

A flexible hyperspectral simulation tool for complex littoral environments

Adam Goodenough^a, Rolando Raqueño^a, Michael Bellandi^a, Scott Brown^a, and John Schott^a

^aDigital Imaging and Remote Sensing Laboratory, Rochester Institute of Technology, 1 Lomb Memorial Drive, Rochester, NY, USA 14623

ABSTRACT

This work describes a visualization tool and sensor testbed that can be used for assessing the performance of both instruments and human observers in support of port and harbor security. Simulation and modeling of littoral environments must take into account the complex interplay of incident light distributions, spatially correlated boundary interfaces, bottom-type variation, and the three-dimensional structure of objects in and out of the water. A general methodology for a two-pass Monte Carlo solution called Photon Mapping has been adopted and developed in the context of littoral hydrologic optics. The resulting tool is an end-to-end technique for simulating spectral radiative transfer in natural waters. A modular design allows arbitrary distributions of optical properties, geometries, and incident radiance to be modeled effectively. This tool has been integrated as part of the Digital Imaging and Remote Sensing Image Generation (DIRSIG) model. DIRSIG has an established history in multi and hyperspectral scene simulation of terrain targets ranging from the visible to the thermal infrared (0.380 - 20.0 microns). This tool extends its capabilities to the domain of hydrologic optics and can be used to simulate and develop active/passive sensors that could be deployed on either aerial or underwater platforms. Applications of this model as a visualization tool for underwater sensors or divers are also demonstrated.

Keywords: Hyperspectral remote sensing, simulation and modeling, hydrologic optics, Monte Carlo, photon mapping

1. INTRODUCTION

We introduce a tool for generating synthetic images of complex littoral scenes for port and harbor security applications integrated with an established hyperspectral remote sensing modeling platform called DIRSIG.¹ In contrast to prior work in this field,² we do not make any assumptions about the sensor location and orientation, the type and form of photon sources, or the spatial distribution of photon accumulation. This results in a much higher degree of computation complexity which is mitigated, in part, by the adaption and further development of an optimized Monte Carlo technique called photon mapping.³ With this tool it is possible to simulate a hyperspectral sensor on an aerial platform over a coastal scene, a radiometric instrument measurement off the side of a boat, a diver under water, or any number of potential collection scenarios – all using the same methodology, scene and, under invariant lighting conditions, the same photon map. This type of multi-scale, flexible approach has many possible applications in the realm of port and harbor security, ranging from sensor design and sensitivity studies to algorithm development.

Send correspondence to: aag7210@rit.edu, www.cis.rit.edu

1.1. DIRSIG

The work discussed here encompasses a major upgrade to the capabilities of the DIRSIG (Digital Imaging and Remote Sensing Image Generation) tool¹ in modeling complex littoral scenes and participating media in general. DIRSIG has an established history in hyperspectral scene simulation of terrain targets ranging from the visible to the thermal infrared (0.380 - 20.0 microns). The current work extends its capabilities into the domain of hydrologic optics by implementing the tools and techniques necessary to propagate and collect light in synthetic media. A summary of the basic DIRSIG components is shown in Figure 1 (the specific contributions discussed in this paper are outlined). The components are combined using a spectral (and polarized) representation and integrated radiance images can be simultaneously produced for an arbitrary number of user-defined bandpasses. The model products can be used to test image system designs, to create test imagery for evaluating image exploitation algorithms, and for creating data for training image analysts.⁴

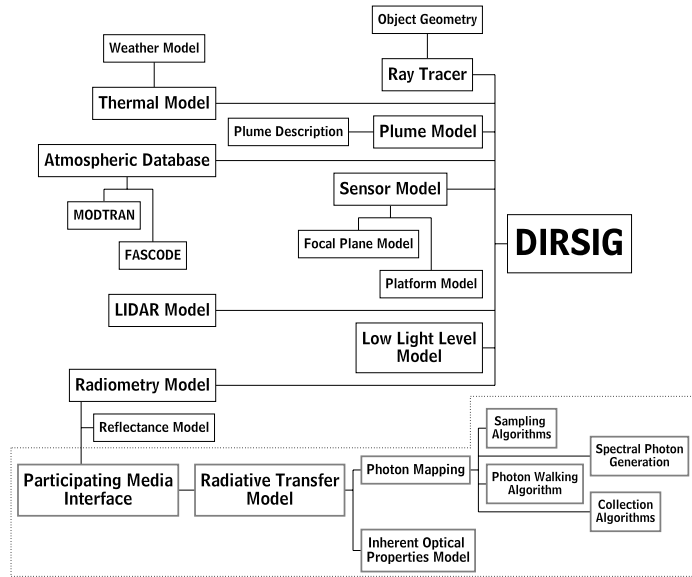


Figure 1. Relationships between DIRSIG components

1.2. Megascene




Figure 2. Selected closeup images from the synthetic Megascene project (left), an overhead view of all five tiles (middle), and a closer view of Tile 5 (right). Measured overflight data has been used to produce the spectral signature of the water

