

Quantitative Visualization of Geophysical Flows Using Low-Cost Oblique Digital Time-Lapse Imaging

Rich Pawlowicz

Abstract—Estuaries and coastal waters are regions where many different important physical processes can be found. Although the physical scale of these processes is often relatively small, their time scales are correspondingly rapid and aliasing is usually a problem in sampling programs. There can be significant spatial variations in mixing and flow patterns, which are usually learned only through long experience in a particular region. Observational and interpretation difficulties might be significantly simplified with a simple remote sensing tool to be used in conjunction with standard techniques. Here, the use of digital time-lapse photography at highly oblique angles as a tool for flow visualization is discussed. The interaction of surface waves with slicks and internal motions can cause apparent changes in the shade and color of water at shallow angles in a way not apparent in downlooking views. The use of time-lapse techniques allows us to isolate time scales of interest and, by “speeding up” low-frequency motions, causes them to become more apparent to the eye. A cheap and portable system based on commercially available equipment is described and various advantages and shortcomings are discussed. Results are shown to illustrate the utility of the observational system.

Index Terms—Internal waves, slicks, turbulence.

I. INTRODUCTION

OCEANOGRAPHIC studies of estuaries and coastal waters are often confounded by the high degree of spatial and temporal variability present. A typical time series contains evidence of processes over a wide variety of time scales, from density-driven estuarine flow, to reasonably deterministic tidal processes, to intermittent turbulent and internal wave features. Many of these are coherent over fairly large spatial scales; that is, they reflect large-scale changes in the pattern of flow (and not merely meter-scale isotropic turbulence). However, identifying the nature of this coherence and optimizing fieldwork to properly observe it is, at best, difficult, requiring a deep knowledge of local waters, and, at worst, impossible. It is generally difficult to know (or even predict) the extent to which measurements taken from a fixed location or a single platform are representative of the surrounding waters.

On the other hand, such regions contain a wealth of interesting and important physical and biological phenomena whose study can be carried out with relatively modest logistical effort if properly applied. A simple flow-visualization technique that

could help identify large-scale patterns of flow would complement and greatly increase the information available from standard ship- or mooring-based observational techniques, allowing them to be placed in a useful context. Fortunately, these waters are often relatively shallow and subsurface flows are many times associated with a strong surface expression. Some kind of surface imaging appears to be an obvious candidate for such a visualization technique.

However, this is not the only motivation underlying oblique surface imaging. Every sea-going oceanographer spends time staring out over the water surface. It is often the case, especially during calm seas, that the water is not a uniform shade, but has a color, perceived roughness, and even sheen that varies from place to place within a field of view [1]. Some of these changes are obviously due to changes in incident lighting (e.g., color of light reflected from hills, clouds, or sky), but even after such differences are accounted for there may be much fine-scale variability. As the ocean surface is the part of the ocean most accessible to our unaided senses and its physics important to air-sea interaction processes, its characteristics and evolution are of interest in their own right. Unhappily, the deck of a ship is not always an ideal place to study such phenomena. Spatial extent and scale is difficult to judge on an otherwise undifferentiated water surface once the view angles are within a few degrees of the horizon and the evolution of these features is also difficult to discern, since they change over time scales of minutes to hours—too long for the eye to perceive, but too short for “synoptic” ship observations. Perhaps surprisingly, down-looking aerial observation is also not an ideal tool except in the direction of the sun’s specular reflection. Although airborne vertical video imagery has been used too, e.g., to study the large boils generated by flow over the bottom of the shallow North Sea [2], many surface features are not visible at near-vertical view angles. Also, whereas overflights at low altitude result in images over only very small regions (with swath widths of no more than a few hundred meters), higher-altitude aerial observations that can cover a wider area are often affected by clouds. High-resolution space-based imaging is also possible, but usually at inconvenient and widely spaced times. Although there are many forms of analysis for which this is not necessarily a problem [3]–[5], studies of feature evolution would in general benefit from time series information. Such information can sometimes be obtained from coastal radars, which are being used to image the surface representation of significant features such as internal waves [6], but there is a significant cost and effort associated with such studies. Simpler visual-imaging systems have been used by a number of workers to study near-shore

Manuscript received April 2, 2002; revised January 13, 2003. This work was supported by the Natural Sciences Engineering Research Council of Canada under Grant OGP0294270.

The author is with the Department of Earth and Ocean Sciences, University of British Columbia, Vancouver, BC V6T 1Z4, Canada (e-mail: rich@eos.ubc.ca).
Digital Object Identifier 10.1109/JOE.2003.819310



Fig. 1. Image of an internal wave packet in the Strait of Georgia, BC, Canada, taken from an aircraft at an altitude of 480 m. These soliton-like internal waves of depression lie on a pycnocline at a depth of about 4 m and are spaced about 100 m apart, with amplitudes of about 5 m and widths of about 30 m. At the top of the photo, the viewing dip angle is about 5° . The internal wave's leading edges appear as thin dark lines and the trailing edges as a brighter shiny patches. At the bottom of the photo, the dip angle is about 35° and the internal waves crests show up as light-colored lines. The light spot at the center right is a vessel of length 29 m.

surf and sediment transport dynamics [7]–[10]. In that work, systems of shore-mounted video cameras were used to produce rectified images of the surf zone over time scales of minutes to years. Generally, the visual contrast between land, water, and breaking waves is large enough that even the poor quality of a video image is sufficient to obtain useful results. An airborne oblique-angle digital imaging system has also been developed recently [11]. This system is designed to image kilometer-scale sections of coastline for periods of a few minutes; the resulting time series of images is processed using the known dispersion of surface waves to deduce water depth. Here, a simple system that has been developed for flow visualization over spatial scales of a few kilometers and time scales of a few tens of minutes to tidal periods is described. It has been found that subtle changes in surface shading, appearing either in response to internal motions or by the advection of surface films, can be monitored in this way. An unexpected variety of physical processes have been observed, including internal wave and frontal propagation, turbulent boils, and the changing patterns of tidal flow.

In this paper, the physical phenomena that cause changes in the appearance of the water surface are first discussed. The underlying “cheap and light” philosophy is explained and its effect on the choice of equipment outlined. Third, the basic mathematics of photogrammetric rectification with specific application to the highly oblique view angles being used are developed. Finally, the utility of this system is illustrated by discussing observations of several oceanographic phenomena.

II. APPEARANCE OF NATURAL WATERS AT OBLIQUE VIEW ANGLES

The surface of coastal waters can appear quite different in many situations; not all (or even much) of this variation is due

to changes in water column currents. In any single image, correct feature identification may be very difficult. However, when viewing a time-lapse animation, the differing speeds of evolution and propagation usually make interpretation of the images much less ambiguous. It is important to emphasize that most of the features of relevance here are rather subtle, consisting of slight changes in shading or brightness, and significant image enhancement is necessary to bring out their details. Many of the points in the following discussion are demonstrated in the image shown in Fig. 1.

Consider first a perfectly flat water surface far from land under a cloudless sky. The appearance of the water as perceived by an observer above the water surface is composed of a mixture of the specularly reflected sky light (which dominates at very shallow angles, i.e., near the horizon) and light that is returned after being scattered within the water column (which dominates at near-vertical angles). The specularly reflected light will come from different parts of the sky, which even in cloudless conditions is not uniformly lit. Instead, it is darkest (and bluest) when looking vertically upward and brightest (and whitest) close to the horizon, due to the increased optical depth in that direction [1]. The appearance of the water will, thus, vary smoothly across the field of view. There also may be a sharp bright patch directly beneath the sun resulting from specular reflection of its light, perhaps surrounded by a fuzzier bright area resulting from light that is weakly scattered by the atmosphere. A smooth variation of brightness is seen clearly in Fig. 1. Upper parts of this figure (at shallower dip angles) are much brighter than the lower parts. This variation can make it difficult to set the correct exposure for the entire image. Often, part of the image will appear very dark while other regions may be washed out.

