Lidar Imaging Through the Ocean Surface

René CARMONA[*], Frederic CEROU[†‡] and Benjamin BLASCO[§]
Department of Operations Research & Financial Engineering,
Princeton University, Princeton, NJ 08544

June 30, 2000

Abstract

This work is concerned with the detection and the identification of immersed objects. Our efforts have been motivated by, and concentrated on, the specific problem of the LIDAR imagery of underwater mines. The difficulties of this problem are illustrated by the experimental data used in this study.

We follow the standard approach which suggests to model the surface elevation as a random field. We combine the statistical analysis needed needed for the model of the roughness of the surface of the ocean with the physics of image forming which include focusing/defocusing, scattering and backscattering of the light beams, and finally the sensor noise, into a large simulation program. The latter is used to produce realistic images and video sequences of moored mines as seen from airborne platforms, making possible the analysis of the statistics of the image degradation process.

1 Introduction

The detection and the removal of sea-mines are of great importance to the Navy. In recent years, the clearance of sea-mines has also gained nationwide support and international attention due to its implications to marine life protection. Motivated by the identification and detection of sea-mines, we propose a complete analysis of the *direct* model which explains the distortion present in gated LIDAR images of immersed objects. We develop systematic simulation tools which take into account the many physical mechanisms involved in the degradation of the images as LIDAR beams propagate through the air and

[*]Partially supported by ONR grant N00178-99-1-9003 and N00178-00-1-9001
[†]Supported by ONR grant N00178-00-1-9001
[‡]Permanent address: IRISA-INRIA, Project Sigma2, Campus de Beaulieu, 35042 Rennes cedex, France
[§]Intern from Ecole Polytechnique, France

the water, and especially through the air-water interface. Our simulation methodology is part of a more general approach to underwater imaging, which has been of particular interest to researchers in fields such as marine science and military defense and we believe that our results will be useful in both military and civil applications in rescue and recovery.

In order to generate random samples for the time evolution of the ocean surface, we use classical models from physical oceanography and standard tools for Monte Carlo simulations. The process taking the expected image (think for example of a large metallic sphere) into the images produced by the LIDAR imaging systems is highly nonlinear, and even in the best case scenarios (see for example Figure 1) very different from the standard linear models of blurring and additive noise used in classical image restoration. This degradation process has an unpredictable component, which needs to be modeled statistically: fortunately, the roughness of the ocean surface can naturally be modeled as the realization of a homogeneous random field with statistics described in spectral terms. As expected, another difficulty comes from the fact that the images suffer from sensor noise, but unexpectedly, this source of distortion is not the main obstacle to a successful analysis. In this study we model this noise as a Poisson shot (quantum) noise. But besides the sensor noise, the main sources of distortion of the images are twofold: First, the strong spatial modulations caused by wave focusing and defocusing (in the present study, we concentrate on the effects of the capillary waves, essentially ignoring the contribution of the gravity waves); Second, the clutter produced by the lack of contrast between the rays reflected by the target and the rays backscattered by the ocean medium. Understanding the image degradation process is only the first part of our research program, and it is clear that the difficult challenge remains: inverting this process. This is addressed in [4] where we report on our results on the detection of the mines of the ML(DC) data set.

The paper is organized as follow. The two data sets used in this study are introduced in the next section. The following Section 3 is devoted to the simulation of random samples of the surface elevation, while Section 4 is concerned with the actual simulation of ray tracing image forming process. We describe the separate components in sequence: the laser, the receiver, the light propagation at the air/water interface, the small angle diffusion approximation and the propagation in the water, the backscattering and finally, the sensor shot noise. Experimental results are reported in Section 5 and the paper ends with a short conclusion explaining some of the shortcomings of our system, and some directions of future research.

Most of the images reproduced in this report are still frames from animations which can be found at the URL `http://chelsea.princeton.edu/~rcarmona`.

2 Description of the Data

This section is devoted to the discussion of the main features of the experimental data, as they were made available to us.

A First Data Set

We first describe the data set, which was at the origin of this study. This first data set consists of images of a fake mine taken during an experiment near Panama City in the Gulf of Mexico. These images were taken with a CCD camera mounted on a crane approximately 50 ft above the ocean surface, the images being taken at angles with the vertical varying from 0^{deg} to 60^{deg}. For the purpose of this experiment, the proxy for the mine was a metallic spherical object with radius 6 in, and it was kept between 1 to 10 m below the surface. Figure 1 shows a couple of generic images from this experiment.

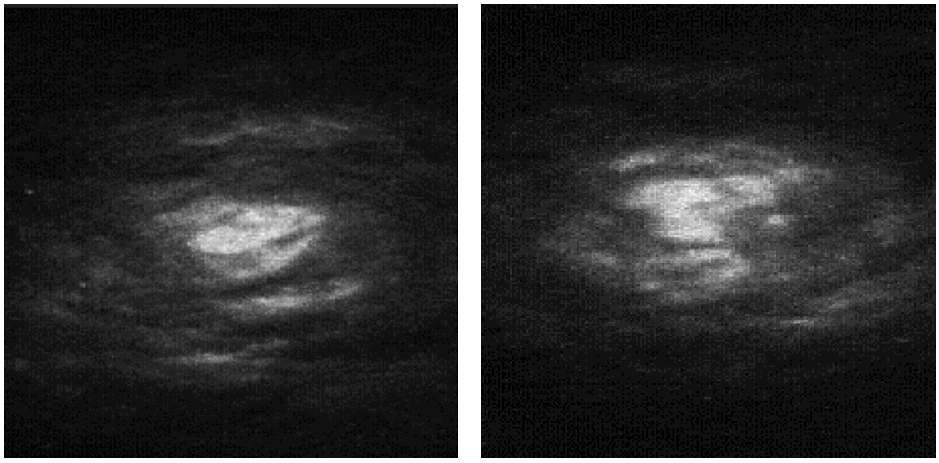

Figure 1: Two typical images from the first data set. The distortion phenomenon described in the introduction is obvious: the spherical object is broken into bright patches.

The variety of images contained in this data set demonstrates the great sensitivity of LIDAR returns to the magnitude of wave slopes as well as to the incident angle of the sensing rays. But it is fair to emphasize that the water was relatively calm during the two days of the experiment.

A Second Data Set

The second data set is from the June 1995 ML(DC) Flight Verification Test conducted in the Gulf of Mexico, off the coast of Panama City. Figure 2 gives a typical example of the image files contained in this data set. Because of the pattern leading to the 5th image from the left, it is possible that a foreign object (possibly a mine) was present when

these images were collected. Unfortunately, direct information on the possible presence of foreign objects in some of these images is not part of the package. In can only be inferred indirectly.

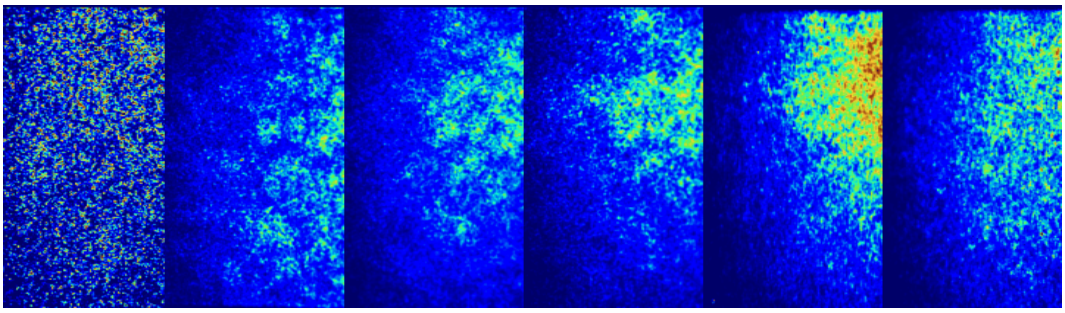

Figure 2: Image returns from 6 different gates (the depth of the gate increases when we go from the left to the right.)

The Magic Lantern Development Contingency (ML(DC)) system is a ranged-gated light detection and ranging (LIDAR) system, consisting of three main optical subsystems: the scanner, the laser and a set of six imaging receivers. The six imaging receivers are locked on six overlapping gates, the depths of these gates being reported in Table 2 below. The images produced by the six receivers are arranged in strips of six images, the depth of the gate increasing from left to right. In this way, each strip represents a water column over a rectangle at the surface of the ocean.

