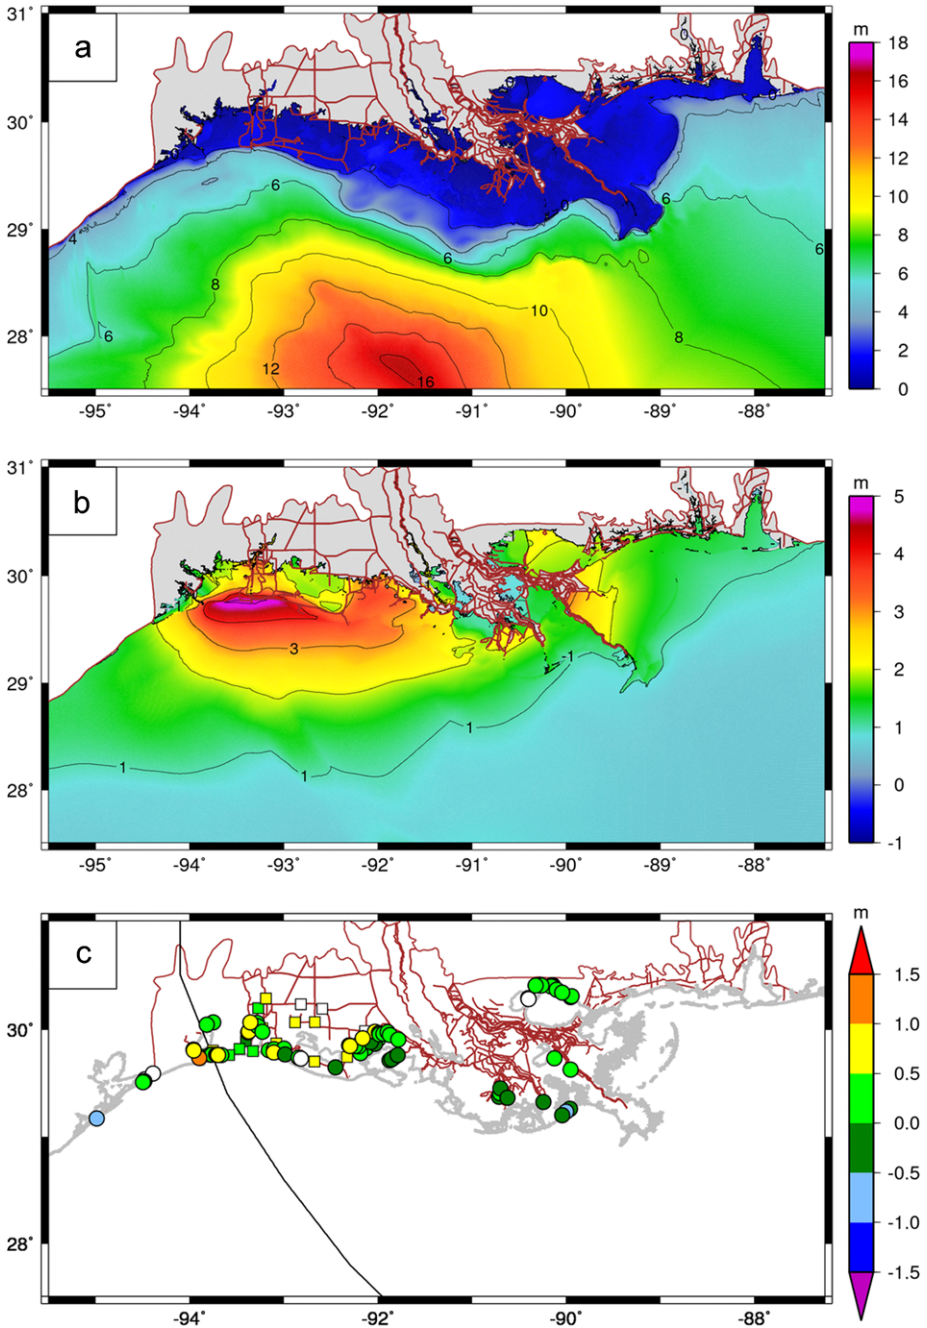


Fig. 10 Contour plots of maximum (a) wave heights (m) and (b) water levels to relative NAVD88 (2004.65) during Rita in Southern Louisiana. Panel (c) indicates locations of the HWMs (circles) and hydrographs (squares) for Rita in Southern Louisiana. The points are color-coded to show the errors between measured and modeled peak water levels; green points indicate matches within 0.5 m. White points indicate locations that were never wetted by ADCIRC

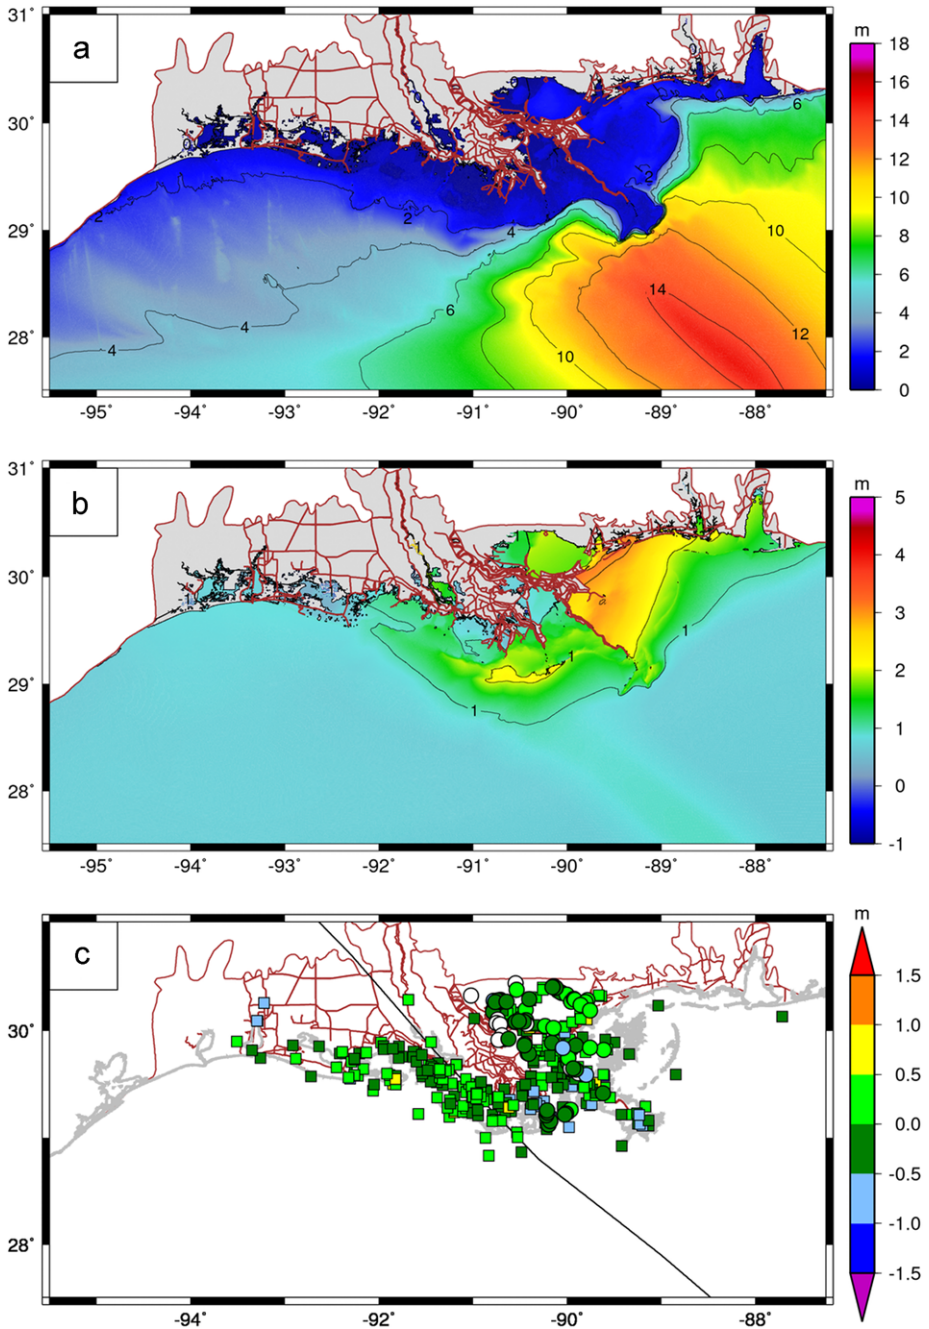


Fig. 11 Contour plots of maximum (a) wave heights (m) and (b) water levels to relative NAVD88 (2004.65) during Gustav in Southern Louisiana. Panel (c) indicates locations of the HWMs (circles) and hydrographs (squares) for Gustav in Southern Louisiana. The points are color-coded to show the errors between measured and modeled peak water levels; green points indicate matches within 0.5 m. White points indicate locations that were never wetted by ADCIRC

