ALGORITHMS FOR SHIP DETECTION AND TRACKING USING SATELLITE IMAGERY

J.K.E. Tunaley

Directorate of Space Development,
National Defence Headquarters,
101 Colonel By Dr., Ottawa. K1A 0K2

ABSTRACT-Ship detection and tracking are important to maintain a Recognized Maritime Picture (RMP). Satellite surveillance provides wide area coverage but is often limited by revisit times and by the time taken to process the imagery, analyze it and send messages in an appropriate format to update the RMP. The paper describes the development of a new ship detection and tracking program that is designed to analyze processed imagery rapidly and inexpensively and to deliver messages automatically by email in the OTH Gold format. The message can include ship position, length, heading and speed.

The program structure is designed to be versatile, flexible and extendable and a RADARSAT image can be loaded and analyzed and a message sent typically in less than one minute.

1. INTRODUCTION

The Polar Epsilon project, which is about to enter its definition phase, will provide a wide area surveillance capability over Canada's maritime approaches out to a distance of 1000 NM. The Arctic Ocean and the northern land mass are also included. Polar Epsilon will provide a capability that includes planning and exploitation tools as well as the means to acquire imagery from commercial and civil satellite-borne sensors. The exploitation tools are being designed to extract useful information so that, unlike the raw data or imagery, this information can be propagated rapidly over low bandwidth channels.

The ship detection software, "OceanWorks" is one of the exploitation tools for the maritime environment and handles radar images though optical images can also be processed. Though R&D continues in the development of advanced ship detection techniques, the present emphasis is on the use of existing sensors and processor software as well as on the need to provide an operational capability in the near term. This constraint implies that normal multi-look radar imagery or pan-chromatic optical imagery is preferred. The satellite agencies or vendors, such as RADARSAT International, routinely distribute such imagery.

Exploitation tools are important for controlling costs. Image analysts are presently an expensive critical component in the analysis chain and must be replaced by automated tools. This implies that much preliminary analysis should be implemented immediately after image processing at the data reception facility to avoid sending high bandwidth products (often with a large time-late penalty) to a Maritime Operations Centre. Thus the data should be reduced to text messages accompanied by image chips.

Polar Epsilon will provide the means to obtain other information about the oceanic environment, such as ocean colour.

The principal sensor in Polar Epsilon will be the RADARSAT-2 Synthetic Aperture Radar (SAR). This is firstly because it is Canadian and provides assured access by the Department of National Defence and secondly because, with a launch in 2005, it will probably offer more flexibility and capability (with respect to resolution and swath width) than any other radar satellite within the next few years. Furthermore, unlike optical sensors that are often limited by small swath widths and are defeated by cloud and fog, both of which are prevalent at northern latitudes, radar satellites offer all-weather operations.

Naval requirements are for precision in the location of a ship target, its heading and speed. Ideally, ships of length greater than 30 m should be detected in sea states up to and including five. Information about the ship, which could be useful in identifying it, is also important so that the information from satellite sensors can fused with other data to form a reliable Recognized Maritime Picture (RMP). Since satellite sensor images are essentially a snap-shot of the ocean surface and other data arrives at different times, it is important to predict where a ship will be when other data arrives so that detections and additional information can be compared and combined. This association of the data will be unreliable in a high-density target environment unless some unique ship features can be extracted to identify a ship. One such feature is ship length and, for accuracy, high resolution is preferred. High resolution also yields information about the ship superstructure that can be used during target association or in classification. It should be emphasized that simple detections or uncorrelated "blobs" on an image are not particularly useful operationally.

The ship heading can be determined from the image of the ship wake, which is often present for large ships moving at service speeds [1]. Ship speed can be ascertained from the azimuthal offset between the apparent position of the ship and the line of the wake

0-7803-8742-2/04/$20.00 (C) 2004 IEEE

Polar Epsilon and the design of OceanWorks involve a number of trade-offs. Perhaps the most important consideration is the desire for high resolution for ship identification, classification and ship speed extraction. Another important consideration is the revisit time, which is closely related to the Area Coverage Rate (ACR). It is necessary to cover the maritime approaches almost entirely in the time interval that a typical ship takes to traverse them. A ship moving at service speed will take about 4 days to cover 1000 NM and this corresponds approximately to the RADARSAT-2 revisit times at mid-latitudes: at high latitudes the revisit times tend to be shorter.

High resolution is fundamentally inconsistent with wide area surveillance because of bandwidth and especially downlink-bandwidth considerations. For radar sensors, dual or full polarimetric modes that might provide enhanced detection or classification performance also tend to require high bandwidths compared to simple non-polarimetric imagery for the same ACR. Moreover, ship motion tends to smear the image in azimuth and these factors reduce the operational value of polarimetry. Therefore polarimetric modes are not currently considered in OceanWorks.