In reality, water is almost never flat. Surface waves can either be large in size relative to the resolution of the eye or camera,

in which case they give rise to obvious wave patterns in the image as the varying slopes change the source location of reflected light, or they can be small, in which case a “smoothed” large-scale effect is perceived. A rough patch of water will usually appear darker than a smooth patch because its tilted slopes reflect light from closer to the vertical (an exception may occur in the vicinity of the sun’s specular point). Rougher patches are seen as curved striations in the upper part of Fig. 1. The degree of smoothing and its spatial extent will obviously vary depending on many factors. In a “perfectly rough” limit under high winds, there may be no noticeable spatial variation, except possibly near the sun’s specular reflection point where a mass of “glint” speckles will appear. As time progresses and the sun moves across the sky, this broad pattern will also change in magnitude and location. In weak winds, more localized spatial variations in surface appearance are often present.

Such smaller scale variability can arise from a variety of sources. First, the lighting source may have sharp variations in intensity and color. Water in front of islands will often appear darker due to the reflection of vegetation, but may also be lighter if islands are rocky or covered in buildings. Topography will affect the pattern of this shading change, but these patterns tend to be relatively constant in time. Clouds can change ambient lighting in ways that vary with time. Although clear days can provide best “viewing” (depending on sun location) with the most vivid blue colors for water, days with a thin overcast have also proved surprisingly useful, as the illumination again varies smoothly from overhead to the horizon. Thick cloud, on the other hand, often provides relatively uniform lighting that reduces contrasts. The water surface becomes uniform in shade. A broken cloud is most problematic because the lighting changes are no longer spatially smooth, causing apparent variations in surface color. Low dark clouds result in correspondingly darker waters in the specular direction, which tend to produce elongated dark areas in photos that appear as “spokes” in rectified images. These spokes may rotate about the camera location in either clockwise or counter-clockwise directions, depending on the cloud motions. Smaller higher clouds can produce shadows on the water. Most cloud effects are usually not “water-like” in photogrammetrically rectified animations (i.e., they produce features that advect across coastlines, or at speeds much greater than can be expected from water column dynamics, or in a fashion obviously related to the viewing geometry) and can be recognized, especially if the original (unrectified) images are examined for evidence of clouds.

The second and more useful source of variability arises from small-scale changes in surface-water properties. These can arise from changes in the nature of the water surface and also in response to flow variations within the water column.

The water surface is often not “clean,” but instead supports surface films of natural or industrial oils [12]. These films damp small-scale capillary waves and have well-defined edges that appear as changes in reflectance at wind speeds below about 5 m/s and can cover up to as much as 40% of the ocean surface at wind speeds less than 2.5 m/s [13]. The smooth water is often brighter than neighboring rippled water, as it reflects the sky light from nearer the horizon. This can change

depending on the location of the sun, nearby land, variations in cloud cover, etc. Such features will tend to be advected with the ambient currents and in animation sequences they produce an easily discernible flow pattern. This can be useful for identifying large-scale changes such as the swirl of water as tidal flow transitions from ebb to flood at different times across part of an estuary.

Internal variability can also change surface properties. Turbulent boils are especially noticeable. These consist of a large central region of upwelling, usually smooth, surrounded by thin sheets of downwelling with steep choppy waves ranging in size from 1–100 cm high, depending on the vigor of the turbulence. Boils will occur in unstratified water in rivers and may also occur in regions of intense mixing of stratified water (flow-separation zones, downstream of sills, etc.). Internal waves are another source of surface variability. This link was recognized in some of the earliest observations of internal waves [14], [15] and is still exploited, usually in a qualitative way [16]. Consider an internal wave of depression at the interface of a thin upper fluid and a thicker lower fluid (as in Fig. 1). At the leading edge of this wave, upper waters are converging and downwelling. At the trailing edge, waters are diverging and upwelling. Water in the convergent region is often rougher. If the internal waves are large, the surface waves in the convergent region may even be breaking, producing white foam. Floating debris may be swept up in the leading wavefront of a large packet. Although clearly visible at highly oblique angles, such internal waves may also be seen at steeper viewing angles, either through specular reflection of the sun or because the water in upper and lower fluid layers is optically different (in either color or transmissivity). In the upper part of Fig. 1, surface reflected light dominates and the leading edge of near-surface internal waves appears as a dark line due to small-scale roughness in the convergence and downwelling that occurs there. Trailing edges are brighter. In this case, the width of leading and trailing edges appears different, although water column observations suggest that the waves themselves are reasonably symmetric. Other observations of these waves show that visibility of this leading-edge roughness is relatively insensitive to the viewing direction as long as the dip angle is small. (In contrast, the large surface waves produced in ship wakes are often invisible when viewed along their crests.) In the lower part of the figure, the water’s appearance is controlled by upwelled rather than surface reflected light and the waves appear as light-colored lines. The color changes arise from the changes in the thickness of the turbid upper layer. A surface wave pattern is just barely resolved in the lower part of Fig. 1, with wave crests aligned from center left to lower right. At this time, there was little wind and only very gentle swell.

Overall, experience suggests that if variability is visibly present, then something interesting can be seen through imaging. However, the physical processes responsible for the variability are often identified only afterward by examining the rectified time-lapse sequences. In addition, the eye is much better at discerning these subtle features than a camera system; in some cases, considerable efforts must be made to adequately extract relevant features from images.

III. DESIGN OF IMAGING SYSTEM

A. Design Philosophy

A primary goal of this work was to develop an imaging system that would be relatively cheap and could be used as an adjunct to conventional observational techniques without a great deal of effort. The ideal would be a system costing only a few thousand dollars at most that could be operated by a single technician or student with little experience. If components are cheap, less effort is needed to ensure adequate security of the installation and to construct shelters in the event of inclement weather. It is easier to justify it as a complement to existing observational techniques rather than as the primary observational tool. It is probably important to note that notwithstanding the discussion of the previous section, various situations can arise in which no useful information can be gathered at all (e.g., when winds are high or fog or thick clouds are present). Another advantage of a cheaper system was less obvious initially, but became more clear once field experience was acquired. A typical camera with a wide-angle lens has a field of view of 50° or so, which may be frustratingly small when looking at expanses of open water (in comparison, the human eye has a field of view of about 100°). This appears not to be a problem in near-shore studies where cameras are aimed along a coastline [8], but is more of an issue when looking offshore. The field of view can be expanded slightly by tilting the camera so that the water surface (which tends to fill only a narrow band over the entire image) is along a diagonal, but even with this simple trick it is difficult in many cases to image the entire area of interest. In consequence, deployment of multiple systems with overlapping or complementary fields of view has been found to be very useful, greatly improving the utility of the results. In general, wider area coverage with a number of relatively cheap instruments appears to be more useful than limited coverage with a more sophisticated package.

B. Spatial Resolution

One limitation on the viewing range possible with such systems is the imaging resolution available. Although actual flow parameters vary considerably depending on the situation, it is reasonable to assume very roughly that typical feature speeds are about $v \approx 1$ m/s in the coastal/estuarine environment and that features (internal wave crests, slicks, etc.) are of order $\delta D \approx 100$ m apart. Consider a camera at a height $H \approx 300$ m imaging water at a distance of $D \approx 4$ km away. Then, the dip angle λ is given by $\lambda = \tan^{-1}(H/D) \approx 4^\circ$, and the vertical angle $\delta\lambda$ subtended by a feature δD across can be estimated from

$$\delta\lambda \approx -\frac{H}{D} \frac{\delta D}{D} \approx 0.1^\circ. \quad (1)$$

(The horizontal angle subtended by such a feature would be about $\delta D/D \approx 1.4^\circ$.) Taking as typical a camera field of view of order 50° (ignoring for now the distinction between horizontal and vertical fields of view), it is clear that at least 1000 pixels, each of which provide average intensity over a region of angular width 0.05° , are required to discern such details. Digital still or video cameras designed for the mass market, producing

images with widths of < 1000 pixels are, therefore, unsatisfactory for this application, but so-called 2–3 “mega-pixel” resolution cameras are adequate.

