

# Real time modelling of multispectral ocean scenes

G. Monnier (1), S. Vince (1), J. Houssay (1)

1. Alyotech Technologies, Rennes, France, {goulven.monnier, sebastien.vince, julien.houssay}@alyotech.fr

## ABSTRACT

In order to address the issue of scene simulation in marine environment with optical sea clutter, Alyotech Technologies recently developed a real-time model of the wind-driven sea surface radiance in the IR and visible spectrum. While IR surveillance from surface ships is among the first considered applications, the model was specifically designed to face the tricky problems related to observation at grazing angles. For this purpose, special effort has been carried out to deal efficiently with the following issues:

- dynamical computing of surface geometry from wave height spectra, including some (limited) nonlinearity,
- representing the surface on optimized multiscale mesh,
- global illumination of the ocean surface by partially cloudy sky-domes,
- dynamical estimate and rendering of unresolved surface rugosity accounting for both capillary waves and distant gravity waves.

Real-time animation and rendering for meshes as large as several  $10^6$  polygons is achieved through massive parallelization on GPU. Full sky domes for both global illumination and sky rendering are precomputed using SKYGEN, a cloudy-sky simulation software developed at Alyotech.



Figure 1 : real time RVB simulation with increasing wind speed

## I. INTRODUCTION

Electro-optical surveillance in the marine environment is of major concern for current defence and security applications, especially when asymmetric threats are considered. Specific issues in threat detection and tracking over sea surface backgrounds mainly own to waves spectral distribution and kinematics, and sea surface illumination and observation conditions. Thus, scene simulations should have the ability to reproduce the temporal behaviour and geometry of sea surface, while accounting accurately for direct solar and global illuminations. Deterministic models of sea-surface 3D geometry and radiometry are very demanding in terms of computational resources, especially when observation at grazing angles is considered, at least for the following reasons:

- wave components over a large range of wavelength have to be taken into account, in order to reproduce both large scale occlusion and small scale slope distribution,
- at grazing angles, a typical surveillance sensor field of view spans extended areas on the surface, inducing large ranges in distance and resolution.

Due to large variations of the sensor pixel footprint on the surface and concentration of a significant part of steepest slopes in the shortest waves scales (which can thus be either resolved or unresolved, even in a single wide FOV), sea surface optical behaviour can range from nearly smooth to very rough within a single scene. This is perhaps the first obstacle one has to overcome in computing realistic synthetic ocean scenes. The current paper aims at describing the way it is treated in our recently developed real time model running on GPU.

## II. SEA-SURFACE MOTION

The generation of the surface is achieved through the following method, which has been extensively used and described in the world of computer graphics:

- a theoretical or measured 2-D wave power spectrum is used as an initial statistical description of the mean sea surface state. The analytical formulations we used are the

JONSWAP spectrum and its improved version developed by Elfouhaily, Chapron and Katsaros [Elfouhaily et al. 1997].

- spatial frequencies are computed from temporal frequencies by use of the appropriate dispersion relationship (the infinite depth limit  $k=\omega^2/g$  hereafter)
- an instantaneous realization of the mean spectrum is simulated through a complex multiplicative Gaussian noise [Tessendorf, 2001] [Thon and Ghazanfarpour, 2002].
- for each simulation time step, the resulting noisy spectrum is shifted in time and converted into surface displacements, using either FFT or direct summation of wave components in the space domain.

Although FFT is the numerically most efficient implementation of DFT, the constraint of computing spectra and displacements on a cartesian grids is a major drawback in that it forces constant resolution in both wave vector and space domains:

- due to uniform spectral resolution, very large grids are necessary to cover the whole spectrum, from the smallest resolved length (defined by the model) to the longest waves (defined by the spectrum),
- uniform grids in space are in contradiction with the need of varying resolution related to wide-FOV observation at grazing angles.

For these reasons, some models make use on direct summation of Fourier components in the space domain, which allows non-uniformly sampled spectra and spatial meshes . However, it has been shown that, in order to achieve a satisfactory number of Fourier components, only FFT-based computations were compatible with real-time requirements, and this conclusion even stands for GPU programming. Thus, we chose to retain the FFT approach, while developing the following scheme for compensating its inherent drawbacks:

- the full 2D-spectrum is divided into two (or more) frequency ranges, from which displacements are computed through FFTs over two (or more) separate grids with high- and low- resolutions.
- A non-uniform spatial mesh is created (see details below), whose vertices are displaced by interpolating and summing the contributions from the high- and low-resolution grids.