Table 1 Summary of the available measurement data for the four historical storms. The values in the columns below the storms indicate the number of available gages or measurement locations

| Type | Source | Katrina | Rita | Gustav | Ike | References |
|--------------|--------------|---------|------|--------|-----|--------------|
| Waves | NDBC | 14 | 11 | 12 | 12 | [26] |
| | CSI | 3 | 2 | 5 | 5 | [42] |
| | AK | | | 16 | | [18] |
| | USACE-CHL | | | 6 | 6 | [13, 35] |
| Water levels | CSI | | | 5 | 5 | [42] |
| | AK | | | 16 | | [18] |
| | NOAA | 3 | | 23 | 23 | [27] |
| | USACE-CHL | | | 6 | 6 | [13, 35] |
| | USACE | | | 39 | 41 | [40] |
| | USGS (Depl.) | | 23 | 24 | 16 | [25, 41] |
| | USGS (Perm.) | 6 | | 18 | 22 | [41] |
| HWMs | CRMS | | | 232 | 240 | [7] |
| | URS/FEMA | 193 | 84 | 82 | 177 | [15, 38, 39] |
| | USACE | 206 | | | | [14] |

Although Ike tracked farther west than the other three hurricanes and made landfall near Galveston, outside the coverage of the SL16 mesh, it impacted all of Southern Louisiana. It created a forerunner surge that traveled down the Louisiana-Texas shelf about 12–24 h before landfall [19], and its waves and storm surge inundated large sections of the marshes and bayou of southwest Louisiana (Fig. 12). The storm was large enough that its outer winds also pushed waves and surge to the east, onto the Louisiana-Mississippi shelf and against the levee protection system near New Orleans.

4.2 Comparisons to Measured Data

Table 1 summarizes the available measured data for all four storms. Waves were measured at buoys, platforms and gages ranging from the deepest water of the open Gulf to the shallow marshes and bayou along southern Louisiana. Water levels were measured at platforms and gages ranging from the continental shelf to the man-made channels near metropolitan New Orleans. Note the significant differences between the available data for the storms of 2005 and 2008. Katrina and Rita are characterized by more than 400 high-water marks but only a few time series of waves and water levels. However, advancements in gage hardening and innovative deployment techniques resulted in expanded time series measurements, so Gustav and Ike are described by a relatively large amount of data.

The importance of the datasets varies by storm. During Rita, the USGS deployed 23 gages to record water levels in southwestern Louisiana, and they are of particular interest because they include the slow recession of surge from the marshes back into the Gulf [4, 11]. During Gustav and Ike, the nearshore wave measurements from [18] and the USACE-CHL are particularly valuable because they represent some of the first measurements taken along the coastline and inside the marshes during a hurricane [13]. Taken together, these measurements describe a wide variety of storm conditions and system responses throughout southern Louisiana.

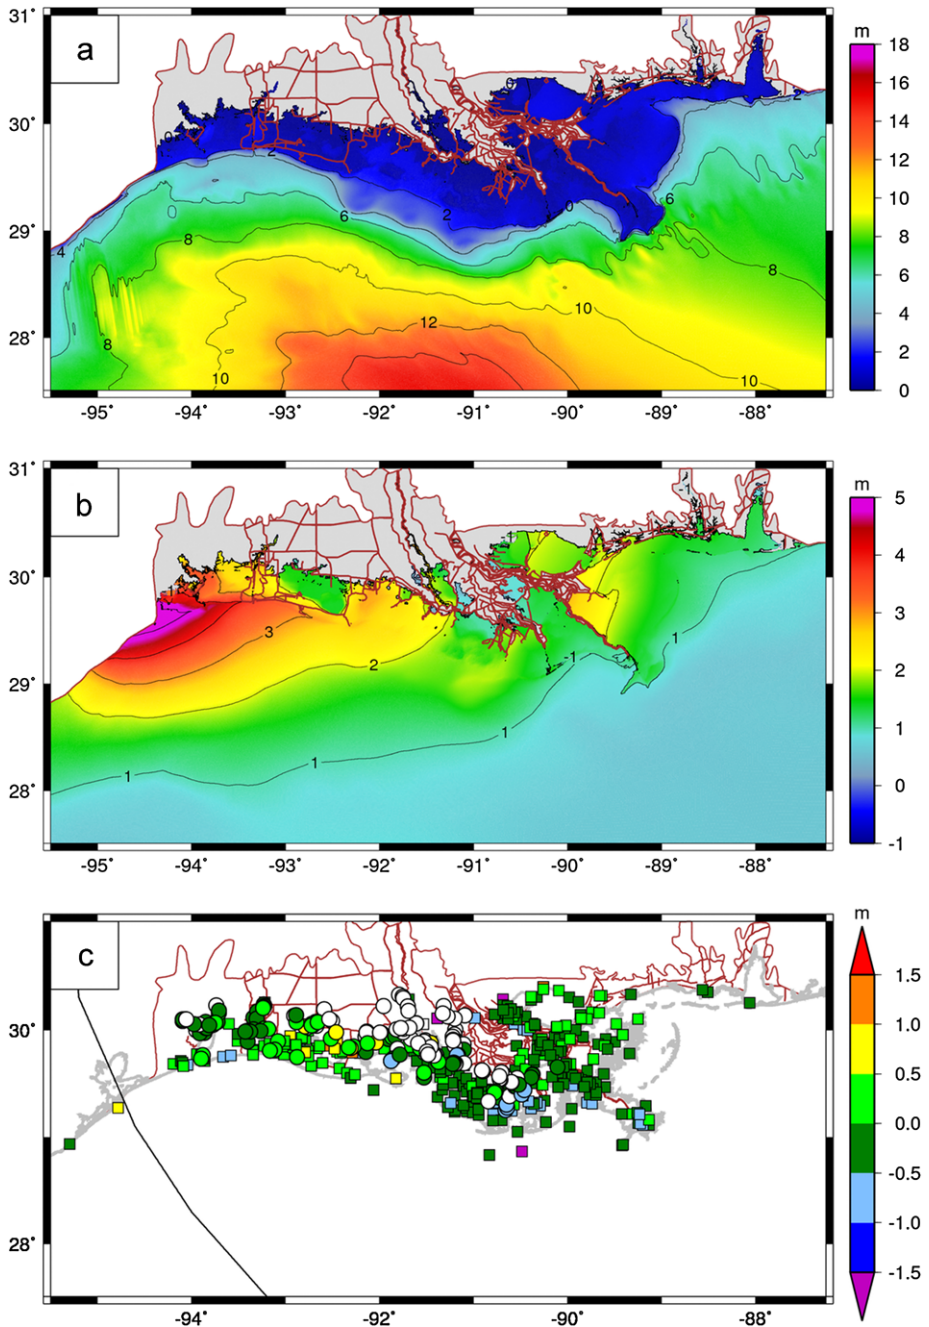


Fig. 12 Contour plots of maximum (a) wave heights (m) and (b) water levels to relative NAVD88 (2004.65) during Ike in Southern Louisiana. Panel (c) indicates locations of the HWMs (circles) and hydrographs (squares) for Ike in Southern Louisiana. The points are color-coded to show the errors between measured and modeled peak water levels; green points indicate matches within 0.5 m. White points indicate locations that were never wetted by ADCIRC