Note that it is possible to improve spatial resolution using a zoom lens. However, the decrease in field of view often more than offsets the advantages of increased resolution. An improved procedure combines a high-spatial resolution but narrow field-of-view (zoomed) image with a lower spatial resolution wide-angle image at each time interval, to provide greater resolution in the central part of the viewing area while still maximizing the size of the imaged region.

C. Image Timing and Time Stamping

In order for motions to appear smooth in animations, the feature displacement between successive images should be smaller than the feature separation, i.e., images should be less than $\delta t = \delta D/v \approx 100$ s or about 2 min apart. Conversely, data-storage requirements increase with the number of images taken and, hence, it is most efficient if images are not obtained too frequently. A typical image spacing of 1–2 min appears to work best as an initial guess for the type of flows described here. These parameters suggest that a high-resolution still camera is an appropriate tool instead of lower resolution video cameras.

A rather unusual degree of knowledge about the timing of the imaging process is required compared with the requirements of most photography, especially when combining them with observations made from another camera system or from various instruments within the water column, e.g., a thermistor chain or an acoustic Doppler current profiler mounted on a ship. Experience suggests that synchronization to better than 15 s is necessary for features such as internal waves and, thus, some kind of time stamping is required. Experience also suggests that having a method for stopping and starting a typical time-lapse sequence without changing the temporal spacing of images greatly reduces the complexity of subsequent image analysis in stationary deployments.

D. Deployment Logistics

It was considered necessary that the system developed could be carried by one or two persons and deployed at a field site without external power. This would provide for maximum flexibility in siting and the independence from available infrastructure would make fieldwork logistics simpler. The system described here has been deployed on hilltops, rooftop decks, embankments, and a ship’s bridge, and has also been used from an aircraft. However, the system does require an operator in reasonably constant attendance, partly for security reasons and partly due to limitations in onboard data storage (described below). This makes it inconvenient for the autonomous long-term deployments that have been described elsewhere [8], [10].

Data-storage requirements vary depending on the camera resolution and the level of compression used, but typically a so-called 2 mega-pixel camera with three 8-bit color planes will require about 6 Mb for an uncompressed image and perhaps 0.5–1 Mb for a JPEG compressed image at least compression. A variety of strategies can be used to decrease file size. In many cases, the water only occupies a narrow strip through the

center of the entire image. Judicious masking of other parts of the image using, e.g., duct tape, can significantly affect compression. Complex terrestrial vegetation (trees and bushes) in particular tends not to compress well; masking this part of an image has resulted in file-size reductions of more than 50% in some cases. A camera deployment requiring two photos every 2 min will then typically generate about 30–360 Mb/h of data. This can be a substantial amount over a tidal period of 6 or 7 h. Although data-storage capacities are quickly increasing, some kind of download capability is probably required. Note that although brightness variations are of primary interest in the appearance of the water surface (implying that a considerable reduction in storage space might be possible by neglecting color information), in practice this option is not available in most commercial cameras. In addition, color information turns out to be useful when attempting to automatically differentiate between land and water pixels during the (usually necessary) image processing required to bring out brightness details.

Two different options for solving storage problems were considered. First, it is possible to control many cameras via a cabled interface to custom software running on a laptop or desktop computer. Such software can then take care of details such as time synchronization and exposure compensation, and images can be downloaded through this interface. This solution requires that the camera be sited in a weatherproof location with adequate power available to support the external computer, as neither camera nor laptop computer batteries are generally capable of an 8-h deployment. Twelve volt lead-acid automobile batteries and inverters were used initially, but their weight is not insignificant and this restricts the available site locations. Alternatively, images can be stored within the camera and extracted at some larger time interval. Currently, many cameras use a form of removable flash memory with maximum available capacities of order 100 Mb (although some “professional” cameras have capacities of up to 1 Gb). Typically, this may suffice for an hour or so of imaging, after which the card must be replaced with another (empty) card. The stored data can then be transferred to a laptop or a number of memory cards can be purchased so that data is transferred at the end of the day (presumably at some base location with available power). Eliminating the need for continuous computer support drastically reduces the power-consumption requirements.

E. System Description

Based on the above discussion and subsequent field experience, the creation of a scientifically useful time-lapse sequence (where “useful” implies quantitative analysis and coordination with other sensors) requires the following:

- 1) relatively fool-proof method for keeping constant image intervals;
- 2) ability to synchronize camera times with other cameras and instruments to a few seconds;
- 3) protocol for starting and stopping the time-lapse sequence for servicing, including emptying onboard storage, while still keeping sequence spacing constant;
- 4) some indicator of when such servicing is required;
- 5) ability to shoot multiple images per interval.

Almost all of these require some kind of programming capability. Although many cameras claim to have a “time-lapse” feature, these are often impractical for scientific purposes. You cannot mix wide-angle and zoom images as discussed above and there is often no indication of available memory filling up until it occurs, so a tediously impractical constant watch must be maintained. Once the memory has been emptied, restarting the time-lapse sequence at even time intervals with the aid of, e.g., a wristwatch can be frustrating (displayed time is usually only given to the nearest minute).

Commercially available “prosumer” camera equipment has the advantage of (relatively) low cost (currently at or below \$1000 US) while providing adequate resolution (image widths of 1000–2000 pixels), but has the disadvantage that relevant performance specifications such as the amount of lens distortion are in many cases not clear or not known at all, so that time must be devoted to understanding and/or determining these parameters. Note that “professional” camera equipment with similar resolution but with better lenses and imaging technology is also available, albeit at considerably greater costs. Such equipment was not considered at this stage of the development.

Product research suggested that one of a group of cameras using the Digita operating system would be of most use since they allow the user to control camera operations by writing programs in a simple scripting language. The programs are stored on the data-storage memory cards and are executed by the camera, thus eliminating the requirement for a controlling computer and considerably reducing the power requirements. The work described here was carried out with a Kodak DC265 and a Kodak DC290. The pixel resolution of these cameras appears adequate to image surface properties out to a distance of about 10–15 times the camera altitude (although landmarks as far as the horizon can be used for determining camera orientation).

Programs have been developed for stationary deployments to ensure that a preset sequence of images are taken at user-selectable intervals, e.g., a combination wide-angle and zoom image every 2 min on the minute, counting from midnight, and to display a count of images taken and an estimate of remaining memory. The risk of large data gaps or irregularly spaced data is greatly decreased (especially when several instruments are deployed simultaneously). In addition, serial communication with other devices including global positioning systems (GPS) and an electronic compass/tilt meter can be used to provide precise time synchronization and estimates of camera orientation, which has proved to be useful in hand-held airborne deployments.

The current system can be powered directly from a rechargeable 6 V sealed lead-acid gell cell (voltage requirements for different cameras vary and documentation usually advises against such procedures). A single cell (weighing less than 1 kg) is sufficient for a typical deployment. A single deployed unit then consists of

- rigid tripod;
- camera;
- external powerpack;
- extra memory cards;
- portable GPS unit (optional).

This can easily be carried by a single person over rough terrain.