The basic mesh is a simple large square grid, in which smaller square grids with twice the initial resolution are iteratively nested (see Figure 2). The number of refining levels (10 in the implementation demonstrated hereafter) as well as the respective size of the nested grids can be freely tuned. This mesh design aims at achieving the highest resolution within a limited region of interest, while covering a much wider sea area. This is of primary importance if observation is made at grazing angle, typically at the horizon either from a ship

or from the shore. Note that the high resolution central patch can be moved freely on the sea surface in order to follow a possible sensor footprint motion.

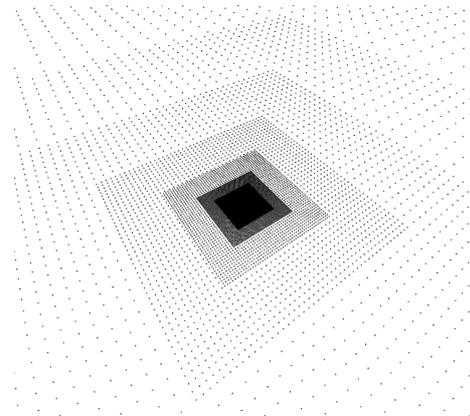


Figure 2 : design of the non-uniform mesh.

In the implementation described here, individual wave components include not only vertical, but also horizontal displacements according to the so-called Gerstner model. This model states that, from a lagrangian point of view, a particle at the sea surface exhibits an orbital motion in a the vertical plane parallel to the dominant wave propagation direction. As such, it is a simple non-linear model which succeeds in reproducing one of the most characteristic features of the sea surface, e.g. the asymmetry between sharp ridges and rounded troughs, which grows with the waves amplitude and ultimately leads to wave breaking. Despite some limitations in terms of hydrodynamics (the Gerstner model is by no way an acceptable approximated solution to the Navier-Stokes equations), this model has been widely used because of its relative simplicity and low cost when compared to more rigorous higher-order non-linear wave models. Various works, still in progress, aim at developing methods for modifying the input spectrum in such a way that the output surface fits high-order models. In this context, the Gerstner approach offers an affordable way to introduce non-linearity in sea-surface models devoted to rendering and imagery (not to hydrodynamics), an will most probably be further improved through spectrum controlled modification.

In our model, the implementation of the Gerstner model is achieved by computing horizontal displacements as a time-shifted and scaled version of vertical displacements, which requires three FFT operations for each grid instead of one. Thus, with two grids (low- and high- resolution), six 512-by-512 FFTs are computed at each time step. The 3D displacements are then interpolated on the vortices of the surface mesh, which typically has  $1.10^6$  to  $3.10^6$  polygones. In order to achieve real time performances, the corresponding code, developed under the CUDA NVIDIA library, is executed on the GPU.

### III. RENDERING

Once the surface mesh is generated, corresponding radiances are rendered as observed by a sensor located at some reasonable height above the sea. In all wavebands, the reflectance of smooth water surfaces reaches nearly-unit values at high incidence observation. While observing the sea at grazing angles, the effective reflectance is reduced because of the slope distribution about the horizontal. However, it remains significant when approaching the horizon, where reflected radiance from the environment (mainly the sky and the sun) thus generally becomes the dominant contribution. To the contrary, radiance observed at nearly normal incidence is often dominated by the water “bulk” contribution, i.e.:

- in LWIR and MWIR : mostly thermal emission from a thin opaque skin layer.
- in the visible range: environment (sky and sun) illumination, transmitted twice through the surface, scattered in the water bulk and possibly on the ocean floor.

The current implementation of the rendering module, which benefits from our multispectral cloudy sky-dome simulator SKYGEN [ref ITBM&S 2007], accounts for:

- Global illumination (cloudy sky and sun)
- Surface reflection (with varying roughness BRDF)
- Atmospheric transmission from sea surface to the observer
- Sea surface thermal emission

#### GLOBAL ILLUMINATION

Global illumination accounts for

- diffuse solar and thermal illumination from the sky, precomputed through SKYGEN
- direct solar illumination from the sun globe, computed in real time.

