Raytheon

TP 13322E

AI Radar Tracker Upgrades for SAR Target Detection

Prepared for:

Transportation Development Centre
Safety and Security
Transport Canada

by:

Raytheon Canada Limited

June 1998

Raytheon

TP 13322E

AI Radar Tracker Upgrades for SAR Target Detection

by:

Peter Scarlett
Graeme Jones
Ambighairajah Yasotharan
Joelle Pineau

Raytheon Canada Limited

June 1998

NOTICES

This report reflects the views of the authors and not necessarily those of the Transportation Development Centre.

The Transportation Development Centre does not endorse products or manufacturers. Trade or manufacturers' names appear in this report only because they are essential to its objectives.

ACKNOWLEDGMENTS

The authors thank the Transportation Development Centre for their whole-hearted support of this project. We also gratefully acknowledge the funding provided by NIF and PERD and the substantial in-kind contributions of the Canadian Coast Guard.

The support and assistance of numerous individuals made the SARAIT a practical reality. Charles Gautier of the Transportation Development Centre believed in and supported the project from its earliest days and contributed greatly to the overall system design. The large recorded data set was only made possible by the patient cooperation of Captain Terry Frost, Captain Drew McNeil and their crews on board CCGS *J. E. Bernier*. Max Johnson and Joe Ryan of Sigma Engineering provided focussed support of the MRI that made the field test and the preliminary analysis a success. The authors would also like to thank Reg Fitzgerald of Oceans Limited, who coordinated the field test, willingly answered our questions and provided all the ground truth data.

Un sommaire français se trouve avant la table des matières.

Transport Canada / Transports Canada

PUBLICATION DATA FORM

1. Transport Canada Publication No.	2. Project No.	3. Recipient's Catalogue No.
TP 13322E	9129	

4. Title and Subtitle	5. Publication Date
AI Radar Tracker Upgrades for SAR Target Detection	June 1998
	6. Performing Organization Document No.

7. Author(s)	8. Transport Canada File No.
P. Scarlett, G. Jones, A. Yasotharan, and J. Pineau	ZCD1460-362-2

9. Performing Organization Name and Address	10. PWGSC File No.
Raytheon Canada Ltd. 400 Phillip Street Waterloo, Ontario N2J 4K6	XSD-6-01871
	11. PWGSC or Transport Canada Contract No. T8200-6-6552

12. Sponsoring Agency Name and Address	13. Type of Publication and Period Covered
Transportation Development Centre (TDC) 800 René Lévesque Blvd. West Suite 600 Montreal, Quebec H3B 1X9	Final
	14. Project Officer Charles Gautier

15. Supplementary Notes (Funding programs, titles of related publications, etc.)

Co-sponsored by Program of Energy Research and Development (PERD), Department of National Defence (DND), the Canadian Coast Guard (CCG), and the New Initiatives Fund (NIF)

16. Abstract

This report describes the development and preliminary testing of the prototype Search and Rescue Artificial Intelligence Tracker (SARAIT). The SARAIT processes plots from any conventional marine radar using M of N (M target detections in N scans) integration and multiple hypothesis tracking (MHT) to detect small, awash, slowly drifting targets such as liferafts, persons in water (PIWs) and wreckage. The SARAIT is implemented on two dual-Pentium Pro single-board computers. It was operated in real time during offshore data-gathering trials and was tested with a small subset of the taped data recorded while sailing at 8 to 10 kn in 3.3 to 3.8 m seas. The SARAIT reliably detected very small PIW-sized targets at 1 to 2 nmi (depending on the clutter intensity) and small liferaft-sized targets at 2 to 3.5 nmi, all with less than 5 false detections per hour. Longer detection ranges are expected to result from the more involved testing planned for early 1999.

17. Key Words	18. Distribution Statement
Artificial intelligence (AI), search and rescue (SAR), M of N (M target detections in N scans) integration, multiple hypothesis tracking (MHT), marine radar	Limited number of copies available from the Transportation Development Centre

19. Security Classification (of this publication)	20. Security Classification (of this page)	21. Declassification (date)	22. No. of Pages	23. Price
Unclassified	Unclassified	—	xxx, 68	—

Transports Canada Transport Canada

FORMULE DE DONNÉES POUR PUBLICATION

1. N° de la publication de Transports Canada	2. N° de l'étude	3. N° de catalogue du destinataire
TP 13322E	9129	

4. Titre et sous-titre	5. Date de la publication
AI Radar Tracker Upgrades for SAR Target Detection	Juin 1998
	6. N° de document de l'organisme exécutant

7. Auteur(s)	8. N° de dossier - Transports Canada
P. Scarlett, G. Jones, A. Yasotharan, et J. Pineau	ZCD1460-362-2

9. Nom et adresse de l'organisme exécutant	10. N° de dossier - TPSGC
Raytheon Canada Ltd. 400 Phillip Street Waterloo, Ontario N2J 4K6	XSD-6-01871
	11. N° de contrat - TPSGC ou Transports Canada
	T8200-6-6552

12. Nom et adresse de l'organisme parrain	13. Genre de publication et période visée
Centre de développement des transports (CDT) 800, boul. René-Lévesque Ouest Bureau 600 Montréal (Québec) H3B 1X9	Final
	14. Agent de projet
	Charles Gautier

15. Remarques additionnelles (programmes de financement, titres de publications connexes, etc.)

Projet coparrainé par le Programme de R&D énergétiques (PRDE), le ministère de la Défense nationale (MDN), la Garde côtière canadienne (GCC) et le Fonds des nouvelles initiatives (FNI)

16. Résumé

Ce rapport décrit le développement et l'essai préliminaire d'un prototype du pisteur à intelligence artificielle de recherche et sauvetage (SARAIT, pour *Search and Rescue Artificial Intelligence Tracker*). Le SARAIT soumet les plots issus d'un radar de bord classique à deux traitements, soit l'intégration de M de N (M détections de cibles en N balayages) et le pistage à hypothèses multiples (MHT - *Multiple Hypothesis Tracking*), le but étant de détecter de petites cibles dérivant lentement à fleur d'eau, tels des radeaux de sauvetage, des personnes à la mer et des épaves. Deux ordinateurs monocarte, chacun doté de deux processeurs Pentium Pro, constituent la configuration matérielle à la base du SARAIT. Celui-ci a été exploité en temps réel pour la collecte de données en mer. Un sous-ensemble restreint de données enregistrées sur bande magnétique, alors que le navire évoluait à une vitesse de 8 à 10 noeuds dans des vagues de 3,3 à 3,8 m de hauteur, a servi à vérifier les performances du système. De fait, le SARAIT a permis de détecter de très petites cibles, de la taille d'une personne à la mer, à des distances de 1 à 2 NM (selon l'intensité du clutter), et des cibles de la grosseur de radeaux de sauvetage, à une distance de 2 à 3,5 NM, le tout avec un taux de fausses détections inférieur à cinq à l'heure. Des distances de détection plus longues sont attendues au terme des essais plus complexes prévus pour le début de 1999.

17. Mots clés	18. Diffusion
Intelligence artificielle (IA), recherche et sauvetage (SAR), intégration de M de N (M détections de cibles en N balayages), pistage à hypothèses multiples (MHT, pour multiple hypothesis tracking), radar de bord	Le Centre de développement des transports dispose d'un nombre limité d'exemplaires.

19. Classification de sécurité (de cette publication)	20. Classification de sécurité (de cette page)	21. Déclassification (date)	22. Nombre de pages	23. Prix
Non classifiée	Non classifiée	—	xxx, 68	—

Executive Summary

The prototype Search and Rescue Artificial Intelligence Tracker (SARAIT) operates in real time to process modular radar interface (MRI)[1] plots, from a re-geared 120 rpm Pathfinder Mark II radar[2], into reliable detections of small, awash, slowly drifting targets such as liferafts, persons in water (PIWs) and wreckage. The SARAIT has been operated during offshore data-gathering trials on the CCGS *J.E.Bernier* and has undergone a limited set of testing on a small subset of the extensive recorded data set. The preliminary tests while sailing at 8 to 10 kn in 3.3 to 3.8 m seas demonstrate that the SARAIT detects very small PIW-sized targets at 1 to 2 nmi (depending on the clutter intensity) and small liferaft-sized targets at 2 to 3.5 nmi, all with less than 5 false detections per hour. These are only the lower bounds and longer detection ranges are expected to result from the more involved testing planned for late 1998. The estimated cost of the SARAIT is only $23,000.

1 Requirements

Marine search and rescue (SAR) usually relies on lookouts to detect survivors because conventional marine radars cannot reliably detect small liferafts and persons in water (PIWs) under typical wave conditions. Lookouts can be very effective but poor visibility and fatigue all too frequently limit their success.

Marine radars *can* detect small SAR targets when operated at maximum sensitivity but in doing so they also detect hundreds of comparably strong reflections from wave crests and other sea or weather clutter. These "false alarm" detections collectively swamp the operator. To compound this problem, awash targets such as PIWs and liferafts are often hidden in wave troughs where they are invisible, and therefore undetectable, for several seconds at a time.

Since developing an entirely new radar for SAR would be prohibitively expensive, the Transportation Development Centre (TDC), has preferred to develop three add-on systems that together enable a normal marine radar to be used for SAR:

- Pathfinder/ST Mark II radar modified for 120 rpm (to illuminate awash targets at most wavecrests),
- Modular radar interface (MRI) (to digitally process 120 rpm radar video into thousands of plots), and
- SAR AI Tracker (SARAIT) (to correlate thousands of MRI plots into high confidence SAR tracks).

The objective of the SARAIT project is to design, build and test a processor that reliably tracks *only* the SAR targets and rejects the plethora of sea clutter detections. Since conventional marine radars do not measure the target Doppler or polarization, the most promising discriminants are that typical SAR targets:

- drift at under 1 kn, unlike sea and weather clutter, which move at 3+ kn, and
- are detectable for several minutes, unlike sea and weather clutter, which last only a few seconds.

Discriminating the SAR targets therefore requires correlating over many scans of radar data to pick out the intermittently recurring detections of the SAR targets as they move along relatively straight lines of drift pushed by the slowly changing currents, winds and waves. Measuring the speed and direction of this drift requires a sophisticated moving target tracker that can account for the usual measurement errors, ship-motion errors, missed detections and many hundreds of false (sea clutter) detections.

[1] The Modular Radar Interface is a product of Sigma Engineering, St. John's, Nfld., Canada.
[2] The Pathfinder Mark II is a product of Raytheon Marine Company, Manchester, N.H., U.S.A.

2 SARAIT Design

The SARAIT is designed to work with high scan-rate radars through Sigma Engineering's new upgraded MRI, which provides plots and displays the processed SAR target tracks. The MRI and SARAIT were developed together and since their processing requirements are similar, they use identical interchangeable hardware. Both the MRI and the SARAIT can be configured to work with most marine radars and in particular with high rotation rate radars such as the modified 120 rpm Pathfinder/ST Mark II.

The SARAIT exploits the slower drift rate of SAR targets compared to the surrounding sea and weather clutter, by correlating over many wave periods. Since SAR targets move so slowly, detections will overlay within the limits of the measurement errors and the target drift. The SAR targets are so small and so infrequently detectable that numerous false radar detections are expected. To correlate the infrequently recurring SAR detections from the much more numerous clutter and noise detections, two stages of correlation are implemented within the SARAIT:

- the M of N detector (MND) correlates scans within stationary windows that are sized to match the actual measurement error plus the growing uncertainty due to the worst-case assumed drift, and
- the multiple hypothesis tracker (MHT) which correlates over many MND updates based on the estimated target drift.

The SARAIT also implements an artificial intelligence (AI) sub-system that is not complete but will:

- identify any non-SAR tracks corresponding to the wake or rain cells, and
- automatically control the MRI and SARAIT processing for maximum performance.

The SARAIT incorporates substantial extensions to the underlying artificial intelligence tracker (AIT) that were designed and implemented in this phase:

- two dual Pentium Pro single-board computers in a single industrial PC enclosure to sustain real-time processing of 1000+ MRI plots per scan at 120 rpm,
- new MND algorithm to correlate over several wave periods and thereby increase the probability of detecting intermittently visible targets such as awash wreckage and PIWs,
- MND software module in C that can correlate more than 160 scans (80 seconds) of 1000 plot MRI data in real time,
- extensions to the MHT software module to initiate and maintain high confidence tracks on only the slow-moving SAR targets, and
- new communications software to enable the SARAIT to receive MRI plots and parameters and to send tracks and control messages.

The SARAIT software is implemented in C++ and Smalltalk and is therefore essentially platform independent.

The SARAIT is configured as two identical independent single-board computers with dedicated peripherals and linked together and to the MRI by separate Ethernet links as shown in Figure 1. Each computer is capable of being operated automatically without its monitor, keyboard and mouse. All hardware is commercial off the shelf (COTS) and interchangeable with equivalent cards, backplanes, chassis and peripherals from other vendors.

The SARAIT incorporates the following minimum hardware set plus optional equipment for use by the SARAIT operator during setup and development:

- dual passive 18 slot ISA/PCI backplane with power supplies,
- two dual Pentium Pro single board computers (Force SRC-620, 200 Mhz, 64 MB EDO RAM),
- two PCI 100T Ethernet cards (one per SBC) (3COM),
- optionally one ISA 10T Ethernet card (for MRI interface),
- optionally up to two PCI SVGA video controllers (one per SBC) (Matrox Millenium 2MB PCI),
- optionally one Exabyte Mammoth tape drive (for playback),
- optionally up to two each 3.5" floppy drives, keyboards, mice and monitors (one per SBC), and
- one rack mountable metal enclosure 0.4 m (H) by 0.5 m (W) by 0.5 m (D) and weighing under 50 kg.

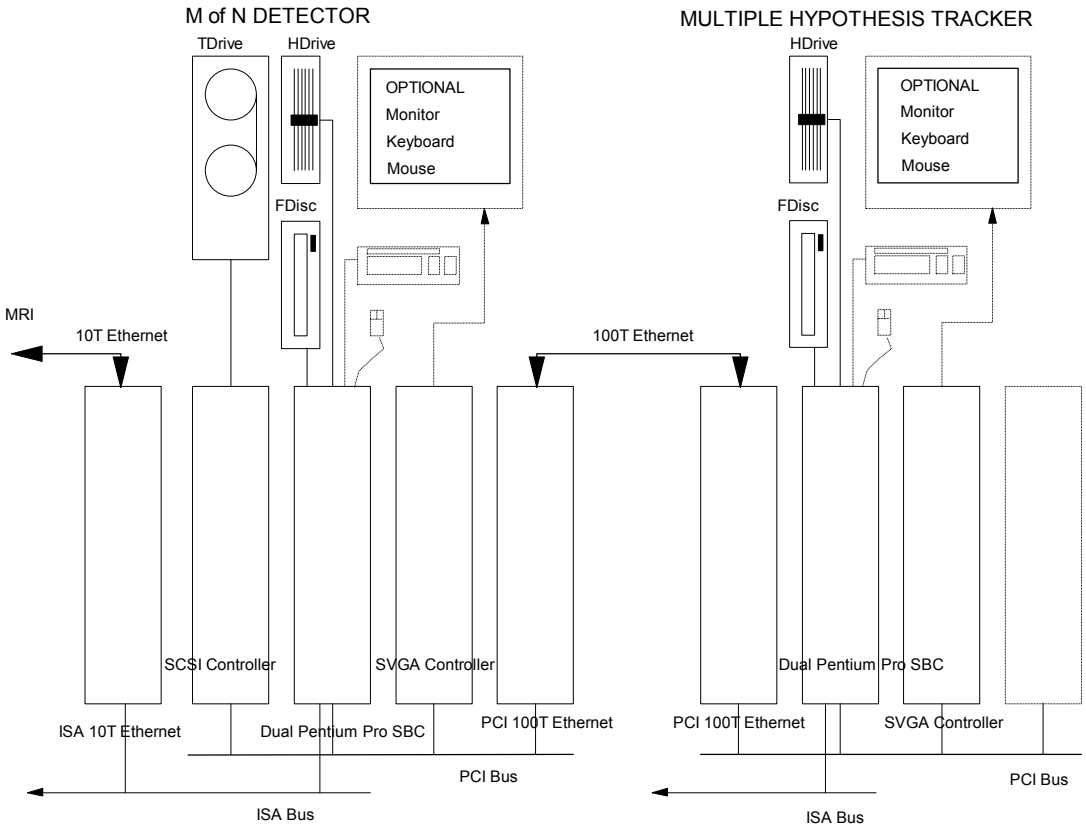

Figure 1 Hardware block diagram

3 Performance Trade-offs

To detect the smallest targets, such as a PIW, at maximum range, the radar detector must be operated at maximum sensitivity, which necessarily entails detecting numerous noise and sea clutter features. The maximum detection range for a given target depends on the target reflectivity, the sea clutter, the underlying swell and the combined sea conditions. The swell is important because SAR targets are often awash and, however high their radar cross section, are frequently hidden by the wave crests. The longer the range, the lower the target visibility. Targets along the direction of the swell will typically be detectable only 10 to 25% of the time. Across the swell, targets are detectable 25 to 40% of the time. The SARAIT correlates across several wave periods using an MND that detects infrequently recurring detections that are more likely to be real targets.

Typically, there are so many false plots that random associations raise the MND false alarm rate. Sea and rain clutter are more important than receiver noise because they are often correlated over a few seconds and therefore, for several scans, appear more like real SAR targets. At maximum sensitivity, the MRI false alarm rate is typically reduced 100-fold by the MND but is still well above the desired rate of 5 to 30 false alarms per hour

The multiple hypothesis tracker (MHT) is therefore used to correlate sequences of MND detections into high confidence tracks. The MHT algorithm is a particularly effective detector of spatial correlation in high false alarm conditions because it retains several candidate interpretations of both pre- and post-detection tracks. Tracks are confirmed as high confidence detections after the evidence has accumulated over many scans. The MND false alarm rate is typically reduced 1000-fold by the MHT to under 20 false alarms per hour. The AI subsystem is not yet complete but will identify non-SAR wake and rain cell tracks to further reduce the false alarm rate and automatically adapt the MRI and SARAIT processing parameters to the

prevailing conditions. The final result when operating at maximum sensitivity is 2 to 10 SARAIT false alarms per hour.

4 Preliminary Tests in 3.8 and 3.3 m Seas

The SARAIT has been tested on two small sets of data gathered in 3.8 and 3.3 m seas, the first data set has very little clutter and the second a great deal. For these preliminary tests, the measure of performance is the maximum detection range for each class of target balanced by an operationally acceptable SARAIT false alarm rate.