IV. PHOTOGRAMMETRIC RECTIFICATION

Although replaying sequences of raw images provides useful information, the true power of this technique lies in the ability to transform photos into photogrammetrically mapped images in ground coordinates with overlaid bathymetry, ship locations, measured currents, etc. Although such a transformation can be carried out by using a general function fitted over an extensive set of “targets” with known locations, the sparsity of recognizable viewpoints over water requires a different approach. This can be broken down as

- 1) conversion from pixel-location indices into a view vector relative to the central axis of the camera (“camera coordinates”);
- 2) transformation of that view vector from camera coordinates into earth-referenced coordinates by applying successive rotations through all three axes;
- 3) determination of the intersection of the view vector with the ground plane;
- 4) conversion from coordinates of the ground point relative to the camera to true ground coordinates.

In order to perform this transformation, various extrinsic and intrinsic camera parameters must be known either *a priori* or through some estimation procedure. Extrinsic parameters include camera latitude LAT and longitude LON , elevation H above water level, and camera orientation (clockwise tilt, dip below horizon, and azimuthal rotation from true north). Intrinsic parameters at minimum include some measure of the field of view (or, equivalently, the calibrated focal length or camera constant) of the camera and may also include a quantification of various mechanical and optical distortions. Although most of the mathematics involved in conversion of pixel images to ground coordinates presented here is classical for photogrammetric applications [17] and has been well studied in the machine vision literature [18], the highly oblique view angles used over large distances require compensation for the sphericity of the earth.

A. Conversion Into Camera Coordinates

It is standard in both the computer vision and photogrammetry literature to begin with an idealized “pinhole” model of the camera. In this model, external objects appear on an image plane a distance f_l (in units of pixels) behind the pinhole, at the intersection of that plane with a straight line ray extending from the external object, through the fixed point or pinhole, and on past the image plane (see Fig. 2). One of the many rays through the pinhole will be perpendicular to the image plane; the intersection of this ray with the image plane at pixel location $\mathbf{p}_{pr} = (i_{pr}, j_{pr})$ is called the principal point. A digital image consists of a two-dimensional (2-D) plane of pixels with w columns and h rows. A pixel is referred to by coordinates $\mathbf{p} = (i, j)$, corresponding to the row and column of the pixel, respectively, numbering from the upper left corner of the image (i.e., lower right corner of image plane). The camera coordinate reference frame is centered on the “pinhole,” with the x direction along the base of the camera, the y direction along the side of the camera, and the z direction or principal axis pointing

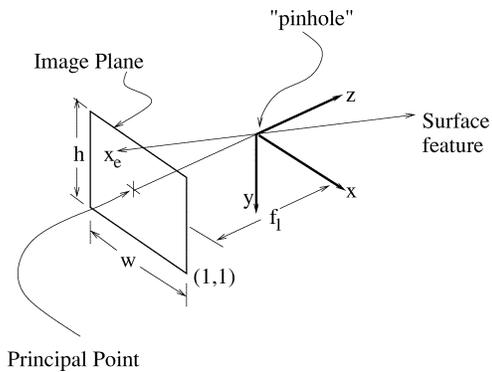


Fig. 2. Camera coordinate system.

outward through the lens. Camera coordinates of points in the image plane $\bar{\mathbf{x}}_c$ are then

$$\bar{\mathbf{x}}_c = \begin{bmatrix} x_c \\ y_c \\ z_c \end{bmatrix} = \begin{bmatrix} j_{pr} - j \\ i_{pr} - i \\ -f_l \end{bmatrix}. \quad (2)$$

This coordinate system has proven to be adequate in relatively simple deployment situations. These include deployments with land points relatively far from the region of interest (although still in the photo image, perhaps near the horizon) with no or few topographic features obscuring nearby waters. However, for more accurate registration of features it is likely that some corrections may have to be made. Two effects are commonly found. First, a radially symmetric lens distortion may be present, especially in lower quality lenses, such that straight lines are bowed outward (“barrel effect”) or inward (“pincushion”). This can be modeled by a multiplicative function containing even powers of radial distance from the principal point. Coefficients of this function will be lens-dependent. Second, imperfect centering of lens elements can give rise to tangential distortions [17], [19].

B. Transformation Into Earth Coordinates

The camera orientation is specified by three extrinsic angle parameters that govern the transformation from an earth-referenced coordinate system to the camera coordinate system. The earth-referenced coordinate system has its origin at the pinhole, with \hat{z} axis pointing north, \hat{x} axis pointing east, and \hat{y} axis vertical downward. The rotation angles represent the azimuthal clockwise rotation of the principal axis from north θ , the dip of this axis below the horizontal λ , and the tilt of the camera in a clockwise direction around the principal axis ϕ . In order to convert an image point in camera coordinates $\bar{\mathbf{x}}_c$ into earth-referenced coordinates $\bar{\mathbf{x}}_e$, rotation matrices are applied as

$$\bar{\mathbf{x}}_e = M_y M_x M_z \bar{\mathbf{x}}_c \quad (3)$$

where

$$M_z = \begin{bmatrix} \cos(\phi) & \sin(\phi) & 0 \\ -\sin(\phi) & \cos(\phi) & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (4)$$

$$M_x = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(\lambda) & \sin(\lambda) \\ 0 & -\sin(\lambda) & \cos(\lambda) \end{bmatrix} \quad (5)$$

$$M_y = \begin{bmatrix} \cos(\theta) & 0 & \sin(\theta) \\ 0 & 1 & 0 \\ -\sin(\theta) & 0 & \cos(\theta) \end{bmatrix}. \quad (6)$$

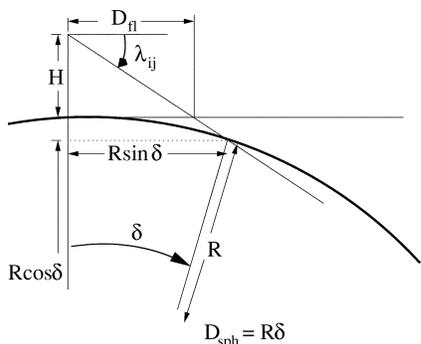


Fig. 3. Sphericity of the earth results in the distance D_{sph} to a ground point at dip angle λ_{ij} from a camera at elevation Z_c being greater than the distance D_{fl} on a flat earth (see text for details).

C. Location of Image Point on the Ground

Assuming the earth's surface to be flat, the ground point in earth coordinates will be the point at which a line extending from the origin in the direction \bar{x}_e intersects the ground plane $y_{gr} = H$

$$x_{gr} = x_e \frac{H}{y_e} \quad (7)$$

$$z_{gr} = z_e \frac{H}{y_e}. \quad (8)$$

If $y_e < 0$, then the image point is in fact above the horizon and can be ignored.

A correction is necessary to account for the curvature of the earth in this mapping. This is especially important in cases when points at or near the horizon are being used to determine camera orientation. Curvature of the earth will result in the distance to the ground point being greater than that predicted using that of (7) and (8) by as much as 100% for points near the horizon. Alternatively, a point near the horizon will appear at a greater dip angle than would be expected based on its distance if the earth was flat. Consider a particular pixel (i, j) . By the formulas given above, this will appear on the ground plane at a distance $D_{fl} = \sqrt{x_{gr}^2 + z_{gr}^2}$ from the camera, at an effective dip angle of λ_{ij} . By geometrical reasoning (Fig. 3), the actual surface distance to the ground point is $D_{sph} = R\delta$ for an earth radius of $R \approx 6370$ km and an angle measured from the center of the earth δ . Then

$$\begin{aligned} \tan \lambda_{ij} &= \frac{H}{D_{fl}} = \frac{H + R(1 - \cos \delta)}{R \sin \delta} \\ &\approx \frac{H}{R\delta} + \frac{R\delta}{2R} \\ &= \frac{H}{D_{sph}} + \frac{D_{sph}}{2R} \end{aligned} \quad (9)$$