Camera#	Receiver Gated Settings
1	12-32 ft
2	16-36 ft
3	20-40 ft
4	25-40 ft
5	35-45 ft
6	40-50 ft

Table 1: Depth Gate Settings for each Camera for ML(DC)

The interested reader is referred to the report [6] for further details on the ML(DC) system. The simulation tools presented in this work were used in a further study devoted to the actual detection of the mines [4].

Comparing Figure 1 and Figure 2, it is obvious that the look of the images from this second data set is very different from the look of the images from the first data set. One of our most difficult challenges is to make sure that our simulation system be able to reproduce realistically both types of images.

3 Simulation of the Surface

As explained in the introduction, our model has a random component. Indeed, instead of entering the physics of the ocean surface via partial differential equations, we choose to model the surface elevation as a random field. This random component is the source as the random realizations of the surface vary, for the statistics of the images. But for the purpose of the present study, once a set of realizations of the surface is generated, it will not be changed.

3.1 Surface Elevation as a Homogeneous Random Field

The first step in the simulation process is to produce realistic realizations of the evolution over time of the ocean surface. Once this is done, we should be able to generate LIDAR images of a target using the successive realizations of the ocean surface, allowing the creation of video sequences of distorted images.

We follow the time-honored approach to ocean surface elevation statistical modeling presented in most (if not all) the text books in physical oceanography. See for example [15] or [21]. In this approach, the deviation of the ocean surface elevation about its mean level is assumed to be a stationary in time and homogeneous in space mean zero Gaussian field specified by its 3-dimensional spectrum. Samples of the ocean surface are simulated using the spectral representation theory. Since we are mostly interested in the effects of capillary waves, we shall ignore that part of the spectrum responsible for the gravity waves. Figure 3 shows a typical frame of the animations of the surface produced by our program.

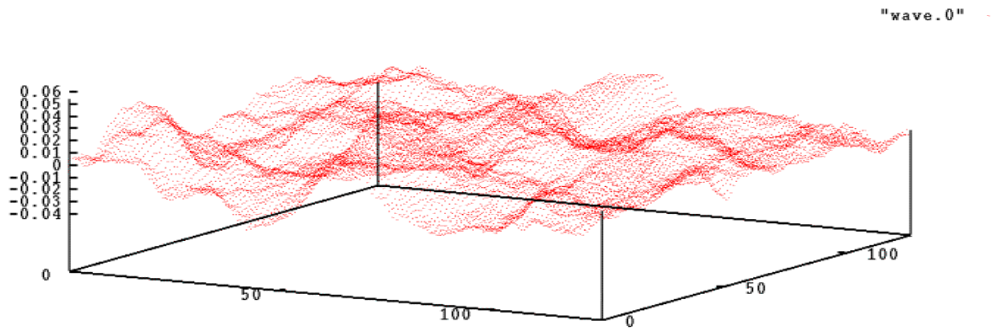

Figure 3: Snapshot of the ocean surface elevation simulated as a time dependent homogeneous Gaussian random field.

The spectral models evolved over time: in 1958, Phillips suggested a spectral model

for gravity waves which had been a mainstay in physical oceanography. Modifications have been made to the Phillip's model. Pierson and Moskowitz incorporated the low-frequency cutoff of a saturated sea. Toba, in 1973, suggested a different analytical form for the gravity wave frequency spectrum. In our simulation we will mainly use the Phillips spectral model. A complete discussion can be found in [15] and [21]. Moreover, further details on these spectra including a discussion of the sampling and random generation issues relative to the spectral representation theorem can be found in [9]. For the sake of completeness, we also mention an alternative approach based on the concept of auto-regression [19].

3.2 Conditional Simulation

We now discuss an interesting twist that came about following a series of experiments, which took place at the ONR Duck facility last October. A proof of concept showed that it was possible to measure the elevation of the ocean surface at a finite number of points, simultaneously and with a great precision. Whether or not it will be possible to implement this information in real time measurements and simulations is still not clear at this stage. But, the accessibility of these measures suggests that, instead of simulating the ocean surface as a homogeneous field with a given spectrum, it should be preferable to simulate a field according to the conditional distribution given the results of these discrete measurements. If we try to include the knowledge of the surface elevation at a finite set of locations, the homogeneity property of the field breaks down, and consequently, the notion of spectrum looses its significance, and the efficient simulation algorithms based on the spectral representation of homogeneous fields cannot be used. In other words, the inclusion of the extra knowledge in the model comes at the price of the loss of the simulation convenience. Attempts have been made (especially in geosciences), and algorithms have been developed to handle conditional simulation. See for example [11], [17], [18] or [8]. Unfortunately the algorithms are only approximations (the most famous one being based on the kriging method) and the computational burden is still prohibitive if one wants this approximation to be reasonable.

We investigated in [3] the possibility to introduce ideas familiar in the simulation of Gibbs random fields. Indeed, since the Gibbs specifications are given in terms of conditional distributions, it appears natural to rely on them in the present situation.

4 The Elements of a First Imaging System

From now on, we assume that simulations of the time evolution of the ocean surface have been performed, and we proceed to the description of the generation of LIDAR images of a target using the successive realizations of the ocean surface, creating in this way, distorted images of the mine. We need a precise analytical model for the target, the laser,

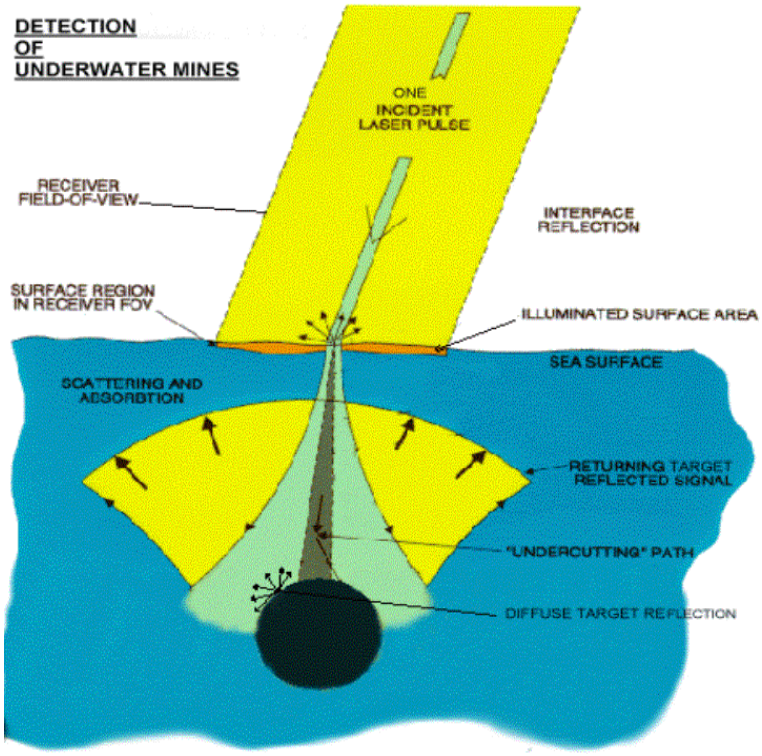

Figure 4: Schematic diagram of our first ray-tracing imaging system.

and the receiver, especially if we intend to include the the imaging platform motion (which will be a helicopter most likely) in the simulations. Throughout this section, we present the various components our ray-tracing simulation program. Computational speed, as well as the possibility to test various models, are also important issues we have to worry about in the development of an accurate real-time direct imaging model.

We first develop a basic ray tracing image forming system based on random surface samples and the laws of Snell and Fresnel from geometric optics. We check that this first simulation tool is capable of reproducing the images and videos of the first data set. But in order to produce images and videos with the features of the files of the second data set, we will need to improve our model. Two new ingredients will have to be included. First, we will add backscattering effects: the photons travelling in the water are scattered by the water particles, and an ensemble of complex interactions produce an overall effective backscattering which can be incorporated in our model. Second, because of the very design of the CCD receiver, and because of the light transport properties, the experimental images and videos are subject to distortion by a photon quantum shot noise which can also be added to our program. These two facts will be crucial in the implementations of the improvements discussed in the second part of this section.

4.1 Properties of the Laser

For modeling purposes we use a divergent laser, in other words, we suppose that the laser beam goes through a lens. We will explain in the next paragraph how the laser is projected onto the water surface. But we need to simulate the real properties of the laser beam after the lens. We use a simulation method from [12]. The initial beam distribution is assumed to be of the form:

$$I_0(P) = Je^{-r^2/a^2}$$

where J is an intensity factor (for simplicity J is assumed to be 1 unless otherwise stated), r is the distance between the point P under consideration and its projection on the central ray of the laser, and a is the 'beam width parameter'. The point P belongs to the region of the water surface illuminated by the divergent laser. From this we can compute the intensity of light received at each point of the surface. Since we do not have enough information about the laser, so we choose a in order to have:

$$0 \leq \frac{r}{a} \leq 1$$

which seems to be a good approximation (see [14]). Using this initial beam distribution produce a good fit to the experimental data.