The testing on these two data sets probed the interrelationships among the many processing parameters but there was insufficient time to explore the performance limits. Determining these limits under the full range of sea states and other conditions must wait for the next phase of focussed testing in early 1999 on the extensive data set gathered off Newfoundland. **The measured performance is therefore a lower bound and does *not* reflect the achievable SARAIT performance.**

4.1 Low Clutter Performance

Tape 22 (Segment 4) was recorded in low clutter conditions with a heavy swell running. The wave rider buoy measured a 3.8 m 9.8 second period swell from 040 degrees. Light 3 kn winds from 140 degrees superimposed small 0.1 m wind driven waves on top of the swell. Most of the clutter arose from small rain cells slowly drifting through the test area and from sea birds alighting on the gentle seas. This data set is ideal for exploring how best to detect distant awash SAR targets in heavy seas that largely obscure them but introduce little competing clutter.

A series of tests were run to explore the effects of the key MRI and SARAIT processing parameters, particularly at long ranges. Figure 2 shows a typical test run with the SAR detections overlaid on the ground truth array positions. All of the array targets are detected by the time the ship has reached the southeast corner of the array. Note that by Tape 22 the array has been in place for over two weeks and several of the larger targets have dragged their moorings.

Figure 3 shows the range and bearing at which the SARAIT initiated tracks on each class of array target. Note the long ranges at which detections occur particularly the wave rider buoy 3.6 nmi to the north and the most distant A0 target 2.5 nmi to the northwest.

Floating birds are detected because of their slow drift speed. Rain cells were also seen as dense clusters of MND detections because of their slow drift in the prevailing 3 kn winds. The MHT eventually initiated several tracks as shown in Figure 2. In practice, the AI sub-system in the SARAIT will be easily adapted to reject such tracks based on the extremely high MND detection density.

Under low clutter conditions most false alarms are due to relatively uncorrelated receiver noise and sea spikes so the MND can detect even the least visible SAR targets by correlating over several swell periods. Because of this, the maximum SARAIT detection range is approximately the same at all bearings. Under these conditions, the SARAIT processing is particularly robust and insensitive to the exact choice of processing.

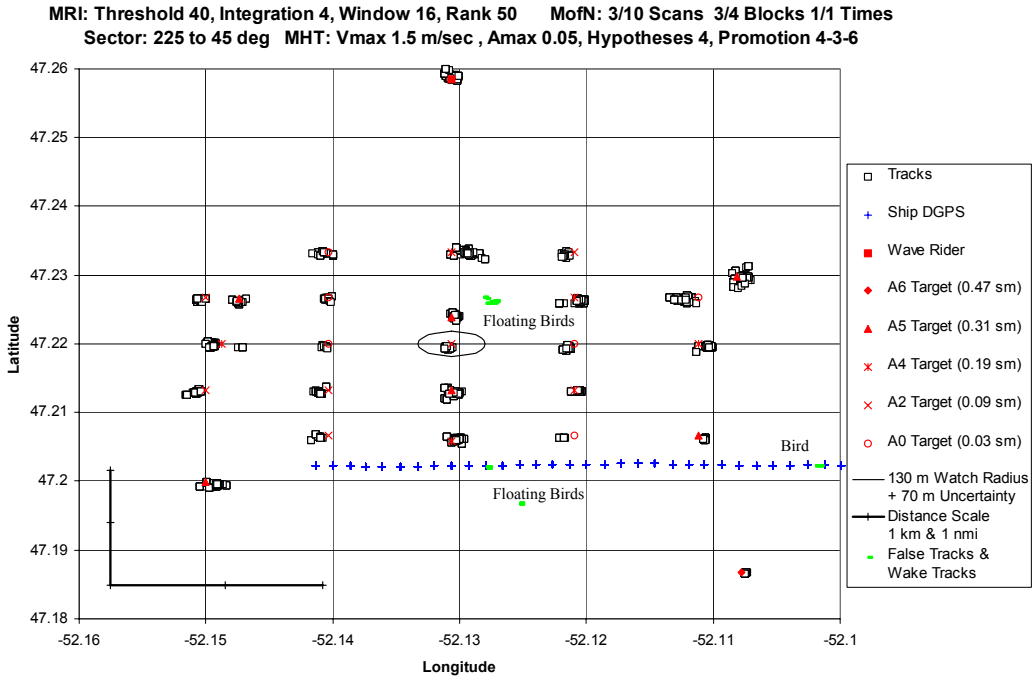

Figure 2 SAR target tracks superimposed on the array ground truth (Low Clutter)

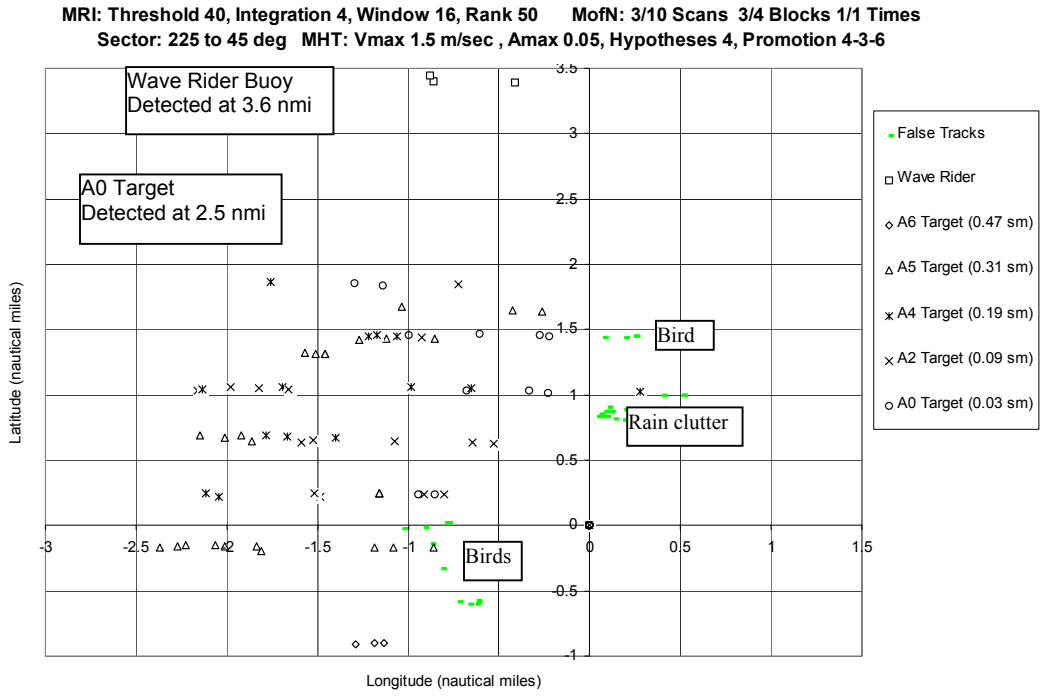

Figure 3 SAR target track initiation, range and bearing from the ship (Low Clutter)

Figure 4 demonstrates that the wave rider buoy is consistently detected at 3.5 nmi whether 2, 4 or 8 scans are integrated in the MRI. Similarly, the SAR array targets are all detected at around 2 to 2.5 nmi

irrespective of the processing or their size. This confirms that target visibility is the main limitation under low clutter conditions such as these.

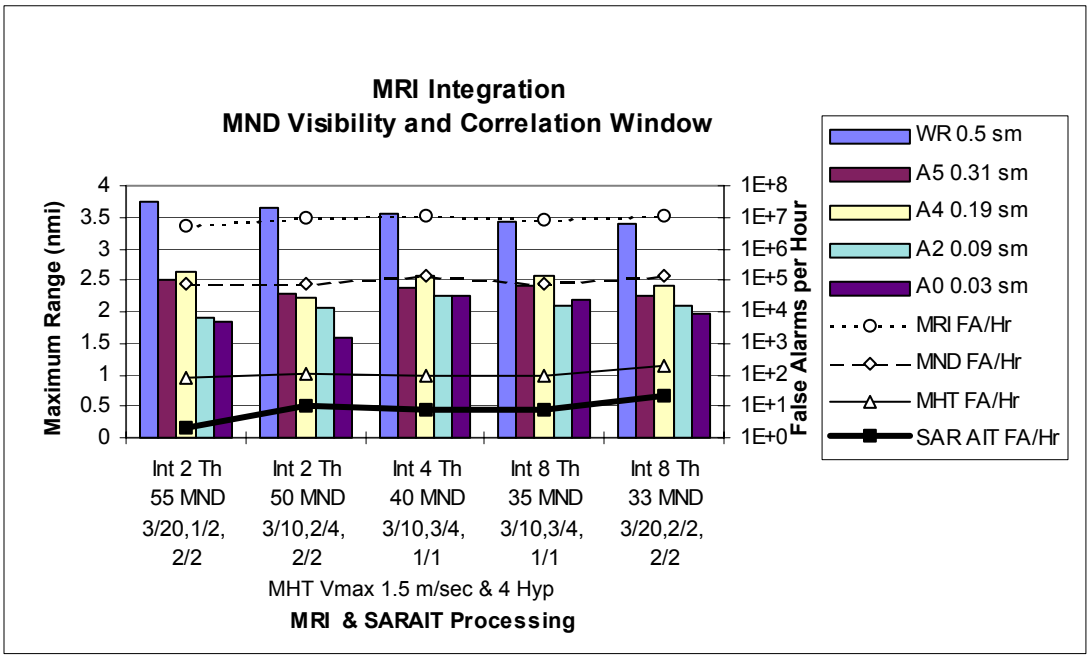

Figure 4 Typical SARAIT performance in 3.8 m seas and low clutter

Reasonable detection ranges with 2 to 20 SARAIT false tracks per hour are therefore achieved under these low clutter conditions by:

- setting the MRI threshold for 750 to 1500 plots per scan (Rfa = 1e7),
- with an MND ratio from 7.5 to 25% giving 200 to 500 MND detections per interval (Rfa = 1e5/hr),
- with an inner MND window less than or equal to the average wave period (10 seconds or 20 scans),
- with an overall MND window spanning several wave periods, and
- with a maximum track velocity of 1.5 m/sec (3 kn) to allow for sudden wave driven motion.

4.2 High Clutter Performance

The SARAIT can also reliably detect awash SAR targets in severe sea clutter and similarly high 3.3 m waves albeit at shorter ranges and not into the face of steep wind-driven waves where the clutter is most intense. In high winds, the sea clutter varies enormously in range and azimuth while the MRI settings are global. The MRI must therefore compromise between dense detections close-in and upwind and sparser detections at longer ranges and downwind. Nonetheless, the wave rider buoy and the A6 and A5 targets are typically detected at 1.5 nmi and the A0 targets at 0.5 to 1 nmi. Longer range detections are expected in the next phase once the clutter dependencies of the MRI and SARAIT processing are explored.

Tape 21 (Segment 1) was recorded three hours earlier, than Tape 22 (Segment 4), when winds were still blowing gustily at 10 kn and veering from around 280 degrees at the start of the record to 0 degrees at the end. The wind-blown seas were therefore somewhat confused and the sea clutter intense. The underlying 3.2 m swell was again from 40 degrees but had a slightly shorter 8.6 sec period.

Detecting the SAR targets under these conditions was particularly difficult because the wind was blowing strongly across the swell and therefore along the troughs where target visibility would normally be highest. Figure 5 shows the SAR target and false alarm tracks for a typical search with the ship positions (for each track update) overlaid on the true target positions. The array targets are observed from most bearings as the

ship sails 3/4 of a complete circuit in this segment.

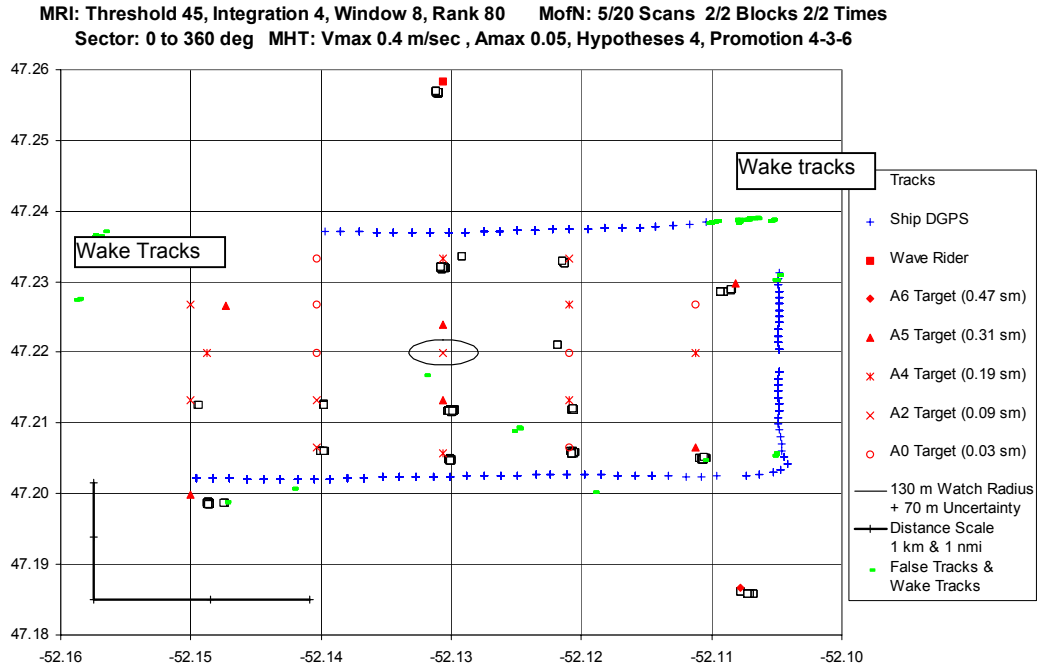

Figure 5 SAR target tracks superimposed on the array ground truth (High Clutter)

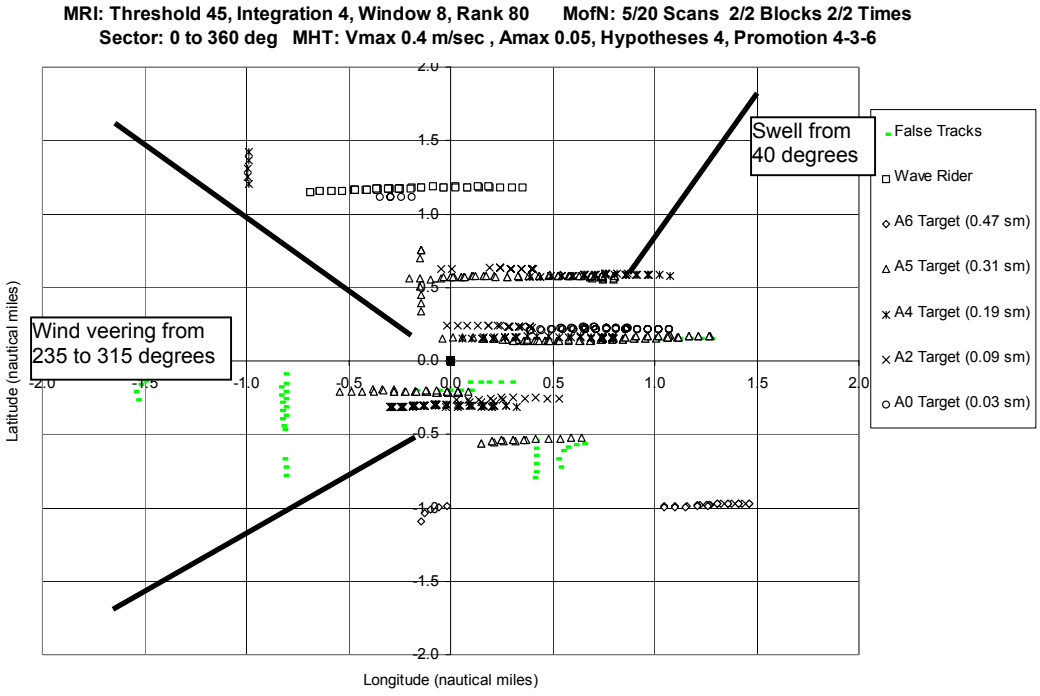

Figure 6 SAR target tracks, range and bearing from the ship (High Clutter)

Figure 6 shows how the detection range varies in azimuth, reaching a maximum of 1.7 nmi midway between the swell and wind direction, a minimum looking into the steep-faced (since still building) wind-driven waves and between 1 and 1.4 nmi elsewhere.

Note also, on Figure 5, the presence of wake tracks, as distant as 1 nmi, along the path of the ship. Figure 6 shows how the (non-wake) false alarms are mostly upwind with the balance downwind.

There are many ways to set up the system to detect the SAR targets in severe clutter with operationally acceptable false alarm rates. Figure 7 illustrates several examples of successful high clutter processing that expose the broad trends that will lead to the best possible performance in the next phase:

- MRI integration over the number of scans targets will be at each wave crest (2 to 4 in this case),
- MRI CFAR window set to 8 for lowest Rfa overall and to 4 for maximum small target range,
- MRI CFAR rank of 30 for nearby small targets to 80 for larger distant targets,
- MND inner window matched to the expected target visibility (10 to 25% depending on range, size and azimuth) and the wave period (20 in this case),
- total MND correlation interval of 2 to 6 wave periods (40 to 120 scans in this case),
- MHT hypotheses set to 4 or more in high clutter environments,
- MHT and MND maximum velocity reduced to 0.5 m/sec (i.e. 1 kn) to limit MHT false alarms, and
- AI rejection of obvious wake tracks.

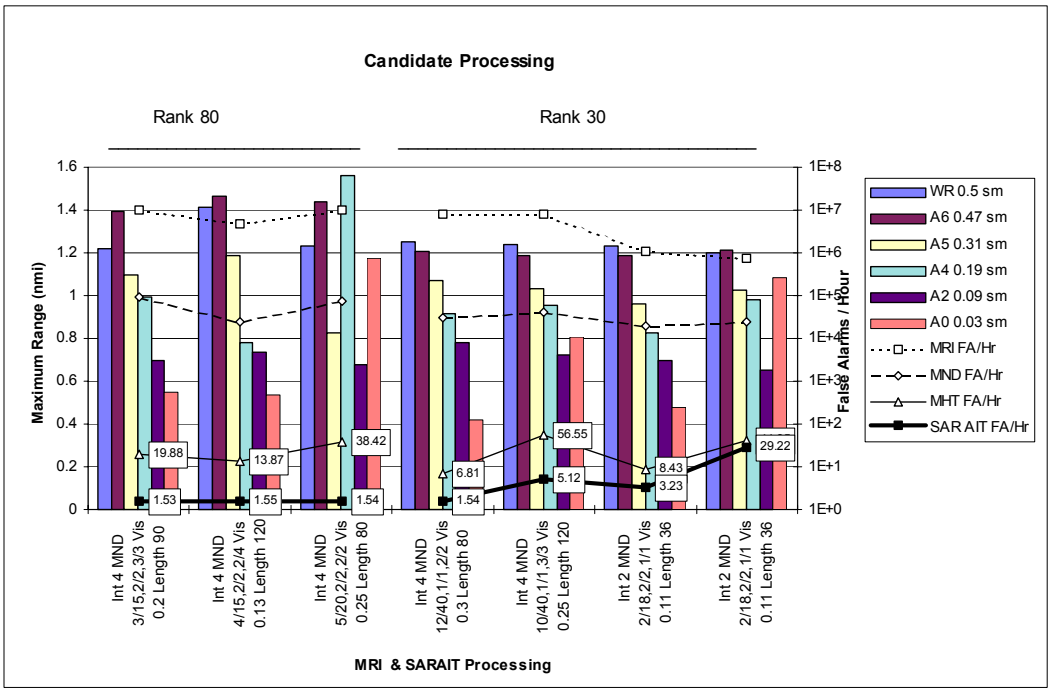

Figure 7 Typical SARAIT performance in 3.3 m seas and severe clutter

5 Conclusions

While the best processing schemes remain to be defined for a diversity of sea and weather conditions in the next phase, an encouraging lower bound on the SARAIT performance in high sea states has been set that can only improve with testing and experience. Moreover, globally setting the processing parameters, as at present, forces compromises that unnecessarily limit the SARAIT performance. For best performance, the MRI and SARAIT processing need to vary with the local clutter conditions.