using $\delta \ll 1$ for low-altitude viewpoints. Equations (7) and (8) are then modified as

$$x_{gr} = x_e \frac{H}{y_e} \frac{D_{sph}}{D_{fl}} \quad (10)$$

$$z_{gr} = z_e \frac{H}{y_e} \frac{D_{sph}}{D_{fl}} \quad (11)$$

where the quadratic equation (9) can be solved exactly for the required correction factor

$$\frac{D_{sph}}{D_{fl}} = 1 + \frac{D_{sph}^2}{D_{hor}^2} \quad (12)$$

$$= 2 \left(\frac{D_{hor}}{2D_{fl}} \right)^2 \left\{ 1 - \sqrt{1 - \left(\frac{2D_{fl}}{D_{hor}} \right)^2} \right\} \quad (13)$$

$$\approx 1 + \frac{D_{fl}^2}{D_{hor}^2}, \quad (14)$$

the last approximation being valid for $D_{fl} < D_{hor}/2$. The distance to the horizon $D_{hor} = \sqrt{2RH}$ at low altitudes. Thus, for a camera mounted 300 m above water level the horizon is 44 km away and the correction at a distance of 10 km will be about 5%, increasing to 20% at 20 km. Note that in the foregoing we have assumed that the earth is spherical for purposes of distance corrections, but not for computing directions. (In essence, the interchangeability of great circles and rhumb lines is assumed.) This is reasonable for typical deployment situations.

A further correction (usually opposed to the sphericity correction) may arise from atmospheric refraction as light rays travel almost parallel to the surface. Downward refraction of light in the horizontally inhomogeneous but stable marine boundary layer during days of light wind has been observed from ships in the region, where most of the camera system testing has occurred. The refraction gives rise to the "hafgerdingar" effect or Fata Morgana [1]. On land, distant objects appear as anvil-shaped spires, but in the coastal ocean the most obvious manifestation of this refraction is an apparent cliff along the edge of distance coastlines. These effects have not yet been observed in deployments, possibly because the effect is of the order of pixel resolution and possibly because they are strongest at near-zero dip angles from observation points quite close to the water. In the observations discussed below, camera locations were usually at elevations of several hundred meters.

D. True Ground Coordinates

Latitudes and longitudes of ground points are most easily found by making a "locally planar" assumption, i.e., assuming that 1° of latitude equals $2\pi R/360$ linear units and that 1° of longitude equals $2\pi R \cos(LAT)/360$ linear units. These values are then added to the known camera latitude and longitude to give true ground coordinates \bar{x}_{true} .

E. Inverse Transformation

The forward mapping

$$\bar{x}_{true} = \mathcal{F}(\mathbf{p}; \mathbf{p}_{pr}, f_l, \theta, \phi, \lambda, LAT, LON, H) \quad (15)$$

from pixel to true ground coordinates using an ideal camera and a spherical earth was described above. For many purposes, it is necessary to also compute the inverse map \mathcal{F}^{-1} . This is useful for mapping coastline data onto the original image to aid in checking that camera-orientation parameters are correct. Also, the most computationally efficient algorithm for image rectification involves conversion of true ground coordinates back into pixel values indexed into the photo image, rather than a simple

transformation of pixel locations to ground coordinates. In the inverse map, one assumes that $y_e = H$ and computes x_e, y_e by inverting (7) and (8) or (10) and (11), finding

$$\bar{\mathbf{x}}_{\hat{c}} = M_z^T M_x^T M_y^T \bar{\mathbf{x}}_{\hat{c}} \quad (16)$$

and, thus, recovers image coordinates through

$$j = j_{pr} + x_{\hat{c}} \frac{f_l}{z_{\hat{c}}} \quad (17)$$

$$i = i_{pr} + y_{\hat{c}} \frac{f_l}{z_{\hat{c}}}. \quad (18)$$

Inversion of lens-distortion parameters involves finding the roots of a high-order polynomial, but this can be done through an iterative technique [19].

F. Determination of Camera Parameters

Carrying out the mapping (15) or its inverse requires knowledge of extrinsic and intrinsic camera parameters. Intrinsic parameters can be estimated from camera specifications or with more precision from laboratory calibration. If the camera is “perfect,” then one would expect $i_{pr} = w/2$ and $j_{pr} = h/2$. Effective focal length f_l may be estimated from manufacturer information. Focal lengths on digital cameras are often specified in equivalent terms to a 35-mm camera, even though the image plane is a different size. A 35-mm camera has a nominal image plane width $w = 35$ mm. Alternatively, the horizontal field of view (FOV) is sometimes specified by the manufacturer. A typical value for a wide-angle lens is about 50° . Focal length and FOV are related through

$$\text{FOV} = 2 \tan^{-1} \left(\frac{w}{2f_l} \right). \quad (19)$$

Note that only the ratio w/f_l appears in this formula so that units are not important. In particular, if FOV is known, then f_l can be deduced in units of pixels since the image width w in pixels is known. In practice there may be some problems with this approach. The camera’s focal length specification is a parameter in a thick-lens camera model; even if considered to be an exact measurement, it will not be the same as the effective focal length considered here, although when focused at infinity they should be the same [8]. It is unlikely that the principal point is actually in the center of the camera image. In fact, the definition of principal point becomes less clear on closer examination [20]. These parameters, as well as other parameters quantifying lens-distortion effects, can be measured by a careful laboratory calibration in advance of field deployments [18]. Software to assist in this process is available [19]. The particular cameras used in the work described here had relatively small amounts of lens distortion and reasonable results were obtained using nominal values only.

An associated problem concerns the repeatability of these parameters as the camera lens is moved in and out. This *stability* of the camera is important, but is generally not quantified except in cameras specifically designed for photogrammetric work. In practice this has also not proved to be a significant issue. For extended deployments, it appears to be more important to ensure that the camera mount is sufficiently rigid so that orientation is not disturbed by periodic maintenance.

Knowledge of the extrinsic parameters can be obtained either by *a priori* measurement or by some estimation technique in which a small number of known locations are mapped into image coordinates and a fit performed by varying parameters until some measure of error is minimized. Various strategies are possible, depending on the amount of information known. For deployments at a fixed location, it is generally possible to determine camera position and altitude through careful use of global positioning system (GPS) receivers or even a large-scale topographic map. Measuring the orientation with the required precision is difficult. At oblique view angles ($\lambda < 10^\circ$) the required precision is less than $1/10^\circ$, beyond the limits of easily purchased compass/tilt meters even if a bore-sighting calibration has been carried out. However, such instruments can provide a useful initial estimate that can be used as a starting point for manual or automatic fitting procedures. The fitting procedure can be formally written as the nonlinear minimization of a cost function. If, for example, we know the intrinsic parameters, as well as the ground locations $\bar{\mathbf{x}}_{\text{true}}^n$, $n = 1, \dots, N$ of N targets, as well as the corresponding pixel indices \mathbf{p}^n of these locations from, say, the first image in a sequence, then we solve for

$$\min_{\theta, \phi, \lambda} \sum_n \|\mathcal{F}(\mathbf{p}^n) - \bar{\mathbf{x}}_{\text{true}}^n\|^2 \quad (20)$$

or

$$\min_{\theta, \phi, \lambda} \sum_n \|\mathbf{p}^n - \mathcal{F}^{-1}(\bar{\mathbf{x}}_{\text{true}}^n)\|^2. \quad (21)$$

Assuming intrinsic parameters are known, identification of two known locations in the photo (or even of one known location and the known orientation of something like a ship wake) suffices to completely determine the orientation angles (with a cost function minimum of zero). Known locations can be fixed features, such as headlands or tracked locations of ships. In the latter case, it becomes important to verify that the various coordinate systems are using similar datums. Alternatively, if images are taken from a moving platform (e.g., an aircraft), then the horizontal location might be known reasonably well but not the altitude; again, identification of two known locations theoretically suffices to determine the missing extrinsic parameters. In practice even “known” parameters are subject to various errors and the minimum value of the cost function is > 0 . Although (20) is preferred in the sense that it deals directly with the ground coordinates of interest, the sensitivity of the map \mathcal{F} (especially for points near the horizon) makes it rather unstable and (21) is usually used, satisfactory images being obtained through either manual or a combination of manual and automatic fitting.