4.2 Description of the Receiver

In this study we use a ray model for temporal pulse analysis. Additionally we reverse the ray-tracing algorithm, which usually proceeds from target to camera. This is in ac-

cordance with the reversibility principle of linear optics, as derived from Fermat's principle. There are significant computational savings in neglecting those rays which after reflection from the target, do not enter the camera's acceptance cone or aperture. In the second data set there are six receivers: they are adjusted to a common boresight but each camera is gated to view different depth slices in the water column, with small overlaps between camera-gated regions. The imaging receivers are Intensified Charge Coupled Device (ICCD) cameras. These cameras are gated to interact with the laser pulse at a specific time. Through the use of a short gate width time and a spectral bandpass filter at 532 nm, each receiver is optimized to amplify the pulsed laser return, and minimize the continuous ambient light. A gated receiver only observes the reflected light, which traveled a distance corresponding to the gate delay time (for our simulation we took 6 m for the "gating".) This "gating" technique is used to minimize (if not completely avoid) surface returns artifacts. Our simulation system also allows imaging the surface and possibly objects immersed under the surface while the helicopter is in motion. See Figures 14 and 15 below. As a motivation for this exercise, we can imagine that the helicopter first flies at low altitude to detect the mines, and then gains altitude to shoot the mine. Our model of the receiver has to be flexible enough to describe all the possible airborne platform elevations. The optical system is assumed to consist of an idealized pinhole camera, and bi-static LIDAR, whose optics is not diffraction-limited. The Lambertian detector exhibits an $N \times N$ - pixel format, addressable in terms of Cartesian coordinates.

Field of View

The present notation and computations are illustrated in the diagram given in Figure 5.

The field of view refers to the angle opening α from the camera onto the part $L1$ of the plane of the average surface elevation imaged by the laser swath. We use the orthogonal projection given by the normal to this average surface to measure the elevation h of the camera. We compute the linear length $L1$ as a function of the elevation h, the opening α and the angle between the vertical and the middle incident ray direction θ. We first notice that:

$$\alpha = \tan^{-1}\left(\frac{L1 - x}{h}\right) + \tan^{-1}\left(\frac{x}{h}\right)$$

or equivalently:

$$\tan\alpha = \frac{\frac{L1}{h}}{1 - \frac{L1-x}{h}\frac{x}{h}},$$

recall that $x < 0$ in the configuration reproduced in Figure 5. Since:

$$\tan\theta = \frac{L1/2 - x}{h}$$

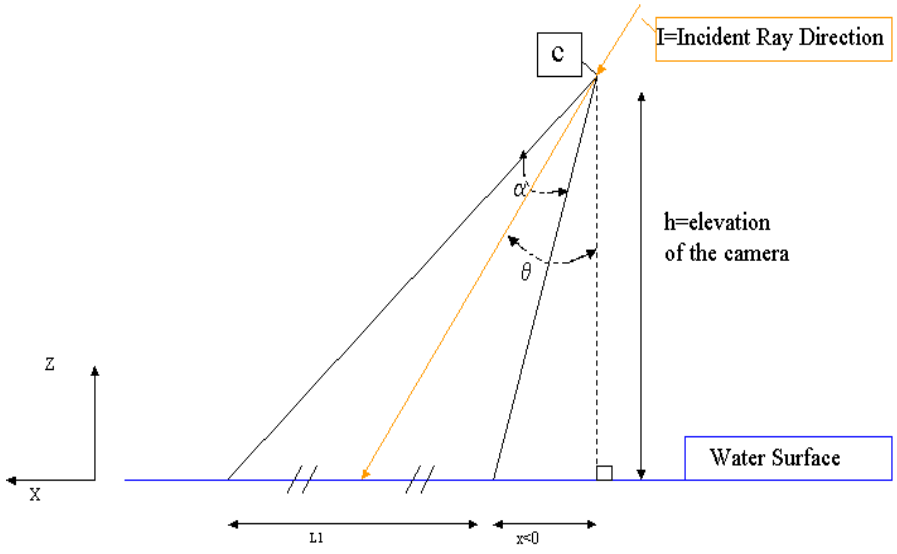

Figure 5: Schematic diagram to illustrate the notation of the field of view discussion.

we can solve for $L1$ and get:

$$L1 = 2\left(-\frac{h}{\tan(\alpha)} + h\sqrt{1 + \frac{1}{\tan^2\alpha} + \tan^2\theta}\right).$$

We perform the same computation in the y-axis direction, using the same opening α, and we compute the corresponding length $L2$, and we keep only the greatest of the two. Our computation has to be refined if the receiver's field of view α is not the same as the laser field of view. In our simulation, we use the same value $\alpha = 8\,\mathrm{deg}$ for both angles. This value should not be too far from the value used in experiments. Next we need to choose the rendering of the surface of the water. We choose a fixed set of grid points at the surface and we only consider the rays (downwelling and upwelling), which go through these points. In fact we assign each cell of the camera to a single point of the water surface grid. In this way, we do not need to project the image at the water level on the plane of the camera sensors.

The camera model

We now review the form of the projection technique from a pinhole camera. This is needed when we consider that the laser is a source of parallel rays. See [10]. A perspective transformation projects 3-d points onto a 2-d plane. It provides an approximation of the way in which a plane image represents a 3-d world. Let us assume that the camera coordinate system (x, y, z) has the image plane as its (x, y) plane, and let the optical axis (established by the center of the lens) be along the z-axis. To be specific we assume that the center of the image plane is at the origin, and that the center of the lens is at the point with coordinates $(0, 0, \lambda)$. If the camera is focused on the far field, λ is the focal length of the lens. Next we assume that the camera coordinate system is aligned with the world coordinate system (X, Y, Z). The coordinates (x, y) of the projection of the point (X, Y, Z) onto the image plane are given by the formulae:

$$x = \lambda\frac{X}{\lambda - Z}, \qquad\qquad y = \lambda\frac{Y}{\lambda - Z}.$$

Remark 4.1 *The homogenous coordinates in the projective space of a point with Cartesian coordinates (X, Y, Z) are defined as (kX, kY, kZ, k), where k is an arbitrary, nonzero constant. Clearly, conversion of homogenous coordinates back to Cartesian coordinates is accomplished by dividing the first three homogenous coordinates by the fourth.*

We define the perspective transformation matrix P as:

$$P = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & -1/\lambda & 1 \end{bmatrix}.$$

The corresponding basic mathematical model is based on the assumption that the camera and the world coordinate systems coincide. We want to consider the more general situation of a world coordinate system (X, Y, Z) used to identify both the camera and 3-d points in a camera coordinate system (x, y, z). The new assumption is that the camera is mounted on gimbals, allowing pan of an angle θ, and tilt of an angle α. Here, pan is the angle between the x and X axes, and tilt is the angle between the z and Z axes.

$$\text{Pan angle:} \quad < Ox, OX >= \theta, \qquad\qquad \text{Tilt angle:} \quad < Oz, OZ >= \alpha.$$

The goal is to bring the camera and the world coordinate systems into alignment by applying a set of linear transformations. After doing so, we simply apply the perspective transformation to obtain the image-plane coordinates for any world point. Translation of the origin of the world coordinate system to the location of the image plane center is accomplished by using the transformation matrix:

$$G = \begin{bmatrix} 1 & 0 & 0 & -X_0 \\ 0 & 1 & 0 & -Y_0 \\ 0 & 0 & 1 & -Z_0 \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

In order to pan the x-axis to the desired direction, we rotate it by θ. We proceed similarly for the tilt angle α, except for the change in the axis of rotation. The composition of the two rotation is given by the single matrix:

$$R = \begin{bmatrix} \cos\theta & \sin\theta & 0 & 0 \\ -\sin\theta\cos\alpha & \cos\theta\cos\alpha & \sin\alpha & 0 \\ \sin\theta\sin\alpha & -\cos\theta\sin\alpha & \cos\alpha & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

So if $\mathcal{C}$ denotes the Cartesian coordinates and $\mathcal{W}$ the world coordinates, we have:

$$\mathcal{C} = PRG\mathcal{W}.$$

We obtain the Cartesian coordinates (x, y) of the imaged point by dividing the first and second components of $\mathcal{C}$ by its fourth. We obtain:

$$\begin{aligned} x &= \lambda \frac{(X - X_0)\cos\theta + (Y - Y_0)\sin\theta}{-(X - X_0)\sin\theta\sin\alpha + (Y - Y_0)\cos\theta\sin\alpha - (Z - Z_0)\cos\alpha + \lambda} \\ y &= \lambda \frac{-(X - X_0)\sin\theta\cos\alpha + (Y - Y_0)\cos\theta\cos\alpha + (Z - Z_0)\sin\alpha}{-(X - X_0)\sin\theta\sin\alpha + (Y - Y_0)\cos\theta\sin\alpha - (Z - Z_0)\cos\alpha + \lambda}. \end{aligned}$$

This camera model is very flexible: we just need to know the coordinates of the helicopter/camera, and the incident ray direction to have all the information we need.