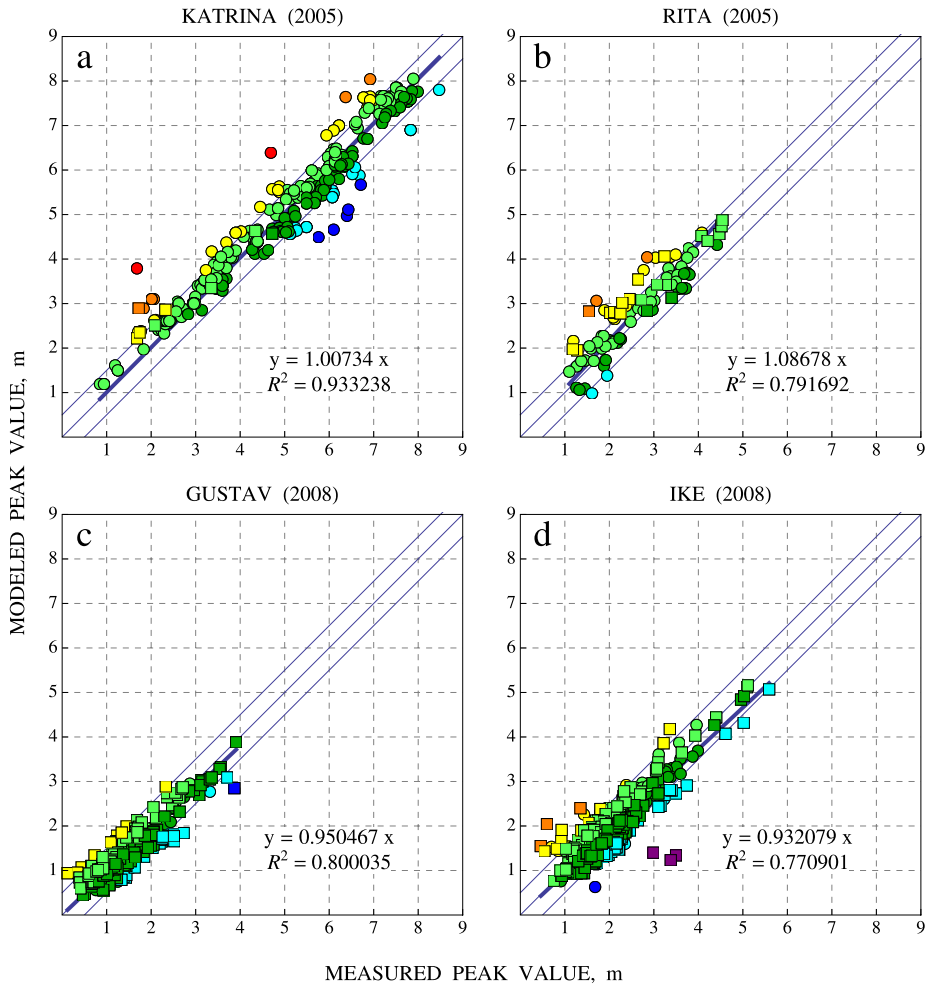


Fig. 13 Scatter plot of USACE and URS/FEMA HWMs (circles) and peak hydrograph water levels (squares) for: (a) Katrina, (b) Rita, (c) Gustav and (d) Ike. Green points indicate a match within 0.5 m. Red, orange, yellow and light green circles indicate overprediction by the model; green, blue, dark blue and purple circles indicate underprediction

These measurement locations are shown spatially for Katrina (Fig. 9c), Rita (Fig. 10c), Gustav (Fig. 11c) and Ike (Fig. 12c), and the locations are color-coded based on the error between the measured and modeled peak water levels for each storm. The matches are excellent overall. The computed SWAN+ADCIRC peak water levels are within 0.5 m at 305 of the 370 (82 percent) locations during Katrina, 71 of the 99 (72 percent) locations during Rita, 353 of the 387 (91 percent) locations during Gustav and 379 of 455 (85 percent) locations during Ike. This performance is also depicted in the scatter plots of measured-to-modeled peak water levels in Fig. 13, with best-fit slopes and correlation coefficients R^2 summarized in Table 2. The best-fit slopes range from 0.93 to 1.09, and the R^2 values range from 0.77 to 0.93, indicating that SWAN+ADCIRC is matching well the peak water levels for a variety of storms, throughout the entire northern Gulf of Mexico.

Table 2 Summary of the error statistics for SWAN+ADCIRC and the four historical storms. The best-fit slope (m) and R^2 values are shown for the water levels (ζ), while the average SI and $Bias$ values are shown for water levels (ζ), significant wave heights (H_s), and peak (T_p) and mean (T_{m-10}) wave periods. For Gustav and Ike, the “All” column includes statistics from an analysis of every available data set in Table 1, while the “Subset” column includes statistics from every source except the USACE-CHL data

| | | | Katrina | Rita | Gustav | | Ike | |
|--------------|------------|--------|---------|------|--------|--------|-------|--------|
| | | | | | All | Subset | All | Subset |
| Water levels | ζ | m | 1.01 | 1.09 | 0.95 | | 0.93 | |
| | | R^2 | 0.93 | 0.79 | 0.80 | | 0.77 | |
| | | SI | 0.19 | 0.28 | 0.24 | | 0.16 | |
| | | $Bias$ | 0.14 | 0.15 | 0.14 | | -0.07 | |
| Waves | H_s | SI | 0.23 | 0.23 | 0.34 | 0.32 | 0.29 | 0.21 |
| | | $Bias$ | 0.05 | 0.43 | 0.35 | 0.22 | 0.09 | -0.05 |
| | T_p | SI | 0.22 | 0.25 | 0.53 | 0.39 | 0.57 | 0.23 |
| | | $Bias$ | 0.07 | 0.25 | -0.03 | -0.01 | 0.02 | 0.04 |
| | T_{m-10} | SI | 0.15 | 0.12 | 0.22 | | 0.16 | |
| | | $Bias$ | 0.09 | 0.18 | -0.03 | | 0.13 | |

To quantify the models’ skill at capturing all aspects of the measured time series, the Scatter Index (SI) is used as:

$$SI = \frac{\sqrt{\frac{1}{N} \sum_{i=1}^N (E_i - \bar{E})^2}}{\frac{1}{N} \sum_{i=1}^N |O_i|}, \tag{1}$$

as well as the Bias:

$$Bias = \frac{\frac{1}{N} \sum_{i=1}^N E_i}{\frac{1}{N} \sum_{i=1}^N |O_i|} \tag{2}$$

where N is the number of observations, $E_i = S_i - O_i$ is the error between the modeled (S_i) and measured (O_i) values, and \bar{E} is the mean error. Thus the SI is the ratio of the standard deviation of the measured-to-modeled errors to the mean absolute measured value. Table 2 summarizes the mean SI and $Bias$ values for water levels, significant wave heights, and peak and mean periods for the four historical storms. Note that this analysis is performed over 3–5 days near the landfall of each storm, and it excludes comparisons at locations where the measurements stopped during the peak of the storm, where there was an obvious datum error, or where SWAN+ADCIRC never wetted due to a lack of sub-mesh-scale features.