It was stated above that one advantage of using a camera with a zoom lens is that for each time interval, images can be taken at different focal lengths to give high resolution in at least part of the imaged region while also including information over a wide field of view. Although the advantage of such a procedure appears obvious in retrospect, it has not (to the author’s knowledge) been exploited previously. However, a new source of error is introduced. It usually is not too difficult to find a few points of known location in the wide-angle views, but is often very difficult to see anything recognizable in zoomed photos, which usually image only the water in the center of the viewed

region. Ideally, the extrinsic parameters would be the same for both images so that determination of these parameters for the wide-angle view would suffice to accurately register the zoomed image as well. However, imperfections in the mechanical construction of the camera can lead to changes in the direction of the principal axis for different focal lengths. These changes are small (around 1°) and can be perfectly stable (and repeatable), but are nevertheless important. In such cases, a further calibration can be carried out to relate the changes in bore-sight direction (or principal axis direction) at different focal lengths. The result of this calibration is a set of correction angles that are applied to extrinsic parameters obtained from the wide-angle view in order to correctly register the zoomed image.

G. Image Enhancement

Once the coordinate transformations described above are carried out, it is generally found that the surface shading features of interest are only very faintly visible in rectified images. This is because the large-scale brightness variations associated with, say, dip angle or specular solar reflection, occupy most of the dynamic range of the image. These tend to swamp the small-scale variations associated with internal features, which are often only just larger than the pixel quantization levels (in this respect, Fig. 1 is somewhat atypical).

A series of image-processing steps are carried out to enhance the detail. Although the operations required vary from deployment to deployment depending on the ambient topography, weather, and viewing conditions, they can generally be carried out automatically for all images within a given deployment. First, vegetated areas (land), if visible, are masked using a threshold based on the ratio of green to blue pixel values, land being relatively green and water relatively blue. Second, only the blue pixel plane is used (results were no different if the average of all three color planes was used). Large-scale variations in brightness are removed by subtracting various best-fit planes or parabolas from the remaining water pixels. The best-fit can be carried out over ground coordinates, or, with more physical basis, as a function of dip angle (to remove the tendency of near-horizon pixels to be brighter). An empirically determined but camera-specific parabolic function of angle from the camera central axis is used to remove so-called “vignetting” effects that result from lens construction. These effects show up as a darkening around the edges of the image. Speckle is removed by applying a 2-D median filter. Finally, a histogram equalization algorithm is used to enhance the remaining small-scale brightness variations. This equalization procedure gives the resulting images a characteristic “grainy” texture.

V. APPLICATION EXAMPLES

One quantitative use of this system in the study of a coastal front is described elsewhere [21]. However, as the technique is still somewhat experimental, it is useful to demonstrate the range of physical processes that have been observed. Two examples are discussed here. First, observations of the radial spreading of a dispersive internal wave packet in which the number of distinct wave crests doubles in about 1.5 h as the

packet propagates 4 km are shown. The existence of such radially spreading wave packets was previously unsuspected in this region. A second example shows the advection of a turbulent field of well-mixed water into deeper regions. In this case, the dominant eddy size is of the same order as water depth (≈ 250 m). Although the existence of these large-scale turbulent features was known, their characteristics were not, in spite of a great deal of standard oceanographic work in this region.

A. Internal Wave Dispersion

One set of field tests were carried out July 8–13, 1999. A single camera was situated just below the peak of a sparsely wooded hillside on San Juan Island, WA, USA, overlooking the southern part of Haro Strait at an altitude of 300 m above water level. Weather was generally good, but the air became increasingly hazy during the week. A coastline about 12 km away was always visible, but a more distant coastline some 40 km away was visible only during the first two days. View directions were to the south and west across Haro Strait and specular sun glint was present in afternoon images. During one afternoon, winds were strong enough that little spatial structure was seen, but in all other cases small-scale brightness variations were visible over the whole region, suggesting a certain robustness of the phenomenon.

During two consecutive days, animation of the rectified and contrast-enhanced sequences showed an arc-like ripple pattern propagating east-northeast across Haro Strait after low slack water. Fig. 4(a) shows a single image from the sequence, covering a region about 7 km^2 . The rectified image forms a west-facing cone, covering only a fraction of the water’s expanse. The slope of the hillside hid the near coastline of San Juan Island and a number of trees also obscured parts of the viewed region. Trees and vegetation have been masked out, resulting in a ragged right-hand edge for the cone. On the water surface, a prominent curved dark leading edge is visible, followed by several less-distinct edges (the fourth, in particular, is difficult to see in single images). Dark features appear to be followed by lighter bands, consistent with the physical processes discussed in Section II. The dark features are, thus, convergent/downwelling leading edge of internal waves. Water depth in the vicinity of the ripple is around 250 m, but the ripple itself appeared to originate just before low slack tide in the gap between two subsurface pinnacles rising to about 70 m below the surface. By extracting pixel data along the propagation axis east-northeast from a point lying between the pinnacles, the time/distance plot shown in Fig. 4(b) can be formed. Fig. 4(c) is a visual interpretation of the important features in Fig. 4(b). A variety of linear features are seen. The most prominent are a set of diagonals from upper left to lower right, which are associated with the curved ripple patterns. From this plot, the propagation speed of the leading wavefront is now easily seen to be about 0.9 m/s (tidal flow is small at this time and roughly perpendicular to the axis used). The second wavefront undergoes an interesting bifurcation near 12:30, splitting into a feature that is moving only slightly slower than the leading wavefront, a faint feature that becomes more clear later, and a third wavefront that moves significantly slower at first, but speeds up slightly after about