A study site north of Rochester, NY, has been instrumental in understanding various components of the remote sensing imaging chain^{5,6}. An extensive effort was made to develop tools and build a representative model of this area (designated Megascene for the sheer size of the endeavor). These products, along with the results from a comprehensive data collection campaign, provide wide area simulations that support algorithm development. To date, the five tiles highlighted in Figure 2 have been simulated in DIRSIG and represent approximately seven square kilometers of real estate in urban, residential, light industrial, and shoreline regions. The intent is not to generate a simulation that exactly reproduces a specific scene; rather, it is to emulate the scene’s spectral character and statistics acquired by sensors, such as those used for data collection over the region.




Figure 3. Tile 5 Topobathymetry

The physics simulating the spectral imaging phenomenology for the terrestrial regions yielded reasonable fidelity.⁵ Implementation of the techniques described in this paper enables us to approach the same level of simulation fidelity for the coastal region (the so-called Tile 5). This area is being modeled with measured topobathymetry (as shown in Figure 3) and characteristic optical properties. The resulting simulations will be compared to collected overflight data.

2. BACKGROUND

2.1. Spectral RTE

We use a standard form⁷ of the spectral radiative transfer equation (RTE) incorporating gains and losses due to the six factors listed in Table 1 and based on the directed differential radiance at point, \vec{x} , time, t , and direction, ω . The notation, $(\omega \cdot \nabla) \frac{L_\lambda(\vec{x}, t, \omega, \lambda)}{n(\vec{x}, t)^2}$, is used to represent this derivative, where ∇ is the gradient operator and we have explicitly shown the dependence on the refractive index. Using symbols for the inherent optical properties listed in Table 3, we can write the individual contributions as shown in Table 2. The superscript e , i , and s stand for elastic, inelastic, and source (emitted), respectively. Solving for the integrated radiance with all six components and reducing the notation to the single dimension, r , along the ray yields the (simplified) equation,

$$L_*(r_b) = \int_{r_a}^{r_b} \left(\boxed{L_*^e(r)} + L_*^i(r) + L_*^s(r) \right) e^{-\int_{r_a}^{r_b} c(r') dr'} dr + L_*(r_a) e^{-\int_{r_a}^{r_b} c(r) dr}, \quad (1)$$

where “ $*$ ” represents implicit refractive index dependence and the integral is along ray section $r_a \rightarrow r_b$.

This work is primarily concerned with efficiently calculating the elastic scattering corresponding to component iv (emphasized in Equation 1). The problem of computing inelastic scattering contributions, (component v), can be solved with similar methods, but the added complexity is beyond the scope of this paper.

i	Loss	Elastic scattering out of the beam
ii	Loss	Inelastic scattering out of the beam
iii	Loss	True absorption
iv	Gain	Elastic scattering into the beam
v	Gain	Inelastic scattering into the beam
vi	Gain	True Emission into the beam

Table 1. Sources of radiance loss and gain along a beam

i	$(\omega \cdot \nabla) \frac{L_\lambda(\vec{x}, t, \omega)}{n(\vec{x}, t)^2} =$	-	$b^e(\vec{x}, t, \lambda) \frac{L_\lambda(\vec{x}, t, \omega)}{n(\vec{x}, t)^2}$
ii	$(\omega \cdot \nabla) \frac{L_\lambda(\vec{x}, t, \omega)}{n(\vec{x}, t)^2} =$	-	$a_*^i(\vec{x}, t, \lambda) \frac{L_\lambda(\vec{x}, t, \omega)}{n(\vec{x}, t)^2}$
iii	$(\omega \cdot \nabla) \frac{L_\lambda(\vec{x}, t, \omega)}{n(\vec{x}, t)^2} =$	-	$a^e(\vec{x}, t, \lambda) \frac{L_\lambda(\vec{x}, t, \omega)}{n(\vec{x}, t)^2}$
iv	$(\omega \cdot \nabla) \frac{L_\lambda(\vec{x}, t, \omega)}{n(\vec{x}, t)^2} =$	+	$b^e(\vec{x}, t, \lambda) \int_\Omega \tilde{\beta}^e(\vec{x}, t, \omega' \rightarrow \omega, \lambda) L_\lambda(\vec{x}, t, \omega') d\omega'$
v	$(\omega \cdot \nabla) \frac{L_\lambda(\vec{x}, t, \omega)}{n(\vec{x}, t)^2} =$	+	$b^i(\vec{x}, t, \lambda' \rightarrow \lambda) \int_\Omega \int_\Lambda \tilde{\beta}^i(\vec{x}, t, \omega' \rightarrow \omega, \lambda' \rightarrow \lambda) L_{\lambda'}(\vec{x}, t, \omega') d\lambda' d\omega'$
vi	$(\omega \cdot \nabla) \frac{L_\lambda(\vec{x}, t, \omega)}{n(\vec{x}, t)^2} =$	+	$S_o(\vec{x}, t, \omega') \tilde{\beta}^s(\vec{x}, t, \omega, \lambda)$

Table 2. Contributions to the directional derivative of radiance for the RTE.