The SI are good, with values of 0.16–0.28 for the water levels and values in the range of 0.12–0.57 for the wave parameters. Note the general agreement between the computed SWAN wave parameters and the measured results, indicating that SWAN simulates well the deep-water wave behavior in the Gulf. The mean SI and $Bias$ values are larger for Gustav and Ike than for the previous storms, especially for the significant wave heights and peak wave periods, but these larger values reflect the difficulties in gathering data in the nearshore environment, as well as the sensitivity of the models to the representation of the fine-scale bathymetry/topography and bottom roughness. When the USACE-CHL gages within the Biloxi and Terrebonne marshes are excluded from the analysis, the mean SI errors decrease to 0.21–0.32 for the significant wave heights and 0.23–0.39 for the peak wave periods, and

the mean *Bias* errors decrease to -0.05 – 0.22 for the significant wave heights and -0.01 – 0.04 for the peak wave periods.

These results indicate that SL16 provides a faithful representation of Southern Louisiana and its response to hurricane waves and storm surge. These four recent hurricanes impacted different parts of the region and in different ways. Katrina and Gustav pushed waves and surge onto the relatively narrow Louisiana-Mississippi shelf and threatened New Orleans, while Rita and Ike made landfall farther west and flooded large portions of the marshes and bayou of southwest Louisiana. SWAN+ADCIRC and the SL16 mesh simulate well these different responses.

5 Performance

5.1 Computing Clusters

SWAN+ADCIRC and its component models were benchmarked on three high-performance computing platforms: Kraken, Ranger and Lonestar.

Kraken is a Cray XT5 cluster located at the National Institute for Computational Sciences (NICS, <http://www.nics.tennessee.edu>). It contains 8,256 compute nodes, each of which has two six-core AMD Opteron processors, and thus the system has a total of 99,072 computational cores. Each node contains 16GB of memory that is shared among the cores, and the overall system has 129TB of memory and a theoretical peak performance of 1030 TFLOPS. The specifications of the nodes are shown in Table 3.

Ranger is a Sun Constellation Linux cluster located at the Texas Advanced Computing Center (TACC, <http://www.tacc.utexas.edu>) at the University of Texas at Austin. It contains 3,936 SMP compute nodes, each of which has four quad-core AMD Opteron processors, and thus the system has a total of 62,976 computational cores. Each node contains 32GB of memory that is shared among the cores, and the overall system has 123TB of memory and a theoretical peak performance of 579 TFLOPS. The specifications of the nodes are shown in Table 3.

Lonestar is a Dell Linux cluster located at the Texas Advanced Computing Center (TACC, <http://www.tacc.utexas.edu>) at the University of Texas at Austin. It contains 1,888 M610 compute nodes, each of which has two six-core Xeon 5680 processors, and thus the

Table 3 Specifications of the compute nodes on the NICS Kraken, TACC Ranger and TACC Lonestar machines

| | NICS Kraken | TACC Ranger | TACC Lonestar |
|--------------|--------------------------------|---------------------------------|-----------------------------|
| Node | Cray XT5 | Sun Blade x6420 | Dell PowerEdge M610 |
| CPU | 2 Six-core AMD Opteron 8435 | 4 Quad-core AMD Opteron 8356 | 2 Six-core Xeon 5680 |
| Core | AMD Opteron 8435 | AMD Opteron 8356 | Xeon 5680 |
| Frequency | 2.6 GHz | 2.3 GHz | 3.33 GHz |
| Architecture | AMD K10 (Istanbul) | AMD K10 (Barcelona) | Intel Nehalem (Westmere-EP) |
| L1-Cache | 64 + 64KB per core | 64 + 64KB per core | 32 + 32KB per core |
| L2-Cache | 512KB per core | 512KB per core | 256KB per core |
| L3-Cache | 6MB on die shared | 2MB on die shared | 6MB on die shared |

system has a total of 22,656 computational cores. Each node contains 24GB of memory that is shared among the cores, and the overall system has 44TB of memory and a theoretical peak performance of 302 TFLOPS. The specifications of the nodes are shown in Table 3.

5.2 Timing Studies

SWAN+ADCIRC and its individual model components were run on all three machines to produce timing statistics. For these simulations, the wind field for Katrina was used as forcing for all models and was either read as an input file every 15 minutes on a regular 0.05° mesh (for all model runs, in the same format as described above for the validation) or generated internally using the Holland vortex wind model (for only the ADCIRC and SWAN+ADCIRC simulations, [17]). Both SWAN and ADCIRC were hot-started at 2005/08/28/0000Z and then run for two days until 2005/08/30/0000Z, thus capturing the peak of the storm as it was making landfall in southeast Louisiana. SWAN+ADCIRC employed the full coupling, with ADCIRC passing winds, water levels, currents and roughness lengths to SWAN, and SWAN passing wave radiation stress gradients to ADCIRC. When ADCIRC was run by itself, it did not receive wave forcing from any source. When SWAN was run by itself, it read and interpolated the wind field to its unstructured sub-meshes, but it did not receive water levels, currents or roughness lengths from any source. File output was disabled in all models.

Figure 14 shows timing results for SWAN and ADCIRC, using both the implicit and explicit solutions of the GWCE, for all three machines. The average time of three runs is reported for each simulation. The solid curves correspond to winds being read as input files, while the dashed curves correspond to wind data being generated internally using the analytical vortex model. We note several aspects of the behavior of the coupled model and its components.

First, SWAN and the explicit version of ADCIRC scale linearly through 7,168 computational cores on Kraken and Ranger, and through 8,196 cores on Lonestar. It is clear that at a very high number of cores, scalability is impacted by the reading of wind data every 15 minutes of simulation time on a regular 0.05° mesh; these wind data are interpolated locally onto the unstructured mesh for each core. However, the fact that each core reads directly the wind file limits scalability. In contrast, the explicit ADCIRC simulation with the winds generated internally scales through 8,192 cores.

As expected, the implicit version of ADCIRC scales linearly as well, but only through about 3,072 cores on Kraken and Ranger. Its required global communication in the matrix solution causes its scaling to tail off at higher numbers of computational cores. When used with the same time step, the explicit version of ADCIRC is much faster than the implicit version; for example, on 1,014 cores on Ranger, the explicit version of ADCIRC requires only 35 min/day of Katrina simulation, compared to 76 min/day for the implicit version. A comparison of the maximum water levels between implicit and explicit solutions in ADCIRC showed minor differences only in regions near wetting and drying, and thus the explicit version of ADCIRC was deemed identical on the SL16 mesh with a time step of 1 s.

When SWAN is coupled to the explicit version, the resulting model also scales linearly, achieving minima of 28 and 32 min/day of Katrina simulation on Ranger and Kraken, respectively, and a minimum of 10 min/day on Lonestar. However, the internally-generated winds are again necessary to achieve scalability through 8,192 cores on Kraken and Ranger. It is noted that the effect of the wind file reading is not as noticeable of a deterrent to scalability for the coupled model. The deterioration of the implicit ADCIRC also affects the coupled model, which loses its linear scaling when it is decomposed on numbers of cores larger than 3,072 for Kraken and Ranger, but maintains scalability on Lonestar.