SKYGEN has been designed for efficiently computing full sky-domes based on MODTRAN radiances, on polar or Mercator projection. It accounts for one heterogeneous 2D cloud layer, which fairly approximates a large range of marine conditions with full or partial cloud cover. Particular attention was given to the modelling of regions of potentially high brightness like cloud borders and solar halo. For the needs of marine scenes simulation, a new function has been added to SKYGEN, preparing and storing all required radiometry and propagation parameters which characterize specified meteorological conditions.

#### REFLEXION AT THE SURFACE

Accounting for diffuse scattering of the sky dome light incoming on the rough sea surface is challenging, especially in the framework of real time constraints. The

computation challenge mainly owns to the following observations:

- any observed sea surface location is characterized by a local BRDF which results from unresolved roughness, which in turns depends on sea surface state and observing distance and direction.
- Radiance from the whole sky dome should be gathered at each observed location at the surface, given its BRDF.
- BRDF and radiance gathering have to be recomputed at each time step

Methods for estimating sea surface local BRDF, given observing and lighting directions and pixel footprint size, have been described in the literature [e.g. Caillaud et al. 2007]. Considering the unresolved wave distribution to be completely realised in a given pixel (though this approximation should be considered with caution), the wave spectrum determines the local distribution of normals. In addition, masking and shadowing in the unresolved surface patch itself have to be accounted for, especially at grazing angles. Very complex functions have been developed to achieve this goal.

Integrating illumination (mainly from sky and the sun) over all incoming directions of the BRDF is the most consuming step in the rendering process, especially because we wish to account for realistic complex sky domes. This problem has been extensively treated in the field of computer graphics; relevant methods recently developed include decomposition of both the so-called “environment map” (here, mapped radiances from the sky dome) and the BRDF over an appropriate basis of angular functions, like spherical harmonics or spatial classes of wavelet functions. This decomposition is typically computed “off-line”, so that the angular integration reduces to dot products of precomputed vectors during the real-time rendering process. Spherical harmonics has proven very efficient provided the BRDF or the environment map can be acceptably approximated by a short series of SH functions. This stands for very diffuse scatterers and/or very isotropic illumination, which is unfortunately nearly never realised in an ocean scene observed at grazing angles. The main reasons are:

- At small distance, high resolution and limited roughness (moderate instantaneous wind), smooth patches of the surface can be observed, thus exhibiting strongly picked (nearly specular) BRDF.
- Unresolved roughness increases rapidly with distance, enlarging the BRDF in the vertical direction. But, to the contrary, increasing incidence thickens the BRDF horizontally. Those contradictory processes lead to BRDFs shaped like vertical blades near the horizon, the most obvious manifestation of which is the sun glitter pillar.

In those two extreme cases, the BRDF includes sharp features which may be reproduced through spherical harmonics at the expense of very large series (typically to the 60th order, meaning more than 3000 coefficients), involving very expensive computations (especially for rotating SH coefficients). Recently proposed wavelets basis may help to overcome this problem.

After unsatisfying testing of SH decomposition, we finally chose another approach, with virtually no restriction on the BRDF shape. Namely, given a set of BRDFs allowing the description of sea surface polygons in the scene, a corresponding set of reflected radiance maps is pre-computed. Each of this maps is in some way an image of the sky dome observed by reflection in the rough surface described by the corresponding BRDF. Its computation involves generating a small 3D piece of the unresolved rough surface, and Monte-Carlo reflective rendering of the initial full-resolution. This method simply intends to perform the required hemispheric integration off-line, and to store the resulting reflected radiances as a series of “blurred” sky domes.

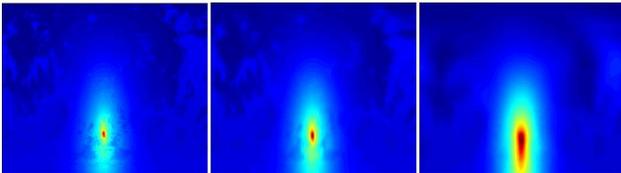


Figure 3 : example of pre-computed blurred sky-domes, with increasing surface roughness.

