GPU based synthetic scene generation for maritime environments

Wayne Keen[a], Michael Tanner[b], Charles Coker[a], Dennis Crow[c]

[a]Air Force Research Laboratory, Munitions Directorate, Eglin AFB FL 32542-6810
[b]L3 GS&ES, 51 3rd Street, Building 9, Shalimar, FL 32579
[c]Kinetics Inc, Cook, Washington 98648

ABSTRACT

Hardware and software in the loop modeling of maritime environments involves a wide variety of complex physical and optical phenomenology and effects. The scale of significant effects to be modeled range from the order of centimeters for capillary type waves and turbulent wake effects up to many meters for rolling waves. In addition, wakes for boats and ships operating at a wide variety of speeds and conditions provide additional levels of scene complexity. Generating synthetic scenes for such a detailed, multi-scaled and dynamic environment in a physically realistic yet computationally tractable fashion represents a significant challenge for scene generation tools. In this paper, next generation scene generation codes utilizing personal computer (PC) graphics processors with programmable shaders as well as CUDA (Compute Unified Device Architecture) and OpenCL (Open Computing Language) implementations will be presented.

Keywords: GPU, KHILS, HWIL, SWIL

1. INTRODUCTION

Small watercraft have demonstrated a capability to pose a significant hazard to modern warships, particularly in the crowded waters in port or in littoral waters. To deal with this threat, while minimizing danger to non-combatants, it is important that capabilities to track such vessels, assess their activities, and, if necessary, engage them be developed and evaluated on an ongoing basis. The development and testing of these capabilities requires that physically realistic models for these craft, the waters they operate in, and their interactions (wakes, dynamics etc.) be developed. To adequately support testing, such models must be computationally tractable and provide data in real- or near real-time rates. In this paper, the elements of the maritime scene will be outlined, as well as the techniques for modeling them. Finally, techniques for applying graphics processing techniques to the problem will be outlined.

2. MARITIME SCENE ELEMENTS

The principle elements of the maritime optical scene are shown in Figure 1. In this section, the components of each of the scene elements will be described, with parameters characterizing their properties and magnitudes presented. In the sections that follow, the issues that are involved with modeling these effects will be presented.

2.1 Sea State Effects

The most obvious dynamic feature of the sea surface at different sea states are the deformations in the surface due to waves. Waves are driven by the interaction of meteorological effects (wind, air pressure, etc.) with currents, tides and, occasionally, geological effects. Waves themselves can be classified by their characteristic frequencies[1], a list of classifications and their associated periods are shown in Table 1.

The combination of waves leads to the sea state. Through the years, a number of systems have been developed to characterize the sea state. These include the Pierson-Moskowitz spectrum and the co-related

Figure 1: Principle Scene Elements

World Meteorological Organization (WMO) sea state scale. Table 2 shows some example WMO sea states and the wave characteristics for each state.

Simple large scale rolling waves have a relatively small impact on the optical appearance of the sea, principally varying the angle at which solar and sky reflections take place. Smaller scale waves, towards the capillary end of the scale can, on the other hand, have impact in the form of small (but bright) glints that lead to the impression the sea surface is sparking (See Figure 2). Such sea states tend to be driven by smaller scale wind fields, and can vary significantly in amplitude and frequency content over short distances. Because of these variable small-scale driving conditions, lower level sea states (0,1) are far more complex to faithfully model optically and hydrodynamically than higher level sea states.
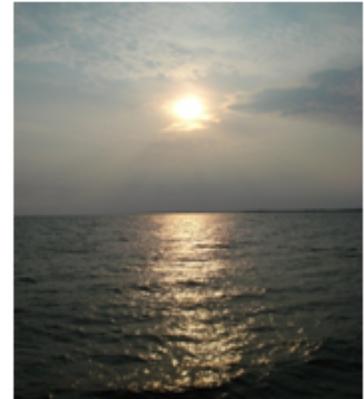

Figure 2: Sea glints

Classification	Period Range
Capillary Waves	Less than 0.1 second
Ultra-Gravity Waves	0.1 – 1 second
Gravity Waves	1 – 30 seconds
Infra-Gravity Waves	30 seconds – 5 minutes
Long Period Waves	5 minutes – 12 hours
Ordinary Tidal Waves	12 hours – 24 hours
Trans-tidal Waves	Greater than 24 Hours