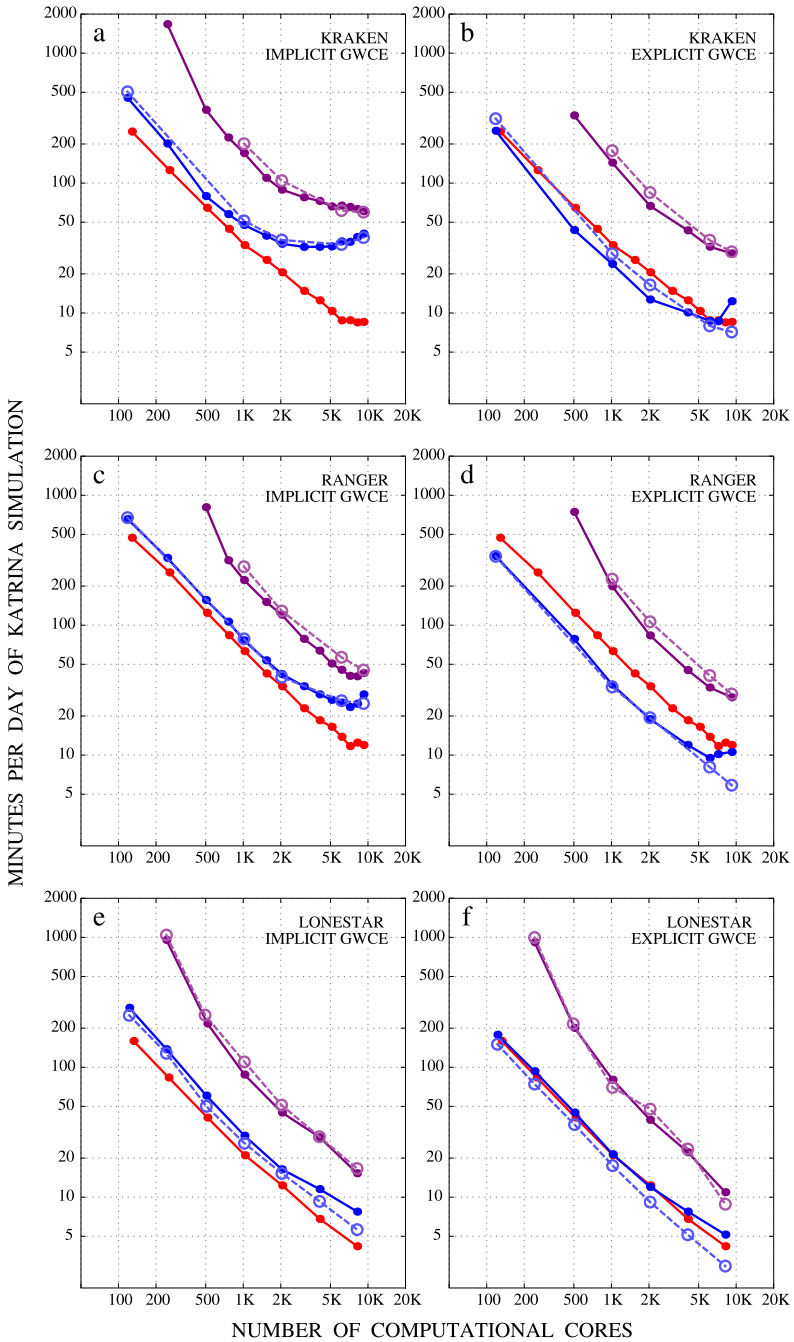


Fig. 14 Benchmarking results for SWAN (red), ADCIRC (blue) and SWAN+ADCIRC (purple). The top row shows results on NICS Kraken, the middle row shows results on TACC Ranger, while the bottom row shows results on TACC Lonestar. The left column shows ADCIRC results using an implicit solution of the GWCE, while the right column shows ADCIRC results using an explicit solution of the GWCE. The dashed lines in all panels show results for ADCIRC with the Holland wind model

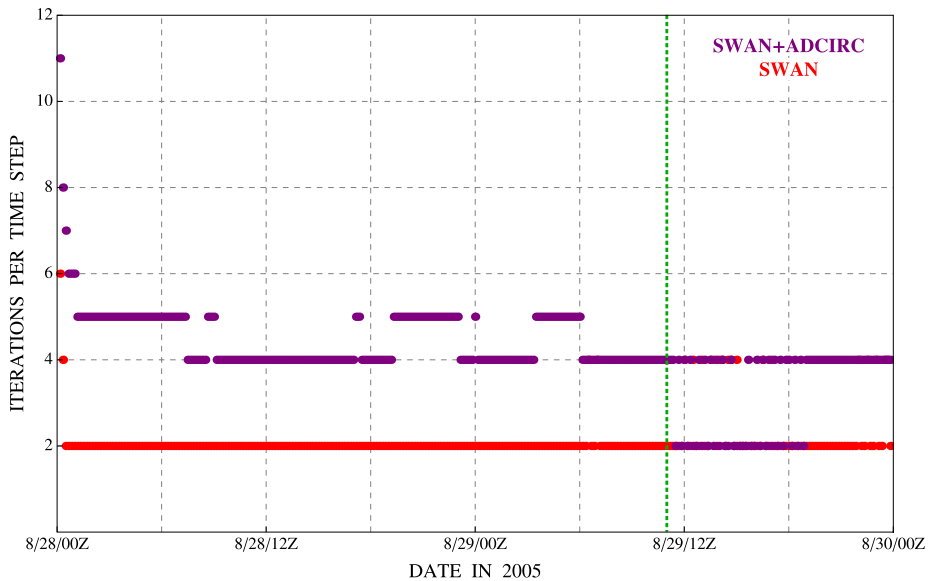


Fig. 15 Iterations per time step for SWAN when it is run individually (red dots) and as part of the coupled model (purple dots), for the two-day Katrina timing simulation on 1,024 cores. The vertical, dotted green line indicates the initial landfall of Katrina along the lower Mississippi River

Second, the individual computational cores are faster on Kraken than Ranger, but its interconnect is slower. The faster speed is evident in the timing curves; for example, on 1,024 cores and using the implicit solution of the GWCE (Fig. 14a, c), the simulation requires 222 min/day on Ranger, but only 170 min/day on Kraken. However, at high numbers of cores, the network on Kraken deteriorates the models' performance, especially that of the implicit version of ADCIRC. The scaling tails off quickly, starting at about 3,072 cores, and prevents the implicit ADCIRC timings from reaching any lower than 32 min/day of Katrina simulation. This behavior deteriorates the performance of the implicit SWAN+ADCIRC, which also does not scale as well on Kraken. Lonestar is significantly faster than both Kraken and Ranger, with a SWAN+ADCIRC simulation on 1,024 cores and using the implicit solution of the GWCE (Fig. 14e) requiring 88 min/day. The SWAN+ADCIRC timings continue to scale to a minimum of 10 min/day on Lonestar.

Third, the timings for the coupled model are greater than the combined total of the timings of its individual components. For example, on 1,024 cores on Ranger and using the implicit solution of the GWCE (Fig. 14a), the SWAN+ADCIRC timing of 222 min/day of Katrina simulation is about 60 percent larger than the combined total of 63 min/day for the stand-alone SWAN and 76 min/day for the stand-alone ADCIRC. This behavior is caused by the SWAN convergence requirements. When SWAN is run by itself, its only forcing is the wind input, and it requires only two iterations per time step. However, when SWAN is run as part of the coupled model, it receives wind velocities, water levels, currents and roughness lengths from ADCIRC, and thus its solution is more complex. Specifically, large portions of the shallow floodplain are now inundated, and significant currents are affecting large portions of both inland regions and the shelf. More iterations per time step are required, as shown in Fig. 15. The convergence requirements remain the same, but the coupled SWAN must iterate more to achieve them, and thus it becomes more expensive computationally. The

mechanics of the SWAN+ADCIRC coupling do not add any overhead to the simulation, but the coupled physics do require additional work from the individual model components.

Fourth, at small numbers of cores, the timings increase significantly for the coupled model. For 512 cores on Ranger and 256 cores on Kraken and Lonestar, the problem size becomes too large to maintain within local cache, and thus the coupled model must access data from memory. This timing increase corresponds to about 10,000 mesh vertices per core on Ranger, and 20,000 mesh vertices per core on Kraken and Lonestar. However, this behavior is understandable when compared to the amount of available cache on each machine. On Kraken and Ranger, the AMD cores have access to the L2 and L3 cache, and thus the cores on Kraken can each access a total of 1.5MB of cache, and the cores on Ranger can each access a total of 1MB of cache. In contrast, Lonestar uses Xeon cores that each have access to 2MB of cache. The problem size becomes too large on all three systems when it reaches 10,000 mesh vertices per 1MB of shared cache on a per-core basis. Thus it can be generalized that SWAN+ADCIRC requires less than 10,000 mesh vertices per 1MB of shared cache, and optimally 5,000–8,000 mesh vertices per 1MB of share cache. When the global mesh is decomposed over smaller numbers of cores, the timings become much slower than expected. However, when the local problem is small enough to be contained on cache, the coupled model does scale linearly, as noted above. The point at which the coupled SWAN+ADCIRC simulations no longer utilize cache is also dependent on the number of frequency and directional bins selected for the SWAN simulation.