The Monte-Carlo computation allows natural account for local masking-shadowing, and could be extended to multiple reflections as well as under-water scattering. The pre-computed blurred sky domes are loaded as texture vectors on the GPU, and accessed through the shaders during real-time rendering according to:

- LOS and local normal directions,
- Local sea surface roughness
- Local mesh resolution

Basically, this scheme allows maintaining a homogeneous roughness over all the observed surface, while accounting for waves which are unresolved even in the finest mesh area (namely capillary waves with wavelengths smaller than 0.2 m) and waves that become unresolved in remote areas of lesser resolution. Moreover, commonly observed spatial variations of the roughness, caused for example by biological or oil slicks or current divergence / convergence.

## WHITE CAPS

Thanks to the Gerstner approach, waves crests tends to get sharper with increasing sea state, as expected from observation. White caps are another dominant feature of moderate to strong winds conditions. Their dynamics and radiative properties are complex, and even not fully understood. Thus, our code currently

includes a simplistic diagnostic white-caps model based on the following principles:

- According to [Sang-Ho Oh et al., 2004], foam location is conveniently determined at each step through a vertical acceleration threshold. This threshold is dynamically computed so that the surface fraction covered with foam obeys the laws given by [Lafon et al., 2005].
- White caps are considered as lambertian scatterers in the visible range, and as perfect blackbodies at the sea-surface temperature in the infrared.

GPU and shaders offer good potential for developing much more realistic white-caps and breaking waves model.

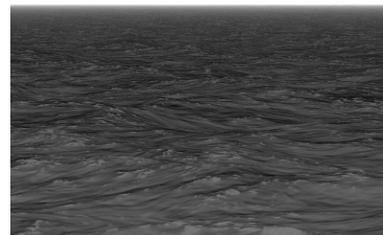
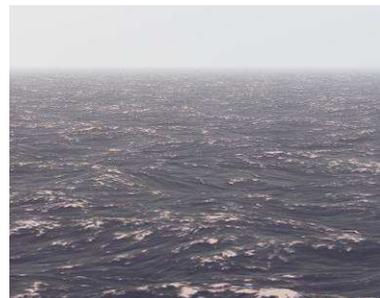


Figure 4 : example of white caps rendering in visible RVB, MWIR and LWIR.

## ATMOSPHERIC PROPAGATION

Propagation of reflected radiances from the surface to the observer accounts for atmospheric attenuation and radiance. Relevant propagation parameters (extinction coefficient and source-functions) are pre-computed, tabulated and stored together with the sky-domes through SKYGEN and MODTRAN.

In infrared wavebands, especially in MWIR, atmospheric absorption exhibits strong variations which must be carefully accounted for while propagating radiances from the sea surface, because of their effect

on band-averaged transmissions. Namely, radiations with distinct spectral distribution undergo different attenuations along a given path. As an example, in the MWIR range (see Figure 5):

- Surface thermal emission follows approximately a continuous blackbody spectrum, with energy radiated regardless of atmospheric absorption lines and bands spectral location, leading to a certain band-averaged attenuation.
- To the contrary, incoming direct solar radiance is filtered by atmospheric absorption before reaching the surface. Thus, while its spectrum is correlated to atmospheric transmission, it undergoes comparatively weak average attenuation.

This is typically the kind of “band-integration” effects a model should reproduce. However, dealing with illuminance and transmittance data with sufficient spectral resolution is clearly not practicable in real time. To overcome this problem, band-averaged transmissions are pre-computed for 4 major spectral contributions to observed radiances: blackbody emission at sea-surface temperature, reflected direct solar illumination, reflected diffuse solar illumination and reflected diffuse thermal illumination. For range R to the surface and a set of sensor heights, transmission is expressed as  $\exp(-k.R^\beta)$ , couples of parameters ( $k, \beta$ ) being computed and stored for each class radiance contribution.

The spectral classification suggested here is somewhat crude, but could easily be extended at nearly no cost (e.g. in order to account for a set of reflectance spectra).