Table 1: Sea wave properties

Wind Speed (Kts)	Sea State	Significant Wave (Ft)	Significant Range of Periods (Sec)	Average Period (Sec)	Average Length of Waves (FT)
3	0	<.5	<.5 - 1	0.5	1.5
4	0	<.5	.5 - 1	1	2
5	1	0.5	1 - 2.5	1.5	9.5
7	1	1	1 - 3.5	2	13
8	1	1	1 - 4	2	16
9	2	1.5	1.5 - 4	2.5	20
10	2	2	1.5 - 5	3	26
11	2.5	2.5	1.5 - 5.5	3	33

Table 2: Sea state properties

2.2 Boat Dynamics

As the boat moves across (and occasionally above) the surface of the water, it can undergo significant translations and rotations including effects such as bouncing, pitching, and splashing over the waves (Figure 3). In addition to impacting the state of the boat, such motions couple with the sea surface to create sprays and wakes that will be discussed in the following sections.

2.3 Sea-Boat Interactions

The principle interactions between the boat and water, beyond the induced boat dynamics, are the wake and spray. The primary components of the wake are shown in Figure 4. The relative magnitudes of these components are a function of the boat hull type and speed, as well as water depth.

The Kelvin wake is the most obvious feature of the wake pattern. It is characterized by a cusp wave that has a half angle of approximately 19.5 degrees with respect to the boats velocity vector. Kelvin wake patterns arise at both the bow and the stern of the boat. Depending on the boats speed, the cusp wave can be broken into a series of divergent waves. The transverse wake is prominent at lower speeds, but its wavelength grows in proportion to boat speed squared, with the effect that it disappears at higher speeds.

The wave patterns of the Kelvin wake provide a modulation in the sea height field. In and of itself, this effect has a minor impact on the optical scene, as it simply modifies the reflective angles of the sea surface. A more significant effect occurs when some of the waves involved are breaking. Breaking waves lead to foam on the sea surface which can be quite prominent optically.

The turbulent wake is generated as the result of the boats hull passing through the water, in combination with the effects of propellers. The turbulence has several effects that impact its optical appearance. Turbulence entrains air bubbles in the wake, as well as engine coolant water and exhaust. Turbulence brings sub-surface waters to the surface, and these waters can have cooler temperatures than the surface water. Finally, turbulence can bring to the surface natural and man-made materials that act as surfactants, leading to a significant dampening of wave action within the region. The net result of these effects is the creation of a region of water that is smoother than the surrounding water, and has an effective temperature that can be lower or higher than the surrounding water. This wake component is unique in that it can persist over a period of several minutes, and can be several kilometers long. The width increases slowly, increasing with distance from the boat, X, approximately as $X^{1/4}$ [2]. The fact that this wake has the effect of tracing the motion of the boat means that it serves as a curving history of boat motions, and this long term, asymmetrical history must be accounted for in any modeling effort.

Figure 3: Boat Dynamics and Spray

Figure 4: Boat wake components

Figure 5: Water Reflections

The boat motion through the water, characterized by both pitch and yaw action, as well as the boat completely exiting the water, leads to a significant amount of splash and spray of water. In addition to direct visibility of such water, its interaction with the water surface leads to phenomenology similar to that of breaking waves, with small scale waves and entrained air disturbing the sea surface. One other effect of the splashed water that must be accounted for is the effect of the cooler water in modifying the temperature of the boat hardbody.

One final interaction between the boat and the water needs to be presented here. The reflection of the boat in the water varies in prominence as a function of the engagement geometry. For relatively still water, and lower look down angles, the boat reflection can be quite pronounced and detailed, as shown in Figure 5.

2.4 Boat Thermal / Optical Behavior

The thermal behavior of watercraft can vary significantly with water and weather conditions, operational conditions, and the structure of the boat itself (Figure 6). It is important to note that some craft in this class feature hulls of coated fiberglass that can have very low emissivity and associated high reflectivity as shown in Figure 7.