The SARAIT has therefore reliably detected PIW-sized 0.03 m^2 targets in 3.2 to 3.5 m waves at 2 nmi in low clutter and 1 nmi in high clutter with 2 to 10 false tracks per hour. Small liferaft-sized 0.5 m^2 targets are detected at 3.5 and 2 nmi depending on the clutter level.

The SARAIT can maintain this performance in real time with a 120 rpm radar and 500 to 1000 MRI plots per scan, depending on the plot density.

The SARAIT processed recorded data from a conventional marine radar (re-geared to rotate at 120 rpm) on a Canadian Coast Guard buoy tender sailing an 8 to 10 kn search pattern as is typical for marine SAR in Canadian waters. Since conventional marine radars cannot usually detect such small awash targets under comparable conditions, the SARAIT has been demonstrated to deliver an operationally significant advance that should greatly improve the effectiveness of marine SAR.

Sommaire

Le prototype du pisteur à intelligence artificielle de recherche et sauvetage (SARAIT, pour *Search and Rescue Artificial Intelligence Tracker*) traite en temps réel les plots générés par l'interface modulaire radar (IMR)[1] à partir des échos reçus par un radar Pathfinder Mark II[2] perfectionné, dont l'antenne tourne à 120 tr/min., pour détecter de manière sûre de petites cibles - radeaux de sauvetage, personnes à la mer, épaves - dérivant lentement à fleur d'eau. Le SARAIT a été mis à l'essai lors d'une campagne de collecte de données en mer, à bord du navire *J.E.Bernier* de la Garde côtière canadienne, et ses performances ont été sommairement contrôlées au moyen d'un sous-ensemble restreint de données tiré de la masse de données enregistrées. Les essais préliminaires, réalisés à une vitesse de 8 à 10 noeuds, dans des vagues de 3,3 à 3,8 m, ont démontré la capacité du SARAIT de détecter de très petites cibles, de la taille d'une personne à la mer, à une distance de 1 à 2 NM (selon l'intensité du clutter) et des cibles de la grosseur de petits radeaux de sauvetage, à une distance de 2 à 3,5 NM , le tout avec moins de cinq fausses détections à l'heure. Il ne faut voir dans ces chiffres qu'une indication des performances minimales du système : des distances de détection plus longues sont attendues au terme des essais plus complexes prévus pour la fin de 1998. À noter que le coût estimatif du SARAIT n'est que de 23 000 dollars.

1 Besoins

Lors d'opérations de recherche-sauvetage en mer (SAR), on compte sur des vigies pour détecter les survivants, parce qu'on ne peut se fier aux radars de bord classiques pour détecter des petits radeaux de sauvetage et des personnes à la mer, en raison des états de mer dans lesquels s'effectuent habituellement ces recherches. Or, tout efficaces que puissent être les vigies, la visibilité réduite et la fatigue nuisent trop souvent à leur travail.

Bien sûr, les radars de bord *peuvent* détecter des petites cibles SAR, lorsqu'ils sont exploités à leur degré de sensibilité maximal. Mais ce mode de fonctionnement les amène à détecter en même temps des centaines d'échos relativement puissants provenant des crêtes de vagues, du clutter de mer et d'autres réflexions parasites. Le radariste est littéralement submergé de détections qui se révèlent être de «fausses alarmes». Et, loin d'arranger les choses, les cibles à fleur d'eau, comme les personnes à la mer et les radeaux de sauvetage, s'enfoncent souvent dans le creux des vagues, où elles deviennent invisibles, donc indétectables, pendant des intervalles de plusieurs secondes.

Compte tenu des coûts prohibitifs qu'exigerait le développement d'un tout nouveau radar expressément conçu pour les opérations SAR, le Centre de développement des transports (CDT) a choisi d'appuyer le développement de trois systèmes complémentaires qui permettent l'utilisation d'un radar de bord ordinaire pour les opérations de recherche-sauvetage. Ces trois systèmes sont :

- un radar Pathfinder/ST Mark II modifié pour tourner à 120 tr/min. (pour éclairer les cibles à fleur d'eau sur la plupart des crêtes de vagues);
- une interface modulaire radar (IMR) (pour traiter et numériser l'imagerie radar à 120 tr/min. de façon à générer des milliers de plots);
- le pisteur IA SAR (SARAIT) (pour établir des correspondances entre les milliers de plots radar produits par l'IMR, de façon à générer des pistes SAR de haute fiabilité).

Le projet SARAIT vise la conception, la construction et la mise à l'essai d'un processeur qui poursuit *uniquement* les cibles SAR, rejetant la pléthore de détections associées au clutter de mer. Comme les radars de bord classiques ne comportent pas de composante Doppler et ne peuvent mettre en oeuvre la

[1] L'interface modulaire radar est un produit de Sigma Engineering, St. John's, Terre-Neuve, Canada.
[2] Le Pathfinder Mark II est un produit de Raytheon Marine Company, Manchester, N.H., É.-U.

discrimination par polarisation, il faut se tourner vers d'autres facteurs discriminants, dont les plus prometteurs sont le fait que les cibles SAR types :

- dérivent à moins de 1 noeud, comparativement aux clutters de mer et d'intempéries, qui se déplacent à des vitesses d'au moins 3 noeuds;
- demeurent détectables pendant plusieurs minutes, contrairement aux clutters de mer et d'intempéries, qui ne durent que quelques secondes.

Pour discerner les cibles SAR, il faut donc établir des correspondances entre les données radar issues de nombreux balayages, afin d'extraire les détections intermittentes mais récurrentes de cibles SAR, lesquelles dérivent selon des trajectoires relativement rectilignes, définies par des courants, des vents et des vagues dont les régimes se modifient lentement. Mesurer la vitesse et la direction de la dérive des cibles nécessite un système complexe de pistage de cibles mobiles qui tient compte de l'erreur de mesure habituelle, de l'erreur due au déplacement du navire, des détections manquées et de plusieurs centaines de fausses détections (clutter de mer).

2 Schéma de principe du SARAIT

Le SARAIT est conçu pour fonctionner en association avec des radars à cadence de balayage élevée. Il traite les plots radar qui lui sont relayés par la nouvelle IMR améliorée de Sigma Engineering et affiche les pistes de cibles SAR. Issus de travaux de développement menés en parallèle, l'IMR et le SARAIT obéissent à des impératifs de traitement semblables, et utilisent donc des composants identiques et interchangeables. Tant l'IMR que le SARAIT peuvent adopter une configuration qui permet de les relier à la plupart des radars de bord et particulièrement aux radars à grande vitesse de rotation, comme le Pathfinder/ST Mark II modifié, qui tourne à 120 tr/min.

Le SARAIT exploite le fait que la vitesse de dérive des cibles SAR est inférieure à la vitesse de déplacement des clutters de mer et d'intempéries dans lesquels elles sont noyées, en établissant des corrélations entre les signaux issus de balayages couvrant plusieurs périodes de vagues. Le mouvement des cibles SAR est tellement lent qu'il en résulte une superposition de leurs échos, une fois qu'on a compensé pour l'erreur de mesure et la dérive de la cible. De plus, les cibles SAR sont tellement petites et si rarement détectables qu'on peut s'attendre à de nombreuses fausses détections radar. Pour distinguer les détections récurrentes mais peu fréquentes de cibles SAR des détections beaucoup plus nombreuses de clutter et d'autres échos parasites, le SARAIT procède en deux étapes de corrélation :

- le détecteur à intégration M sur N (DMN) établit des corrélations entre les balayages effectués à l'intérieur de fenêtres de distance fixes, déterminées en fonction de l'erreur de mesure réelle et de l'incertitude croissante due à la dérive supposée dans le pire cas;
- le pisteur à hypothèses multiples (MHT) établit à son tour des corrélations entre les nombreuses mises à jour du détecteur DMN en vue de prévoir la dérive de la cible.

Le SARAIT utilise également un sous-système à intelligence artificielle (encore en développement) conçu pour :

- repérer toutes les «fausses» pistes correspondant à un sillage ou à une cellule de pluie et non à une cible SAR;
- commander automatiquement les paramètres de traitement de l'IMR et du SARAIT de façon à en maximiser les performances.

La présente phase de développement du SARAIT a été l'occasion de procéder à des ajouts majeurs au pisteur à intelligence artificielle (pisteur IA) :

- deux ordinateurs monocarte, chacun doté de deux processeurs Pentium Pro, réunis dans un boîtier de PC industriel, pour le traitement en temps réel des plus de 1 000 plots radar fournis par l'IMR à chacun des 120 balayages effectués chaque minute;
- un nouvel algorithme DMN, pour la corrélation des signaux correspondant à plusieurs périodes de vagues, ce qui permet d'augmenter la probabilité de détection de cibles visibles par intermittence, telles des épaves et des personnes à la mer dérivant à fleur d'eau;

- un module logiciel DMN en langage C, capable d'associer en temps réel les plus de 160 balayages (en 80 secondes) de 1 000 plots de l'IMR;
- un module logiciel MHT (pistage à hypothèses multiples) perfectionné, pour la reconnaissance et la poursuite de pistes fiables, correspondant uniquement à des cibles SAR à dérive lente;
- un nouveau logiciel de communications, pour la réception par le SARAIT des plots de l'IMR et des paramètres de traitement, et la transmission par ce dernier des pistes initiées et des messages de commande.

Le logiciel du SARAIT utilise les langages C++ et Smalltalk et est donc indépendant, pour l'essentiel, du système d'exploitation.

Le SARAIT est mis en oeuvre sous la forme de deux ordinateurs monocarte identiques et autonomes, reliés à des périphériques spécialisés et communiquant entre eux et avec l'IMR via des liaisons Ethernet distinctes (voir la Figure 1). Chaque ordinateur peut être exploité automatiquement, sans moniteur, clavier ni souris.

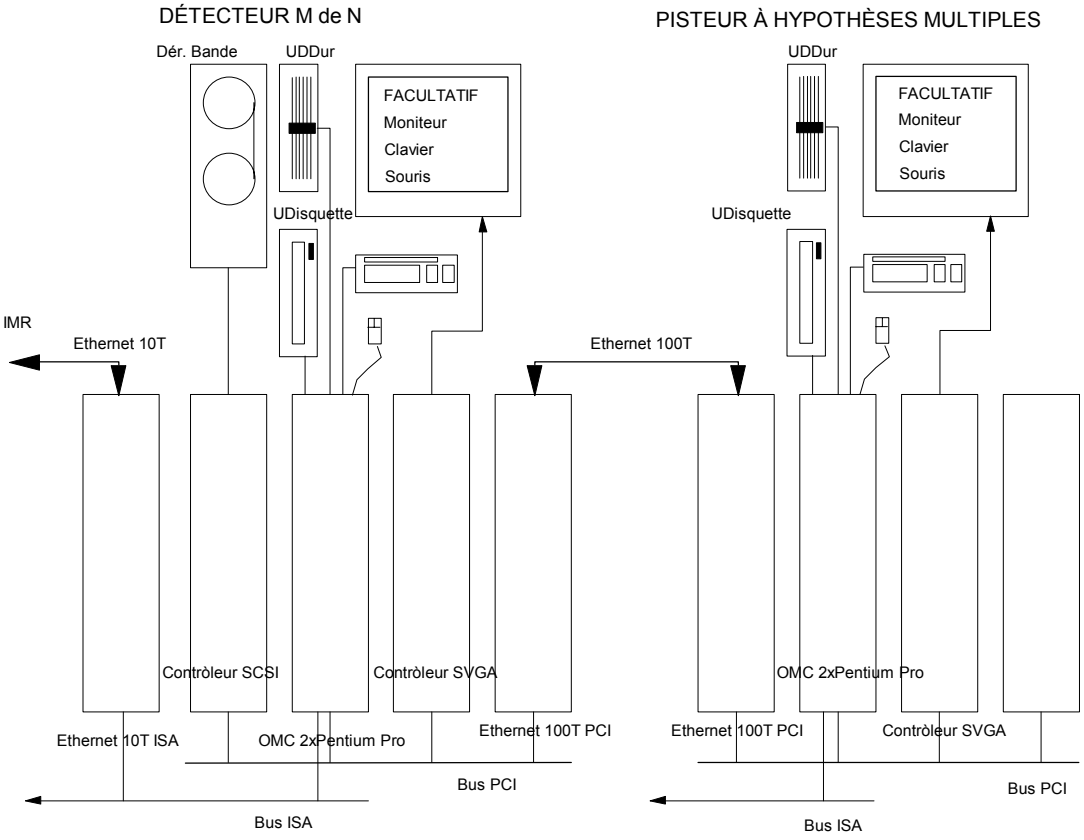

Figure 1 Schéma des composants

Le système utilise uniquement des composants du commerce qui sont interchangeables, c'est-à-dire qu'ils peuvent être remplacés par des cartes, fonds de panier, châssis et périphériques d'autres fournisseurs.

L'ensemble de composants ci-après représente la version minimale du SARAIT, complétée d'équipements facultatifs utilisés pour le montage et le développement du SARAIT :

- fond de panier passif double à 18 fentes ISA/PCI avec deux blocs d'alimentation;
- deux ordinateurs monocarte, chacun doté de deux processeurs Pentium Pro (Force SRC-620, 200 Mhz, mémoire vive EDO de 64 Mo);
- deux cartes Ethernet 100T PCI (une par ordinateur monocarte) (3COM);
- facultatif - une carte Ethernet 10T ISA (pour interfaçage avec l'IMR);
- facultatif - une ou deux cartes vidéo SVGA PCI (une par ordinateur monocarte) (Matrox Millenium 2 Mo PCI);
- facultatif - un dérouleur de bande Exabyte Mammoth (pour lecture en différé);
- facultatif - un ou deux de chacun des éléments suivants : lecteur de disquettes 3,5 po, clavier, souris et moniteur (un par ordinateur monocarte);
- un boîtier métallique pouvant être monté sur bâti de 0,4 m de hauteur par 0,5 m de largeur par 0,5 m de profondeur et pesant moins de 50 kg.

3 Arbitrage des performances

Pour détecter de très petites cibles, comme des personnes à la mer, à une distance maximale, le détecteur radar doit être exploité à son degré de sensibilité maximal, ce qui entraîne nécessairement la détection du clutter de mer et de nombreuses autres réflexions parasites. La distance de détection maximale d'une cible donnée dépend de la réflexibilité de celle-ci, du clutter de mer, de la houle de fond et des conditions de mer. Le facteur houle est important du fait que les cibles SAR sont souvent à fleur d'eau et que, peu importe qu'elles aient une grande surface équivalente, elles sont souvent cachées par les vagues. Et plus la cible est éloignée, moins elle est visible. Le taux de détection des cibles qui se déplacent dans le sens de la houle n'est habituellement que de 10 à 25 %. Dans une houle de travers, le taux de détection atteint de 25 à 40 %. Le SARAIT traite des signaux couvrant plusieurs périodes de vagues, grâce à un détecteur dit «DMN», qui détecte les images récurrentes mais peu fréquentes, qui ont plus de chances de constituer des cibles réelles.

Généralement, les «faux» plots radar sont si nombreux que les associations aléatoires entraînent un taux élevé de fausses alarmes lors du traitement DMN. Les échos parasites provenant des clutters de mer et de pluie sont plus importants que le bruit du récepteur, du fait qu'ils sont associés plusieurs fois en quelques secondes : au bout de plusieurs balayages, ils se confondent davantage (que le bruit du récepteur) à des cibles SAR réelles. Lorsque le radar est réglé à sa sensibilité maximale, la détection DMN diminue de cent fois, habituellement, le taux de fausses alarmes de l'IMR, mais ce taux reste supérieur au taux visé de 5 à 30 fausses alarmes à l'heure.

D'où l'intervention du pisteur à hypothèses multiples (MHT), qui établit des correspondances entre les séquences de détections du DMN, de façon à générer des pistes haute fiabilité. L'algorithme MHT est un détecteur de corrélations spatiales particulièrement efficace dans un milieu où le taux de fausses alarmes est élevé, parce qu'il garde ouvertes plusieurs hypothèses de pistes, même après en avoir retenu une. Ainsi, les pistes haute fiabilité ne sont confirmées comme telles qu'après que leur bien-fondé ait été appuyé par de nombreux balayages. Le MHT réduit généralement de mille fois le taux de fausses alarmes associé au DMN, pour l'amener à moins de 20 à l'heure. Le développement du sous-système à intelligence artificielle n'est pas encore terminé. À terme, il sera en mesure de repérer les pistes correspondant à des retours de sillage et de pluie et de diminuer encore le taux de fausses alarmes, et d'adapter automatiquement les paramètres de traitement de l'IMR et du SARAIT aux conditions environnementales, notamment de clutter. Le taux de fausses alarmes du SARAIT, lorsque exploité avec un radar réglé à sa sensibilité maximale, sera alors de 2 à 10 à l'heure.

4 Essais préliminaires dans des mers de 3,8 et 3,3 m

Le SARAIT a été mis à l'essai sur deux petits ensembles de données recueillies dans des mers de 3,8 et 3,3 m. Le premier correspondait à des conditions de très faible clutter, et le second, à des conditions de clutter intense. Lors de ces essais, la performance du système était mesurée sous l'angle de la distance de détection maximale pour chaque catégorie de cible, conjuguée à un taux de fausses alarmes acceptable du point de vue opérationnel.