2.2. Photon Mapping

Equation 1 introduced a form of the RTE that can only be solved numerically except in special, and not particularly useful, cases. Stepwise numerical integration takes the form of “ray marching”; that is, the ray along which we are integrating is broken into segments within which we assume that the local optical properties and light field influences are invariable. These segments can be of arbitrary length and have a center point, r_c , which will be the “sample” point from which measurements are made (as demonstrated in the accompanying figure). In practice, we use an adaptive algorithm to choose segments based on changes in the local volume properties and choose a sample point close to the center point, but with a random offset in order to avoid sampling artifacts. Integration in a highly complex and variable medium requires many ray segments and computations, while relatively homogeneous volumes can be traversed rather quickly.




Figure 4. Stepwise integration along a ray from r_a to r_b

The integrated form of the in-scattered radiance from a segment centered at point r_c and of length Δr_c can be written as

$$L_\lambda \left(r_c + \frac{\Delta r_c}{2}, \omega \right) = \left(b(r_c, \lambda) \int_{\Omega} \tilde{\beta}(r_c, \omega' \rightarrow \omega, \lambda) L_\lambda(r_c, \omega') d\omega' \right) \Delta r_c, \quad (2)$$

where we have assumed that the index of refraction does not change significantly within the segment. Given that the inherent optical properties are known at point r_c (see Section 3.4) at each wavelength λ , the only unknown in the equation is the radiance coming from all directions ω' . Conceptually, the easiest way to calculate $L_\lambda(r_c, \omega')$ would be to use the same backwards ray tracing process that has been used thus far to compute the contribution from the ray. That is, for every direction, ω' , a new ray could be sent out and the radiance contribution could be calculated. Since there are an infinite number of directions in Ω , the integral would either have to be solved using standard numerical techniques or using a Monte Carlo approach. Within a volume, every new ray that would be produced to calculate the integral would also need to do a ray marching integration along its own path.

Diminishing returns eventually limits the effectiveness of sending out numerous new rays and the contributions of subsequent generations could be approximated or neglected. Nonetheless, the number of operations grows approximately exponentially as each new generation of rays attempts to calculate the in-scattered radiance. Even with a moderate number of samples, a limited number of scattering generations, and large integration step sizes, the calculation is infeasible for any practical application – and we have not even discussed the problem of making sure that rays eventually “find” the important sources in a scene.

It is therefore impractical to do purely forward (rays from sources) or purely backwards (rays from the detector) calculations for multiple scattering situations. Instead, we adapt a hybrid approach called photon mapping that computes the integral along the ray using a two-pass technique. The photon mapping approach was developed and popularized by Henrik Wan Jensen³ and integrates existing Monte Carlo techniques, an efficient search algorithm, and the concept of geometry independent photon storage. It has been used to efficiently produce synthetic, ray-traced images for computer graphics applications. The primary contribution of this work is to adapt this method in order to generate spectral (rather than three band) synthetic imagery driven by bio-optical models of water properties and integrate it into the active and passive sensor testbed already described.

The “photon map” is a collection of discrete bundles of energy (the “photons”) that are organized in a way that conveniently expresses the spatial relationships between energy bundles (the “map”). Within a particular scene, the photon map is a static entity which contains information about energy in the volume. When ray-tracing a scene, scattering contributions are measured using local density estimates obtained by querying the photon map. These calculations are very efficient, especially compared to Monte Carlo sampling techniques for evaluating in-scattered radiance.

The first pass of the photon mapping technique consists of generating photons at sources and propagating them through the scene. The propagation process is exactly equivalent to established forward Monte Carlo techniques²⁷. As each photon interacts with the volume (through scattering or absorption), a record of the photon is stored in the map (its position, direction, flux, and wavelength – note that the wavelength cannot be derived from the flux since these are not true photons, only bundles of energy). The photons are stored in a k-d tree,⁸ independent from any geometry or abstract concept of layers or voxels. The k-d tree is essentially a binary tree for higher-dimensional data and allows for efficient searching in the next pass. Since the photon map is only dependent on sources and scene geometry, it is possible to re-use the same photon map for different sensor locations and orientations.