4.3 Light Propagation through the Air/Water Interface

Refraction at the air-sea interface results in focusing and defocusing of both downwelling and upwelling radiations. This effect is commonly observed in every day life, a typical instance being the spatial patterns formed at the bottom of a swimming pool on a sunny day. For upwelling radiation, the image of a submerged object is strongly distorted by the refractive artifacts produced by of the waves.

In this work, we do not consider scattering in the air: we assume that the rays arriving at, or leaving the air/water interface, travel along straight lines. Also, for this part of the research, we assume that each time t, we have a way to compute the surface elevation and the normal to this surface at each point of the surface above a given grid. We assume that the illumination of the surface is such that one ray arrives at each point of the surface grid. Reflection and refraction at a surface can be specular or diffuse. In our case, the process is specular: the energy of the incident ray is divided solely between reflected and refracted daughter rays.

Snell's Law

As explained in [15], the properties of reflected and refracted rays from a level surface are summarized in Figure 6.

Following Mobley's book's convention, ξ' represents the unit vector along the direction of the incident ray, while ξ_r and ξ_t denote unit vectors in the directions of the reflected and transmitted (refracted) rays respectively. The unit normal vector to the interface is denoted by n. These notation are illustrated in Figure 6.

The relevant formulae are reported in Table 2 below. The first row of equations follows from Snell's Law, where $n_w \approx 1.34$ is the refractive index of water. Subsequent rows present equations governing properties of the reflected and refracted daughter rays, specifying the directional vector as well as the incident angles. These equations can be easily derived using Snell's law.

AIR-INCIDENT RAYS	WATER-INCIDENT RAYS
$\sin\theta' = n_w \sin\theta$	$n_w \sin\theta' = \sin\theta$
$\xi_r = \xi' - 2(\xi' \cdot n)n$	$\xi_r = \xi' - 2(\xi' \cdot n)n$
$\xi_t = (\xi' - cn)/n_w$ where $c = \xi' \cdot n - \sqrt{(\xi' \cdot n)^2 + n_w^2 - 1}$	$\xi_t = n_w\xi' - cn$ where $c = n_w(\xi' \cdot n) - \sqrt{n_w(\xi' \cdot n)^2 - n_w^2 + 1}$
$\theta_r = \theta' = \cos^{-1}(\xi' \cdot n)$	$\theta_r = \theta' = \cos^{-1}(\xi' \cdot n)$
$\theta_t = \sin^{-1}(\sin\theta'/n_w)$	$\theta_t = \sin^{-1}(\sin\theta'/n_w)$

Table 2: Snell's law formulae

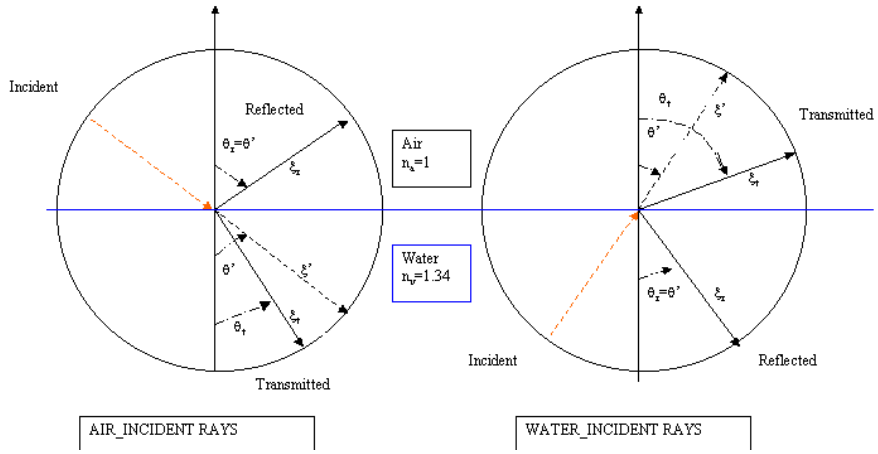

Figure 6: Schematic diagram to illustrate the discussion of Snell's law.

Fresnel's Reflectance

Assuming perfect specular returns, the Fresnel's reflectance r is the fraction of the incident ray's energy retained by the reflected daughter ray. If we let Φ denote the radiant energy carried by a single ray, the following equations hold for the radiant energies Φ_r and Φ_t of the reflected and transmitted rays:

$$\Phi_r = r\Phi \qquad \text{and} \qquad \Phi_t = t\Phi$$

if $t = 1 - r$ denotes the transmittance, i.e. the fraction of energy passed to the transmitted daughter ray (recall that the Fresnel's reflectance r is a number between 0 and 1.) The reflectance r is a function of the incident angle θ'. This function is given, for both the air-incident and the water-incident cases, by Fresnel's formula:

$$\begin{aligned} r(\theta') &= \frac{1}{2}\left(\frac{\sin^2(\theta' - \theta_t)}{\sin^2(\theta' + \theta_t)} + \frac{\tan^2(\theta' - \theta_t)}{\tan^2(\theta' + \theta_t)}\right) \\ r(0) &= \left(\frac{n_w - 1}{n_w + 1}\right)^2 \end{aligned}$$

As the incident angle θ' increases from $0\,\mathrm{deg}$ to $90\,\mathrm{deg}$, Fresnel's reflectance becomes more significant until it eventually reaches the maximum value of 1. The difference in

refractive indices leads to a higher reflectance for water-incident rays as compared to air- incident rays. In other words, it is much easier for light to enter than to leave the ocean. Fresnel's reflectance equals to 1 for a water-incident ray at or beyond the critical incident angle of $48\,\mathrm{deg}$. This phenomenon of total internal reflection greatly influences the passage of radiant energy across an air-water. When this critical point is reached, no energy from the ray can be transmitted from water to air.

4.4 Small Angle Diffusion Approximation in the Water

Like many authors, see for example [13] or [14], we use the small angle diffusion approximation introduced by Wells in [22].

4.4.1 The scattering Explanation

Not only is scattering a complex phenomenon, but its effect on imaging also depends greatly on the biological and chemical conditions of the ocean. While scattering and absorption due to natural water are by and large well understood, distortions due to underwater organisms and other particles are extremely hard to quantify. Refractive effects often appear as spatial feature dissociations (breakup), while multiple scattering tends to produce resolution degradations and contrast reduction in the received image. Analysis of the propagation of light through water is complicated by multiple scattering and by the strong forward directivity of each scattering. Only relatively simple imaging systems operating over short water paths have been analyzed successfully using the conventional optical oceanographic parameters (the attenuation coefficient, the absorption coefficient and the volume scattering function). This is an active research topic in the field of oceanography, and has led to the development of various forms of scattering theories. The small angle scattering approximation is one of them. It is of great importance in many imaging applications. In this approximation, one circumvents these practical and theoretical difficulties in obtaining analytic solutions to imaging problems, by using a technique well known in the analysis of linear electrical and optical systems. It involves the characterization of the system (optics and water paths) in terms of its response to an impulse of scene reflectance. The output can be expressed as a convolution of the input reflectance distribution with the system impulse response function. Alternatively, these operations may be performed in the Fourier domain. The system spread function is computed as the product of the transmitter spread function and the receiver spread function. Each of these spread functions describes the spreading produced by the optics on a one-way water path. Such spread functions have been used to predict successfully the performance of several types of imaging systems over long water paths. We define the two essential water scattering functions: the point spread function PSF and the beam-spread function.

The Beam Spread Function (BSF for short) $BSF(\theta, \Phi, R)$ is used for the normalized irradiance distribution on a spherical surface of radius R centered on the transmitter. The

beam B travels from a source S at the origin to its focal point at the range R. Irradiance is measured on a small area of the spherical surface at range R. See Figure 7 for details.