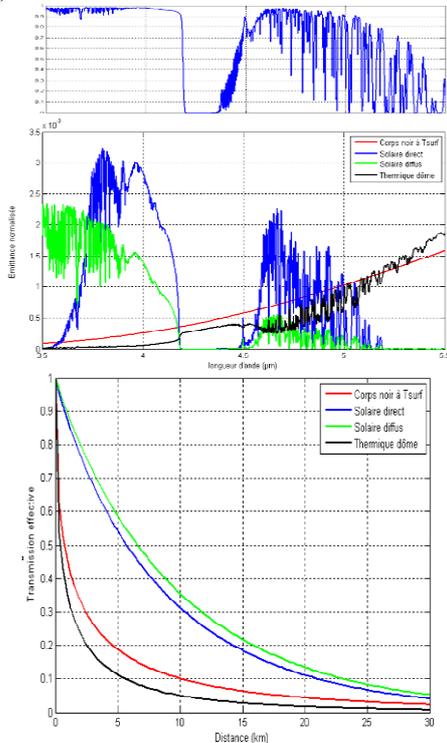


Figure 5 : left: correlation between atmospheric transmittance (above) and contributions to observed radiances (below). Right: corresponding band-averaged transmissions as a function of range

#### IV. EXAMPLE RESULTS

The current application, developed for testing and demonstration, includes a GUI (Figure 6) which allows the user to:

- Select precomputed illumination, sky-domes and waveband data,
- Modify the sensor height, direction and field of view,
- Interactively modify several parameters like wind speed, fetch, air-sea temperature difference, atmospheric attenuation,
- Set numerous rendering options (including multiple rendering for simultaneous visualization of various fields)

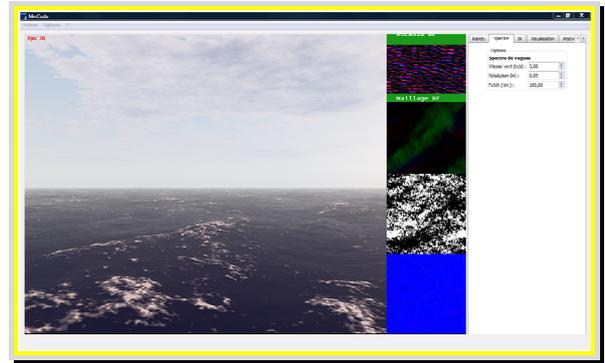


Figure 6 : snapshot of the GUI with the main rendering window, additional rendered fields and the control tabs.

Figure 7 illustrate successive computing and rendering steps in a case of low sun (20° elevation). The last image demonstrate the model ability to account for roughness spatial variability (provoked by currents or slicks).

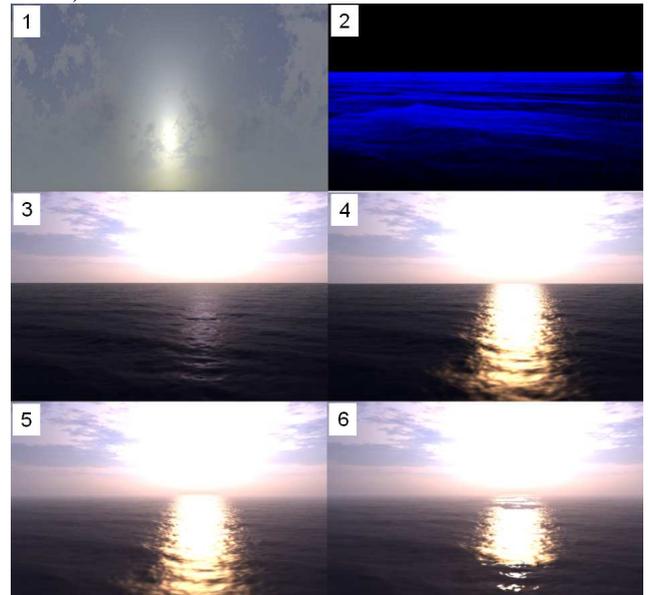


Figure 7 : successive modelling steps (U10= 4 m/s, fetch=100 km). 1: SKYGEN RVB Mercator sky dome with altocumulus layer. 2: multi-resolution mesh with sea surface displacements. Rendering with, 3: sky diffuse contribution only, 4: diffuse + direct solar contribution, 5 : propagation effects, 6: smooth zones (slicks).

Figure 8 illustrates the significance of unresolved roughness effects in VIS (RVB), MWIR and LWIR wavebands. Images in the first column were rendered without any account for sub-grid roughness (rendering #1). In the second column, the roughness loss of resolution in distant areas due to the mesh design is taken into account (rendering #2). Finally, images in the last column also account for additional roughness due to short waves unresolved everywhere (mainly capillary waves, rendering #3).