The detection of ships and of wakes involves another trade-off. Ship detection is best achieved against a low clutter background but wake structures are most visible if the return from the sea is much higher than the thermal noise. This affects the choice of polarization and the beam modes, the range of useful incidence angles and the radar resolution. For example, HH polarization is preferred for ship detection because the sea clutter tends to be low. On the other hand VV polarization is preferred for sea and wake visibility.

The effect of thermal noise is described by the Noise Equivalent Sigma Zero (NESZ). To obtain useful ship and wake detection performance in the far range requires a low NESZ. Beam spoiling may be desirable to implement ScanSAR modes, which are important for wide swaths, but spoiling tends to increase the NESZ.

The requirements and trade-offs imply that the optimal operational configuration is obtained by RADARSAT-2 without polarimetry, in a ScanSAR beam mode and with HH polarization. Because imagery from other satellites such as Envisat will also be used, OceanWorks will accept various standard image formats. OceanWorks is being developed using RADARSAT-1 imagery from the archive and from the ongoing Integrated Satellite Tracking of Oil Polluters (ISTOP) program of Environment Canada. By May 2005 it will form an operational package that will detect ships, oil spills and provide environmental data on wind and waves. Most algorithms have been developed but some need integrating into the OceanWorks software package.

2. SHIP DETECTION

Ships can be detected in Synthetic Aperture Radar (SAR) images because their returns tend to be much stronger than the returns from the sea surface. It is well-known that the sea clutter statistics are closely approximated by the K-distribution, which is described by a mean and shape parameter [2]. The mean and shape parameter can be estimated from the clutter surrounding the ship and, if a Probability of False Alarm (PFA) is specified, the distribution yields a threshold for detection. Like its predecessor, the Ocean Monitoring Workstation [3], this forms the basis of the OceanWorks detection system but, on its own, the statistical method results in far too many false alarms. Even small rates of false alarms cannot be tolerated operationally in part because large areas with swath widths of 300 km or more and lengths of up to 4000 km are contemplated. Servicing even a few false alarms using aircraft patrols is unacceptably expensive.

False alarms occur for many reasons. First, the sea surface is not homogeneous and the sea clutter statistics vary from point to point. Statistical parameters can be estimated over small patches or "tiles" of sea but, as the size of the tiles is reduced, statistical stability is impaired and bias is introduced. This limits the minimum tile size but significant changes in the sea surface may occasionally be present over a single tile. Thus the statistics are only approximate and residual bias typically results in false alarms.

Land returns are another source of false alarms, especially small islands. Most of these can be removed using a land mask but the masks are often incomplete or inaccurate. Moreover, at the present time, the geolocation of the image may not be accurate (particularly with old RADARSAT-1 images). Geolocation of RADARSAT-2 images should be excellent since this satellite will carry a GPS receiver.

False alarms can be generated by nadir ambiguities. In RADARSAT images, these occur as lines in the azimuthal direction at nearly the same ranges in every image from a given beam mode. The false alarms can be suppressed by increasing the detection threshold in the neighborhood of the ambiguity.

Fixed targets, such as oil rigs, should usually be suppressed. This is achieved by entering their coordinates into a list of targets that are to be ignored.

To reduce the actual PFA, the statistics-based ship detections are modified by the result of a segmentation process. The areas around potential targets are segmented using a fast "quad-tree" algorithm. This utilizes a top-down approach to grow segments from successive square blocks of pixels of finer and finer resolution.

The segmentation process collects and identifies all the pixels associated with the ship and their intensities provide an estimate of Radar Cross Section (RCS), as well as estimates of the ship length and beam. In the absence of a wake, the ship heading is estimated from the orientation of the ship in the image relative to true north. This estimate tends to be quite inaccurate for the standard ScanSAR resolutions, which are about 50 m. This is because there are typically only a few pixels in the ship image and because of the effect of ship motion in distorting the ship image. Segmentation also allows the clutter background to be separated from the ship and provides more estimates of the sea scattering coefficient and its variability.

To remove false targets, a set of rules is applied using fuzzy logic. For example, if the target is too large or has the incorrect shape, it tends to be rejected. The performance goal is an average of three false alarms in a complete swath of data of length 4000 km. With improved land masks, this is believed to be achievable.

The approach is different from the Analyst's Detection Support System (ADSS), which is used in the Australian Coastal Watch program. ADSS uses statistical, segmentation and other methods separately with a final decision based on voting [4].

3. HEADING AND SPEED EXTRACTION

Wakes are extracted from the image using a high pass spatial filter and implementing a Radon transform to detect bright and dark lines in the image [5,6]. The Radon transform is applied only to areas including a detected ship. Initially, transforms are implemented to determine in which of four small quadrants the wake lies. Then the area of the transform is increased to obtain high sensitivity and accuracy. Rules are applied to reduce false alarms due to naturally occurring internal waves or other linear structures on the ocean surface. For example, the azimuthal distance between wake lines and the ship must be less than a maximum.