The second pass consists of backward ray tracing from the detector where the contribution from a ray in a volume is calculated using the ray marching technique discussed previously. Rather than attempting to calculate the in-scattered radiance directly, we now use the photon map to estimate the in-scattered radiance. Given that we know the direction, ω_i of each photon in the map when it was stored, we can apply the local scattering phase function to estimate the amount of flux that is scattering in the direction of the integral. Thus, we transform Equation 2 into,³

$$L_\lambda \left(r_c + \frac{\Delta r_c}{2}, \omega \right) \approx \left(\frac{1}{V_{search}} \sum_{i=1}^{k_\lambda} \Phi_i \tilde{\beta}(r_c, \omega_i \rightarrow \omega, \lambda) \right) \Delta r_c, \quad (3)$$

where k_λ is the number of photons found, V_{search} is the volume of the search, and the scattering coefficient has been incorporated into the density estimate. The search volume is usually a sphere that encompasses all of the photons that were found by a search by radius or number. We will modify this method somewhat in Section 3.5 in order to minimize the errors inherent in practical density estimation.

3. APPROACH

3.1. Sampling Algorithms

Any Monte Carlo technique can only be successful if the code that drives it produces samples that are representative of the underlying probability density functions in the model. We divide the process of generating uniform samples into two steps. During the first step we generate pseudo-random samples from a uniform distribution using a hybrid approach that combines stratified and Latin hypercube sampling⁹ in order to guarantee a high level of uniformity for any number of samples (see Figure 5).

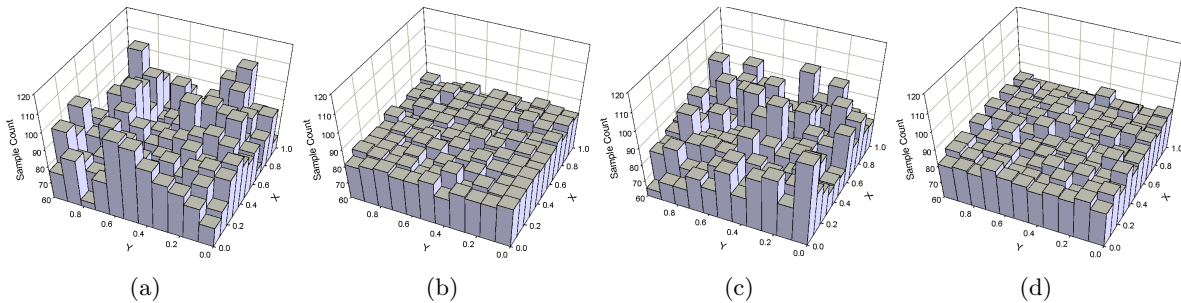


Figure 5. Example histograms of two-dimensional samples in the unit square generated from (a) random (by a standard linear congruential method), (b) stratified, (c) Latin hypercube, and (d) hybrid sampling. Note that while no significant improvement over stratified sampling can be seen, the hybrid approach gains beneficial uniformity properties from the inclusion of Latin hypercube constraints that are not apparent in these histograms.

Using a uniform sampling of the unit square as a basis, we now project the samples onto a 2-manifold using an area preserving algorithm that correctly distributes the samples onto the new surface. The projection operators are derived using a process compiled from the literature on the subject¹⁰ that consists of finding a parametrization of the manifold, linking the parametrization to the surface area and applying an optional weighting function,

defining and inverting cumulative distribution functions, and forming a new parametrization that takes the unit square samples as input. The result is a set of samples on the manifold that have the same uniform distribution properties as the original samples.

In practice, we use a wide variety of projections to sample anything from arbitrary constructed geometry (using triangular facet projections) to hemispheres that are weighted by a bi-directional reflectance distribution function (BRDF). In the next two sections we state, without derivation, two of the simpler sampling projection algorithms that are relevant to atmospheric photon generation.

3.1.1. Sphere section sampler projection




Figure 6.

A sphere section projection is used to sample a sky dome quadrant by generating a vector that points to a location within that quadrant for a sphere centered at $[0, 0, 0]$. Angles θ_1 , θ_2 , ϕ_1 , and ϕ_2 define the range of zenith and azimuth angles in the quadrant, respectively (in the equation, $\mu = \cos(\theta)$). The projected samples are computed from the two-dimensional uniform random deviates (ξ_1, ξ_2) as

$$\psi(\xi_1, \xi_2) = \begin{bmatrix} \sin(\cos^{-1}(\mu_1(1 - \xi_1) + \mu_2\xi_1)) \cos((\phi_2 - \phi_1) \cdot \xi_2 + \phi_1) \\ \sin(\cos^{-1}(\mu_1(1 - \xi_1) + \mu_2\xi_1)) \sin((\phi_2 - \phi_1) \cdot \xi_2 + \phi_1) \\ \mu_1(1 - \xi_1) + \mu_2\xi_1 \end{bmatrix} \quad (4)$$

3.1.2. Disk sampler projection




Figure 7.