Ces essais ont permis d'approfondir les relations réciproques entre les nombreux paramètres de traitement, mais le temps a manqué pour définir des limites de performances. La définition de ces limites, dans toute la gamme des états de mer et autres conditions opérationnelles, devra attendre les prochains essais prévus pour la debut de 1999 sur le vaste ensemble de données recueillies au large de Terre-Neuve. **La performance mesurée représente donc une limite inférieure et n'est *pas* indicatrice de la performance du SARAIT dans des conditions optimales.**

4.1 Performances dans des conditions de faible clutter

La bande 22 (segment 4) a été enregistrée dans des conditions de faible clutter et de forte houle. La bouée Waverider a enregistré une houle de 3,8 m de hauteur et d'une période de 9,8 secondes, venant à 40 degrés. Un vent faible, de 3 noeuds, soufflant de 140 degrés, ajoutait des vagues de 0,1 m de hauteur à la houle de fond. La plupart du clutter était produit par de petites cellules de pluie qui traversaient lentement la zone d'essai et par des oiseaux de mer se posant sur l'eau. Cet ensemble de données est idéal pour étudier la meilleure façon de détecter des cibles SAR éloignées flottant à fleur d'eau dans une forte houle qui les rend souvent invisibles, mais génère peu de clutter.

Une série d'essais ont été réalisés pour étudier les effets des principaux paramètres de traitement de l'IMR et du SARAIT, en particulier lorsque les cibles sont éloignées. La Figure 2 illustre le résultat d'un essai type : les «détections» de cibles SAR sont montrées en surimposition sur les emplacements vrais des cibles. Toutes les cibles d'essai ont été détectées avant que le navire atteigne l'angle sud-est de la grille. À noter qu'au moment où a été enregistrée la bande 22, les cibles étaient en place depuis plus de deux semaines et plusieurs des grosses cibles avaient chassé sur leurs ancres.

La Figure 3 montre la distance et le gisement du point de départ des pistes initiées par le SARAIT, pour chaque catégorie de cibles faisant partie de la grille. À noter cette fois les grandes distances de détection, notamment de la bouée Waverider, à 3,6 NM au nord, et de la cible A0, la plus éloignée, à 2,5 NM au nord-ouest.

Des oiseaux flottants ont été détectés en raison de leur faible vitesse de dérive. Les cellules de pluie sont également apparues comme des groupes denses de détections DMN, en raison de leur faible dérive, au milieu d'un vent dominant de 3 noeuds. Le MHT a fini par initier plusieurs pistes, comme on peut le voir à la Figure 2. Dans la pratique, il sera facile d'adapter le sous-système IA du SARAIT pour qu'il rejette les pistes issues de densités de détections DMN extrêmement élevées.

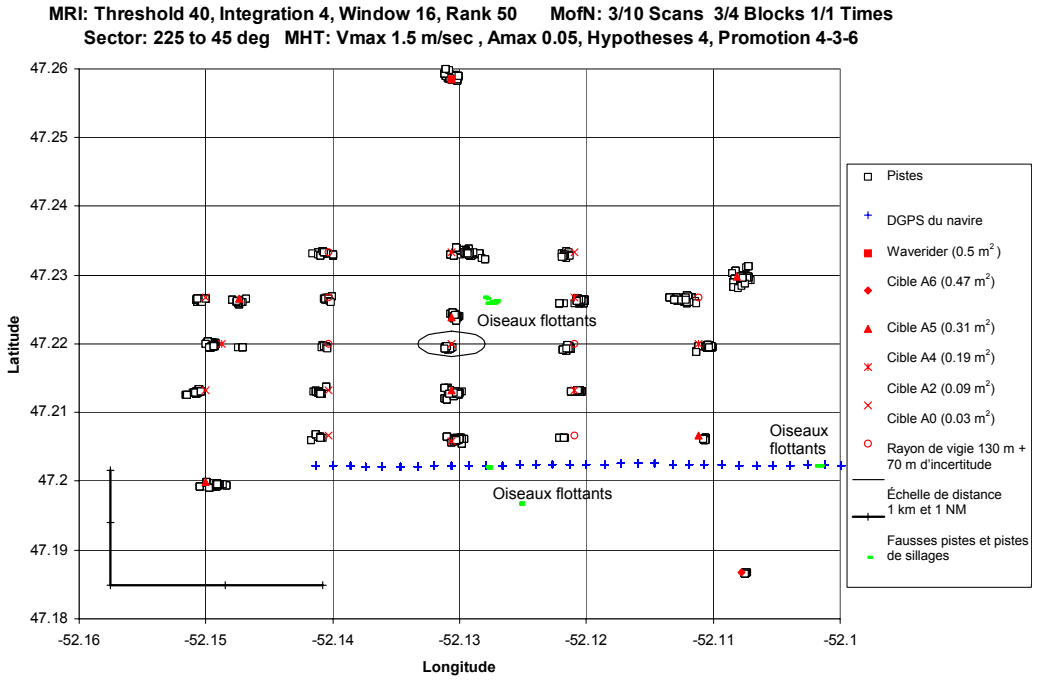

Figure 2 Pistes de cibles SAR en surimposition sur les emplacements vrais des cibles (faible clutter)

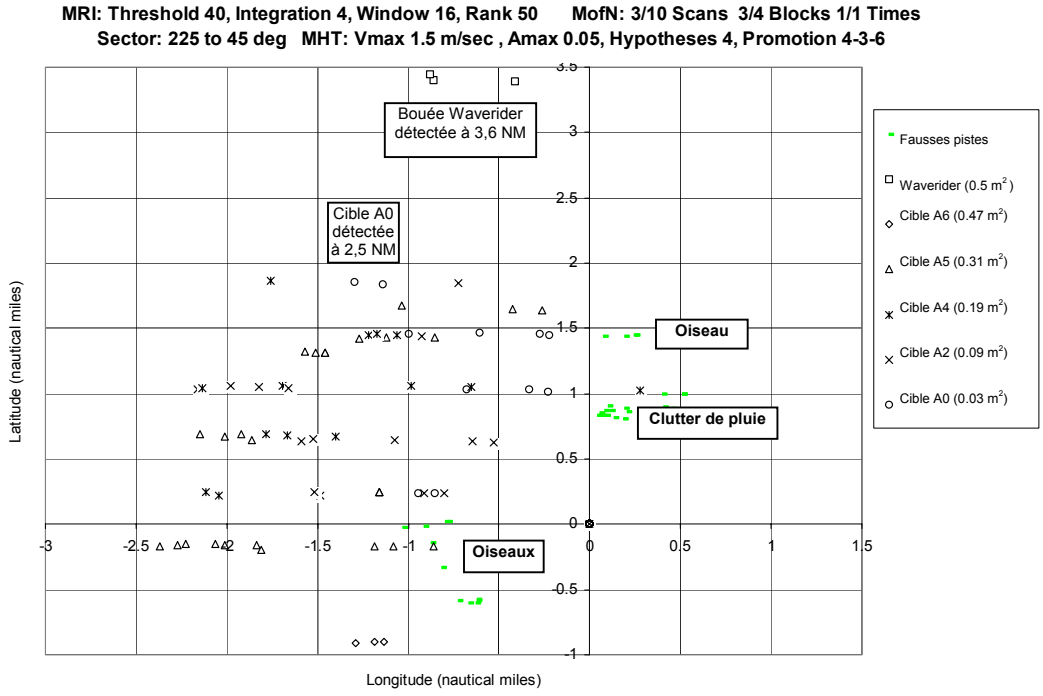

Figure 3 Distance et gisement du point de départ des pistes de cibles SAR initiées depuis le navire (faible clutter)

Dans des conditions de faible clutter, la plupart des fausses alarmes sont dues au bruit du récepteur relativement peu corrélé et à des perturbations transitoires de la surface de la mer, de sorte que le DMN

peut détecter même les moins visibles des cibles SAR en établissant des correspondances entre plusieurs périodes de houle. C'est pourquoi la portée de détection maximale du SARAIT est à peu près la même, quel que soit le gisement de la cible. Dans ces conditions, le traitement du SARAIT est particulièrement robuste et indifférent aux paramètres de traitement.

La Figure 4 montre que la bouée enregistreuse de vagues est toujours détectée à une distance de 3,5 NM, que l'IMR intègre 2, 4 ou 8 balayages. De même, les cibles de la grille sont toutes détectées à une distance approximative de 2 à 2,5 NM, peu importe le type de traitement effectué ou les dimensions de celles-ci. Cela confirme que la visibilité des cibles constitue le principal obstacle à leur détection dans des conditions de faible clutter.

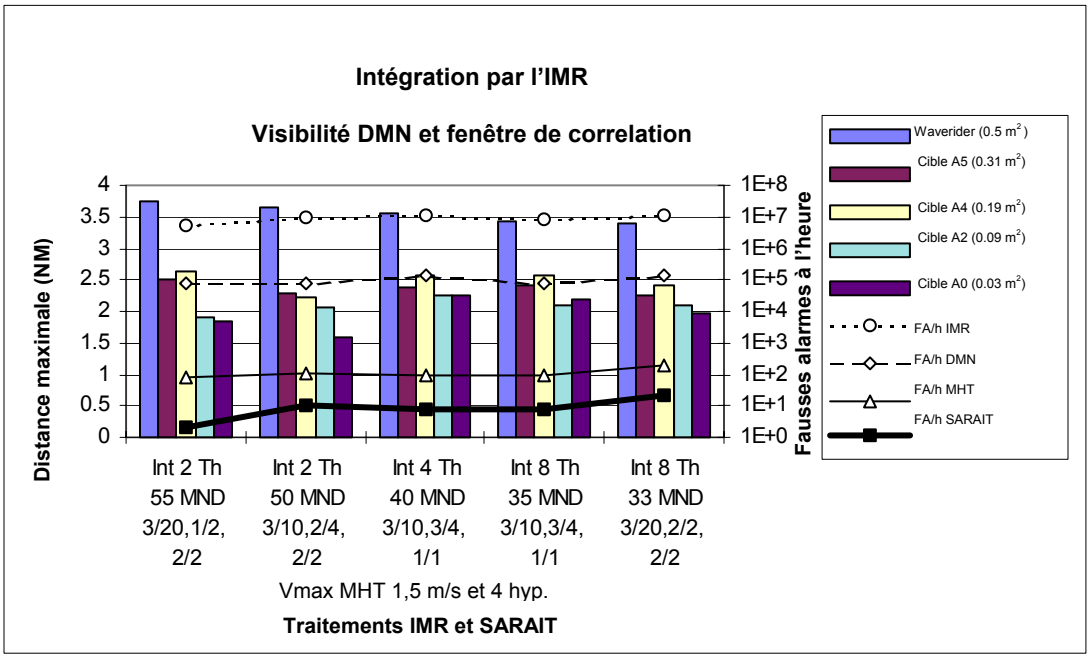

Figure 4 Performances types du SARAIT dans une mer de 3,8 m avec faible clutter

Donc, pour obtenir des distances de détection raisonnables dans un faible clutter, compte tenu d'un taux d'initiation de fausses pistes de 2 à 20 à l'heure, il faut réunir les conditions suivantes :

- réglage du seuil de l'IMR à 750 à 1 500 plots par balayage (TFA = 1e7);
- rapport du DMN de 7,5 % à 25 %, donnant 200 à 500 détections DMN par intervalle (TFA = 1e5/h);
- fenêtre interne du DMN égale ou inférieure à la période de vague moyenne (10 secondes ou 20 balayages);
- fenêtre globale du DMN couvrant plusieurs périodes de vague;
- vitesse maximale de la poursuite de la piste de 1,5 m/s (3 noeuds), de façon à pouvoir compenser une dérive soudaine due aux vagues.

4.2 Performances dans des conditions de clutter intense

Le SARAIT peut également détecter de manière sûre des cibles SAR dérivant à fleur d'eau dans des conditions de clutter de mer intense et des vagues de hauteur équivalente de 3,3 m, mais à des distances plus faibles, et pour autant que l'antenne ne soit pas positionnée pour balayer le versant abrupt des vagues générées par le vent, là où le clutter est le plus intense. Par vent fort, le clutter de mer varie énormément en distance et en azimut, tandis que les réglages de l'IMR sont globaux. L'IMR doit donc trouver un équilibre entre des détections denses à faible distance et dans le vent, et des détections plus rares à de plus grandes distances et sous le vent. La bouée enregistreuse de vagues et les cibles A6 et A5 ont quand même été détectées à une distance de 1,5 NM, et les cibles A0, à des distances de 0,5 à 1 NM. Des distances de

détection supérieures sont attendues à l'issue de la prochaine phase des essais, du fait que les facteurs de dépendance entre le clutter et le traitement de l'IMR et du SARAIT seront alors mieux compris.

La bande 21 (segment 1) a été enregistrée trois heures avant la bande 22 (segment 4). Le vent soufflait alors par rafales à 10 noeuds, dans une direction d'environ 280 degrés, au début de l'enregistrement, avant de tourner pour souffler de 0 degré à la fin de l'enregistrement. Ces vents ont produit une mer agitée associée à un clutter intense. La houle de fond, de 3,2 m, venait toujours de 40 degrés, mais sa période était légèrement plus courte, de 8,6 secondes.

La détection des cibles SAR dans ces conditions s'est avérée particulièrement difficile, car le vent violent soufflait en transvers de la houle, et donc parallèlement aux creux de vague, là où les cibles sont normalement le plus visibles. La Figure 5 montre les pistes de cibles SAR et les fausses pistes pour une recherche type, ainsi que les positions du navire (à chaque mise à jour de piste), en surimposition sur les emplacements réels des cibles. Les cibles sont observées dans la plupart des gisements, le navire ayant effectué les trois quarts d'un circuit complet au cours de ce segment d'enregistrement.

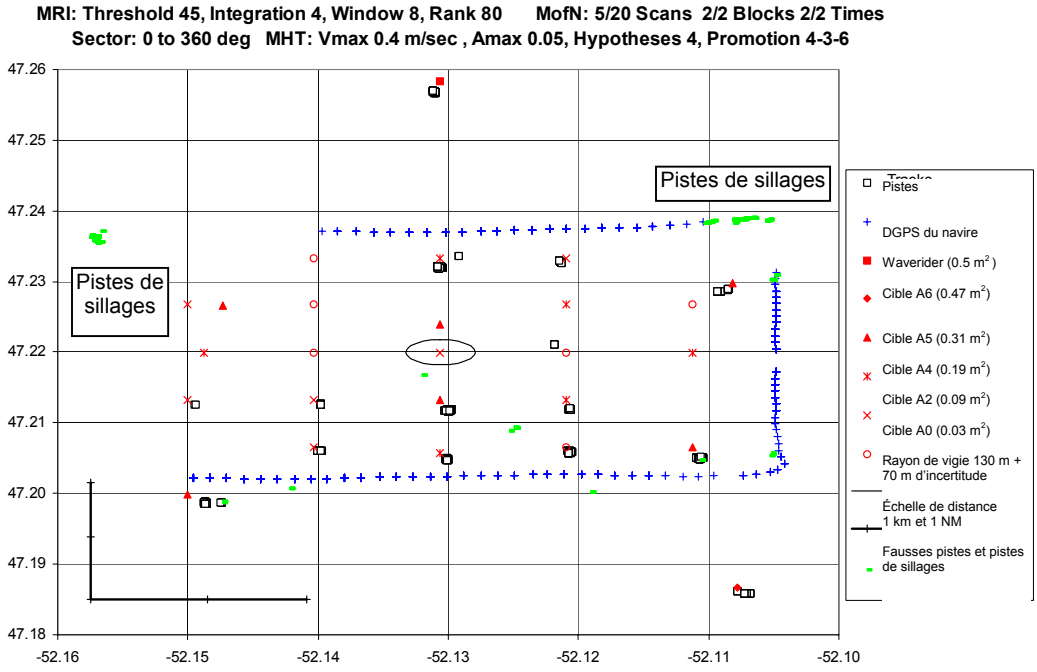

Figure 5 Pistes de cibles SAR surimposées sur la grille vérité-sol (clutter intense)

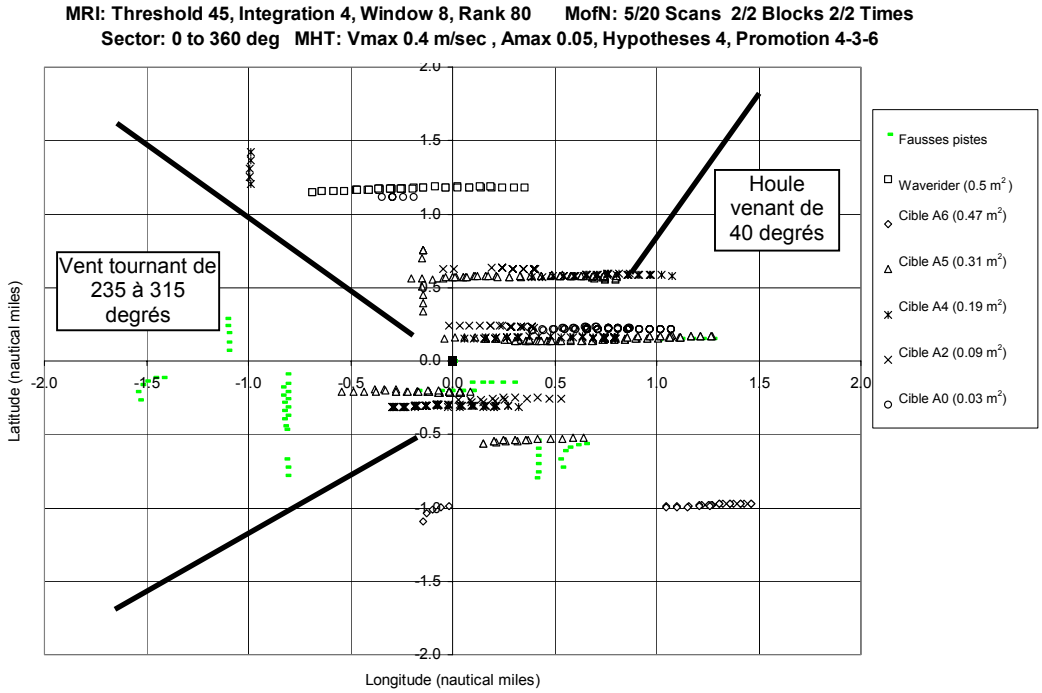

Figure 6 Pistes, distance et gisement des cibles SAR depuis le navire (clutter intense)

La Figure 6 illustre la variation de la distance de détection selon l'azimut, celle-ci atteignant un maximum de 1,7 NM lorsque le navire se situe au milieu de l'arc limité par la direction de la houle et celle du vent, un minimum lorsque le radar éclaire le versant abrupt des vagues (encore en formation) générées par le vent, et entre 1 et 1,4 NM, partout ailleurs.

À noter également, sur la Figure 5, la présence de pistes de sillage, à des distances aussi grandes que 1 NM, le long de la trajectoire du navire. On peut constater, à la Figure 6, que les fausses alarmes (autres que celles dues aux sillages) sont surtout dans le vent, et les autres sous le vent.