Radar wakes are associated with different types of hydrodynamic wake. For example, the turbulent wake usually appears as a dark streak astern of the ship in a direction opposite to the heading and is very common in some radar imagery. However, there are often bright lines associated with internal waves or with other types of crest pattern, some of which are related to time varying excitations. In some cases, wake arms, which tend to lie symmetrically about the heading vector, may not be visible on one side of the wake and the rule base must be able to detect this by performing a limited classification to identify the turbulent wake and choosing this preferentially if it is sufficiently strong.

Visual assessment suggests that the wake extraction algorithm provides the ship heading to better than two degrees. The ship speed is derived from the offset in azimuth between the center-line of the wake and the centroid of the ship image. The goal is to obtain ship speed to an accuracy of about 2 kt; for RADARSAT ScanSAR Narrow (SCN) images, this corresponds to estimating the azimuth separation to an accuracy of about one pixel and should be achievable.

4. OCEANWORKS

The software for OceanWorks is written in C++ using the Microsoft Foundation Classes and the Visual.NET framework. It should be supportable for many years. An object-oriented design is used for maximum flexibility, extensibility and reliability. Several Dynamic Link Libraries have been created to allow easy updating for new satellite image formats and software expansion in general.

The program has been designed for speed of operation to minimize the end-to-end time from satellite illumination of the area of interest to the sending of target text messages automatically by email to specified recipients. A large fast Random Access Memory (e.g. 3GB) is desirable for optimal performance: this minimizes overall data loading times by keeping the entire image as well as its bitmap representations in memory.

To minimize operator involvement and the time taken to run OceanWorks, the only operations are file selection and clicking on three buttons. The algorithms do not require the operator to enter or choose any parameters. Eventually operator validation of ship detections should not be needed.

The software was recently exercised at the RADARSAT International reception facility at Gatineau. A RADARSAT-1 image was taken of the Grand Banks area; the data were received at Gatineau, processed into two SCNB image frames, which were input to OceanWorks. The first frame was processed and a target text message sent in less than 40 minutes from the illumination time. The second frame could have been processed and the message sent in less than one hour but this was not done. Further more comprehensive trials will be carried out in the next few months.

The OceanWorks screen from the second frame of the Grand Banks trial is shown in Fig. 1. The entire image is shown in a small pane on the upper right, while a part of the image is shown at full resolution on the left. The user may navigate by clicking in the right pane or by dragging the scroll bars. At the bottom of the screen, the OTH Gold text message is displayed. (The OTH Gold format will soon be replaced by the Extended OTH Gold format or "XCTC".) These messages are designed to be received by the Maritime Operations Centre Trinity in Halifax, Nova Scotia, where they will be integrated into the RMP and be accessible on CANMARNET. This is the unclassified version of the RMP and is available to Canadian Government Departments.

5. OPTIMIZING THE SYSTEM

The standard beams modes on RADARSAT-2 are expected to be similar to those on RADARSAT-1 but augmented by polarimetric modes and other experimental modes. The additional modes are not of immediate interest for wide area surveillance because of downlink bandwidth and ACR considerations. Therefore the optimal standard mode for ship detection will probably be SCNB because of the wide 300 km swath width with an acceptable, though not ideal, resolution of 50 m. SCNB is likely to be better than SCNA for ship detection because the grazing angle is less and this reduces the relative importance of clutter. However, the visibility of the sea surface is better in SCNA. Fig.1 shows a bright ship wake that extends for a distance of about 38 km behind the ship.

The beam modes for RADARSAT-1 and -2 have not been optimized for maritime surveillance. It appears to be possible to develop a new beam mode that results in greater swath width, higher resolution and lower NESZ. The new mode would be of the ScanSAR type but with more individual beams to avoid beam spoiling and to reduce the NESZ. For example, it could be based on the current fine beams augmented by additional fine beams in the far range. The downlink bandwidth would be fully utilized to maximize resolution.

Cross-cueing to other sensors is important for obtaining or confirming classification and track information. Optical satellites can provide high resolution but are unreliable because of the likelihood of cloud cover, especially at northern latitudes. Cross-cueing experiments are planned for 2004 using RADARSAT-1 to cue other satellites that can be tasked at short notice. Cross-cueing using optical satellite sensors requires reliable, accurate tracking information because the sensor field of regard is often quite small, at least in one dimension. The orbits of RADARSAT and many optical satellites are sun-synchronous but the RADARSAT orbit is about 4 hours ahead of most of the optical satellites. Therefore the tracking accuracy must be sufficient to place a ship in a box of side about 10 km after 4 hours. OceanWorks includes a capability for predicting the position of a target after some specified time interval (see dialog box in Fig. 1).