A disk projection is used to sample a solar or lunar disk with radius, ρ . It is simpler to derive the projection on the x - y plane for a uniform disk and then transform the samples to the proper position defined by the solar/lunar vector and solid angle as retrieved from appropriate ephemeris data for the scene. The projected samples are computed as

$$\begin{aligned} \psi_1(\xi_1, \xi_2) &= \rho\sqrt{\xi_1} \cdot \cos(2\pi\xi_2), \\ \psi_2(\xi_1, \xi_2) &= \rho\sqrt{\xi_1} \cdot \sin(2\pi\xi_2) \end{aligned} \quad (5)$$

3.1.3. One-dimensional importance sampling

In addition to two-dimensional projections, we often need a method to randomly select a particular element from a set when each element has a weight. Given an arbitrary number of elements, e_i , with associated weights, $W(e_i)$, the weights are converted to probabilities, $P(e_i)$, by dividing by the sum of all of the weights (① and ②),

$$P(e_i) = \frac{W(e_i)}{\sum_i W(e_i)}. \quad (6)$$

The probability elements are ordered so that the largest probability comes first in the probability vector (③), which helps optimize the next step when the probabilities vary. In step ④, a uniformly distributed random number $r \in [0, 1]$ is pulled from a generator and the value is iteratively compared – this is why we put the larger probabilities first – to the cumulative probability of the vector elements. The sampled element is the one corresponding to the location of the random number in the vector (element E in the example).

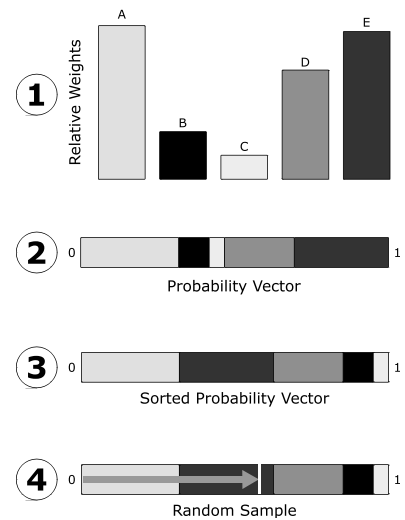


Figure 8.

3.2. Photon Generation

For the purposes of this discussion, the source of photons entering the scene is the sky (both the sky dome and the sun/moon). Nonetheless, the same sampling techniques could be applied to any type of source and used for low-light level or active sensor modeling.

The sky dome (which is a source of down-welled radiance) is broken into quads, each of which is assumed to be homogeneous (i.e. the radiance coming from any point within a quad is the same as any other point in the same quad). Quads are implemented via a sphere section sampler and the weight is equal to the integrated flux coming from the quad (using the area of the current section). All of the quads have equal area and, therefore, define equal solid angles. The sun and moon disks are constructed using a disk sampler and oriented according to ephemeris tables for the current date and time. The weight of the solar/lunar disk(s) is the total integrated flux coming from the solid angle defined by the disk. Figure 9 demonstrates how the atmosphere is divided. The zenith, θ , and azimuth, ϕ , angles are shown.

The entire process of generating photons is summarized in the list below.




Figure 9. Sampling of atmospheric illumination

Atmospheric Photon Generation

- ❶ Initialize a count of “shot photons” to zero
- ❷ Initialize the photon map with the pre-determined number of photons to be stored *The number of photons defined in the preceding two steps are independent from each other*
- ❸ For each photon, until the photon map is filled, perform the following steps:
 - ❶ Uniformly sample spatially within a pre-determined horizontal extent of the scene that defines the **PropagationArea**
 - ❷ Record the sampled point, **HorizontalSample**
 - ❸ Use one-dimensional importance sampling to select an element of the atmosphere (a sky quad or a solar/lunar disk) *Weighting is based on the amount of flux produced by each element*
 - ❹ Sample spatially within the element using a sphere section or disk sampler (see Sections 3.1.1 and 3.1.2, respectively) and record the sample position (**AtmosphereSample**)
 - ❺ Calculate the initial direction of the photon from point to point, **PhotonDirection = AtmosphereSample** → **HorizontalSample**
 - ❻ Use one-dimensional importance sampling to select the wavelength of the photon. *Weighting is based on the amount of flux contributed by each bandpass for the flux associated with the atmosphere element*
 - ❼ Propagate the photon through the scene and store it in the photon map as it is either absorbed or scattered
 - ❽ Increment the count of shot photons *The photon count must be incremented regardless of whether the photon is stored*
- ❹ Calculate the flux associated with each photon by dividing the total flux passing through the **PropagationArea** by the number of photons that were “shot” *This is not the number of photons stored in the map(s)*

3.3. Surface Modeling

In addition to general scene geometry considerations, we must also create an air/water interface that is representative of a wind-driven wave surface. The primary approach to generating such surfaces via spectral wave synthesis and the Fresnel theory that covers propagation across the interface has already been introduced in detail to this field.² We have used an analogous approach to generate height fields from spectral wave models (a surface generated by a JONSWAP¹¹ spectrum is shown in Figure 10). We are also investigating the use of windowed Fourier transforms and hybrid wave train models to allow for spatially changing the spectral properties of the wave surface as waves are propagated into shallower regions. We have also implemented a microfacet-based bi-directional reflectance/transmittance distribution model¹² to handle cases where the spatial wave frequency is higher than that resolvable by the modeled sensor.