Lastly, on Kraken and Ranger, all of the models experience a hard lower limit on the timings at high numbers of cores. On Ranger, even the models that otherwise scale linearly reach a minimum of about 10 min/day, and on Kraken, the minimum is about 8 min/day. This behavior is caused by the reading of the wind fields, which are made available to the models from a single set of files containing pressures and wind velocities. These two files must be read by all of the cores. This input method becomes increasingly expensive when large numbers of cores are accessing the same files. The models cannot continue to scale linearly because the cost of reading the wind fields begins to dominate the overall computational time. Note that this behavior is not as pronounced on Lonestar, which handles better the input of the wind fields from the single set of files. In contrast, when ADCIRC uses the Holland vortex wind model [17], which does not require large amounts of file input because it generates internally the wind field, the scaling continues linearly to 9,216 cores, reaching a minimum of 6 min/day of Katrina simulation.

The coupled SWAN+ADCIRC maintains the already-excellent scalability of its components, and it scales linearly to large numbers of cores, especially when ADCIRC solves explicitly the GWCE. Hurricane waves and surge can be modeled efficiently, so that results can be provided in a timely manner.

5.3 File Output

As noted above, the coupled model can output as many as 10 global output files containing the wind fields, water levels, currents and several wave parameters. These global files can be large; for example, a file containing water levels at all mesh vertices at 30-min intervals during the two-day Katrina timing simulation would be about 11GB, while a similar file containing wind velocities at all mesh vertices would be about 21GB. These file sizes would increase for longer simulations, or if data is written more frequently. This large amount of file output can slow down the computations if it is not handled properly.

The coupled SWAN+ADCIRC utilizes the stand-alone ADCIRC's capability to dedicate cores for file output [37]. Figure 16 shows timing results for file output during the two-day

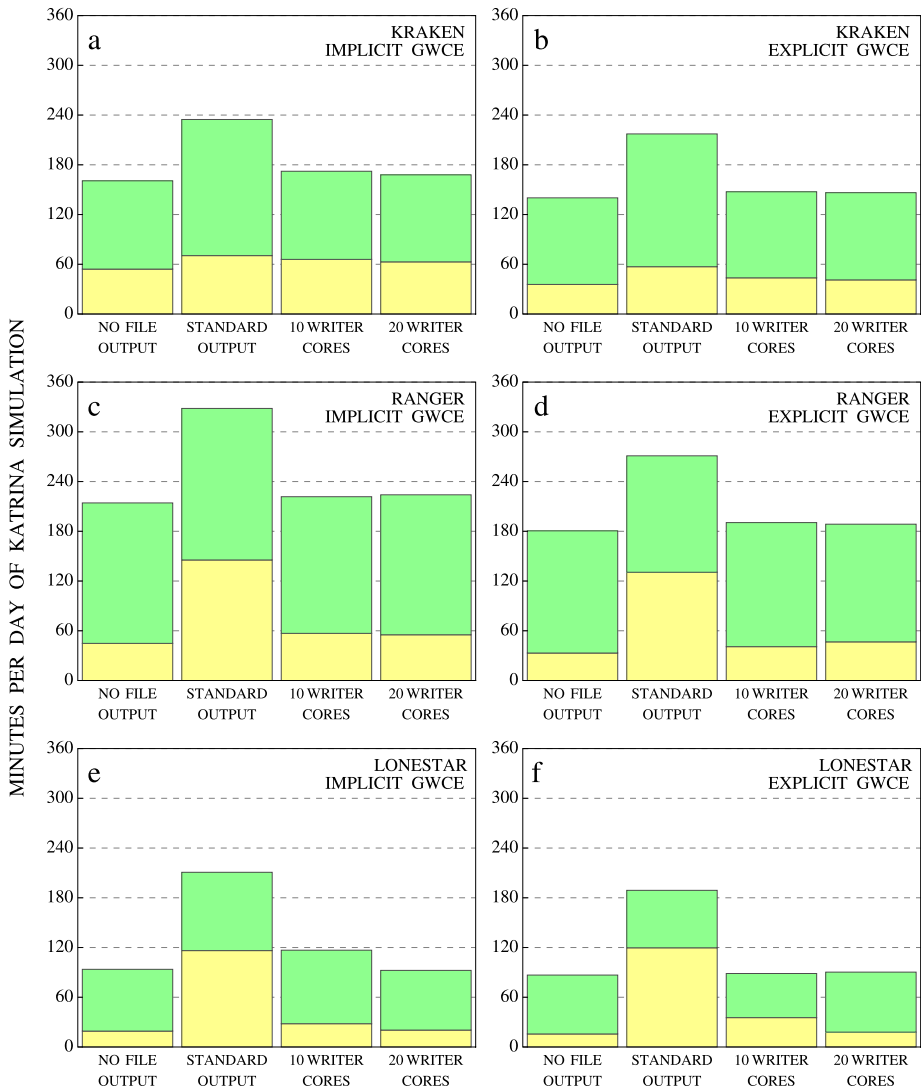


Fig. 16 Timing results for SWAN+ADCIRC on 1,024 cores (green) and 6,144 cores (yellow). The top row shows results on NICS Kraken, the middle row shows results on TACC Ranger, while the bottom row shows results on TACC Lonestar. The left column shows ADCIRC results using an implicit solution of the GWCE, while the right column shows ADCIRC results using an explicit solution of the GWCE

Katrina simulation at 30-min intervals. Note the increase in wall-clock timings when the files are written via standard output, i.e., when the computational cores must also manage the writing of the global files. On Ranger with 1,024 cores, the timings increase by 53 percent when ADCIRC solves implicitly the GWCE, and 50 percent for the explicit solution of the GWCE. On Kraken, the timings increase similarly, indicating that a significant portion of the overall simulation is devoted to file output.

When the global mesh is decomposed over 1,024 cores, the results indicate that 10 writer cores are sufficient to manage the file output without slowing the simulation. Each writer

core is able to write its time snap of data before the next time snap is available, and thus the computational cores do not have to wait to pass their information to the writer cores. The timings with 10 writer cores are similar to the timings without any file output, for all machines and both GWCE solution methods. Nothing is gained by employing an additional 10 writer cores (or 20 total).

When the global mesh is decomposed over 6,144 cores, the timings decrease accordingly; note the decrease of 79 percent in the wall-clock time for the case without file output in Fig. 16d. These faster timings are more influenced by the file output, and thus a larger increase is experienced when standard output is employed. For example, when the GWCE is solved explicitly on Ranger with 6,144 cores, the increase is 294 percent when files are written with standard output. The increases are not as acute on Kraken, indicating that machine manages better its file output.

This increase is not mitigated completely when 10 cores are dedicated for file output. The use of these 10 writer cores does decrease the wall-clock timings, but not to the same level as when file output is disabled. Whereas the writer cores increase the wall-clock timings by 3–7 percent when the global mesh is decomposed over 1,024 cores, they cause an increase of 15–40 percent on 6,144 cores. In this case, the computational cores have finished the next 30 simulation minutes before the writer cores have finished writing the last time snap of data to the output files, and thus the computational cores must wait to pass information and resume their work. When 20 writer cores are employed, the increase in the wall-clock timings is not as severe, but it does still exist.

Thus, it is important to use a sufficient number of dedicated writer cores for file output. Without them, the simulation time can be increased by 50–300 percent, depending on the domain decomposition, ADCIRC solution algorithm, and machine specifications. And, even when the writer cores are employed, it is important to employ them in a manner that does not slow down the computations. On small numbers of cores, it may be sufficient to utilize one writer core per global output file. However, on large numbers of cores, it is necessary to increase the number of writer cores by using the sequential approach, to ensure that writer cores are always available to receive information from the computational cores.