Il y a plusieurs façons de configurer le système pour détecter les cibles SAR dans un clutter intense, tout en limitant le taux de fausses alarmes à un minimum acceptable. La Figure 7 illustre plusieurs exemples probants de traitement dans des conditions de clutter intense, dessinant les grandes tendances qui déboucheront sur les performances optimales attendues au cours de la prochaine phase :

- réglage de l'intégration par l'IMR à un nombre de balayages tel que les cibles se trouvent à chaque fois sur la crête des vagues (2 à 4 dans le cas présent);
- réglage de la fenêtre TFAC IMR à 8, pour un TFA global minimal, et à 4, pour une distance de détection maximale des petites cibles;
- réglage du rang TFAC IMR à 30, pour des petites cibles à faible distance, et à 80, pour des cibles éloignées de plus grandes dimensions;
- réglage de la fenêtre interne du DMN en fonction de la visibilité attendue des cibles (10 % à 25 % selon la distance, les dimensions et l'azimut) et la période des vagues (20 dans le cas présent);
- réglage de l'intervalle de corrélation total du DMN à 2 à 6 périodes de vague (40 à 120 balayages dans le cas présent);
- réglage du MHT pour qu'il élabore 4 hypothèses ou plus dans des conditions de clutter intense;
- réduction à 0,5 m/s (1 noeud) de la vitesse maximale de poursuite des pistes MHT et DMN, de façon à limiter le nombre de fausses alarmes du MHT;
- rejet par le sous-système IA des pistes de sillages évidentes.

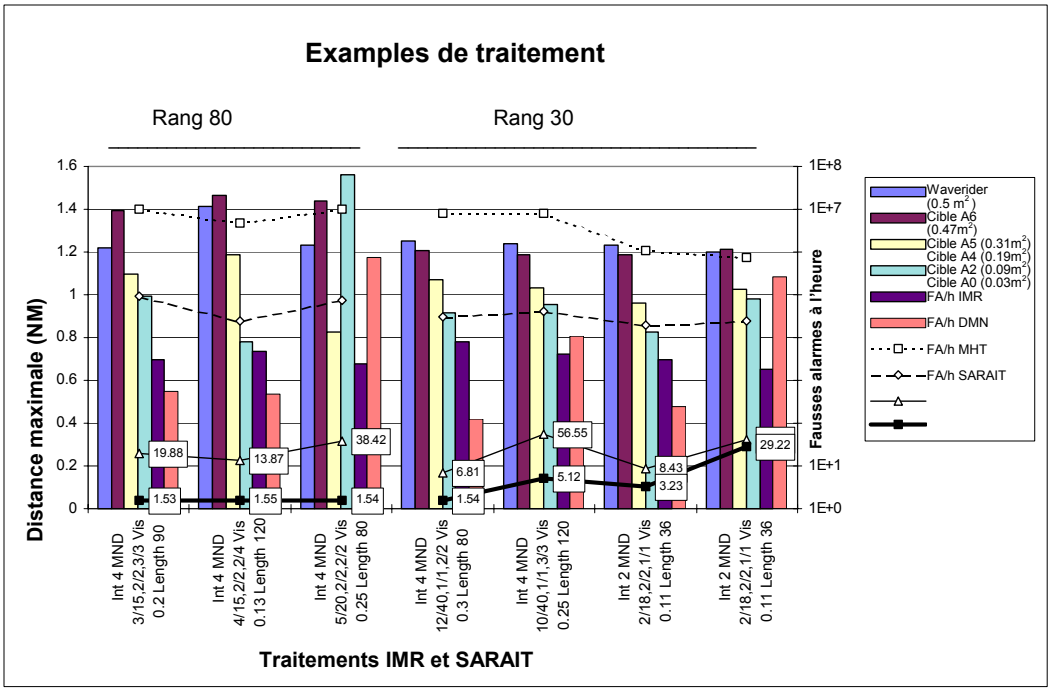

Figure 7 Performances types du SARAIT dans des vagues de 3,3 m et un clutter intense

5 Conclusions

Les méthodes de traitement optimales, pour une diversité d'états de mer et de conditions météorologiques, restent à définir, mais l'essai du SARAIT dans des mers agitées a débouché sur un seuil de performances prometteur, qui ne pourra que s'élever à mesure que l'on acquerra davantage d'expérience à la faveur d'autres essais. De plus, le fait de fixer globalement les paramètres de traitement, comme on le fait présentement, oblige à des compromis qui limitent indûment les performances du SARAIT. Pour des performances optimales, les traitements de l'IMR et du SARAIT doivent varier en fonction des conditions locales de clutter.

Le SARAIT a donc détecté de manière sûre des cibles de 0,03 m^2 (représentant des personnes à la mer) dans des vagues de 3,2 à 3,5 m, à une distance de 2 NM, dans des conditions de faible clutter, et de 1 NM dans des conditions de clutter intense, avec 2 à 10 fausses pistes à l'heure. Les cibles de la grosseur de petits radeaux de sauvetage (0,5 m^2) sont détectées à des distances variant entre 3,5 et 2 NM, selon l'intensité du clutter.

Le SARAIT peut soutenir ces performances en temps réel, pourvu qu'il soit appuyé par un radar tournant à 120 tr/min., une IMR qui l'alimente en plots radar à une cadence de 500 à 1 000 plots par balayage, selon la densité des détections.

Le SARAIT a servi à traiter des données enregistrées provenant d'un radar de bord classique (modifié pour tourner à 120 tr/min.) qui avait été installé à bord d'un baliseur de la Garde côtière canadienne exécutant un circuit de recherche à une vitesse de 8 à 10 noeuds, ce qui représente les conditions types des opérations de recherche-sauvetage en mer dans les eaux canadiennes. Comme les radars de bord classiques sont habituellement incapables de détecter des petites cibles à fleur d'eau dans de pareilles conditions, le SARAIT marque un important pas en avant, qui devrait améliorer considérablement l'efficacité des opérations de recherche-sauvetage en mer.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

GLOSSARY

AI	Artificial Intelligence
AIT	Artificial Intelligence Tracker
CCGS	Canadian Coast Guard Ship
CFAR	Constant False Alarm Rate
CNR	Clutter to Noise Ratio
DGPS	Differential GPS
DLL	Dynamic Link Library
GPS	Global Positioning System
M of N	M target detections in N scans
MHT	Multiple Hypothesis Tracking
MND	M of N Detector
MRI	Modular Radar Interface
OS-CFAR	Ordered Statistic CFAR
PC	Personal Computer
PCI	PC Interface
Pd	Probability of Detection
Pfa	Probability of False Alarm
Ptd	Probability of Track Detection
Ptfa	Probability of Track False Alarm
PIW	Person in Water
PPI	Plan Position Indication
PRF	Pulse Repetition Frequency
PRI	Pulse Repetition Interval
Ptd	Probability of Track Detection
Ptfa	Probability of Track False Alarm
RAM	Random Access Memory
RCS	Radar Cross Section
Rfa	Rate of False Alarm
Rtfa	Rate of Track False Alarm
SAR	Search and Rescue
SARAIT	Search and Rescue Artificial Intelligence Tracker
SBC	Single Board Computer
SIR	Signal to Interference Ratio
SLOC	Source Line of Code
SNR	Signal to Noise Ratio
SVGA	Super Video Graphic Adapter
TDC	Transportation Development Centre
UTC	Universal Time Coordinates
VTS	Vessel Traffic Services

1. Introduction

The Search and Rescue Artificial Intelligence Tracker (SARAIT) processes plots from any conventional marine radar to reliably detect and track small, slowly drifting targets such as liferafts, persons in water (PIW) and wreckage. This report describes the development and preliminary testing of the SARAIT.

1.1 Requirements

Marine search and rescue (SAR) usually relies on lookouts to detect survivors because conventional marine radars cannot reliably detect small liferafts and persons in water (PIW) under typical wave conditions. Lookouts can be very effective but poor visibility and fatigue all too frequently limit their success.

Marine radars ***can*** detect small SAR targets when operated at maximum sensitivity but in doing so they also detect hundreds of comparably strong reflections from wave crests and other sea or weather clutter. These "false alarm" detections collectively swamp the operator. To compound this problem, awash targets such as PIWs and liferafts are often hidden in wave troughs where they are invisible, and therefore undetectable, for several seconds at a time.

Since developing an entirely new radar for SAR would be prohibitively expensive, the Transportation Development Centre (TDC), has preferred to develop three add-on systems that together enable a normal marine radar to be used for SAR:

- Pathfinder/ST Mark II radar modified for 120 rpm (to illuminate awash targets at most wavecrests)
- Modular Radar Interface (MRI) (to digitally process 120 rpm radar video into thousands of plots)
- SAR AI Tracker (SARAIT) (to correlate thousands of plots into high confidence SAR tracks)

The objective of the SARAIT project is to design, build and test a processor that reliably tracks ONLY the SAR targets and rejects the plethora of sea clutter detections. Since conventional marine radars do not measure the target Doppler or polarization, the most promising discriminants are that typical SAR targets:

- drift at under 1 kn unlike sea and weather clutter which moves at 3+ kn, and
- are detectable for several minutes unlike sea and weather clutter which last only a few seconds.

Discriminating the SAR targets therefore requires correlating over many scans of radar data to pick out the intermittently recurring detections of the SAR targets as they move along relatively straight lines of drift pushed by the slowly changing currents, winds and waves. Measuring the speed and direction of this drift requires a sophisticated moving target tracker that can account for the usual measurement errors, ship-motion errors, missed detections and many hundreds of false (sea clutter) detections.

Fortunately, TDC funded Raytheon Canada from 1994 to 1996 to develop just such a tracker for Vessel Traffic Service (VTS) applications. Extending this Artificial Intelligence Tracker (AIT) for the Search and Rescue problem is the subject of this report.

1.2 Previous Phase

The Artificial Intelligence Tracker (AIT) was developed to track VTS targets through severe clutter and blind zones. To this end, the AIT incorporates a sophisticated Multiple Hypothesis Tracker (MHT), an Expert System and a Geographic Information System (GIS) Following extensive trials at the Vancouver VTS Centre, the potential use of the AIT to track SAR targets was recognised .[1] The AIT was therefore used in 1997 to process recorded marine radar data from the CCGS *Gilbert.* This test demonstrated that in 1.5 m seas, the AIT reliably detected PIW targets at 0.5 nmi and small liferafts at more than 1.3 nmi.[2]

Although serious deficiencies in the radar data and the GPS ship position data were shown to be limiting the AIT performance, this early work provided convincing proof that an advanced tracker could reliably pick out slowly drifting SAR targets in the presence of numerous sea clutter false alarms.

1.3 Objectives of this Phase

The key requirement of the current phase was to upgrade the proof of concept AIT processor into a prototype Search and Rescue Artificial Intelligence Tracker (SARAIT) that could operate in real time for use in field trials. To this end substantial extensions to the underlying AIT were designed and implemented:

- new multi-processor computer hardware to sustain real-time processing of 120 rpm MRI plots,
- new M of N Detector (MND) software module to increase the probability of detecting intermittently visible targets such as awash wreckage and PIWs,
- extensions to the Multiple Hypothesis Tracking (MHT) software module to initiate and maintain high confidence tracks on only the slow moving SAR targets, and
- new communications software to enable the SARAIT to receive MRI plots and param and to send tracks and control messages.

In parallel with the development of the SARAIT, Oceans Limited of St. John's was separately contracted by the Transportation Development Centre to record radar data off the east coast of Newfoundland. To this end, the modified 120 rpm Pathfinder/ST Mark II was mounted on the CCGS *J.E.Bernier* and search

[1] Raytheon Canada, "Artificial Intelligence Applied to Target Tracking Phase III: Tracker Testing and Evaluation", TDC Report TP 12972E, February 1997.
[2] Raytheon Canada, "Radar Processing for Detection of SAR Targets Using the AI Tracker", TDC Report TP 12987E, March 1997.

patterns sailed around both drifting and moored calibrated targets. Nine data sets were recorded on 22 tapes in seas ranging from 1.7 to 3.8 m. Staff from Sigma Engineering were on board to operate the MRI and from Raytheon Canada to test the SARAIT with live and recorded data. Preliminary testing of the SARAIT on two data tapes was undertaken to explore the complex tradeoffs between the different MRI and SARAIT processing param.

The SARAIT is designed to work with high scan-rate radars through Sigma Engineering's new upgraded Modular Radar Interface (MRI), which provides plots and displays the processed SAR target tracks. The MRI and SARAIT were developed together and since their processing requirements are similar, they use identical interchangeable hardware. Both the MRI and the SARAIT can be configured to work with most marine radars and in particular with high rotation rate radars such as the modified 120 rpm Pathfinder/ST Mark II.

2. Requirements Analysis

The overall requirement is to separate the numerous sea and weather clutter plots from the few SAR target plots. Although their radar reflectivities are all comparable, both waves (which dominate the sea clutter detections) and weather clutter move rapidly (usually above 3 kn and usually down-wind) while PIWs and life rafts drift slowly in response to the wind, ocean currents and waves (usually below 1 kn).

Another potential discriminant is that SAR targets are physically small while much (but not all) sea clutter spans several range azimuth cells. However, the apparent extent of SAR targets (and the associated position error) is often increased at low detection thresholds when nearby clutter of comparable amplitude is clustered with it. Rejecting the largest extent plots can therefore remove blatantly non-SAR plots from consideration but cannot discriminate most clutter from most SAR targets. Target extent is therefore expected to be a less reliable discriminant than drift speed.

Table 2-1 shows the typical range of wave speeds in fully developed seas arising from 10, 20 and 40 kn winds. Fully developed seas are generated by steady winds blowing for many hours over many miles of unobstructed ocean (the fetch); the higher the winds, the longer and further they must blow to impart their full energy to the seas below. The weather in Canadian waters is so changeable and upwind coastlines so common in SAR work that fully developed seas cannot be assumed. In practice, most seas are limited by the available fetch or the duration of the wind. This changes the wave distribution from the fully developed case as follows:

- fetch or duration limited seas never develop the longest (largest) waves,
- rising winds create more of the shorter (smaller) waves and falling winds less, and

- long period swells are often present from distant storms and these propagate independently of the wind direction modulating the locally generated wind waves.

Table 2-1 Significant wave heights[3], wave speeds, wavelengths and periods for fully developed seas[4]

Wind Speed kn		10			20			40		
Min. Duration hrs		2.4			10			42		
Min. Fetch nmi		10			75			710		
Wave Characteristic:		Min	Avg	Max	Min	Avg	Max	Min	Avg	Max
Sig. Wave Height $H_{1/3}$ m			0.5			2.5			11.5	
Wave Speed	kn	3.0	8.7	18.2	9.1	17.3	33.7	19.7	34.6	65.8
	m/sec	1.6	4.5	9.4	4.7	8.9	17.3	10.1	17.8	33.8
Wave Length	m	2	13	56	14	51	192	66	203	735
Wave Period	sec	1	2.9	6	3	5.7	11.1	6.5	11.4	21.7

Most ocean waves are quite modest; on average, 45% are below 1.2 m, 65% are below 2.1 m, 80% are below 3.5 m and 90% below 6 m[5]. Unfortunately, most calls for SAR arise in foul weather and high seas. However, rapid sailing to the estimated location of the survivors and efficient searching limits practical marine SAR operations to somewhat lower seas. Reliably detecting SAR targets in 3.5 m waves (maximum wave period of 15 seconds) in falling winds after a storm or in fully developed 20 kn seas therefore appear to be reasonable objectives that cover most cases and are consistent with safe ship-handling.

Many SAR targets are awash and therefore only visible at wave crests, typically for about 20% of the wave period. Except during storms, most winds are below 15 kn and generate wind-driven waves with periods ranging from 1 to 5 seconds. Although such waves are usually under 2 m high, targets are unlikely to be visible for more than 1 second at a time. Conventional 20 rpm marine radars scan every 3 seconds and may therefore be looking elsewhere as the target becomes detectable at a wave crest. During proof of concept testing, targets were seen to fall out of phase with the radar rotation and thereby remain undetected for over 1 minute, which naturally degraded the performance.

The trials Pathfinder/ST Mark II radar was therefore modified to scan at 120 rpm (i.e. every 0.5 seconds) to be more certain of detecting each target as it reaches its wave crest. Note that even the shortest waves, with a 1 second period, will still be illuminated twice per wave period at 120 rpm with a correspondingly high likelihood of detecting those targets that are only visible at wave-crests.

Three effects of high-speed scanning must be recognised:

- There are many intervening scans when an awash SAR target is hidden in the trough. On average, an awash target in 3 second waves will only be detectable for 1 or 2 scans out of 6.
- Slowly moving targets will be detected at the same location (within an error radius) for many scans. At 1 kn, the target will remain in a 10 m range cell for about 20 seconds or 40 scans at 120 rpm (versus 7 at 20 rpm).
- Compared to a 20 rpm radar, there are 1/6 the number of pulses per beamwidth so the single scan radar performance will be somewhat degraded *assuming of course that the target was at a wave crest or otherwise detectable* when the beam moved by. For awash targets, this loss is the necessary cost of exploiting each detection opportunity. The amount of loss will depend on the clutter and the processing but integrating over several scans can be used to reclaim some of this lost performance for the larger, more visible, targets.

[3] Significant wave height is the average of the highest 1/3 of waves.

[4] Pierson, Neumann and James, "Practical Methods for Observing and Forecasting Ocean Waves by Means of Wave Spectra and Statistics", U.S. Navy Hydrographic Office Pub. No. 603, 1955.

[5] B. Kinsman, "Wind Waves", Prentice Hall, N.J., 1965. p. 9.

The characteristic behaviour of awash targets is therefore to be obscured in the wave trough for several seconds and to then be briefly detectable at the wave crest. An awash target in a typical 2 m wind-driven wave of 5 second period will only be detectable 15% of the time by a 10 m elevation radar at 1 km[6]. The average Pd at ranges beyond 1 km is therefore unlikely to exceed 0.15 for even the largest awash targets. Moreover, to detect the smaller PIW targets, marine radars must be operated at very low CFAR thresholds to achieve even a modest 0.7 Pd at wave crests. To achieve 0.7 Pd against PIWs at wave crests the radar must detect every other small wave feature and therefore typically suffer a very high 0.001 Pfa.

Trackers rapidly increase in efficiency and performance as the quality of the input detections improves. Since measurement errors are largely dictated by the radar and the characteristics of the targets and of the sea are beyond our control, this leaves improving the input probabilities of detection and of false alarm as the key SARAIT objectives.The MHT handles poor detection data particularly well but it rapidly degrades when the probability of detection (Pd) is low AND the probability of false alarm (Pfa) is high. The SARAIT therefore needs to either increase the probability of detection (Pd) or reduce the probability of false alarm (Pfa). Ideally, the SARAIT would do both.