The conditions inherent in operation in the maritime environment lead to wide variations in the conditions of the surface of the boat hulls, even for the same model of boat. It is important therefore that the models for standard boat materials allow for variations to be applied to the standard definitions.

Finally, it should be noted that the boat hull is not the only element contributing to the optical signature. The crew members form a significant contribution to that signature. In addition to their signature, boat crew can and do move about the boat, which can change its optical profile significantly.

Figure 6: Boat surface reflection

Figure 7: Thermal Instrumentation

3. MARITIME SCENE COMPONENT MODELING ELEMENTS

The elements of the marine scene can be divided into four main components:

- Boat hardbody and dynamics
- Sea surface geometry
- Sea surface properties
- Spray

The boat hardbody signature is generally the easiest to render optically. Optical models for this sort of hardbody target are supported by a variety of tools at AFRL/RWGGS. Thermal models for such craft, particularly those representing the target in motion, are less readily available. Fortunately, the AFRL/RWGGS team has participated in a number of data collections in which test watercraft were instrumented with a variety of temperature measurement equipment, both direct contact and optical (Figures 7 and 8). Data was collected over a variety of seasonal, meteorological and diurnal conditions. This data has allowed for the construction of baseline thermal database which can be used to support a variety of test programs.

Figure 8: Optical Instrumentation

As part of the thermal measurements outlined above, a number of measurements to better characterize realistic marine materials have been performed. Measurements of hemispherical reflectivity were also performed to characterize a variety of hull materials, in a variety of conditions of weathering (Figure 8).

The boat dynamics model consists of a simplified rigid grid model of the boat that has sampling locations along the boat's centerline and in a lateral line through the boat's center-of-gravity (Figure 9). The sea heights are computed at each sampling location and are used to generate forces and torques for the 6-DOF model. The motion model is based on using a spring-mass-damper system for each degree-of-freedom (boat height, yaw, pitch, and roll) as shown in Figure 10.

The sea surface geometry is a combination of the sea state with the disturbance generated by the passing of the boat. The height field created by the sea surface constitutes a driving function for the boat dynamics as well as a source for obscuration for lower sensor observing angles. Descriptions of sea surface geometries have been obtained from a variety of sources, ranging from relatively simple physics models generated in-house to databases generated for specific watercraft using commercial fluid dynamics tools. Future efforts requiring generation of this data will include tools such as Wave Watch III[3].

One complication when modeling sea surface details is the varying scale of the features involved. The capillary wave components, as well as the turbulent effects of breaking waves and spray on the sea surface leads to effects whose spatial extent is typically measured in centimeters. For many modeling tasks, such small scale effects could in fact be safely ignored. However, for maritime scenes, such effects as the glints produced can be quite bright. To preserve the fine scale information, while avoiding having the sea surface data generated at such fine levels, a method has been developed to interpolate the geometry data and encode the relevant surface data, such as surface normals, into textures. This technique allows optimized computation, with the reflection and emission operations performed directly on the Graphics Processing Unit (GPU). One advantage of such an approach, which will increase going forward with GPU implementations, is that it lends itself to frame buffer copy operations which enable scene features, such as noising, to be implemented in an optimized fashion.

Figure 9: Rigid boat model grid

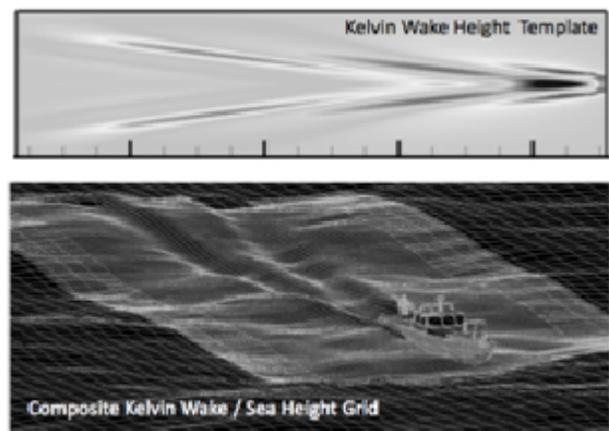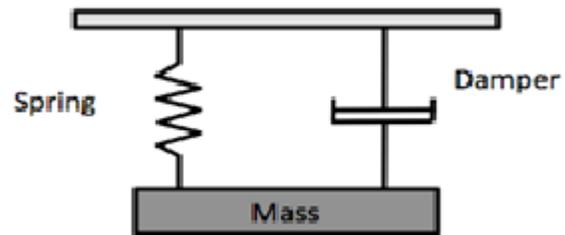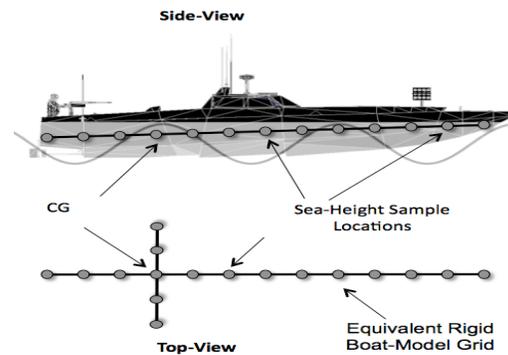Figure 10: Sprint dynamics model