6 Conclusions

The recent coupling of SWAN+ADCIRC allows hurricane waves and circulation to be modeled on unstructured meshes that resolve their generation in deep water and on the continental shelf, their propagation into shallow water, and their dissipation in the complex nearshore environment. In a parallel computing environment, the component models alternate simulation on the same local sub-meshes. Information can be passed between models through local cache or memory, without the necessity of interpolation, thus ensuring that inter-model communication is intra-core and highly scalable. Intra-model communication is untouched by the coupling, and the component models communicate locally and/or globally, depending on their solution algorithm. The coupling paradigm adds no overhead to the computations, but it does allow waves and circulation to interact in complex problems.

This work validated the coupled model through comprehensive hindcasts of four recent, well-documented Gulf hurricanes on the new SL16 mesh, which offers unprecedented resolution in the Gulf of Mexico (4–6 km), on the continental shelf (500–1000 m), within the wave-dissipation zones (200 m), and inside the fine-scale channels and features in Southern Louisiana (down to 20 m). These hurricanes are represented by a wealth of measurement data, for winds, waves and water levels, and comparisons show that SWAN+ADCIRC simulates well the evolution of wave and circulation responses for these storms. Comparisons

to high-water marks and peak hydrograph values show the computed ADCIRC water levels to be a faithful match, with best-fit slopes near unity and correlation coefficients R^2 of 0.77–0.93. The scatter indices (SI) for the hydrographs were in the range of 0.16–0.28 for the four storms, and they indicate that ADCIRC is capturing well the hurricane forerunner before landfall, the peak water levels, and the recession afterward. When outliers are removed from the analysis, the SI for significant wave heights, peak and mean wave periods are also in the range of 0.12–0.34, again indicating a faithful simulation by SWAN of the wave generation, propagation and dissipation. It is an ongoing process to understand the required mesh resolution for both the wave and circulation models, as the resolution is dependent on storm, other physical forcings and geography.

SWAN+ADCIRC is also efficient and scales very well. The coupled model was benchmarked on the high-performance computing clusters Kraken, Ranger and Lonestar, by applying a two-day hindcast near landfall of Hurricane Katrina on the SL16 mesh. The parallel, unstructured version of SWAN scales linearly through 8,196 computational cores, with tailing off of scalability on Kraken and Ranger related to the reading of the wind input files. The scalability of ADCIRC depends on the machine and the time-stepping algorithm employed in its solution of the GWCE. The implicit solution scales linearly through about 3,072 cores on Kraken and Ranger before scalability degrades. This behavior is related to the global *MPI_ALLREDUCE* needed to solve the system of simultaneous equations. The implicit solution on Lonestar appears to degrade only due to the wind file reading. When wind data are generated internally, for either the implicit or explicit solutions of the GWCE, the simulations scale linearly through at least 8,192 cores.

The coupling paradigm adds no overhead to the computations, but the coupled physics do lengthen the run-times relative to those of the component models. The updated wind velocities, water levels, currents and roughness lengths throughout the domain require SWAN to iterate more per time step to achieve the same level of convergence, and thus its computations are more expensive within the coupled model than the sum of the component models run individually. In addition, the increased memory requirements when SWAN and ADCIRC are coupled causes the combined code to use memory outside cache earlier than the component models, resulting in a significant slowdown when an insufficient number of cores is assigned. However, SWAN+ADCIRC scales as well as its components, including linear scaling through at least 8,196 cores, on all three machines, when the GWCE is solved explicitly. The dedicated writer cores improve the wall-clock timings by 50–300 percent, depending on the domain decomposition, ADCIRC solution algorithm, and machine specifications. Future work will implement dedicated reader cores for the hurricane wind forcing data, to further improve the scalability of SWAN+ADCIRC and its model components, as well as the implementation of OpenMP, to allow the model components to run on separate cores and share information through local memory. The coupled model produces hindcasts in as little as 10 min/day of simulation on Lonestar with 8,192 cores. These results indicate that SWAN+ADCIRC is well-positioned for hurricane forecasting applications, which require rapid dissemination of results to emergency managers.

Acknowledgements This work was supported by awards from the National Science Foundation (DMS-0915223, OCI-0749015 and OCI-0746232).

References

1. Atkinson, J.H., Westerink, J.J., Hervouet, J.M.: Similarities between the wave equation and the quasi-bubble solutions to the shallow water equations. *Int. J. Numer. Methods Fluids* **45**, 689–714 (2004)