The Pathfinder/ST Mark II is a typical marine radar in having 10 m range resolution on short pulse and 40 m on medium pulse. The azimuth resolution with a 1 degree beam (2.5 m antenna aperture) is 40 m at 2 km and this increases proportionally with range. The radar typically detects small weak targets with an error of about one quarter of a resolution cell to which must be added scan-to-scan motion compensation errors and North mark errors. Typically, the net radar errors for SAR targets will not be less than half a range resolution cell for nearby targets and may be much more for longer range targets. To the radar, these targets are therefore stationary while the range/azimuth errors exceed the actual target drift.

Targets can therefore be usefully integrated or correlated over the time it takes to drift the error distance. Note that the absolute error is less important than the scan-to-scan errors during the integration or correlation period. A SAR target drifting at 1 kn (0.54 m/sec) will drift outside a typical short pulse 5 m error in 10 seconds and outside a medium pulse 20 m error in 40 seconds. The maximum integration or correlation interval for 1 kn targets is therefore 40 seconds or 80 scans at 120 rpm. Two approaches are well documented for increasing the Pd and/or reducing the Pfa of such quasi-stationary targets:

- Linear Integration of the radar video increases the signal-to-noise ratio (SNR) prior to detection, and
- Binary M of N Detection declares a tentative detection if M centroided radar detections (called plots) are coincident (within the known errors) in N radar scans.

The MRI can linearly integrate up to 8 scans before detection and plot extraction but this is only effective in special cases:

- long linear integration over several wave periods often works in low clutter conditions or when targets are very reflective or highly visible, and
- short linear integration can increase the detectability of targets that are consistently visible at wave crests for more than one scan at a time.

Linear integration is ineffective for the usual SAR case of small awash targets in high clutter conditions because the occasional weak target returns are swamped by the much more numerous clutter returns. Integrating many clutter returns with occasional weak target returns actually reduces the output target power and is counter-productive.

Binary M of N Detection (MND) is much more effective in the high clutter conditions typically encountered in SAR work because strong clutter detections are only detected if they recur sufficiently often to pass the MND criterion. The MND must be able to correlate plots from the MRI over at least one wave period, and preferably over several, for two key reasons:

- to improve the probability of detecting the target at one or more wave crests, and
- to reject a useful proportion of the sea clutter plots before entering the numerically intensive MHT.

[6] Raytheon Canada, “Radar Processing for Detection of SAR Targets Using The AI Tracker”, TDC Report TP 12987E, March 1997. p. 3-2.

For the challenging case of 20 kn fully developed seas, most wave periods are between 3 and 11 seconds. Since target drift limits the MND integration interval to 10 seconds in short pulse mode, only 1 to 3 wave periods can be integrated. The situation is somewhat improved in medium pulse mode where a 40 second integration is possible but this still limits the MND to between 4 and 13 wave periods. In 15 second period waves, as would arise after a storm, about one third fewer wave periods could be correlated without target drift. Note, however, that many targets will drift at well below 1 kn and that the azimuth errors at longer ranges, where detection is most difficult, are proportionally larger than the worst case and therefore permit significantly longer correlation periods for targets drifting other than radially.

The number of coincident detections needed to declare a detection will depend on the scan rate, the target's visibility and the Pd/Pfa of the particular target. To achieve a high Pd, the ratio M/N must be somewhat smaller than the target visibility and N sufficiently long to span several waves. For example, to detect 20% visible 0.8 Pd targets in 10 second (maximum) waves, the MND might be set up to span two of the longest waves with a 5 of 40 detection criterion. Such an MND should increase the average Pd from approximately 0.16 to 0.8 as illustrated in Table 2-2. However, typical input Pfa of 0.001 and 0.002 (i.e. 1000 and 2000 false plots in each scan of 1,000,000 resolution cells), approximately 30 and 736 clutter plots will randomly recur and be detected by the MND. Doubling the input Pfa can be seen to dramatically increase the MND false alarms against low visibility targets. Increasing the MND criterion to 7 of 40 drops the Pd to 0.46 but removes most of the false alarms. Unfortunately, the practical input Pfa from the MRI CFAR detector will often approach 0.01 in the vicinity of the radar where clutter is most severe which will severely limit the MND performance there. The MND cannot therefore reliably detect the least visible or lowest Pd targets in severe clutter without false alarms.

Table 2-2 Approximate MND performance in moderate uniformly distributed clutter

MRI Inputs	***Pd (MRI) = 0.8 @ 20% Visible***	***Pfa (MRI) = 1e-3***	***MRI FA = 1000***	***Pfa (MRI) = 2e-3***	***MRI FA = 2000***
M of 40 MND	**Pd (MND)**	**Pfa (MND) =**	**MND FA =**	**Pfa (MND) =**	**MND FA =**
4	0.90	4.6e-3	463	5.7e-3	5725
5	0.79	3.0e-5	30	7.4e-4	736
6	0.63	1.6e-6	2	7.7e-5	78
7	0.46	6.9e-8	0	6.8e-6	7
8	0.30	2.6e-9	0	5.1e-7	1

For this reason, the Multiple Hypothesis Tracker (MHT) is used to correlate the MND output plots that are output every N scans to detect the SAR targets over much longer periods along their course of drift. This gives the system more time to accumulate corroborative evidence of SAR targets while rejecting the more numerous but faster or more random wave detections.

The overall SARAIT objective is to be operationally useful. To this end, the SARAIT must reliably detect slowly drifting SAR targets including liferafts, Persons in Water and debris in typical North Atlantic and North Pacific sea conditions. The reliability of the SARAIT is described by the probabilities of track detection (Ptd) and track false alarm (Ptfa) at specified ranges and sea conditions. A more intuitive equivalent to the Ptfa is the SARAIT track false alarm rate (Rtfa) per hour.

Since most SAR detections must be investigated at least visually by a lookout and probably by a change of course, they must be strictly limited. This in turn limits the achievable probability of detection. High SARAIT false alarm rates up to 30 per hour would probably be acceptable in good visibility when lookouts could quickly be cued to assess whether each detection is real or not. At night or in poor visibility, with only the radar to go on, no more than 5 false alarms per hour are likely to be acceptable. Until the field trial data is completely analysed, the achievable balance between Ptd and Ptfa (or Rtfa) can only be estimated.

The SARAIT, being constrained to use commercially available processing hardware and high level languages such as C++ and Smalltalk, is not expected to run in real time with the full range of operating param. It must nonetheless operate in real time with at least a useful subset of the operationally useful MND and MHT param. A reasonable real-time objective is therefore to operate the inner MND sliding window at N1*N2 up to 40 with 1000 MRI plots per scan. Multiple N1*N2 blocks can then be chained N3 times to correlate over much longer intervals up to 160 scans long.

3. SARAIT Design Description

The overall Search and Rescue Radar Processing System comprises a Pathfinder/ST Mark II modified for 120 rpm, a Modular Radar Interface (MRI) and a Search and Rescue Artificial Intelligence Tracker (SARAIT). The complete system detects and tracks small awash targets by processing the radar video from any conventional marine radar.

The MRI processes the radar video into constant false alarm ratio (CFAR) detections and clusters them into plots. The MRI also displays raw or processed video, CFAR detections, clustered plots or tracks.

The SARAIT processes plots from the MRI into motion compensated SAR target tracks and sends them back to the MRI. The SARAIT can also display the tracks directly.

3.1 Design Approach

The SARAIT exploits the slower drift rate of SAR targets compared to the surrounding sea and weather clutter, by correlating over many wave periods. Since SAR targets move so slowly, detections will overlay within the limits of the measurement errors and the target drift. The SAR targets are so small and so infrequently detectable that numerous false radar detections are expected. To correlate the infrequently recurring SAR detections from the much more numerous clutter and noise detections, two stages of correlation are implemented within the SARAIT:

- the M of N Detector (MND) correlates scans within stationary windows that are sized to match the actual measurement error plus the growing uncertainty due to the worst-case assumed drift, and

- the Multiple Hypothesis Tracker (MHT) which correlates over many MND updates based on the estimated target drift.

The SARAIT also implements an Artificial Intelligence sub-system that is not complete but will:
- identify any non-SAR tracks corresponding to the wake or rain cells, and
- automatically control the MRI and SARAIT processing for maximum performance.

3.2 Operational Context

The SARAIT interacts with the MRI and the SARAIT Operator as shown in the context diagram below.

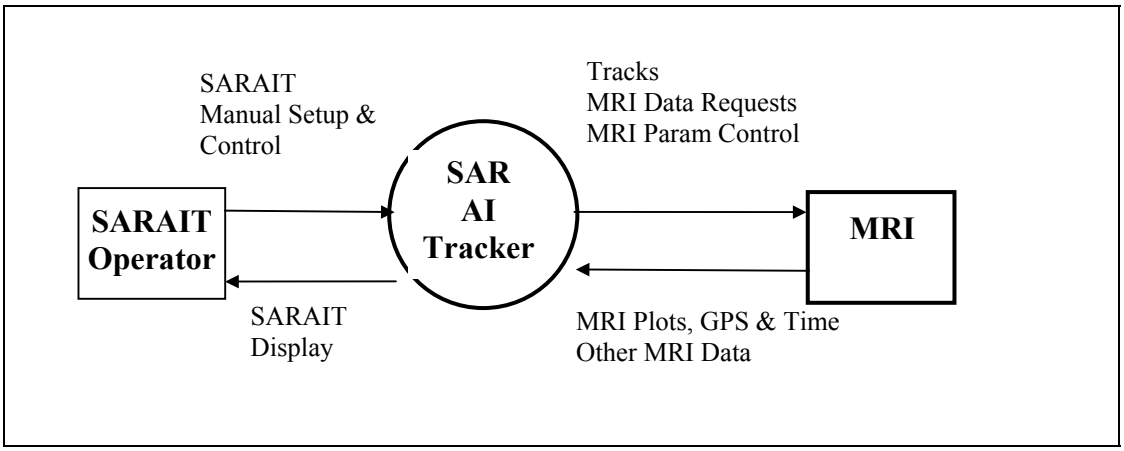

Figure 3-1 Context diagram

The SARAIT automatically requests and receives plot, GPS and time data from the MRI. The SARAIT automatically requests and receives param and other data from the MRI and automatically controls selected MRI param. The (optional) SARAIT Operator has complete access to all SARAIT setup and control param via the SARAIT Operator Interface.

3.3 Hardware Implementation

The SARAIT is configured as two identical independent single-board computers with dedicated peripherals and linked together and to the MRI by separate Ethernet links as shown in Figure 3-2. Each computer is capable of being operated automatically without its monitor, keyboard and mouse.

All hardware is Commercial Off The Shelf (COTS) and interchangeable with equivalent cards, backplanes, chassis and peripherals from other vendors.

The SARAIT incorporates the following minimum hardware set plus optional equipment for use by the SARAIT Operator during setup and development:
- dual passive 18-slot ISA/PCI backplane with power supplies
- two dual Pentium Pro single board computers (Force SRC-620, 200 Mhz, 64 MB EDO RAM)
- two PCI 100T Ethernet cards (one per SBC) (3COM)
- optionally one ISA 10T Ethernet card (for MRI interface)
- optionally up to two PCI SVGA video controllers (one per SBC) (Matrox Millenium 2MB PCI)

- optionally up to two 3.5 " floppy drives (one per SBC)
- optionally one Exabyte Mammoth tape drive (for playback)
- optionally up to two keyboards (one per SBC)
- optionally up to two mice (one per SBC)
- optionally up to two monitors (one per SBC)

The SARAIT is implemented in one rack mountable metal enclosure 0.4 m (H) by 0.5 m (W) by 0.5 m (D) and weighing under 50 kg.

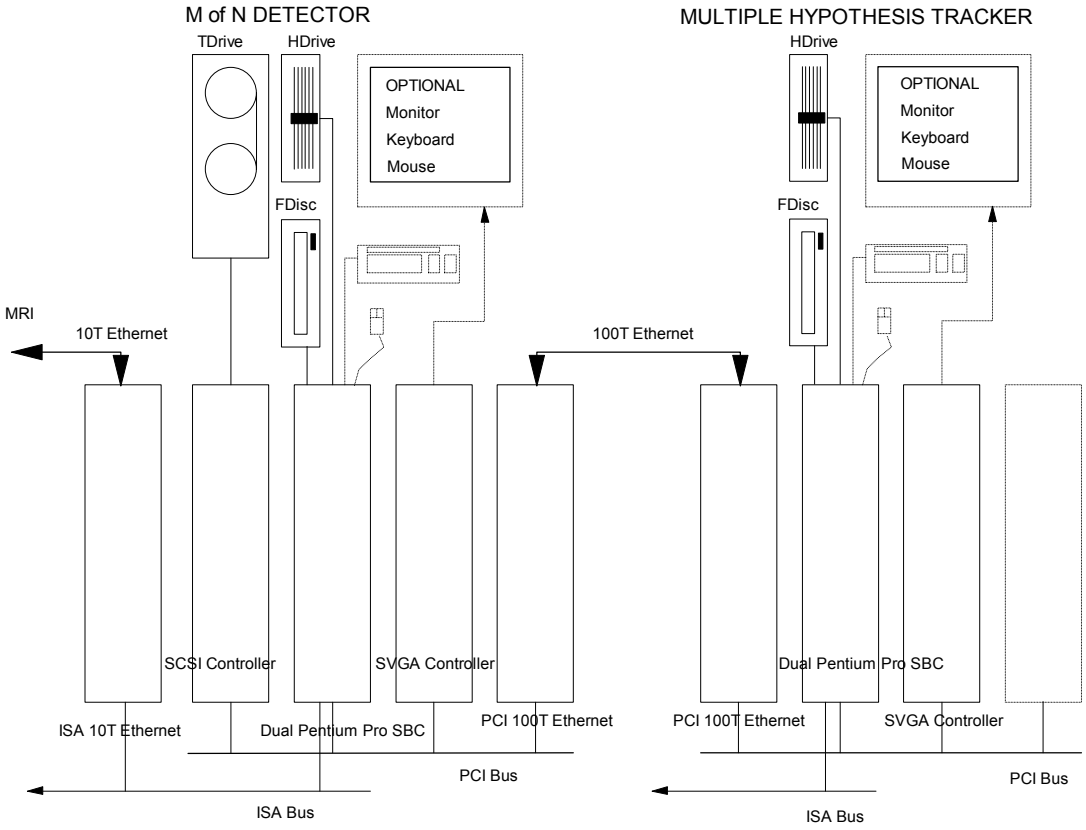

Figure 3-2 Hardware block diagram

3.4 Physical Interfaces

The MRI-SARAIT External Interface is a single 10T Ethernet link as shown in Figure 3-2, for all data and control exchanges between the MRI and the SARAIT. To minimize data exchanges, the first computer receives hundreds of MRI plots every half second over the MRI-SARAIT 10T Ethernet link and implements the MND algorithm which significantly reduces the data rate. MND detections are passed to the second (MHT) computer over the MND-MHT SARAIT Internal Interface (100T Ethernet link) every N1*N2 scans. The small number of SAR tracks are then passed back up the chain to the MRI for display.

3.5 Functional Description

The SARAIT receives plots from, and sends target tracks to, the Modular Radar Interface (MRI). The MRI processes the radar video and applies an ordered statistic (OS) Constant False Alarm Rate (CFAR) detector that is robust in the face of non-Gaussian clutter features such as sea spikes. Detections are clustered and plots generated at the centroid of each detection cluster. The detection and plot extraction is manually controlled by setting the following MRI Plot Extraction parameters:

- detection threshold (above amplitude of CFAR rank plus CFAR offset),
- minimum and maximum plot extent,
- number of scans to average, and
- OS-CFAR window, offset and rank.

Each plot is defined by the following measurements:

- range,
- bearing,
- major axis run,
- minor axis rise,
- maximum intensity, and
- orientation.

Since most targets are smaller than one radar resolution cell and extremely variable in intensity, only the plot range and bearing are used by the SARAIT. Each scan of plot data also includes:

- radar and waveform parameters,
- time of day,
- ship latitude and longitude, and
- number of plots.

The SARAIT calculates the range and azimuth variance from the MRI parameters and uses them to set the size of the correlation area for each MRI plot. Motion compensation is first applied to all incoming MRI plots which are then passed through the M of N Detection (MND) algorithm which has the following characteristics:

- implements up to three levels of nested M of N integration to maximize target detections and limit false alarms,
- processes motion compensated range and bearing plot data,
- the M of N parameters are operator selectable,
- the initial radius of the association gate and its growth over time are automatically determined by the MHT Maximum Velocity,
- plots that fall into an association gate are clustered and a least squares estimate of the target position, its motion vector and the estimated errors updated,
- Levels 1 and 2 M of N normally process a sliding window of N1*N2 MRI scans,
- Level 2 is optionally configured to process N2 successive blocks of N1 windowed detections (to process larger windows in real time),
- Level 1 and 2 detections are accumulated, clustered and the centroids output as detections every N1*N2 MRI scans (where N1*N2 is sized to span sufficient wave periods that most targets are expected to be visible at least M1*M2 times within the window),
- Level 3 M of N processes a sliding window of N3 successive blocks, each of N1*N2 MRI scans, and outputs detections comprising a least squares estimate of the target position, its estimated motion vector and the estimated errors, and
- These M of N Integrated Detections are passed on to the MHT every N1*N2 MRI scans.

The Multiple Hypothesis Tracking (MHT) algorithm has the following characteristics:

- implements full Multiple Hypothesis algorithm that assesses Bayesian probability of each possible association within a controllable hard gate,
- processes motion compensated M of N Integrated Detections every N1*N2 scans,
- all MHT processing parameters are operator selectable,
- tracks are promoted from Preliminary, to Tentative and to Confirmed as MND detections are successfully associated, and
- track estimates are smoothed by a Kalman filter and updated every N1*N2 scans.

A motion compensated display of tracks in Latitude and Longitude is operator-selectable and updated every N1*N2 scans.

Ship-centred tracks are output after each block of N1*N2 scans by removing motion compensation prior to passing them to the MRI (with typical N1*N2 of 30 scans, MRI tracks will only jump by 7.5 m for targets moving at 0.5 m/sec).

Real-time operation is manually maintained by:
1. controlling the MRI Extraction Parameters to limit the number of plots entering the M of N Detector
2. controlling the M of N Detector parameters to limit the number of detections entering the MHT

NOT YET IMPLEMENTED: Adaptive parameter control to maintain optimum performance.

NOT YET IMPLEMENTED: On power-up, the SARAIT waits a set time for operator inputs to the GUI and in their absence automatically sets up and starts the MRI and starts processing MRI data using default parameters stored in an operator editable script file. The SARAIT operates whether or not any keyboard, monitor or mouse are connected.

The SARAIT executes long scripted sessions of **repeated** play back and processing. The following functions are provided:
- control of the MRI tape drive by client level **Tape Function DLL Calls** to find and playback a specified number of scans from a specified segment and time interval,
- NOT YET IMPLEMENTED: control of the MRI plot extractor by client level **Plot Extraction Function Calls** to setup the processing param in above, and
- control of the SARAIT MND, and
- NOT YET IMPLEMENTED: control of the SARAIT MHT.