Figure 10.

3.4. IOP Models

Definition of the inherent optical properties of the water (as listed in Table 3) has been made as flexible as possible to allow for complicated constituent models and spatial distributions. We take the approach of defining the IOPs by starting with a base medium (eg. pure water using a hybrid of data for absorption^{13,14}, scattering,¹⁵ a roughly Rayleigh form scattering phase function,⁷ and the IAPWS standard for the refractive index¹⁶). Additional constituent properties are then added to the base in linear combination using a variety of common parameterized model forms.⁷ Accordingly, the linear models for the absorption coefficient, scattering coefficient, and phase function are given by:

Symbol	Units	Description
$a(\vec{x}, t, \lambda)$	$\left[\frac{1}{m}\right]$	absorption coefficient
$b(\vec{x}, t, \lambda)$	$\left[\frac{1}{m}\right]$	scattering coefficient
$c(\vec{x}, t, \lambda)$	$\left[\frac{1}{m}\right]$	attenuation coefficient
$\tilde{\beta}(\vec{x}, t, \omega, \omega' \rightarrow \omega, \lambda)$	$\left[\frac{1}{sr}\right]$	scattering phase function
$n(\vec{x}, t, \lambda)$	<i>dimensionless</i>	index of refraction

Table 3. Summary of inherent optical properties

by:

$$a(\vec{x}, t, \lambda) = a_0(\vec{x}, t, \lambda) + \sum_{i=1}^{N_a} a_i(\vec{x}, t, \lambda), \quad (7)$$

$$b(\vec{x}, t, \lambda) = b_0(\vec{x}, t, \lambda) + \sum_{i=1}^{N_b} b_i(\vec{x}, t, \lambda), \quad (8)$$

$$\tilde{\beta}(\vec{x}, t, \omega, \omega' \rightarrow \omega, \lambda) = w_0 \cdot \tilde{\beta}_0(\vec{x}, t, \omega, \omega' \rightarrow \omega, \lambda) + \sum_{i=1}^{N_b} w_i \cdot \tilde{\beta}_i(\vec{x}, t, \omega, \omega' \rightarrow \omega, \lambda), \quad (9)$$

where we have made the dependence on position and time explicit and N is equal to the number of additional components. The decomposition of the scattering phase function required weighting the component scattering phase functions by the effective contribution of the corresponding scattering coefficient to the whole. Each weight, w_i , is given by:

$$w_i = \frac{b_i(\vec{x}, t, \lambda)}{b(\vec{x}, t, \lambda)}, \quad 0 \leq i \leq N_b. \quad (10)$$

Note that we do not include the index of refraction in the linear combination models. We will assume that a single refractive index model is used, which is usually defined by the base medium, but can be overridden by the user if necessary.

3.5. Collection Algorithms

As the number of photons being searched for (or equivalently, the search radius) is increased, we increase the risk of finding photons uncharacteristic of the local light field. We want to grow the search so we can drive down the variance inherent in the photon map due to the random, sampled nature of construction (despite using very uniform sampling strategies). The objective of the collection algorithms that follow is to find an efficient means of extending the search volume (to reduce noise) while maintaining the assumption of locality. There are many plausible and potentially more accurate approaches to compensating for the two sources of error presented here (boundary and spectral). Our goal in developing these methods is to make the process as efficient as possible while still compensating for the bulk of the error.

3.5.1. Boundary collection algorithm (Boundary map)

By far, the largest errors are likely to occur near boundaries of the volume where the search radius can extend beyond the bounds of the volume. This is also the region where, at least for turbid media, the most scattering will occur. Since the photon map has no internal concept of boundaries, it will treat anything beyond a boundary as part of the local volume. In most cases, this ends up causing the density estimate to be biased lower than it should be due to the effective inclusion of a photon “void” in the density estimate. To compensate for this error, we construct a “boundary map” (analogous to a photon map) composed of nodes containing the distance to the boundaries of the volume at that particular point. These nodes are computed and stored during the first pass of the algorithm. The boundary map is constructed using a sparsely populated k-d tree which can be searched to estimate the location of boundaries from the local boundary nodes. Once these distances are known, an effective volume can be calculated. Within a complex scene it is possible for boundaries to be arbitrarily oriented, however, we assume that it is sufficient to know the boundary distances in the six axial directions and to be conservative in our estimates of the constraints.