2. Battjes, J.A., Janssen, J.P.F.M.: Energy loss and set-up due to breaking of random waves. In: Proceedings of the 16th International Conference on Coastal Engineering, pp. 569–587. ASCE, Reston (1978)
3. Booij, N., Ris, R.C., Holthuijsen, L.H.: A third-generation wave model for coastal regions, Part I, Model description and validation. *J. Geophys. Res.* **104**, 7649–7666 (1999)
4. Bunya, S., Dietrich, J.C., Westerink, J.J., Ebersole, B.A., Smith, J.M., Atkinson, J.H., Jensen, R.E., Resio, D.T., Luettich, R.A. Jr., Dawson, C.N., Cardone, V.J., Cox, A.T., Powell, M.D., Westerink, H.J., Roberts, H.J.: A high resolution coupled riverine flow, tide, wind, wind wave and storm surge model for Southern Louisiana and Mississippi: Part I—Model development and validation. *Mon. Weather Rev.* **138**(2), 345–377 (2010)
5. Cardone, V.J., Cox, A.T.: Tropical cyclone wind forcing for surge models: critical issues and sensitivities. In: First JCOMM International Storm Surge Symposium. Seoul, Korea, October, 2007. Submitted to *Nat. Hazards* (2007). doi:[10.1007/s11069-009-9369-0](https://doi.org/10.1007/s11069-009-9369-0)
6. Cox, A.T., Greenwood, J.A., Cardone, V.J., Swail, V.R.: An interactive objective kinematic analysis system. In: Fourth International Workshop on Wave Hindcasting and Forecasting, Atmospheric Environment Service, Banff, Alberta, Canada, pp. 109–118 (1995)
7. Coastal Wetlands Planning: Protection and Restoration Act, <http://www.lacoast.gov>
8. Dawson, C.N., Westerink, J.J., Feyen, J.C., Pothina, D.: Continuous, discontinuous and coupled discontinuous-continuous Galerkin finite element methods for the shallow water equations. *Int. J. Numer. Methods Fluids* **52**, 63–88 (2006)
9. Dietrich, J.C., Kolar, R.L., Luettich, R.A. Jr.: Assessment of ADCIRC’s wetting and drying algorithm. In: Miller, C.T., Farthing, M.W., Gray, W.G., Pinder, G.F. (eds.) Proceedings of the XV International Conference on Computational Methods in Water Resources (CMWR), vol. 2, pp. 1767–1778 (2004)
10. Dietrich, J.C., Kolar, R.L., Westerink, J.J.: Refinements in continuous Galerkin wetting and drying algorithms. In: Proceedings of the Ninth International Conference on Estuarine and Coastal Modeling, pp. 637–656 (2006)
11. Dietrich, J.C., Bunya, S., Westerink, J.J., Ebersole, B.A., Smith, J.M., Atkinson, J.H., Jensen, R.E., Resio, D.T., Luettich, R.A. Jr., Dawson, C.N., Cardone, V.J., Cox, A.T., Powell, M.D., Westerink, H.J., Roberts, H.J.: A high resolution coupled riverine flow, tide, wind, wind wave and storm surge model for Southern Louisiana and Mississippi: Part II—Synoptic description and analyses of Hurricanes Katrina and Rita. *Mon. Weather Rev.* **138**, 378–404 (2010)
12. Dietrich, J.C., Zijlema, M., Westerink, J.J., Holthuijsen, L.H., Dawson, C.N., Luettich, R.A. Jr., Jensen, R.E., Smith, J.M., Stelling, G.S., Stone, G.W.: Modeling hurricane waves and storm surge using integrally-coupled, scalable computations. *Coast. Eng.* **58**, 45–65 (2011)
13. Dietrich, J.C., Westerink, J.J., Kennedy, A.B., Smith, J.M., Jensen, R.E., Zijlema, M., Holthuijsen, L.H., Dawson, C.N., Luettich, R.A. Jr., Powell, M.D., Cardone, V.J., Cox, A.T., Stone, G.W., Pourtaheri, H., Hope, M.E., Tanaka, S., Westerink, L.G., Westerink, H.J., Cobell, Z.: Hurricane Gustav (2008) waves, storm surge and currents: hindcast and synoptic analysis in Southern Louisiana. *Mon. Weather Rev.* (2011, in press). doi:[10.1175/2011MWR3611.1](https://doi.org/10.1175/2011MWR3611.1)
14. Ebersole, B.A., Westerink, J.J., Resio, D.T., Dean, R.G.: Performance Evaluation of the New Orleans and Southeast Louisiana Hurricane Protection System, Vol. IV—The Storm. Final Report of the Interagency Performance Evaluation Task Force. U.S. Army Corps of Engineers, Washington (2007)
15. Federal Emergency Management Agency: Louisiana Hurricane Ike coastal high water mark data collection. FEMA-1792-DR-Louisiana, Draft Report, February 2009
16. Hasselmann, S., Hasselmann, K., Allender, J.H., Barnett, T.P.: Computations and parameterizations of the nonlinear energy transfer in a gravity wave spectrum, Part II: Parameterizations of the nonlinear transfer for application in wave models. *J. Phys. Oceanogr.* **15**(11), 1378–1391 (1985)
17. Holland, G.J.: An analytical model of the wind and pressure profiles in hurricanes. *Mon. Weather Rev.* **108**, 1212–1218 (1980)
18. Kennedy, A.B., Gravois, U., Zachry, B., Luettich, R.A. Jr., Whipple, T., Weaver, R., Fleming, J., Chen, Q.J., Avissar, R.: Rapidly installed temporary gauging for waves and surge during Hurricane Gustav. *Cont. Shelf Res.* **30**, 1743–1752 (2010)
19. Kennedy, A.B., Gravois, U., Zachry, B., Westerink, J.J., Hope, M.E., Dietrich, J.C., Powell, M.D., Cox, A.T., Luettich, R.A. Jr., Dean, R.G.: Origin of the Hurricane Ike forerunner surge. *Geophys. Res. Lett.* (2011, in press). doi:[10.1029/2011GL047090](https://doi.org/10.1029/2011GL047090)
20. Kolar, R.L., Westerink, J.J., Cantekin, M.E., Blain, C.A.: Aspects of nonlinear simulations using shallow water models based on the wave continuity equations. *Comput. Fluids* **23**(3), 1–24 (1994)
21. Luettich, R.A. Jr., Westerink, J.J.: An assessment of flooding and drying techniques for use in the ADCIRC hydrodynamic model: implementation and performance in one-dimensional flows. Report for the Department of the Army, Contract Number DACW39-94-M-5869 (1995)
22. Luettich, R.A. Jr., Westerink, J.J.: Implementation and testing of elemental flooding and drying in the ADCIRC hydrodynamic model. Final Report for the Department of the Army, Contract Number DACW39-94-M-5869 (1995)

23. Luettich, R.A. Jr., Westerink, J.J.: Formulation and Numerical Implementation of the 2D/3D ADCIRC Finite Element Model Version 44.XX (2004). http://adcirc.org/adcirc_theory_2004_12_08.pdf
24. Madsen, O.S., Poon, Y.-K., Graber, H.C.: Spectral wave attenuation by bottom friction: theory. In: Proceedings of the 21st International Conference on Coastal Engineering, pp. 492–504. ASCE, Reston (1988)
25. McGee, B.D., Goree, B.B., Tollett, R.W., Woodward, B.K., Kress, W.H.: Hurricane Rita Surge Data, Southwestern Louisiana and Southeastern Texas, September to November 2005. U.S. Geological Survey Data Series, vol. 220 (2006)
26. National Data Buoy Center: <http://www.ndbc.noaa.gov>
27. NOAA Tides Currents: <http://www.tidesandcurrents.noaa.gov>
28. Powell, M.D., Houston, S., Reinhold, T.: Hurricane Andrew's landfall in South Florida. Part I: Standardizing measurements for documentation of surface windfields. *Weather Forecast.* **11**, 304–328 (1996)
29. Powell, M.D., Houston, S., Amat, L., Morrisseau-Leroy, N.: The HRD real-time hurricane wind analysis system. *J. Wind Eng. Ind. Aerodyn.* **77–78**, 53–64 (1998)
30. Powell, M.D., Vickery, P.J., Reinhold, T.A.: Reduced drag coefficient for high wind speeds in tropical cyclones. *Nature* **422**, 279–283 (2003)
31. Powell, M.D.: Drag coefficient distribution and wind speed dependence in tropical cyclones. Final Report to the National Oceanic and Atmospheric Administration (NOAA) Joint Hurricane Testbed (JHT) Program (2006)
32. Powell, M.D., Murillo, S., Dodge, P., Uhlhorn, E., Gamache, J., Cardone, V.J., Cox, A.T., Otero, S., Carrasco, N., Annane, B., St. Fleur, R.: Reconstruction of Hurricane Katrina's wind field for storm surge and wave hindcasting. *Ocean Eng.* **37**, 26–36 (2008)
33. Resio, D.T., Westerink, J.J.: Modeling the physics of storm surges. *Phys. Today* **61**(9), 33–38 (2008)
34. Rogers, W.E., Hwang, P.A., Wang, D.W.: Investigation of wave growth and decay in the SWAN model: three regional-scale applications. *J. Phys. Oceanogr.* **33**, 366–389 (2003)
35. Smith, J.M., Jensen, R.E., Kennedy, A.B., Dietrich, J.C., Westerink, J.J.: Waves in wetlands: Hurricane Gustav. In: Proceedings of the International Conference on Coastal Engineering 2010, Shanghai, China, vol. 32 (2010). Retrieved from <http://journals.tdl.org/ICCE/article/view/1027>
36. Snyder, R.L., Dobson, F.W., Elliott, J.A., Long, R.B.: Array measurements of atmospheric pressure fluctuations above surface gravity waves. *J. Fluid Mech.* **102**, 1–59 (1981)
37. Tanaka, S., Bunya, S., Westerink, J.J., Dawson, C.N., Luettich, R.A. Jr.: Scalability of an unstructured grid continuous Galerkin based hurricane storm surge model. *J. Sci. Comput.* **46**, 329–358 (2011)
38. URS: Final coastal and riverine high-water marks collection for Hurricane Katrina in Louisiana. Final Report for the Federal Emergency Management Agency (2006)
39. URS: Final coastal and riverine high-water marks collection for Hurricane Rita in Louisiana. Final Report for the Federal Emergency Management Agency (2006)
40. USACE New Orleans District, Personal communication (2009)
41. Walters, D.J.: Personal communication (2009)
42. Wave-Current-Surge Information System for Southern Louisiana: <http://www.wavcis.lsu.edu>
43. Westerink, J.J., Luettich, R.A. Jr., Feyen, J.C., Atkinson, J.H., Dawson, C.N., Roberts, H.J., Powell, M.D., Dunion, J.P., Kubatko, E.J., Pourtaheri, H.: A basin to channel scale unstructured grid hurricane storm surge model applied to Southern Louisiana. *Mon. Weather Rev.* **136**(3), 833–864 (2008)
44. Zijlema, M.: Computation of wind-wave spectra in coastal waters with SWAN on unstructured grids. *Coast. Eng.* **57**, 267–277 (2010)