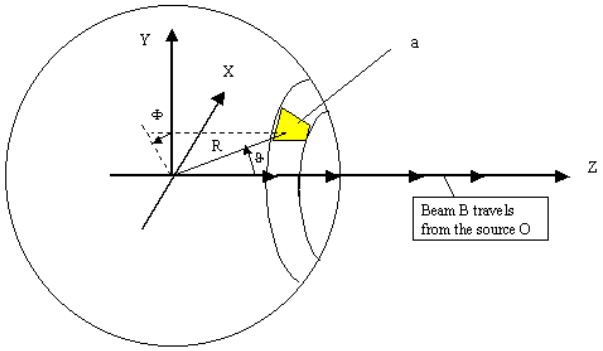

Figure 7: Beam Spread Function Diagram.

The Point Spread Function (PSF for short) $PSF(\theta, \Phi, R)$ is used for the apparent normalized radiance of an unresolved Lambertian source at the position $(0, 0, R)$. The receiver is located at the origin and it points to a small area a of the spherical surface at range R. In other words, $PSF(\theta, \Phi, R)$ is the apparent radiance in the direction (θ, Φ) resulting from scattering in the medium of radiance from an unresolved Lambertian source at $(0, 0, R)$.

In the practical applications we are interested in, the angle Φ is very small, and it will be convenient to redefine $BSF(\theta, R)$ and $PSF(\theta, R)$ as the irradiance and radiance values present in planes tangent to the spheres at $(0, 0, R)$ for small θ's.

Scattering Computations

Because of the symmetric roles played by the PSF and the BSF, we shall only need to compute the PSF (see Wells' fundamental paper [22] for a discussion of the reciprocity between the PSF and the BSF.) Wells also suggested a simple algebraic form for the scattering phase function $s(\theta)$, expression which can be analytically integrated, thus allowing for a straightforward computation of the PSF. The system can also be described in terms of the Mean Transfer Function (MTF for sort) which is obtained from the PSF by Fourier transform, the latter being implemented with the Fast Fourier Transform since we are

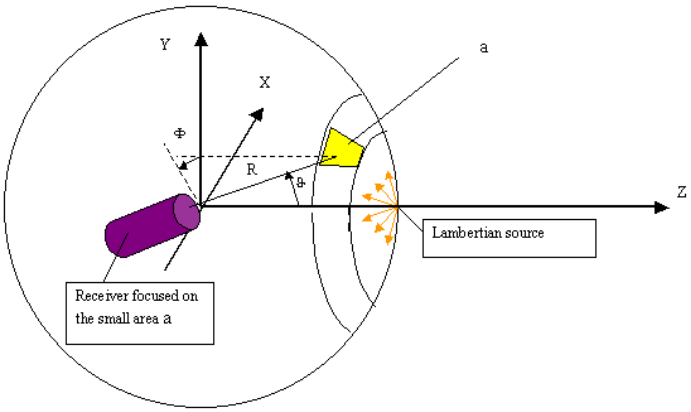

Figure 8: Point Spread Function Diagram.

working on a finite grid. Accordingly:

$$MTF(\psi) = e^{-ct} e^{b_f z D(\psi)},$$

where we follow Wells' notation with the exception that the range dependence z is factored out of the MTF exponent $D(\psi)$, so that D depends only on the scattering phase ψ. This exponent is given by the formula:

$$D(\psi) = \frac{1}{\psi} \int_0^{\psi} \Sigma(\psi')\, d\psi'$$

where:

$$\Sigma(\psi) = 2\pi \int_0^{2\pi} s(\theta) J_0(2\pi\theta\psi)\theta\, d\theta.$$

In our simulation system, we use the following alternative algebraic form of the scattering function:

$$s(\theta) = \frac{s_0}{\theta^{3/2}\sqrt{\theta_0^2 + \theta^2}}$$

for $0 \leq \theta \leq \pi$, the constant s_0 being chosen so that $D(0) = \Sigma(0) = 1$, yielding:

$$MTF(0) = e^{-(c-b_f)z} = e^{-K_d z}$$

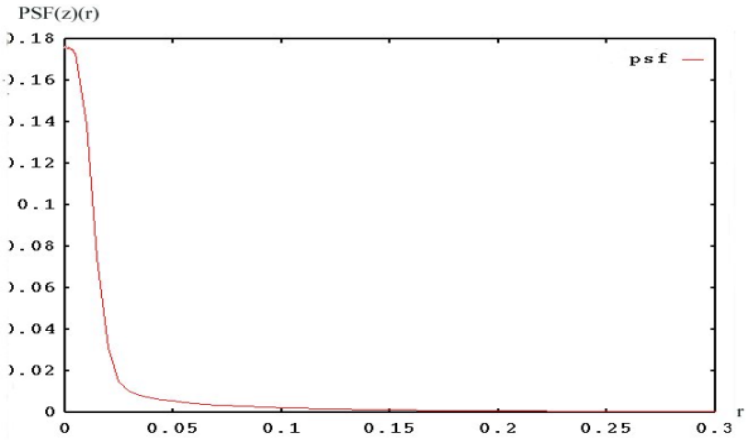

Figure 9: Graph of the PSF used in our simulation system.

for diffuse attenuation, where c is the beam attenuation coefficient, b_f is the forward scattering coefficient, and K_d is the diffuse attenuation coefficient. $\theta_0 = 0.12$. See [13] for details.

In summary, the radiant energy is distributed according to the PSF computed from the small angle scattering approximation. It describes the vertical attenuation and the radial spreading of the ray energy as it propagates through the water. Summing over all rays yields the downwelling irradiance distribution at a given depth. This downwelling irradiance distribution is multiplied pointwise by the objet reflectivity or water backscatter to obtain the upwelling radiance distribution. For the receiver considered here, this is equivalent to the downward propagation of the transmitted rays (see Wells for a discussion of the reciprocity of the PSF). It is customary to truncate the PSF in order to speed up the computations. To be specific, viewing the PSF as a probability density function, we truncate it at its 90 percentile. As it can be seen from Figure 9, this approximation will not change significantly the numerical results because of the fast decay of the PSF.

Empirical Evidence

The small angle approximation to the radiative transport equation is used extensively in imaging models where the transport medium is thick. Its popularity stems from a sound physical rationale and the computational convenience that it offers. It is generally considered valid when the particles in suspension are not very large compared to the wavelength, the refractive index ratio is close to one, and the optical thickness of the medium is not

too large. But the limits beyond which this approximation becomes unreasonable are not well understood. We use as a guideline the recent empirical study [20] which is devoted to the quantification of the limits of this approximation validity.

4.5 Target Shape/Reflectivity Adjustment

Our basic assumption is that mines are spheres whose reflectance are diffuse: this means that all the incident energy coming into a small area of the surface of a mine, multiplied by the mine reflectance is equal to the overall reflected energy leaving this area. Moreover, we assume that the reflected energy is the same in all the directions. We shall come back later to the discussion of the target reflectance.

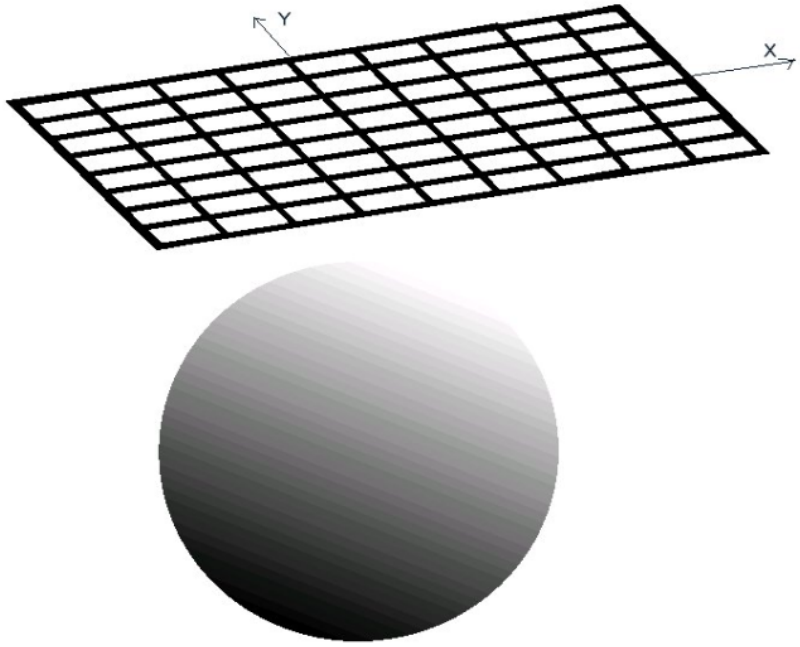

Figure 10: Projection of the spherical target on a rectangular grid.