3.6 User Interface

The SARAIT processing is highly paramized and therefore user controllable. Two interfaces are provided for operator control:
- MHT Graphical User Interface for major operational settings that could change frequently during a mission or playback run and
- MND Setup and Playback File for MND settings and for scripting repeated playbacks of tape sections.

All parameter changes require stopping the SARAIT processing, entering the new parameters and restarting the SARAIT. This will clear the MND and MHT memory which must refilled with new plots before any tracks can be initiated.

3.6.1 Graphical User Interface

The GUI allows the SARAIT operator to change the following key parameter using drop-down menus and user entered or selected parameters:

- Maximum Track Velocity
- Maximum Track Acceleration
- Input Script File
- Input Channel (TCP/IP or File)
- Output Log File
- Detection Probability
- False Alarm Density
- New Target Density
- Diagnostic Display (Tracks & Detections, Tracks or None)
- Display Min/Max X and Y position
- MHT Display: Tracks & Detections, Detections, None
- Start SARAIT

The GUI is controllable by either mouse with keyboard or keyboard alone as follows:

- Numerical entries are typed into the appropriate box and
- Options (e.g. for display) are by drop-down menu (indicated by a down arrow at right side of entry box) and selected by “Enter” or a “Left Mouse Click”.

3.6.2 Setup File Scripted Playback and Processing

Numerous MRI data tapes have already been recorded during the trial period. Rapid playback, processing and analysis are critical to deliver timely conclusions inexpensively. To this end, the MHT and MRI parameters are set manually at the start of each scripted run and are not themselves scripted at present.

4. M of N Detection Algorithm

An M of N Detection algorithm was developed first in Matlab and then ported to multi-threaded C++.

4.1 Problem Statement

Given a set of radar plot images, $P = \{ P_i : i = 1, 2, ...,U\}$, each corresponding to a single radar scan containing a set of distinct extracted plots $P_i = \{P_{ij} : j=1, 2, ..., V_i\}$, find which elements P_{ij} are correlated through M of N consecutive scans from set P. Correlation is obtained when an element $P_{kl} \exists P_k$ is found within range R of element $P_{ij} \exists P_i$, where $k > i$ (i.e. element P_{kl} occurs in a later data scan and is therefore considered an update of element P_{ij}).

Extending this problem, one can consider two-block logic, where elements P_{ij} must be correlated over $M_2(M_1/N_1)/N_2$ scans. That is, the M_1/N_1 correlation must be repeated M_2/N_2 times. (From this point on, the $M_2(M_1/N_1)/N_2$ criteria will be referred to, more simply, as the M/N criteria.)

Solving this problem will generate results which can be used for track initiation in a Multiple Hypothesis Tracker (MHT). It significantly reduces the number of hypothesis, and by consequence, reduces the processing time and resources required by the MHT. The radar tracking system is intended for ship-based search-and-rescue (SAR) operations.

4.2 Data I/O

Radar plot images, either extracted from real radar detections, or generated through simulation, are stored in *.mat* or *.txt* files. Each file contains radar return data corresponding to a single radar scan, which includes both returns from SAR targets and false alarms.

Simulated data was generated and processed into data files each with upwards of 1000 extracted plots $P_i = \{P_{ij} : j = 1, 2, ..., 1000\}$, distributed over $2x10^6$ resolution cells, which corresponds to a false alarm probability of $0.5x10^{-3}$. Anywhere between 1 and 36 SAR targets were placed in the 11.2x11.2 km^2 search area. Scan rates of 20 and 120 rpm were investigated. The total number of scans P_U considered for the MofN track initiation varied, but in every case: $P_U >= N_1xN_2$.

Each element P_{ij} satisfying the M/N criteria is output to a text file, which includes the original location (x-y or lat-long coordinates) of the plot, as well as the time at which the plot was initialized. This file is used directly by the MHT to initiate tracks.

4.3 Data Storage

At any given time memory storage is required for all elements P_{ij} from scans i = [a, b], where the number of scans in the interval [a, b] is equal to N_1xN_2. Whenever a new set P_b is input, every data point in this set must be compared to all remaining points from the N_1xN_2-1 previous scans {i.e. P_{ij} : i = [a, b-1], j = [1, V_i] }. This comparison step represents the major bottleneck in data processing for this M/N problem. To accelerate processing, rather than use a simple brute force procedure, in which every new point is compared to every previous point sequentially, a tree structure is used, which will reduce the number of comparisons needed by limiting to comparisons to most similar elements.

A 2-D recursive binary tree with fixed-location nodes is defined (see document on Gating Problem for further details). Nodes located at the very base of the tree point to linked-lists in which are stored the elements P_{bj} from the last incoming data set, P_b.

Computing of the correlation between elements is accelerated by the use of the tree structure. Since the most recent data elements P_{bj} are placed in the tree, elements from previous data sets can be passed down the tree structure one by one, and compared only to the most similar new elements, which are linked at the lower nodes.

Data elements from previous sets P_i, i = [a, b-1] are stored in a linked-list in sequential input order. Data elements are removed from the linked-list once they reach the N1xN2 time limit. If, by that time, they satisfy the M/N criteria, they are written to an output file before being deleted, to be used later by the MHT.

4.4 Tree Structure

A key feature of the tree structure used is that it is spanned using fixed-label nodes. This implies that none of the actual data elements are placed directly in the nodes. Rather, a center location is specified for the search zone, along with the half-width of the area covered (i.e. ~ = radar range), and from that, subsequent lower nodes are specified by progressively sub-dividing the area in both dimensions alternately, dividing it in half every time.

Due to the nature of the radar scanning, it is likely that the data will become available for processing in a somewhat ordered manner. This precludes the use of a tree where data points are used as node labels as it would result in an unbalanced tree. The data, however, will not necessarily be in perfect order, precluding

the use of a simple linear storage system. A tree with fixed-label nodes was chosen to accommodate partly ordered data input.

When all data points from a given set P_i have been processed, the centroid of these points is calculated. This centroid is used to reset the location of nodes in the tree for the next scan, to ensure that the tree does not become unbalanced if the co-ordinate system is not fixed. Using the new centroid, the entire tree is traversed, resetting location of all nodes accordingly. Furthermore, after a given set P_i has been processed and compared with existing points, the linked lists attached to the lower nodes of the tree are cleared before any new data is entered, both to release memory and to prevent overlap in data when the next scan set is processed.

4.5 Comparison Criterion

The comparison itself is a gating problem, where one attempts to determine which of the elements in the latest set P_i are, in fact, updates of elements observed in previous sets. The motivation behind this is that radar detection due to false alarms will not be correlated through multiple scans, though they may have short birth-death behaviour. SAR targets, on the other hand, should be visible in a relatively fixed position (with some drift allowance), over a significantly longer time period, though not necessarily at every scan. The visibility rate of the targets will vary, depending on target range and radar rotation rate. Furthermore, the portion of the target visible above water will often be less than typical wave heights. Targets at closer range from the radar will have significantly higher visibility rates than those at farther range.

The gating problem assumes that there is a circular gate of radius R_{ij} around each element P_{ij}. The radius has a fixed minimum, to accommodate radar resolution, and grows over time, as long as no correlated elements from subsequent scans are found within it. Initial gate sizes of 5-15 m were used, and gates were assumed to grow at 0-0.5 m/s (depending on trials), until a correlated point was found, at which time the gate returned to its initial size.

A correlation is declared if the following expression is true:

$$(X_{a1} - X_{b1})^2 + (Y_{a1} - Y_{b1})^2 < R_{a1}^2$$

where P_{a1} is a gate with center (X_{a1}, Y_{a1}) and radius R_{ai}

P_{b1} is an element from the latest data set, located at (X_{a1}, Y_{a1})

All correlated data points are assumed to be updates of the same target. The location and time of the update are stored to allow calculation of a track centroid and covariance matrix at the end of the $N_1 x N_2$ time period for those tracks that satisfied the M/N criteria.

4.6 M/N Criteria

Each of the gates is associated with a counter used to keep track of correlations and verify if the gate satisfies the M/N criteria. The counter comprises four sections, accounting respectively for M_1, N_1, M_2 and N_2. When a new element is created, counters M_1 and N_1 are initialized to 1, and counters M_2 and N_2 are initialized to zero.

The N_1 counter is incremented with the input of every new data set P_i. Whenever a correlated element is found, the M_1 counter is incremented. Whenever the N_1 counter reaches the N_1 criterion, or the M_1 counter reaches the M_1 criterion (whichever happens first), the M1 and N_1 counters are reset to zero, N_2 is incremented, and if the M_1 criterion had been reached, the M_2 counter is incremented. Similarly, whenever the N_2 counter reaches the N_2 criterion, or the M_2 counter reaches the M_2 criterion (whichever happens first), the gate is deleted from the list, and if the M_2 criterion had been reached, the details of that gate are placed in the output data file for further processing by the MHT.

The algorithm described thus far uses a simple criterion, requiring a certain re-occurrence in the data before declaring a track initiation. The role of the M/N algorithm was then extended to include some elements of tracking, in an attempt to improve track formation in the tracker. New developments include calculating the centroid of original point P_{ij} and its updates, and sending this information to the MHT, rather than simply the original location of P_{ij}. This may improve track smoothing and reduce jitter during track formation. Furthermore, in addition to the centroid, the covariance in position and velocity is also calculated, and provided to the tracker whenever a plot passes the M/N criteria.

5. MHT Algorithm

The greatest problem associated with tracking multiple targets, based on large numbers of radar detections, is the possibility of incorrect association. In other words, if an established target track is assigned an incorrect detection, the update of its trajectory could send it in the wrong direction, resulting in a false or deceiving target track. Since such tracking is inherently a causal system, these problems can not be recovered. In the last twenty years, and with the advent of more powerful computers, new techniques have been developed which allow better association decisions to be made, even allowing for the possibility of incorrect assignments. One of the most advanced is the multiple hypothesis tracking (MHT) algorithm.

MHT was originally conceived in its most complete form by Reid[7]. It is a statistical approach, incorporating false targets, new tracks, missed detections and finite track lifetimes. The basic premise is that, through the application of Bayes' rule, the probability of any track/detection combination, over a given number of radar updates, is solely dependent on the probability of the combination from the previous scan, and the probability of the current track-detection updated association. The algorithm thus does not make any "hard" assignments at each step; instead it keeps all possible track/detection associations, ranking them by their probabilities (i.e. how likely a given association actually is). Such hypotheses may be efficiently updated at each step merely by calculating the current probabilities for association. Thus, a combination of tracks and detections that looks very likely at one stage, may at a later time be revealed to be less feasible as its updated probability decreases. The correct (or more likely) association hypotheses will then predominate, allowing incorrect decisions to be prevented.

[7] D.B.Reid, "A Multiple Hypothesis Filter for Tracking Multiple Targets in a Cluttered Environment", Tech. Report LMSC-D560254, Lockheed Palo Alto Research Lab., Palo Alto, CA, 1977.

The MHT processing assumes that each new radar report is either an extension of an existing track, a new target or a false alarm (these last two options are combined).These possibilities, together with a missed detection scenario for each of the tracks, account for the additions to the hypothesis set at each update time. The combination and extension of tracks and hypotheses implies that there is exponential growth as new hypotheses are formed at each update. The AIT implementation of the MHT propagates several of the most likely possible track/detection scenarios forward, thus still allowing for the likelihood of missed detections, crossing tracks and false alarms.

Efficiency is further maintained by clustering the data, whereby the multiple hypotheses are considered only for groups of tracks and detections that are close to each other. This enables the gross complexity of the problem to be reduced. These clusters encompass hypotheses that share common reports, and are separate from those in other clusters. In this way, clusters may be processed independently and in parallel, preventing unconstrained growth of the hypothesis tree. The best track/detection hypotheses are determined as solutions to a linear assignment problem, where the elements or "costs" are determined (probabilistically) by the closeness of targets and detections. The auction algorithm is used (repeatedly) to solve these problems. Once under track, the target trajectories are propagated using a Kalman filter.

A number of changes were introduced to the basic MHT used in the baseline Artificial Intelligence Tracker (AIT). These changes were initially implemented during the requirements specification and simulation phase of the project and then to improve the performance, several other features were added for the search and rescue trials that occurred off the coast of St. John's in November 1997. The major changes to the MHT are listed below.

5.1 Algorithmic Modifications

In regards to the MHT, the major changes affected the gating, data association and logic. The data association was sped up through changing the existing N-best hypothesis algorithm to remove redundancies in the search tree. In addition, potential tracks were kept for longer (2-3 updates) to maximize the chances of track initiation. As a result of the slow target velocity (compared to the available resolution), many problems were encountered with velocity gating, since the primary estimates could often be inaccurate and lead to large velocity initial values in the Kalman filter. Resetting the velocity to zero improved performance.

5.2 Interfacing/Communications

Considerable programming effort was devoted to making all processes successfully communicate with each other. RPC calls were used by the MHT to spawn the MND, and data was transferred to the MHT through Windows NT sockets. Many difficult timing issues were resolved to ensure that the MHT received current detection lists. Compatibility was also the focus of much work, to make sure that timing and coordinate bases were consistent across all processes. The MRI's DLL was successfully linked into the MHT, with only the display track option untested.

5.3 Display

A new diagnostic display was designed for the MHT, which showed latitude/longitude information, together with detections, tracks and the instantaneous ship position.

5.4 C/C++ Algorithms

Substantial new C code (summarized in Table 5-1) was written to implement the MND and handle the interfaces to the MRI and between the two SARAIT computers. The new MRI client calls (e.g. EnableClientWait) are now employed so we only receive the plots when processing may continue (i.e. after the MHT has completed its update). Similarly, the M/N code now waits until the MHT is ready.

Table 5-1 C and C++ software blocks

File Name	*Description*	*SLOCs*
mofnanal.cpp	Analysis (offline) version of M/N Detector	1546
mofnol.cpp	Online version of the above	1454
extras.cpp	Addition routines (e.g. coordinate conversions) for M/N Detector	268
mncldll.c	Client dll to handle RPC calls from tracker to M/N Detector	598
mnctrl.c	Master RPC code (for operation of M/N and MRI through tracker)	57
mnctrl.h	Header of above (generated by compiler)	96
mnctrl.idl	Visual C++ setup file for RPC	22
mnctrl.acf	As above	5
mnctrl_s.c	Server stub for RPC calls (generated by compiler)	478
mnctrl_c.c	Client stub for RPC calls (generated by compiler)	487
mnctrlp.c	Procedures called by RPC processes (including MRI calls)	256

5.5 MHT Algorithmic Modifications

It is not always an easy task to document the changes in such a large and modular implementation such as the Smalltalk MHT system. With this in mind, Table 5-2 shows new or extensive modifications to the system, indicating whether just new methods have been added or whole new classes (denoted by *) have been added.

Table 5-2 Modifications to the core MHT software

Class Name	*New/changed Methods*	*Description*	*SLOCs*
MofNInterface*	all	Handles all calls to M/N Detector	71
MRI_Interface	all	Expanded MRI Controller	111
CustomViewExample	showDetectionsASR: with: with: with: with:	Displays current and old detections on MHT diagnostic display	44
CustomViewExample	showRange2: with:	Displays latitude and longitude lines on MHT diagnostic display	64
CustomViewExample	showShip: with:	Displays ship position and history on MHT diagnostic display	18
SwothrMHT	processDataSAR	Initializes and runs MHT and display	138
SwothrMHT	processDataSARAnal	Above for analysis phase	132
SwothrMHT	getDetections	Gets the detections sent over a socket from the M/N detector	65

The MHT has also been modified so that the scoring is once again consistent. In addition, some extra additions have been incorporated to provide efficient operation. Most notable among these is upgrading to the latest 2.5.2 (c) VisualWorks executable and associated DLLs, and a “no-processing” option that splits a cluster that has become too large to run the association algorithm efficiently.

6. SAR Radar Plot Simulator

The design of the SARAIT has motivated the development of a radar plot simulator. The tracker system needed to be tested under a variety of conditions to ensure proper performance. The lack of available measured data representing the full range of test conditions placed severe restrictions on testing possibilities. A simulation package was therefore developed to address that need.

The objective of the simulation is to generate a set of radar scans showing returns from drifting targets, as well as false alarms due to sea surface. The search is assumed to be done by a ship based non-coherent scanning radar whose range, azimuth, and signal strength outputs are fed into a tracking algorithm.

6.1 Approach

A suite of Matlab routines was developed to address the problem. SAR target and false alarm simulation were treated as separate problems, but subject to the same radar and sea conditions.

The radar is assumed to be capable of measuring range and azimuth, which are then assumed to be further processed by a CFAR detector and a plot extractor. The MATLAB code directly simulates the output of the plot extractor on a scan-by-scan basis.

The simulation of SAR targets allows generation of any number of targets distributed throughout the scan area. These targets are assumed to drift over time, and therefore their position is calculated at every scan. Their visibility is governed by a set of fully developed waves.

The false targets are divided into two types of populations. The first set has a birth-death behaviour over a small number of consecutive scans, and therefore some of the false alarms are correlated over short time periods. These go through a birth-detection-motion-death cycle, where the detection/motion states are repeated through multiple scans. The second set of false alarms is generated randomly at every scan. If a false target is alive, it is assumed to be visible, and therefore visibility coefficients need not be calculated for the false alarms.

6.2 Simulation Components

The overall structure of the simulator is shown in Figure 6-1, below.