Figure 11: Wake height field

Modeling of boat wakes, as outlined in the previous section, is composed of two main elements, the Kelvin wake, and the turbulent wake. For boats operating within the planning regime, the Kelvin wake simplifies significantly. The cusp waves take the form of a simple "V" shape, with the transverse component disappearing.

The Kelvin wake disturbance is input into the maritime scene generator as a height field (Figure 11). To date, this height field has been generated using a range of modeling complexity, including simplistic numeric models, moderate level of fidelity tools such as the Michlet[4] code, and high level CFD codes.

Both the Kelvin and turbulent wakes are computed for steady-state conditions on fixed, axi-symmetric wireframe grids. To account for boat heading changes during the simulation, these grids are convectively recomputed based on the velocity and turn-rate of the boat. This allows boat maneuvers and serpentine-like boat motion to be considered since the method provides a time-history of the wake evolution. An example of the wake from a small boat turning during the simulation is shown in Figure 12.

Figure 12: Curving wake

The turbulent wake is represented by a semi-empirical model that applies an offset temperature-field using an off-line model that includes Perlin[5] noise variations (Figure 13). Like the Kelvin wake, this temperature-field is towed behind the boat and is used to modify the uniform temperature of the sea.
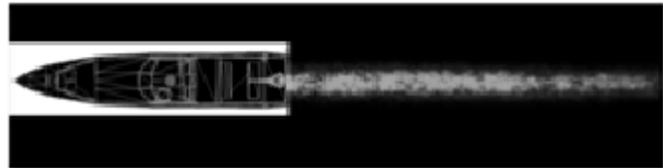

Figure 13: Turbulent wake representation

The sea-spray generated by a small boat moving through a dynamic sea is an important part of the scene radiometrics since it varies considerably over time as the boat encounters waves and swells in the sea. The boat sea-spray is modeled by using a particle system to decompose the spray into individual particles that are ejected around the boat's waterline. Each particle follows a 3-DOF trajectory that is subject to gravitational and drag forces as the particle moves through the air (Figure 14). Each particle has a finite area and is modeled in a billboard fashion that always faces the observer. The finite area is textured with a semi-translucent cloud so that a smoky appearance is generated as many individual particles overlap as the particle spray is ejected.
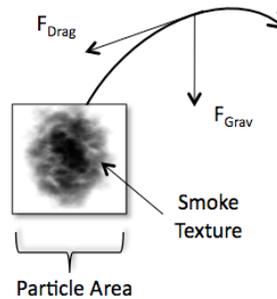

Figure 14: Spray Particles

The boat sea-spray model is used to represent the steady-state spray in the side-aft areas as the boat displaces water moving through the sea. The boat sea-spray model is also used to model the dynamic nature of the bow-spray as the boat encounters waves and swells. Particle ejection locations are arranged around the boat waterline and have different source conditions depending on where they are located (see Figure 13). The dynamic ejectors, for example, are triggered by the boats pitch-rate as computed by an internal 6-DOF model.
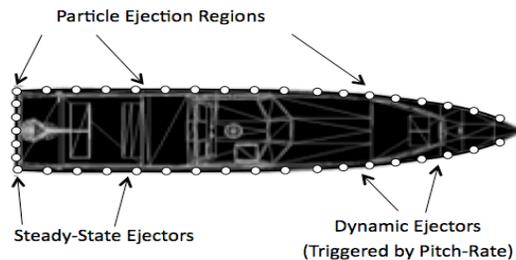