In visible and without sun glitter (line 1), rendering #1 causes a strong excess of sky reflexion while approaching the horizon, giving the sea an unusual silver-like appearance. This is fairly corrected by rendering #2, rendering #3 bringing limited additional change. In visible again but while observing towards the sun (line 2), the effect of small scale roughness is even more obvious: only rendering #3 succeeds in reproducing a realistic glittering area. In infrared bands (lines 3 and 4) with significant thermal glow from the atmosphere (respectively MWIR from 3.7 to 5.1  $\mu\text{m}$  and LWIR), the improvement brought by rendering #2 over rendering #1 is very strong: rendering #1 artificially increases the reflective contribution of near-horizon thermal glow at grazing angles, leading to underestimated contrast between ocean and sky. Rendering #2 acceptably corrects this trade-off, except in the MWIR sun glitter zone where, as in visible, small scale roughness added by rendering #3 proves critical.

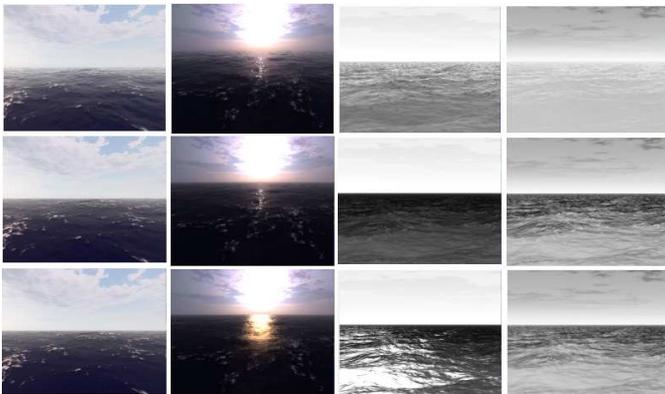


Figure 8 : effect of accounting for unresolved roughness. Lines : 1, RVB, 2, RVB towards the sun ; 3, MWIR; 4, LWIR. Columns: 1, without roughness; 2, with roughness from mesh resolution decreasing only; 3, with additional capillary-waves roughness.

## V. FUTURE WORKS

The model described in this paper is still currently under development, with special efforts in the following topics:

- Improving and extending the sea surface model (output spectrum tuning, advanced white caps modeling, inclusion of swell and measured wave spectra...),
- Improvement of sea surface rendering (parametric model for bulk water scattering in

the visible range),

- Validation against observations and measurements,
- Introduction of boats with their wakes.

In relation to the validation tasks, we plan to take advantage of real-time capacities to develop fast automated tuning methods, based on inversion algorithms. This approach could help finding the best fit of a necessarily imperfect and approximated model to available observations.

In addition, we are working at project of real-time modeling of sea-surface radar signal and imagery, which will benefit from the approach and techniques involved in the optical model presented here.

## VI. ACKNOWLEDGMENTS

We want to thank Dr. Bertrand Chapron (IFREMER Brest / CERSAT) for his precious guidance and communicable passion for sea surface dynamics and radiometry.

## BIBLIOGRAPHY

- [1] K. Caillault, S. Fauqueux, C. Bourlier, P. Simoneau, and L. Labarre: Multiresolution optical characteristics of rough sea surface in the infrared, *Applied Optics*, Vol. 46, Issue 22, pp. 5471-5481 (2007).
- [2] Elfouhaily T., B. Chapron, and K. Katsaros (1997): "A unified directional spectrum for long and short wind-driven waves" *JOURNAL OF GEOPHYSICAL RESEARCH*, VOL. 102, NO. C7
- [3] Lafon, C., Piazzola, J, Forget, P. & Despiau, S. (2005). Modelling of whitecap coverage in coastal environment. 37th International Liège Colloquium on Ocean Dynamics, Gas Transfer at Water Surfaces.
- [4] Sang-Ho Oh, Natsuki Mizutani, Kyung-Duck Suh, Noriaki Hashimoto (2005). Experimental investigation of breaking criteria of deepwater wind waves under strong wind action. *Applied Ocean Research* 27.
- [5] Tessendorf, J. (2001). Simulating ocean water. ACM SIGGRAPH course notes. Updated in 2004.
- [6] Thon, S. and Ghazanfarpour, D. (2002). Ocean waves synthesis and animation using real world information. *Computers and Graphics*, 26(1):99.108.