We will consider the volume of a sphere that has been constrained in any combination of six specified directions. These six directions correspond to the positive and negative orthogonal axial directions centered at the sphere origin (the search site). The axes would presumably be aligned with the global coordinate system though this is not necessary. If the original search sphere does not hit a boundary in a particular direction then the constraint is zero and the search sphere was able to expand to its full radius in that direction (whether or not it found nodes at that distance). Each constraint is also less than the search radius, R , since the search site must be within the boundary. This gives each constraint a possible range of $[0, R)$.

In order to calculate the effective search volume, the volumes outside of the boundaries (the occluded volumes) must be removed from the volume of the unconstrained sphere. In the simplified case of a single constraint, the effective volume can be found by removing a spherical cap corresponding to the constraint. When there is more than one constraint and the spherical caps intersect we need to account for removing the intersected volumes (a conceptual rendering of six intersecting constraints is shown). Given the level of symmetry in the types of intersections possible we can pre-compute the intersecting volumes at selected intervals in order to optimize the volume calculation.



Figure 11.

3.5.2. Spectral collection algorithm

Accounting for spectral photons for multi/hyperspectral modeling introduces an effective bias error incurred by searching for enough photons in each spectral band to drive down the spectral noise. Under ideal conditions, we expect that the total variance in the estimate is linearly related to the number of photons that are used to form the estimate. Accordingly, under optimal circumstances where the photons are uniformly distributed across all of the spectral bands, we expect that the spectral variance will increase by a factor of s , where s is the number of spectral bands. In other words, to maintain the same level of variance as we would have without considering spectral photons, we need to search for $s \cdot N$ rather than N . This makes us more susceptible to errors as we violate the locality assumption. The approach we take to compensate for this source of error is to perform a two-step search.

The first step is to perform a local core search by radius across all spectral bands. The radius is set so that the assumption of a local volume is reasonably met. The density that results from this search is used to set the *mean* density of the final spectral density (no spectral shape exists at the moment). Thus, we ensure that the final integrated density will be consistent with the local light field intensity.

The second step performs an expanded spectral search within each spectral band. The number of spectral photons to find is set such that we ensure a certain level of fidelity. We also calculate the relative weight of each spectral density and apply this weight to the core density calculated in the previous step (a variety of methods can be used here). The spectral density that results is our final estimate.

4. VALIDATION

For any model of this size and complexity it is necessary to consider both the theoretical and procedural validity of the model. We have adopted, and are currently carrying out, an intensive validation process for the model consisting of five phases of increasing complexity. The validation process is summarized in Table 4.

Phase	Type	Description
I	Code	Extensive numerical algorithm testing, step-by-step ray tracing, and optimization
II	Intuitive	Design and rendering of a series of simple test scenes which can be judged based on visual inspection
III	IOP \leftrightarrow AOP	Examines apparent optical properties for which an approximate relationship to IOPs is known ^{17,18}
IV	Peer-to-peer	Comparison against Hydrolight for fully modeled illumination, IOP, and surface models ¹⁹
V	Sub-space	Examines the correspondence between synthetic invariant target sub-spaces ²⁰ and measured targets in the Megascene Tile 5 region

Table 4. Validation phases

5. APPLICATIONS

The potential applications of this tool are multiple and varied and include:

- Passive broadband/multispectral/hyperspectral sensor design and sensitivity studies
- Radiometric instrument design and testing
- Evaluation of existing sensors under various environmental conditions to determine operational limits and optimum conditions for data acquisition
- Studies of sensor design tradeoffs with respect to cost per sensor, coverage, and data quality
- Thermal modeling (using DIRSIG's built-in thermal model and weather histories)
- Surface, partially submerged, and fully submerged target detection algorithm design and evaluation
- Spatially variant bio-luminescence studies
- LIDAR simulation (using DIRSIG's built-in LIDAR model along with photon mapping capabilities)
- Low-light level simulation using above-water or submerged illumination sources
- Geometric obscuration studies coupled with temporal illumination study
- Studies on the impact of adjacency effects on target detection

- Multi-scale collection scenario design and mission planning
- As a sensor fusion testbed where a simulated host of sensors across the EM spectrum provide a baseline of image data sets to best present unique imaging phenomenologies to a user or algorithm
- Background signature characterization
- Invariant sub-space construction for target detection algorithms²⁰

6. SUMMARY

The photon-mapping technique and its extension into the hyperspectral domain shows great potential in addressing basic hydrologic optical phenomenology in the littoral region including embedded targets and platform/instrument obscuration of the scene. Prior to this implementation, approaches that could simulate littoral scenes relied on assumptions that limited the model's capabilities and usability. The stochastic nature of a Monte Carlo solution coupled with the extreme scales between scene and sensor renders classic techniques prohibitively computationally demanding when considering scattering media. Photon mapping optimizes the computation of in-scattered radiance through an efficient photon searching and radiance estimation and brings such simulation tools closer to operational viability for flexible, multi-scale usage.