In order to make the computation easier, we project the sphere on a rectangular grid (see Figure 10.) We call this grid the target plane. The matrix of the grid points coordinates is necessary huge. This is unfortunate because it will be needed later when we include the backscattering effects in our simulation. The sharpness of the definition is free: this is a common technique in computer vision. On the other hand, the discretization is regular. We will study the incoming light on each square plaquette of the grid. But in order to give to this plane the appearance of a sphere, we need to give to each square element of the plane grid a weight, which is function of the part of the surface of

the sphere, which corresponds to this square plaquette. In order to do so, we perform the projection as indicated in Figure 11:

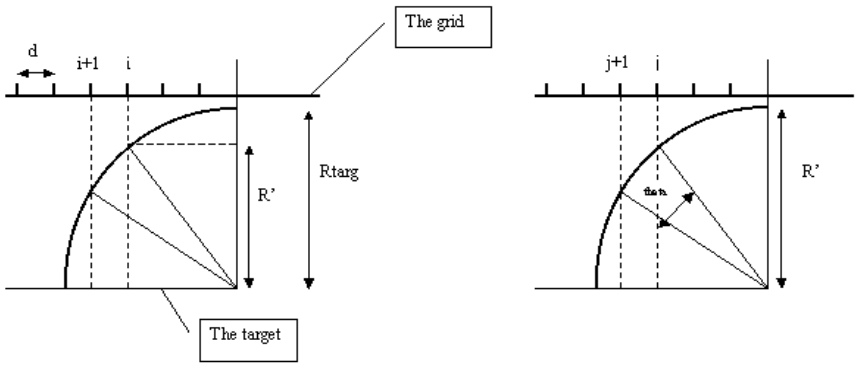

Figure 11: 2-d details of the projection procedure.

As before we work with an orthogonal set of axes. Working first in the x-direction we get the angle θ and the length R'. Doing the same with the other axis, we get R'' and the angle ψ. Then we are able to compute the area A, which is the projection of the square $[i, i+1] \times [j, j+1]$ on the sphere. We will call this area ponder.

The ponder will be multiplied by the target reflectance. Since all the incident energy (coming into a small area on the surface of the 3-d target) multiplied by the target reflectance is equal to all the reflected energy leaving this area, which is the same in all the directions, our plane representation takes into account the real area of each element of our matrix.

$$\begin{aligned} A = \text{ponder}(i,j) &\approx (\theta R')(\psi R'') \\ &\approx R'R'' \left| \sin^{-1} \frac{(i+1)d}{R''} - \sin^{-1} \frac{jd}{R''} \right| \left| \sin^{-1} \frac{(j+1)d}{R'} - \sin^{-1} \frac{jd}{R'} \right| \end{aligned}$$

Using this weighting technique, our effective target does look like a 3-d sphere even though it is only a 2-d plane disk. To complete the illusion we need to modify the rules governing the reflection of the incident rays. Consequently, if the incident ray direction is not vertical (vertical rays are the only ones reflected in the same way whether the target is

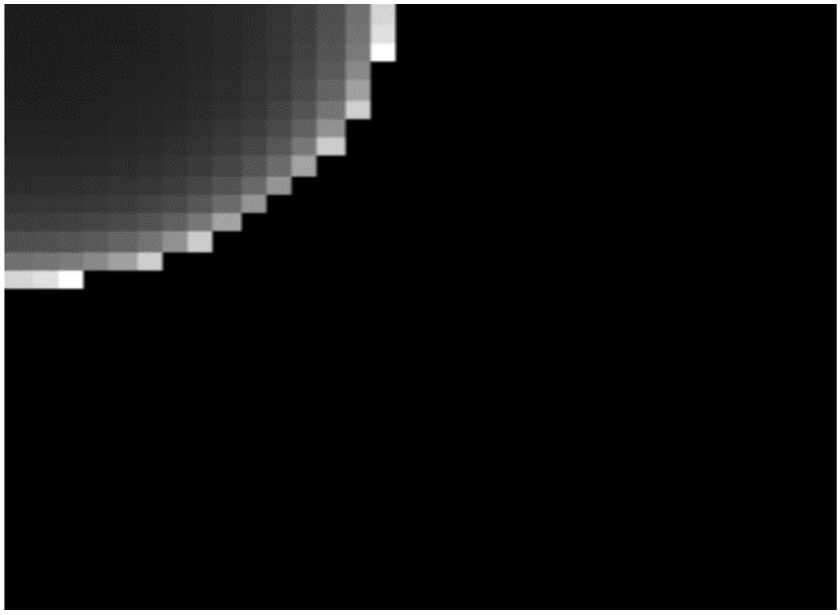

Figure 12: Ponder function for a quarter of a circular mine.

a sphere or a horizontal plane), we change artificially the target plane to fool the ray into believing that it hit a spherical target. A simple 3-d-geometric transformation can be used to do the trick. The illusion is satisfactory and the computations are still much lighter than those needed to consider a real 3-d spherical model for the target.

Remark 4.2 *At this stage it is useful to make sure that there is no ambiguity in the terminology used throughout this report. We call incident ray direction the direction of the central ray of the laser. When this ray is refracted at the air/water interface, we call it the refracted incident ray direction. It gives us the average direction of the downwelling rays in the water, and their intersection with the target plane gives us the average position of the center of the spot of light at this level. The quality of this approximation depends heavily on the specific realization of the water surface, and unfortunately, it sometimes gives wrong results. Despite these occasional problems, it will be good enough for our purposes. The alternative was to consider the incident ray direction after being refracted by a quiet water surface. This alternative was considered inadequate because of the presence of big gravity waves, which shift horizontally the location of the light spot in the target plane level.*

Our first simulation system was based on the considerations presented so far. It was quite adequate for the simulation of images and video sequences of the type of the first data set. The first two examples presented in the section on experimental results in Section 5, were generated with this version of our simulator. The last two sections describe

improvements needed to make our system capable of producing realistic simulations of experiments such as those leading to the second data set.

4.6 Backscattering in the Gated Region

We now explain how we included backscattering effects in our imaging system. Since we compute only once the PSF for each image, we cannot use the model of [13] and [14], which would require too much computing time. Instead we choose to use an effective reflectance R:

$$R = \frac{s_b}{2K}(1 - e^{-2K_d \Delta Z_g})$$

on the target plane, where $s_b \approx \pi s(\pi)$ and s is the scattering function defined earlier, K_d is as before the diffuse attenuation coefficient, and ΔZ_g is the width of the gate.

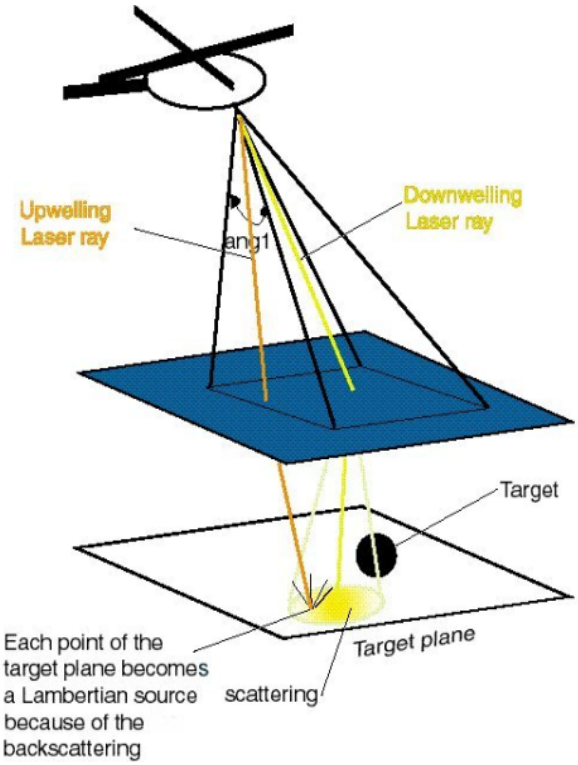

Figure 13: Diagram explaining the computation of the effective reflectance.