Figure 15: Spray ejector locations

Boat self-reflections off the sea are a result of the boat thermal emission being reflected off the nearly specular sea and onto the sensor. This is a classic multi-bounce reflection phenomenon and it is currently being addressed using a semi-empirical planar reflection model commonly used in computer graphics[6]. To do this calculation the model renders a

second copy of the boat while scaling the heights coordinates with a negative one. This process essentially inverts the boat. The radiance of the inverted boat is then scaled based on empirical data to account for the scattering losses, as the original source radiance is re-reflected off the sea towards the observer.

The user then enables object blending when the object is rendered so that the graphics hardware adds the self-reflection of the boat to the sea radiance it covers (Figure 16).

When considering blending, one subtle but important point is the order of rendering within a scene. Because a portion of the boat is actually underwater, it is possible to have water appear inside the boat. For mostly closed craft, such as the one in Figure 16, this is not a significant issue, but for most watercraft such water intrusion will be visible due to the open nature of their geometry. To deal with this issue, it is necessary to render a "hole" in the water corresponding to the boat's hull prior to actually rendering the hull within the water.
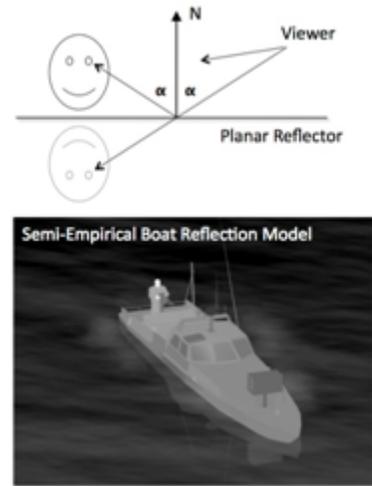
Figure 16: Boat reflection

4. GPU HISTORY AND APPLICATION TO SCENE GENERATION

Utilization of graphics hardware to perform scene generation calculations for hardware- and software-in-the-loop testing has a long history at RWGG. Initial work in this area centered on high end graphical work stations, in large part due to bit limitations of commercial PC hardware. The cost of such work stations, both from a purchase and a maintenance standpoint, led to the adaptation of workarounds which allowed PC hardware to be initially utilized[7]. The gaming industry drove a significant wave of development of the numerical capabilities in the 2004 – 2010 time frame, leading to a growing revolution in the form of the General Purpose computing on Graphics Processing Units (GPGPU). The relatively low cost commercial graphics cards have added hardware capabilities in the area of programmable stages, with higher arithmetic precision in the rendering pipeline. This has enabled stream processing to be performed in a highly parallel and efficient manner for both graphical and non-graphical data.

Current generation modeling capabilities developed at AFRL/RWGGS utilize programmable shaders to perform radiance calculations on surfaces. For data structures such as sea surfaces, that might have height field databases requiring gigabytes of memory, tiling techniques, sea state interpolation, and memory paging techniques have been developed to avoid issues with graphics memory limitations.

Continuing to use the sea surface radiance model as an example, for each pixel within the seeker FOV, a radiance computation is performed on a GPU to determine the apparent radiance at the sensor (Figure 17). The radiance computation includes thermal emission from the sea, reflected solar irradiance, reflected sky-shine irradiance, atmospheric path radiance between the sensor and sea-pixel, and atmospheric path transmission loses. The sea surface normals are calculated "on-the-fly" from the sea height-field and Bidirectional Reflectance Distribution Function (BRDF) is used in the reflection computations.
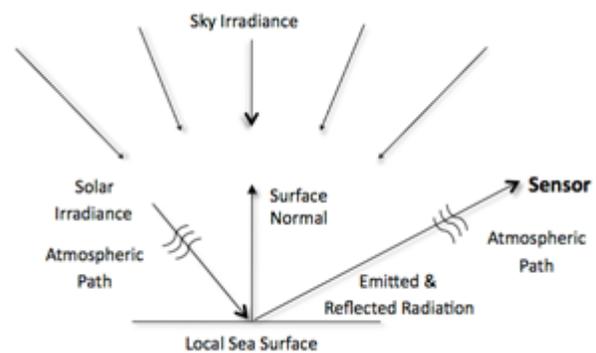
The atmospheric path effects are modeled using a three-dimensional lookup table of pre-computed MODTRAN[8] runs.