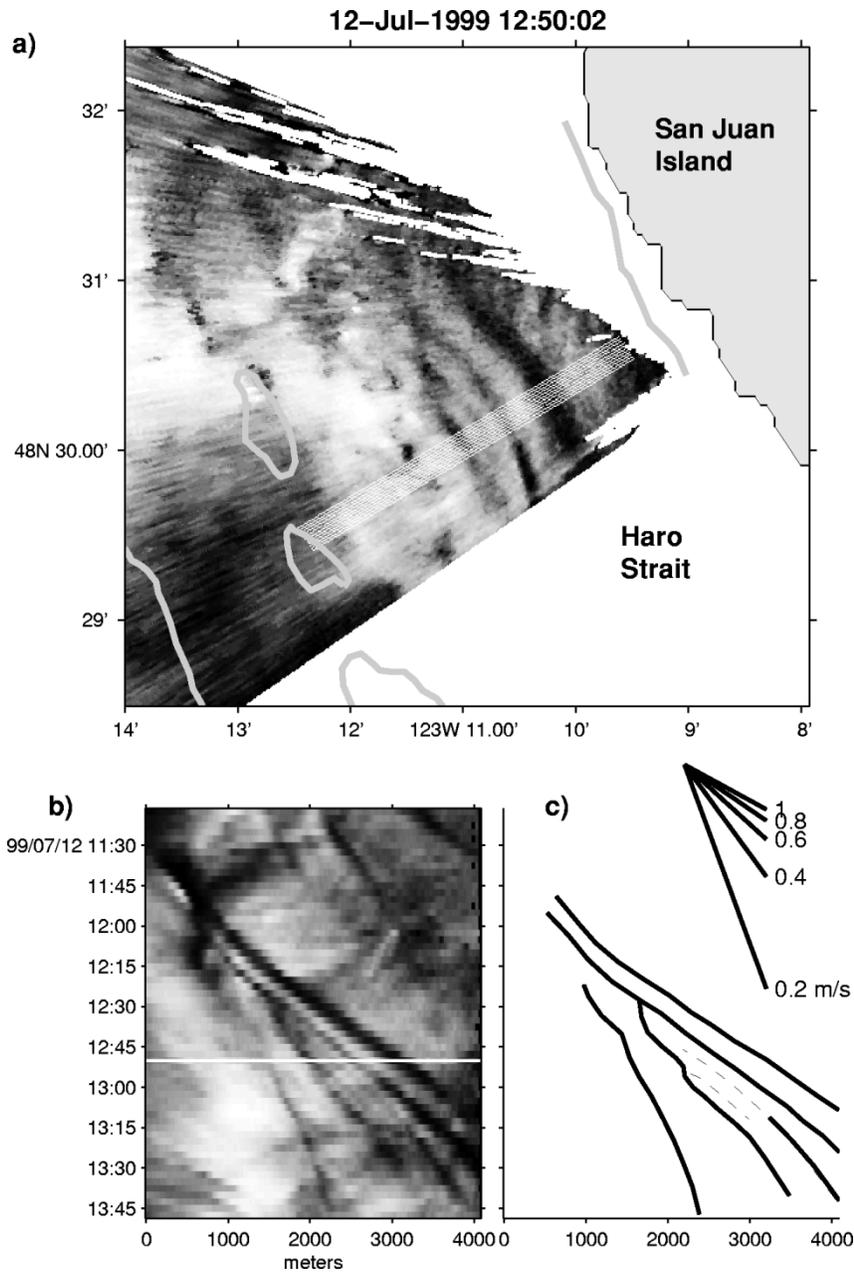


Fig. 4. (a) Image of an internal wave packet propagating east-northeast toward the coastline of San Juan Island. The area shown is slightly larger than 7×7 km. A 120-m isobath is shown in gray. (b) Time/distance plot for wave propagation along the line marked by white pixels in (a). The time of the image in (a) is marked by the thin white line. (c) Visually extracted wavefronts from (b). Note the splitting of a trailing wavefront at 12:30 and again at 13:00. Slopes corresponding to different speeds are shown at right.

12:50, bifurcating again near 13:00. A final trailing wavefront does not move uniformly, but appears to slow with time. The wavefronts are separated by 300–500 m in a water depth of about 270 m. Although no water-column measurements were made at the time, the typical summer stratification in this area can support a first baroclinic mode with a linear wave speed of about 0.7 m/s [22]. Higher order modes will have slower speeds. As nonlinear waves often travel at speeds slightly faster than linear waves, the obvious interpretation is that the first few wavefronts are nonlinear (and, hence, large amplitude) first-mode internal waves. The bifurcation thus shows the decay of a large topographically generated internal perturbation into a series of soliton-like wave crests.

B. Large Eddy Turbulence

Another set of field tests were carried out June 29–July 3, 2000, from the same location as in the previous year. To provide better spatial coverage, two cameras with overlapping fields of view were used; each camera also taking both wide and narrow-angle images at 2-min intervals. The spatial coverage is significantly better than for the sequence discussed previously [Fig. 5(a)]. Contrast enhancement was carried out for each set of images before they were combined. Although this simplifies some aspects of the processing, changes in brightness levels between processed images result in visible “cones” showing the narrow-angle images superimposed on wide-angle images.

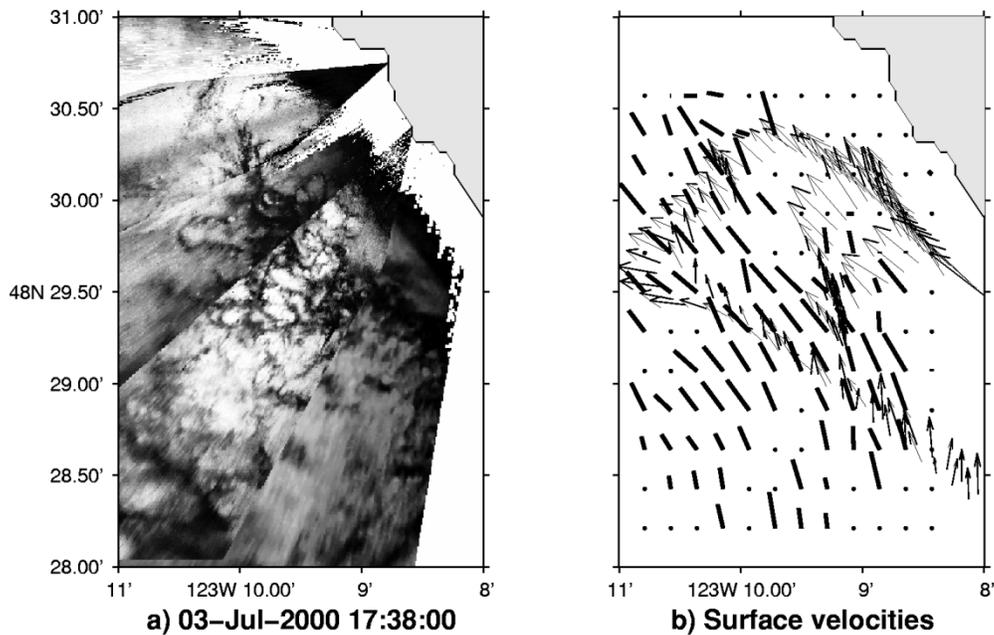


Fig. 5. (a) Image of turbulent boils propagating northward at the height of the spring flood tide. Area shown is 5.5×3.6 km. (b) Surface velocity field over 6 min, estimated by correlation analysis applied to images (thick lines) and from 2 h of ship-mounted acoustic Doppler current profiler observations (thin arrows).

Remaining uncertainty in the bore-sight calibration at different focal lengths are visible as offsets in features on either side of the cone edges. During this set of tests, water-column measurements of current were made using a 150-kHz acoustic Doppler current profiler (ADCP) mounted on a ship repeating a large rectangular track. At the height of the spring flood tide on July 3, the entire region became turbulent with large boils covering the water surface. The appearance of the water surface over a 5.5×3.6 -km region is shown in Fig. 5(a). Although the clearest view of the boils, which appear as elongated light-colored patches, is seen within a narrow cone pointing at the lower left corner (containing a high-resolution narrow-angle image from one of the cameras), they can be seen over much of the area. A typical dimension for the patches is about 200–300 m, which is similar to the water depth. It is likely that these represent large eddies that fill the entire water column. Other evidence suggests that the water column is unstratified at this time [21]. Although the eddies evolve on relatively short time scales, the animated sequences clearly show a general movement of features in a curving pattern entering from the south and leaving to the northwest. A quantitative estimate of the surface flow field can be performed using 2-D correlation analysis applied to small subblocks of the entire image. Fig. 5(b) shows estimated surface flow vectors after removal of artifacts near the edges of individual images and averaged over 6 min (i.e., three subsequent pairs of images). Also shown in the image are velocity vectors from a depth of 8 m measured from the ship during a 2-h period of steaming centered on the image time. Note that ship-measured velocities change considerably over this time span in the northwestern portion of the track. Both image-derived and ship-measured velocities show northward flow in the southern part of the region, turning northwestward in the northern region. Flow magnitudes are also very similar (in this case, speeds at maximum flood are about 1.5 m/s).