In this way, the reflectance coefficient R depends only upon the gating (i.e. the permissible travel time of the rays from the source and to the sensor) and on the characteristics of the water turbidity. In words, we combine all the backscattering in the gated region into a cumulative effect, which is incorporated in the model as the effective reflectance of the plane at the bottom of the gate. We add this reflectance to the reflectance of the

target, and we apply the same tools as earlier: we use the PSF to describe the scattering and the diffusion for each ray, even those that are far from the target: backscattering is everywhere. We compute first the incoming light on the target plane (i.e. the plane at the bottom of the gate), and then we multiply the incoming luminance of each element of our discretization of the target plane by the reflectance distribution we just computed. We can then produce an image of the target plane and the target itself seen from the water just above this plane. We use again the PSF to compute the upwelling rays, and we take into account the refraction at the water-air interface as before. Finally we obtained the desired (distorted) image.

Notice that the equation giving the effective reflectance due to backscattering does not depend upon the actual travel time of each ray but that it only depends upon the thickness of the gate. It does not enter into the computation of the PSF.

4.7 The Sensor Poisson Shot Noise

The receiver comprises CCD sensors. The emission of photons from any source is a random process and the number of photogenerated carriers collected in a potential well in a given time interval of length T, is a random variable. The standard deviation of this random variable characterizes the (quantum) photon noise. Since the number of photons emitted follows a Poisson distribution, the standard deviation of the number of charges collected equals the square root of the mean and thus $N_{\text{noise}} = \sqrt{N_s}$. This noise source is of course a fundamental limitation in most applications since it is a property of light itself, not the image sensor. This limitation tends to become serious only under conditions of low scene contrast, as it is often the case in imaging at low light levels and like in our case. We have:

$$N_s = \frac{HSAT}{q}$$

where H is the image irradiance in W/m^2, S is the effective responsivity in A/W (the typical for an ideal receiver being $238\ mA/W$), A is the geometric area of each cell, and q the charge of the electron. In order to get a feel for this irradiance, we reproduce in Table 3 the figures computed for some astronomical objects:

	Irradiance (in W/m^2)
Starlight	5.10^{-5}
Full moon	5.10^{-3}
Twilight	5.10^{-3}

Table 3: Examples of Physical Irradiances

Since we do not have enough information on the physical characteristics of the laser

and the receiver used in the experiments, we do not have exact values for image irradiance. For the numerical experiments we performed with N_s, we chose N_s from 0 to 13. The covariance and the mean of the random variable are the number N_s of charges collected in this pixel. The new pixel value is n, whose probability is:

$$p(n) = \frac{N_s^n}{n!} e^{-N_s}.$$

We will apply this noise to the final images. For the computation, we just need a random number generator.

5 Simulation Results

We present some of the results obtained with the simulation system described in the previous section, our main goal being the comparison with the experimental data introduced in the first section. Recall that the fixed images shown below are in fact frames from animations (animated GIF files) which can be viewed at the URL:

`http://www.princeton.edu/~rcarmona`.

Example 1

For this first example, the ocean surface was generated with a Phillips spectrum and the following parameters:

- ◇ The wind speed was set to 5.0 m/s
- ◇ The maximum wave elevation was setto 0.1 m
- ◇ The helicopter flies up from 16 m to 66 m
- ◇ A spherical mine of radius 0.3 m was present 5 m below the surface
- ◇ The incident rays were vertical
- ◇ The truncation of the PSF was at its 95 percentile

The following sequence of images was produced by a simulation of the elevation of the helicopter, without any backscattering or shot noise. We also consider the laser as a group of parallel incident rays. This result shows that our simulation system (even without the backscattering and shot noise effects) is adequate to produce images similar to the images of the first data set. Moreover, it also shows that it can handle motion of the imaging platform.

Example 2

The parameters used to produce the next example are very similar, except for the fact that the helicopter motion is not necessarily vertical. The output includes the surface of the sea

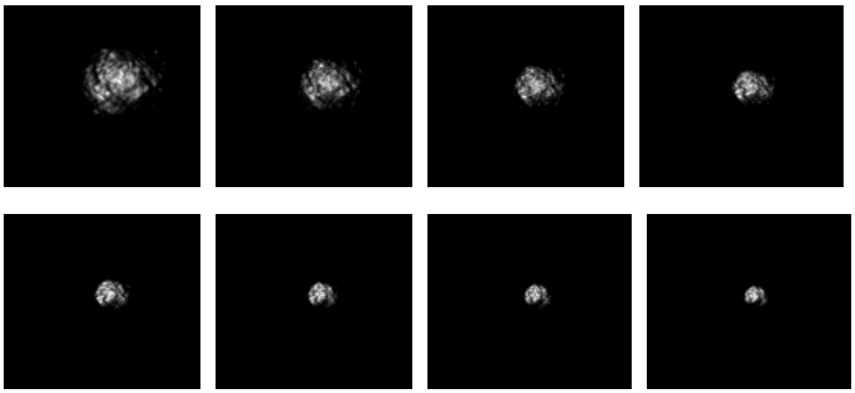

Figure 14: Mine view from an helicopter moving up.

evolving with time, together with the position of the helicopter (cross above the surface.) The cross below the surface gives the position of the target.

For the remaining examples included in this study, we shall only produce still frames (i.e. instantaneous snapshots), and as a consequence, there is no point in trying to move the imaging platform. This effect can be added if needed.

Example 3

The next example includes shot noise, backscattering, and it uses the laser beam distribution given in the text. The simulated conditions are essentially the same as in Example 1, except of the fact that:

◇ The grid has 256×256 separated by a mesh of 0.016 m

◇ The field of view (for the receiver and for the laser) was set at $\alpha = 6.7\,\mathrm{deg}$

◇ The laser illuminates a square of 4 m by 4 m.

◇ Reflectance of the target = 20

◇ The bottom of the gate was 6 m below the surface

◇ The backscattering coefficient was set at 7.9% (it is not arbitrary, but the result of the computation)

◇ The elevation of the helicopter was fixed and equal to 34 m

In order to understand better the effects of the backscattering, we produced images

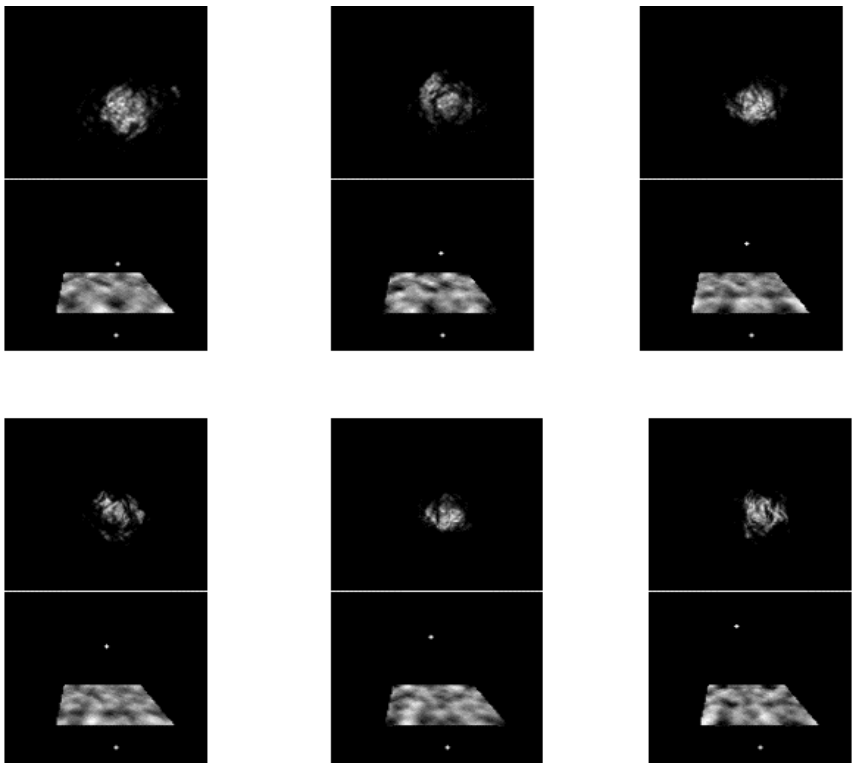

Figure 15: The bottom images show the ocean surface together position of the helicopter, and the top panes show the image of the mine as viewed from the helicopter moving up.

with a mine (left column of Figure 16), and without a mine (right column of Figure 16). One can clearly see the differences due to the presence of a mine.