The integration of these techniques into the DIRSIG environment makes it possible to address many issues in port and harbor security by using an extensive collection of sensor, environmental, geometric, and radiative transfer models and tools. Extensive validation of both overhead, surface, and underwater simulations against both modeled and measured data will be used to verify subjective phenomenology as well as radiometric accuracy and fidelity. Finally, simulation of the the Megascene Tile 5 littoral zone will provide a large scale project for which expansive ground truth data has been collected.

ACKNOWLEDGMENTS

Support for this work has been provided by the Department of Defense Office of Naval Research under a Multidisciplinary University Research Initiative (MURI) (N00014-01-1-0867): Model-based Hyperspectral Exploitation Algorithm Development.

REFERENCES

1. E. Ientilucci and S. Brown, "Advances in wide area hyperspectral image simulation," in *Proceedings of the SPIE conference on Targets and Backgrounds IX: Characterization and Representation*, **5075**, pp. 110–121, (Orlando, FL), April 2003.
2. C. R. Bostater, "Hyperspectral simulation and recovery of submerged targets in turbid waters," in *Proceedings of the SPIE*, **5780**.
3. H. W. Jensen, *Realistic Image Synthesis Using Photon Mapping*, A K Peters, 2001.
4. Digital Imaging and Remote Sensing Laboratory, *DIRSIG User's Manual*, February 2006.
5. N. G. Raqueño, L. E. Smith, D. W. Messinger, C. Salvaggio, R. V. Raqueño, and J. R. Schott, "Megacollect 2004: Hyperspectral collection experiment of terre strial targets and backgrounds of the RIT Megascene and surrounding area (Rochester, New York)," in *Proceedings of the SPIE conference on Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XI*, (Orlando, FL), March 2005.
6. R. Raqueño, N. Raqueño, A. Vodacek, J. Schott, A. Weidemann, S. Effler, M. Perkins, W. Philpot, and M. Kim, "Megacollect 2004: Hyperspectral collection experiment over the waters of the rochester embayment," in *Proceedings of the SPIE conference on Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XI*, (Orlando, FL), March 2005.
7. C. Mobley, *Light and Water: Radiative Transfer in Natural Waters*, Academic Press, Boston, MA, 1994.
8. J. L. Bentley, "Multidimensional search trees used for associative searching," *Communications of the ACM* **18**, pp. 509–517, September 1975.
9. Chiu, Shirley, and Wang, *Multi-jittered Sampling*. Graphics Gems IV, Academic Press, 1993.

10. J. Arvo, "Stratified sampling of 2-manifolds," *Monte Carlo Ray Tracing, SIGGRAPH 2003 Course Notes*, pp. 39–61, 2003.
11. D. E. Hasselmann, M. Dunckel, and J. A. Ewing, "Directional wave spectra observed during JONSWAP 1973," *Journal of Physical Oceanography*, pp. 1264–1280, 1980.
12. M. Ashikhmin, S. Premoze, and P. Shirley, "A microfacet-based BRDF generator," in *Siggraph 2000, Computer Graphics Proceedings*, K. Akeley, ed., pp. 65–74, ACM Press / ACM SIGGRAPH / Addison Wesley Longman, 2000.
13. R. M. Pope and E. S. Fry, "Absorption spectrum (380-700 nm) of pure water. ii. integrating cavity measurements," *Appl. Opt.* **36**, pp. 8710–8723, 1997.
14. R. C. Smith and K. Baker, "Optical properties of the clearest natural waters," *Applied Optics* **20**(2), pp. 177–184, 1981.
15. H. Buiteveld, J. M. H. Hakvoort, and M. Donze, "The optical properties of pure water," *SPIE Proceedings on Ocean Optics XII* **2258**, pp. 174–183, 1994.
16. International Association for the Properties of Water and Steam, *Release on the Refractive Index of Ordinary Water Substance as a Function of Wavelength, Temperature and Pressure*, September 1997.
17. H. R. Gordon, O. B. Brown, and M. M. Jacobs, "Computed relationships between the inherent and apparent optical properties of a flat homogeneous ocean," *Applied Optics* **14**, pp. 417–427, 1975.
18. J. T. O. Kirk, *The relationship between inherent and optical properties*. Ocean Optics, Oxford University Press, New York, 1994.
19. C. D. Mobley and L. K. Sundman, *Hydrolight 4.1 Technical Documentation*. Redmond, WA 98052, Aug. 2000.
20. G. Healey and D. Slater, "Models and methods for automated material identification in hyperspectral imagery acquired under unknown illumination and atmospheric conditions," *IEEE Trans. Geosci. Remote Sensing* **97**(6), pp. 2706–2717, 1999.