6. CONCLUSIONS

OceanWorks should provide reliable ship detection with minimal operator intervention and minimal false alarms. Optimal performance of a wide area satellite surveillance system requires not only reliable algorithms for ship detection and ship heading and speed extraction but also a consideration of the satellite radar hardware, especially its ability to image wide swaths, obtain high resolution and exhibit low NESZ. OceanWorks will soon be operational and be a significant part of Project Polar Epsilon.

REFERENCES

[1] J.K.E. Tunaley, "The estimation of ship velocity from SAR imagery", Proc. IGARSS-03, Toulouse, July 21-25, 2003.

[2] D. Blacknell, "Comparison of parameter estimators for K-distribution", IEE Proc.-Radar, sonar Navig., vol 141(1), pp 45-52, 1994.

[3] P.W. Vachon, S.J. Thomas, J. Cranton, H.R. Edel and M.D. Henschel, " Validation of Ship Detection by the RADARSAT Synthetic Aperture and the Ocean Monitoring Workstation", Canadian Journal of Remote Sensing, Vol. 26(3), pp. 200-212, June 2000.

[4] " Faster, surer and cheaper detection in wide area surveillance", Australian Defence Science, vol. 1(1), 2003.

[5] M.T. Rey, J.K. Tunaley, J.T. Folinsbee, P.A. Jahans, J.A. Dixon and M.R. Vant, "Application of Radon transform techniques to wake detection in Seasat-A SAR images", IEEE Trans. Geosci. Remote Sensing, vol 28, pp 553-560, July, 1990.

[6] K. Eldhuset, "an automatic ship and wake detection system for spaceborne SAR images in coastal regions", ", IEEE Trans. Geosci. Remote Sensing, vol 34(4), pp 1010-1019, 1996.

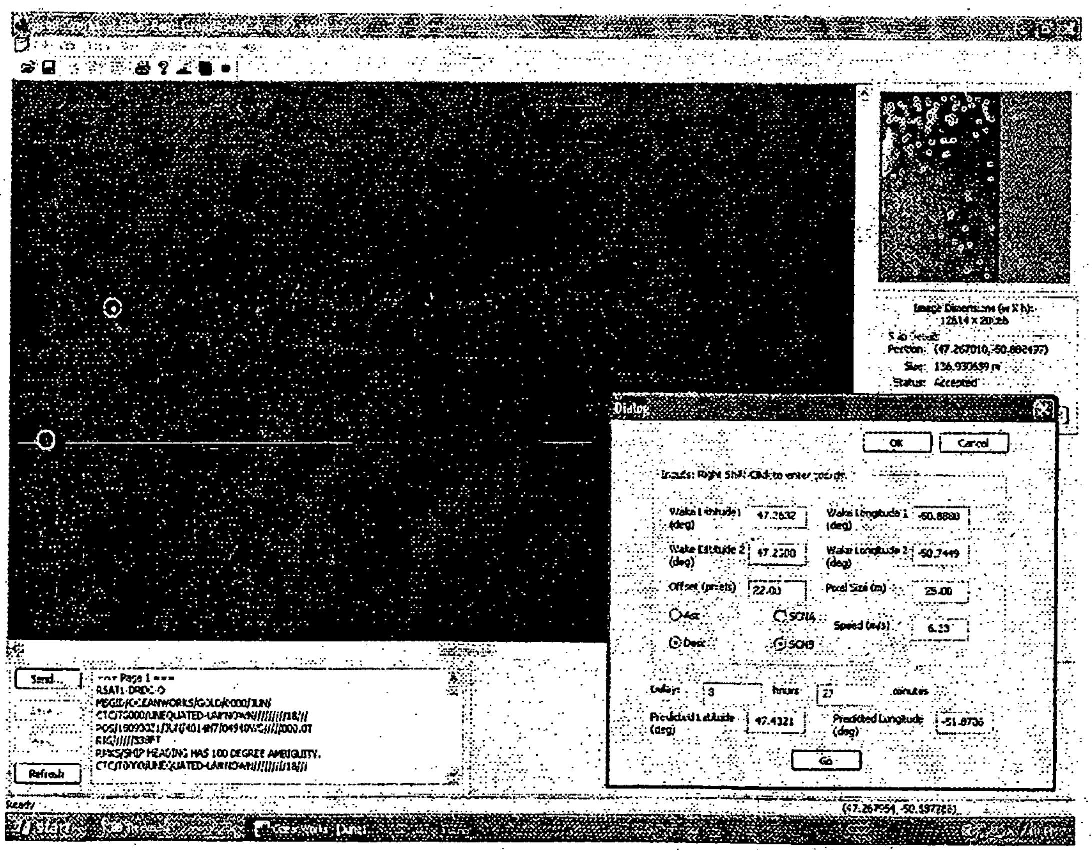

Figure 1

527374

CA 029137