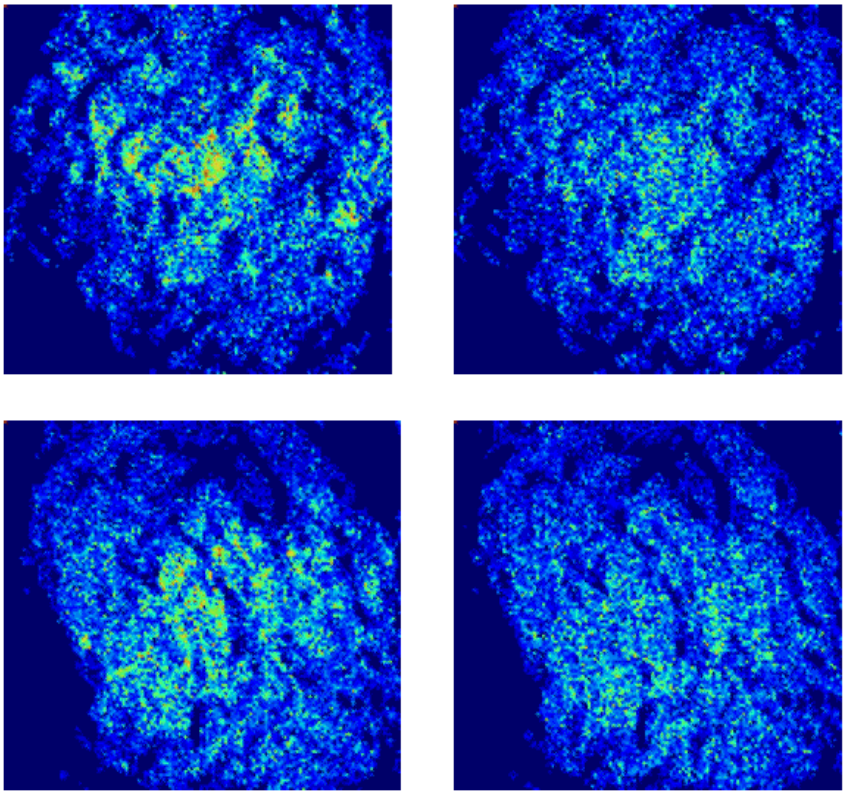

Figure 16: The images on the left contain a mine while the images on the right don't.

More images generated in this way shows that how apparent the mine is in the image depends strongly on the relative values of the target reflectance and the water turbidity as encapsulated by the backscattering coefficient. The mines can easily be seen in Figure 16 because the target reflectance is 20% and the backscattering coefficient is 7.9%. Real experimental data suffer from the same problems: the backscattering affects significantly the detection.

Example 4

Finally for this last example, we choose to illustrate the effects of the "gating" just below the target. The goal of this experiment is to see the shadow of the target. We take the same ocean surface, and the same target properties as above. The depth of the target plane is 3 m and the gating depth 9 m and we produce images as seen from the helicopter. The

image on the left of Figure 17 was produced with a target while the image on the right was produced without a target.

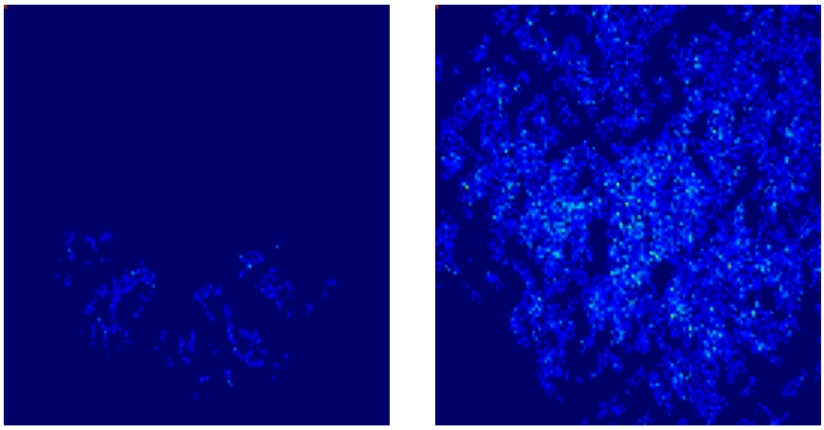

Figure 17: Simulation of the gating effect.

We can see the shadow and the influence of the target for the upwelling rays and downwelling rays in the left image. The right image is a reference of the same gating image (with backscattering....) but without the target above. This example shows the limits of our methodology and presumably of the computation of the PSF and the effective backscattering. Indeed, these images are not very realistic. The shadow is too big, with too much contrast: the multiple backscattering must have an influence despite the presence of the target.

6 Conclusion

This report records our first attempt at deriving a real-time imaging system for the detection and identification of underwater objects with an airborne LIDAR system. Implementations of a wave simulation model together with an imaging algorithm have been completed and experimental results have been presented and analyzed. The lack of objective metrics to quantify the quality of the results is the Achilles heel of simulation.

Nevertheless, the images produced by our system are very satisfactory. Moreover, the detection results obtained in [4] confirm its usefulness. In fact, because of its modularity, its usefulness should not be limited to the detection of moored mines. Many other applications in search and rescue, whether they take place in the ocean, in a lake, a pond or a pool, should benefit from the tools developed in this study.

Future research should aim at further developing this imaging tool in several directions, robustness being an obvious one.

But the ultimate challenge remains the derivation of a stable inversion algorithm, which will lead to a reliable underwater object detection and identification device.

Acknowledgments:

This work would not have been possible without the contributions of several individuals who have been instrumental in speeding up our learning curve. They include Dick Lau and Nancy Swanson for providing access to the data and to funding sources, and for many enlightening conversations on the real life experimental conditions. We are also grateful to Steve Moran from Kaman Aerospace for several discussions of the same vein. Finally, we would like to thank Annette Fung for her contribution at a very early stage of the project.

References

[1] R. J. Adler (1981): The Geometry of Random Fields. John Wiley & Sons. New York, NY

[2] J.D.E.Beynon and D.R.Lamb: Charge-Coupled devices and their applications, Mc Graw Hill

[3] R. Carmona and Y. Picard (1999): Gibbs Fields for Conditional Simulation of Random Fields. Tech.Rep. ORFE Dept. Princeton University.

[4] R. Carmona and L.Wang (2000): Lidar Detection of Moored Mines. Tech. Rep. ORFE Dept, Princeton University.

[5] R. Cipolla and P. Giblin (2000): Visual Motion of Curves and Surfaces. Cambridge Univ. Press. Cambridge UK

[6] Costal Systems Station: *Magic Lantern (ML) Automatic Recognition (ATR) Database Report*, Jan. 1996.

[7] N.A. Cressie (1993): Statistics for Spatial Data. John Wiley & Sons, New York, NY

[8] S. Frimpong and P.K Achireko (1998): Conditional LAS stochastic simulation of regionalized variables in random fields. *Computational Geosciences* **2**, 37-45.

[9] A.N.Fung (1998): Imaging through an ocean surface: a simulation approach. Senior Thesis, ORFE Dept, Princeton University

[10] R.C. Gonzales and R.E. Woods (1992): Digital Image Processing, Addison Wesley Pub. Cpy

[11] A.G. Journel and C.G. Huijbregts (1978): Mining Geostatistics. *Academic Press* New Yourk, NY.

[12] P. Liu and R.A. Kruger (1994): Semianalytical approach to pulsed light propagation in scattering and absorbing media. *Optical Engineering,* **33**

[13] J.W. McLean and J.D.Freeman (1996): Effects of ocean waves on airborn lidar imaging. *Applied Optics* **35**#118 pp. 3261-3269.

[14] L.E. Mertens and F.S. Replogle, Jr (1977): Use of point spread and beam spread functions for analysis of imaging systems in water

[15] C.D. Mobley (1994): Light and Water. Academic Press, New York, NY

[16] M.S. Schmalz (1990): Rectification of refractively-distorted imagery acquired through the sea surface: an Image Algebra formulation, *SPIE Image Algebra and Morphological image Processing*

[17] M. Shinozuka and R. Zhang (1996): Equivalence between kriging and CPDF methods for conditional simulation. *J. of Eng. Mechanics* **122**(6), 530-538.

[18] M. Shinozuka, R. Zhang and M. Oshiya (1996): Conditional Simulation Techniques: Recent Developments and Applications. *in IASSAR report on Computational Stochastic Mechanics* Sect. 1.6 pp. 224-226.

[19] P-T. D. Spanos (1983): ARMA Algorithms for Ocean Wave Modeling. *J. Energy Resources Technology* **105**,300-309.

[20] N.L. Swanson, V.M. Gehman, B.D. Billard and T.L. Gennaro (2000): Limits of the Small Angle Approximation to the Radiative Transport Equation. *NSWC preprint* Dahlgren VA.

[21] R.E. Walker (1994): Marine Light field Statistics. Jon Wiley & Sons, New York, NY.

[22] H.G. Wells (1973): theory of small angle scattering, *AGRAD Lect. Ser.* **61**