VI. CONCLUSION

A technique for observing the changing characteristics of the water surface over spatial scales of several kilometers and time scales of minutes to hours has been presented. The particular hardware configuration described has been motivated by the need for a cheap, relatively simple, robust, and portable visualization system that can be operated independently of any infrastructure with little effort. It is thus ideally suited to exploratory or opportunistic efforts or as an adjunct to more conventional observational programs. The simplicity is made possible partly by the recent availability of high-resolution digital still cameras with low-distortion lenses and a high degree of configurability. The mathematical processing of photos into rectified images is largely straightforward, although at low angles and large distances a correction must be added for the sphericity of the earth. The rectified images must themselves undergo considerable image enhancement in order to be useful. Results found without calibration of the camera (other than relative boresight calibrations at different zoom levels) to correct for lens and mechanical imperfections have been found to be adequate for many purposes. Uncertainties in the oceanographic state (e.g., the actual density profile) are often of more concern than errors in the image rectification, but this will obviously not be the case in all situations and some kind of calibration may be required in more elaborate observational programs. The system described has been used in a variety of situations, from hillsides, ship's bridges, and aircraft, and in areas accessible only by extended foot travel. Two examples of physical phenomena have been described here. In both cases, the results showed surprising large-scale coherence of relatively compact flow features. Not only can these processes be observed, but quantitative estimates can be made of parameters such as propagation/advection speed and direction. It is likely that understanding of many other coastal processes can be enhanced by use of this technique.

ACKNOWLEDGMENT

Development and testing of this system over the past few years could not have occurred without the cooperation and help of a number of people. The author would first like to thank P. Ballenger, from whose backyard the image examples discussed were recorded in 1999 and 2000. He would also like to acknowledge B. Neufeld, M. Steininger, A. Waterhouse, P. Baines, and Yuval, who helped with programming and/or manned camera stations at various times. C. Wang suggested the use of wakes to find a direction and D. Farmer first demonstrated how a camera could be used scientifically. Calibration tests of the cameras were greatly simplified by software written by J.-Y. Bouguet [19].

REFERENCES

- [1] D. K. Lynch and W. Livingston, *Color and Light in Nature*. Cambridge, U.K.: Cambridge Univ. Press, 1995.
- [2] W. A. M. Nimmo-Smith, S. A. Thorpe, and A. Graham, "Surface effects of bottom-generated turbulence in a shallow tidal sea," *Nature*, vol. 400, pp. 251–254, 1999.
- [3] W. Munk, L. Armi, K. Fischer, and F. Zachariassen, "Spirals on the sea," in *Proc. R. Soc. Lond. A, Math Phys. Sci.*, vol. 456, 2000, pp. 1217–1280.
- [4] V. Artale, D. Levi, S. Marullo, and R. Santoleri, "Analysis of nonlinear internal waves observed by landsat thematic mapper," *J. Geophys. Res.*, vol. 95, no. C9, pp. 16 065–16 073, 1990.
- [5] A. K. Liu, Y. S. Chang, M. K. Hsu, and N. K. Liang, "Evolution of nonlinear internal waves in the East and South China Seas," *J. Geophys. Res.*, vol. 103, no. C4, pp. 7995–8008, 1998.
- [6] R. A. Kropfli, L. A. Ostrovski, T. P. Stanton, E. A. Skirta, A. N. Keane, and V. Irisov, "Relationships between strong internal waves in the coastal zone and their radar and radiometric signatures," *J. Geophys. Res.*, vol. 104, no. C2, pp. 3133–3148, 1999.
- [7] T. C. Lippman and R. A. Holman, "Quantification of sand bar morphology: A video technique based on wave dissipation," *J. Geophys. Res.*, vol. 94, no. C1, pp. 995–1011, 1989.
- [8] K. T. Holland, R. A. Holman, T. C. Lippman, J. Stanley, and N. Plant, "Practical use of video imagery in nearshore oceanographic field studies," *IEEE J. Oceanic Eng.*, vol. 22, pp. 81–92, Jan. 1997.
- [9] H. F. Stockdon and R. A. Holman, "Estimation of wave phase speed and nearshore bathymetry from video imagery," *J. Geophys. Res.*, vol. 105, no. C9, pp. 22 015–22 033, 2000.
- [10] B. D. Morris, M. A. Davidson, and D. A. Huntley, "Measurements of the response of a coastal inlet using video monitoring techniques," *Mar. Geol.*, vol. 175, pp. 251–272, 2001.
- [11] J. P. Dugan, G. J. Fetzer, J. Bowden, G. J. Farrugia, J. Z. Williams, C. C. Piotrowski, K. Vierra, D. Campior, and D. N. Sitter, "Airborne optical system for remote sensing of ocean waves," *J. Atmos. Ocean. Technol.*, vol. 18, no. 7, pp. 1267–1276, 2001.
- [12] R. D. Pfeltzer, O. M. Griffin, W. R. Barger, and J. A. C. Kaiser, "High-resolution measurement of surface-active film redistribution in ship wakes," *J. Geophys. Res.*, vol. 97, no. C4, pp. 5231–5252, 1992.
- [13] H. A. Espedal, O. M. Johannessen, J. A. Johannessen, E. Dano, D. R. Lyzenga, and J. C. Knulst, "COASTWATCH'95: ERS 1/2 SAR detection of natural film on the ocean surface," *J. Geophys. Res.*, vol. 103, no. C11, pp. 24 969–24 982, 1998.
- [14] G. Ewing, "Relationship between band slicks at the surface and internal waves in the sea," *Sci.*, vol. 111, pp. 91–94, 1950.
- [15] J. A. Shand, "Internal waves in Georgia Strait," *Trans. Amer. Geophys. Union*, vol. 34, no. 6, pp. 849–856, 1953.
- [16] D. M. Farmer and L. Armi, "The generation and trapping of solitary waves over topography," *Sci.*, vol. 283, pp. 188–190, 1999.
- [17] *Manual of Photogrammetry*, 4th ed., C. C. Slama, C. Theurer, and S. W. Henriksen, Eds., Amer. Soc. Photogrammetry, Falls Church, VA, 1980.
- [18] R. Tsai, "A versatile camera calibration technique for high-accuracy 3D machine vision metrology using off-the-shelf TV cameras and lenses," *IEEE J. Robot. Automat.*, vol. RA-3, no. 4, pp. 323–344, 1987.
- [19] J. Bouguet. (2002) Camera Calibration Toolbox for MATLAB. [Online]. Available: http://www.vision.caltech.edu/bouguetj/calib_doc/
- [20] R. G. Wilson and S. A. Shafer, "What is the center of the image?," *J. Opt. Soc. Amer.*, vol. 11, no. 11, pp. 2946–2955, 1994.
- [21] R. Pawlowicz and C. Wang, "Observations of the strongly nonlinear internal tide generated by a tall sill," *J. Geophys. Res.*, 2002, submitted for publication.
- [22] R. Pawlowicz, "Observations and linear analysis of sill-generated internal tides and estuarine flow in Haro Strait," *J. Geophys. Res.*, vol. 107, 2002, C6, 3056, doi: 10.1029/2000JC000504.



Rich Pawlowicz received the B.Sc. degree in mathematics and electrical engineering from Queen's University, Kingston, ON, Canada in 1987 and the Ph.D. degree from the Massachusetts Institute of Technology, Cambridge, MA/Woods Hole Oceanographic Institute Joint Program, Woods Hole, MA, in oceanographic engineering in 1994. His Ph.D. research involved acoustic tomography in the Greenland Sea.

Between degrees, he was an Engineer for Computing Devices Canada, Ottawa, Canada. Since 1996, he has been an Assistant Professor of Oceanography with the Department of Earth and Ocean Sciences, University of British Columbia, Vancouver, Canada. His research interests include various aspects of coastal oceanography and internal wave processes. Frustration with trying to understand dynamics in these regions using standard techniques has led him to explore oblique photography as a scientific tool.