Figure 17: Sea surface optical model

5. PERFORMANCE CHARACTERIZATION AND PATHS FORWARD

A profile of code operation on a typical marine scene is shown in Figure 18. The primary bottlenecks in the scene generator performance are due to the sea state and bow-spray computations. The sea state requires significant time to performing the temporal interpolation on the CPU and could be sped up considerably by representing it as a 3D texture and moving the database pages to the GPU for processing.

The bow-spray takes most of its time computing the 3-DOF trajectories for the thousands of particles being ejected. This calculation is done on the CPU and could be done much faster if moved to the GPU.

It should be noted that the bulk of the development work for the maritime modeling capability was performed in a timeframe in which the G80 family of GPU's represented the state of the art. The current generation of hardware offers nearly 5 times as many processing units as the development hardware, as well as offering significantly improved numerical precision.

Support for GPGPU programming languages such as CUDA and OpenCL, as well as growing momentum to use this hardware as massively parallel computing devices from a variety of industries, are providing dramatic upgrades to the techniques described here. In addition, the vendor community has shown an increased interest in providing improved support for access to data directly from the cards themselves, rather than having to transfer through the CPU. Such direct memory access (DMA) would provide dramatically improved support for hardware-in-the-loop applications.

Figure 18: Performance profile

6. SUMMARY

The development program for maritime modeling at AFRL/RWGG has demonstrated the capability to model the key phenomenology of the maritime scene, including small boat thermal signatures, sea state, wakes, and optical environments. The graphical methods applied to this problem have broad implications for scene generation software. The capability of the hardware to perform massively parallel calculations on streams of data can be exploited to perform a wide range of optical and thermal phenomenology modeling, including thermal solvers, plume flowfield chemistry and radiance calculations, and environmental effects. Such capability can provide support for increased scene fidelity, as well as improved frame rates to support hardware-in-the-loop and software–in-the-loop testing applications.

7. ACKNOWLEDGEMENTS

8. REFERENCES

[1] Reddy, M.P.M., Descriptive Physical Oceanography, Swets & Zeitlinger Publishers, India, page 166, 2001
[2] Buller, E.H., Tunaley, J.K.E., The effect of the ships screws on the ship wake and its implication for the radar image of the wake. Proc. IGARSS 89, Victoria, B.C. Canada, pp. 362-365,1989
[3] http://polar.ncep.noaa.gov/waves/wavewatch/wavewatch.shtml#documentation
[4] http://www.cyberiad.net/michlet.htm
[5] http://http.developer.nvidia.com/GPUGems2/gpugems2_chapter26.html
[6] http://www.cse.ohio-state.edu/~holewins/781/implementation.html
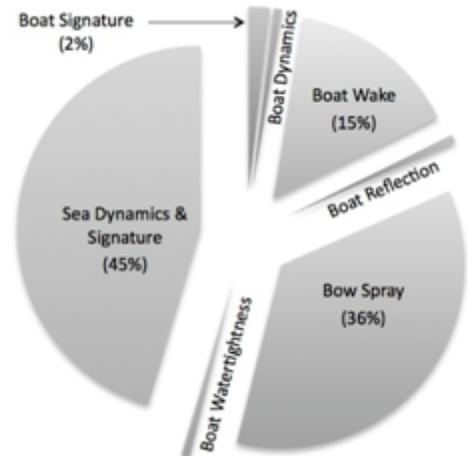
[7]   Bergin T., Thompson R., "Application of Innovative Rendering Techniques for Hardware-In-the-Loop (HILS) Scene Generation", Technologies for Synthetic Environments, Hardware-In-The-Loop Testing VIII, Proc of SPIE, Vol. 5092, 2003

[8]   http://www.kirtland.af.mil/library/factsheets/factsheet.asp?id=7915