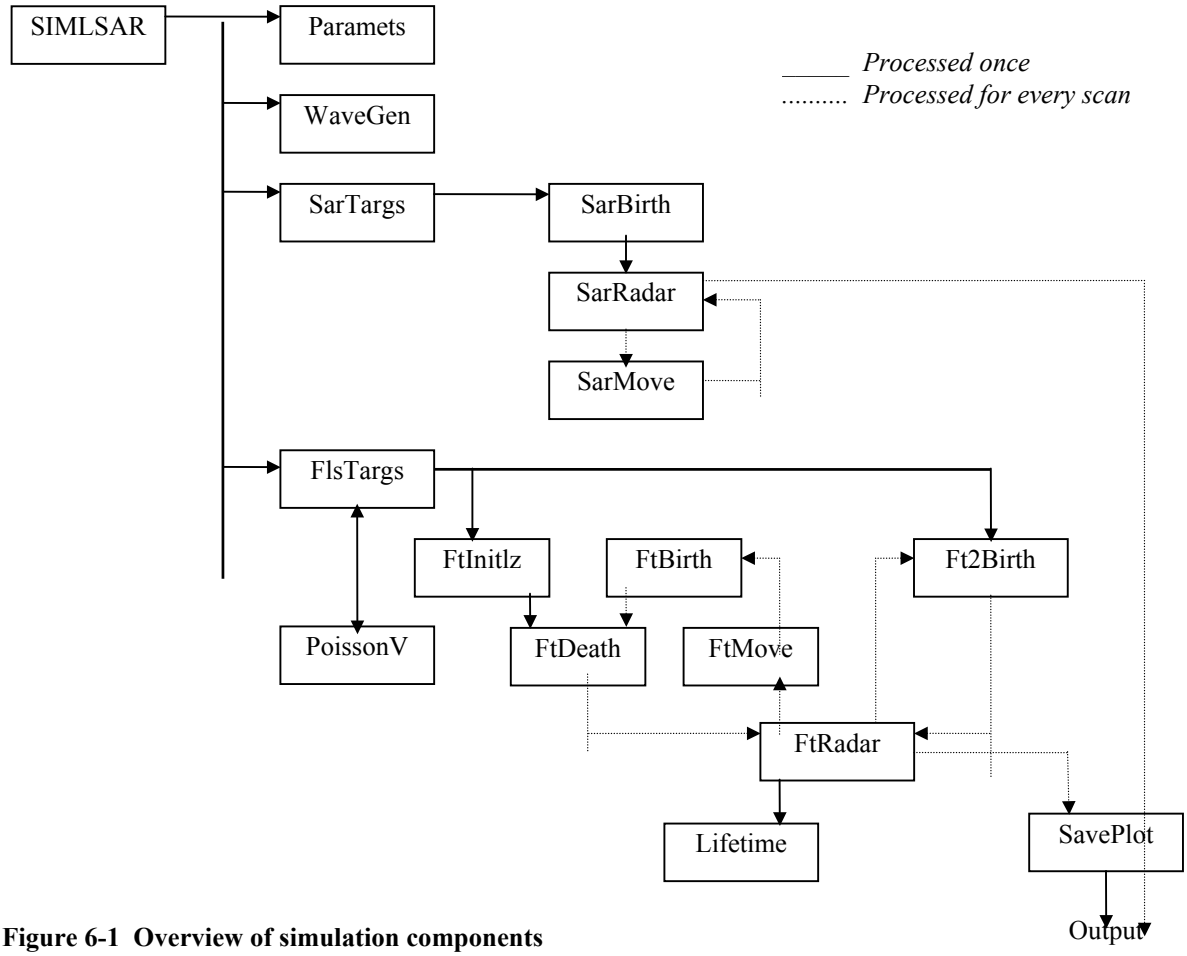

Figure 6-1 Overview of simulation components

6.2.1 Simulation Main - *Simlsar.m*

This is the main routine of the SAR Radar Plot Simulator. Its purpose is to call the principal subroutines, first of all loading the param file, then initiating the SAR targets simulation and finally generating the false alarm detections.

6.2.2 Param File - *Paramets.m*

This routine contains param for all components of the simulation package. It assigns values to the param used in simulating a SAR ocean scenario using a non-coherent radar. Positions of the radar and the targets are described by their rectangular (x, y, z) co-ordinates in metres and, because of the short ranges involved,

assume a flat-earth model. Similarly, velocities are described by their components (Vx, Vy, Vz) in m/second, and so on.

The simulation package can be set to any of the following three modes:

- simulate only the false targets
- simulate only the real (SAR) targets
- simulate both false and real targets

The simulation assumes a fixed-radar location and search zone. Measurement errors are set by the operator to reflect the sum of the radar range/azimuth error, DGPS motion compensation error and North-mark errors. The radar has two range modes (1.5 and 3 nmi), which define the dimensions of the search zone.

The false targets represent radar detections from the ocean surface. The false alarm density is determined by dividing the search area into a given number of resolution cells, and specifying a probability of false alarm for those cells. The heading and speed of false targets are assumed to follow a Gaussian distribution $G\sim(\mu,\sigma^2)$. The mean lifetime of targets having a life-death behaviour, as well as the relative population size of the two types of false targets, must be user-defined.

Parameters of the wave spectrum include the mean and standard deviation of the wave speed, azimuth and amplitude. A spatial range for the generation of the wave spectrum is chosen, which can be smaller than the actual range of the radar, simply because target obstruction due to waves is significantly more likely to be caused by wave crests closer to the target, than by waves further away. Reducing the range of generated waves significantly accelerates computing of the Fourier transforms required to generate the wave spectrum.

Characteristics of the SAR targets considered include the target height and cross section, as well as the target velocity both parallel, and perpendicular, to the direction of wave propagation. Both nominal values and variations of each of these measurements must be specified. The probability of detection is also set as a parameters, and is applied to targets passing the visibility requirement. This implies that visible targets will not be automatically detected.

Both SAR and false targets are assumed to have a certain position error; the standard deviation of that error is user-defined.

6.2.3 Wave Spectrum Generation - *WaveGen.m*

This routine generates wave samples in time and space to model target visibility. The total simulation observation interval usually determines the wave temporal sampling interval. The spatial wave sampling occurs in only one dimension, assumed to be the direction of propagation of the waves. The total wave spatial sampling interval is assumed to be smaller than the search zone, but greater than any distance a SAR target could be reasonably expected to travel over the simulation period. Furthermore, the spatial wave spectrum available should be greater than that from which a target could reasonably expect visibility obstruction.

A one-dimensional temporal wave spectrum can be generated from the Pierson-Moskowitz equation for a fully developed sea-state [1]:

$$E_\infty(w) = (\alpha\, g^2 / w^5) \exp[-B (w_o / w)^4]$$

where $w = 2\pi f$ ($f = 1 / T$)

$\alpha = 8.1 \times 10^{-3}$

g = gravity

$B = 0.74$

$w_o = g / U_{19.5}$

$U_{19.5}$ = wind-speed measured 19.5 m above sea surface

$(H_{1/3})_\infty = 2.12 \times 10^{-2}\, U_{19.5}^2$

($H_{1/3}$ is defined as the average of the 1/3 highest waves)

A power spectrum is obtained from this equation, to which a frequency threshold is applied, expressed as a percentage of the power spectrum area. This threshold defines a maximum sampling rate. The number of samples, however, should correspond to a power of 2 to significantly accelerate computation of the Fourier transforms needed to obtain the time-domain representation of the wave spectrum. This implies that the duration of the temporal wave spectrum may have to be extended beyond the observation period. This causes no problems, as targets may simply not go through the entire cycle of generated wave spectrum.

The second phase is to extend the wave spectrum from a one-dimensional time-varying wave to a two-dimensional spectrum, to reflect spatial variation. Spatial wave generation is only performed in one dimension in space, which is assumed to be in the direction of wave propagation. Spatial wave intervals, therefore, simply describe a wave range in the distance between two points as projected onto the direction of wave propagation.

6.2.4 SAR Targets - *SARAITargs.m*

This routine generates the output of the radar due to SAR targets. It sets initial conditions for the radar and SAR targets, and proceeds to call upon the *SARBirth.m* sub-routine which initiates SAR targets. Following this, it alternately calls upon the *SARRadar.m* and *SARMove.m* subroutines, repeating the sequence once for every scan in the specified simulation observation period. At the end of the SAR targets simulation, the radar detection results are saved to a file (*sarplot.mat*) which includes the following:

- all SAR targets' xy-position at every scan,
- standard deviation of position errors,
- number of SAR targets detected at every scan,
- xy-position of detected targets, and
- visibility rate of each target (percentage of scan visible).

6.2.4.1 SAR Target Birth - SARBirth.m

This routine initially generates SAR targets, and is invoked by *SARAITargs.m*. It generates as many distinct SAR targets as required in the parameters file. For each target, a height, original position and velocity are randomly generated according to the specified parameters.

Targets' positions are uniformly distributed in both x-y dimensions, within the defined search zone. Pre-determined target locations might be required for testing purposes. If this is the case, desired locations can be specified at the beginning of the *SARBirth.m* file. Target height, as well as velocity components (both parallel and perpendicular to wave propagation) are uniformly distributed around given nominal values, and within limits established by the variability param.

6.2.4.2 SAR Target Radar Detection - SARRadar.m

This routine determines the visibility and detection of SAR targets at every scan. It combines the generated wave spectrum with knowledge of a target's location to calculate the visible fractional height of the target.

The distance between a given target and the radar is projected onto the direction of propagation of the waves. The resulting relative distance corresponds to the spatial interval of wave spectrum to be considered for visibility.

The spatial wave spectrum originates from the radar location, and samples are taken in the positive direction, with azimuth corresponding to the direction of wave propagation. If the projected relative distance of a target is negative with respect to that origin, samples of the spatial wave spectrum are taken in the opposite direction.

The spatial dimension of the generated wave spectrum can be smaller than the projected distance between the radar and a target. In such cases, the full spectrum is considered, and it is assumed that visibility interference of the target can only be caused by these waves, which is adequate assuming the wave spectrum contains at least 5-10 wave periods. It is unlikely that a target would be obstructed from a peak

at farther range than this. The spatial dimension of the generated wave spectrum can be set to ensure a sufficient minimum range.

The actual visibility is calculated by measuring the portion of a target which will not be obstructed from the radar by intervening wave crests. If desired, a diffraction factor can be included to account for the diffraction effect due to wave crests. The diffraction factor is incorporated by simply lowering the given radar height, resulting in an effective height. The minimum visible portion after calculating the obstruction by all the intervening wave samples defines the percentage visibility. Visibility, however, does not necessarily imply detection. The user-defined probability of detection is applied in addition to the visibility ratio, to determine whether or not a target will be detected at a given scan.

6.2.4.3 SAR Target Motion - SARMove.m

This routine governs the motion of SAR targets on a scan-to-scan basis. A target's new position is calculated based on that target's velocity and drift azimuth. The calculated velocity is applied directly to the present position to generate the next scanned position. The drift due to velocity is not randomized in any way, and therefore the distance travelled by a given target between each scan is constant.

6.2.5 False Targets - *FlsTargs.m*

This routine simulates the radar output due to water surface features which appear as false targets. In this case, despite the fact that false alarms are usually attributable to water surface features, the simulated false targets are not extracted from the wave spectrum discussed above. Rather, false alarms are generated randomly, according to specified conditions. The purpose of these false targets is to simulate the water surface as seen by the radar after CFAR detection and plot extraction.

Two types of false alarms were considered. The first population exhibits a birth-death behaviour, with a lifetime of more than one scan, which is consistent with radar detections of water surfaces at high rotation rates. Average population size and birth/death rates can be selected through the parameters file; however, the actual number of births and deaths for each scan is calculated based on a Poisson distribution.

The second population is randomly generated at every scan and is consistent with radar detection of sea spikes. A Poisson distribution also describes its population size.

Sub-routines governing the behaviour of false targets are all called from within this routine. Sequentially, sub-routines for the birth, death, observation and motion of false targets are each activated once every scan, and the information is stored on a scan-by-scan basis in individual output files.

At the end of this sequence, a diagnostic can be obtained from the lifetime routine, which shows distribution of lifetime duration among those targets with a birth-death behaviour.

6.2.5.1 Poisson Variable Generation - Poissonv.m

This function generates a set of Poisson random variables with the same mean value, from a set of random variables that are uniformly distributed in [0, 1]. It takes as argument the mean of the Poisson random variables and the number of random variables to generate. This number is equal to three times the number of scans in the simulation period - one Poisson variable for each of population 1 birth, population 1 death and population 2 birth is needed for every scan.

6.2.5.2 False Targets Initialization (Population 1) - FTInitlz.m

This routine creates an initial population of false targets with birth/death behaviour and assigns kinematics and other parameters. The size of the false target population is specified as a user-defined expected value. At each scan, a number of targets will be born and others will die. The birth and death rates are determined according to a Poisson distribution, but are generated independently, so that the net population does not necessarily equal the expected value. Over a number of scans, however, the mean population size will approach the specified expected value.

In the initialization routine, the number of false targets initialized is equal to the sum of the expected population size and Poisson-sized birth sample. Since a birth sample is added directly to the initial population, it is necessary that the death routine be called immediately, before any radar processing is done (see Figure 1), to ensure a population of the expected size when outputting false targets radar detection results for this scan.

Once the size of the population has been determined, the position of these initial targets is determined randomly according to a uniform distribution over the search zone.

6.2.5.3 False Targets Birth (Population 1) - FTBirth.m

This routine creates new false targets and assigns kinematics and other param. For the first scan, the *FTInitlz.m* routine is called however, on subsequent scans, the *FTBirth.m* is called instead. The two routines are almost identical, with the difference that the *FTBirth.m* routine only creates new targets according to a Poisson variable, since a base population is already alive and has been conserved, following some motion allocation, from the previous scan.

6.2.5.4 False Targets Death (Population 1) - FTDeath.m

This routine kills the longest living false targets. The number of deaths is determined by a Poisson random variable. As false targets are created, they are placed in order of appearance in an array. Therefore, to ensure that in fact, the longest living false targets are being killed, targets are removed from the end of the array, where longer living targets are stored. The array is used as a queue, containing all targets alive at any given scan.

6.2.5.5 False Targets Birth (Population 2) - FT2Birth.m

This routine creates an initial population of false targets with no life/death behaviour corresponding to sea spikes. These targets are purely random, uniformly distributed, and a new set is generated at every scan. The size of the population at every scan is determined from a Poisson variable. The mean of the Poisson distribution is obtained from the param file, where is specified the expected value of the total false target population size, as well as the ratio between the transient and the persistent populations. Location of these targets is determined randomly following a uniform distribution over the search zone.

6.2.5.6 False Targets Radar Detection - FTRadar.m

This routine computes the radar output due to the false targets in both populations. A measurement error is added to the position of all false targets alive. The measurement error is added to both x and y dimensions independently, according to a normal distribution. Eventually, computation of the signal power output by the radar could also be calculated as part of this function, if deemed necessary.

6.2.5.7 False Targets Motion (Population 1) - FTMove.m

This routine updates the positions and velocities of the false targets in population 1. The current positions of live false targets are determined in *FTRadar.m*, and these positions are used as a starting point for motion. Since a velocity with (x, y) components is generated at each false target's birth, the motion effect can be calculated by adding the product of the velocity and scan interval to the present position of the target. This yields the false target position for the next scan.

6.2.5.8 *False Targets Statistics* - Lifetime.m

This routine computes the distribution of lifetime duration of the false targets in population 1. It is only a diagnostic tool, and has no effect on the simulation itself. Based on the birth/death behaviour of targets over the full simulation period, it determines the lifetime distribution for false targets in population 1, and can output the results in graphical format.

6.2.6 Result Output - *SavePlot.m*

This routine saves the plot extractor output due to real (SAR) and false targets in a file. Both SAR targets and false alarms are placed in the same array, leaving it up to the tracker to determine which detections should be discarded, and which should be considered for full tracking. The (x, y) position and scan index of each target is written to the output file. Each scan of data generates a separate binary file (e.g., *plot1.mat*, *plot2.mat*, etc.) to enable the tracker to process data scans discretely and sequentially.

In the event that the user selects a mode 1 simulation (i.e. false targets only), an alternate output routine named *SavFTPlt.m* is invoked instead of *SavePlot.m,* and similarly it outputs scan-by-scan data in separate files (e.g., *ftplt1.mat*, *ftplt2.mat*, etc.). If a mode 2 simulation is selected, the results of the SAR detections can simply be found in the *sarplot.mat* file (see above *SARAITargs.m* description for more details on *sarplot.mat*). In this case, no scan-by-scan plot file is generated.

6.3 Data Generation

The param file allows a user to control the characteristics of the simulation. This section describes typical configurations used to generate data for testing of the SAR tracker system. Typically, a full simulation was generated (Mode 3), where both real (SAR) targets and false alarms were produced.

6.3.1 Radar Configuration

A radar in 3 nmi mode was assumed, corresponding to a maximum range of 5600m in both x and y directions. This yields a square search area of 11200 x11200 m^2, over which the targets were uniformly distributed. The radar height was varied between 10 m and 25 m, and as expected, an increase in visibility was observed when increasing height.

Antenna rotation rates of both 20 rpm and 120 rpm were investigated; however, most of the later testing focused on 120 rpm radar data processing. The faster rotation rate emphasizes the effects of correlated false alarms from the population with birth/death behaviour (population 1).

6.3.2 False Alarm Simulation Characteristics

The false alarm density was specified by dividing the search zone in $2x10^6$ cells, and assuming a $0.5x10^3$ probability of false alarm for each. On average, this resulted in 1000 false alarms per scan. They were assumed to have a mean speed of 5.5 kn and mean heading of 0°. The population distribution was usually composed of 25% of population 1 targets (birth/death behaviour) and 50% of population 2 targets (lifetime = 0). The mean position error was 15 m to reflect the expected Global Positioning System (GPS) and radar measurement error.

6.3.3 Wave Characteristics

Most of the wave characteristics are defined by the equation of the power spectrum described above. The nominal wave amplitude was set to 1.5 m. A Gaussian distributed wave azimuth with mean 0° and standard deviation 45° was specified and generated randomly for each SAR target.

6.3.4 SAR Targets Simulation Characteristics

As many as 36 SAR targets were simulated simultaneously. It was assumed that the search zone was sufficiently large to prevent crossing of distinct tracks. In fact, unless targets happened to be initially generated very close to each other, their drift speed was such that it was highly unlikely that they would ever approach each other. As such, they each had the potential to spawn a distinct track.

The SAR target nominal height was typically set to 0.1 m, with a 0.01 m variability. A nominal cross section of 0.1 m^2, with 0.01 m^2 variability was also assumed, though this parameters was not used in the visibility calculations. It could be used eventually to simulate the signal power of a target detection. The probability of detection, assuming the visibility requirement is satisfied, was typically set to 0.8.

The target velocity was expressed in terms of its (U, V) components, where U is the direction given by the SAR target drift azimuth, and V is perpendicular to U. Nominal velocities were (U, V) = (0.5 kn, 0 kn), with variability of (U_{var}, V_{var}) = (0.04 kn, 0 kn).

Even though not accessible through the param file, the location of SAR targets was modified through the *SARBirth.m* script, where one can specify the desired location of targets. By default, targets are placed randomly throughout the search zone following a uniform distribution. For testing purposes, however, it was useful to look at target visibility, detection and tracking potential for targets as a function of range and azimuth in the direction of wave propagation. Typically, targets were specifically placed at 2, 4 and 6 km from the radar. At each of these three ranges, a ring of 12 targets was placed such that there was a 30° intervals between targets. Since the wave direction was given an azimuth of 0°, this configuration allowed for the study of the effects of range and wind direction on target visibility and tracker performance.

7. Basic Parameter Trade-offs

Preliminary testing of the SARAIT used synthetic data from the SAR Radar Plot Simulator described in Chapter 6 to test the newly developed algorithms, assess processing loads and determine appropriate processing parameters for representative target scenarios. The 36 targets, of constant drift rates (0.5 ms^{-1}), were spread around three range rings (2000, 4000 and 6000 m) at 30^{o} intervals with appropriate simulated return strength and visibility profiles. All plots are scaled in m.

It should be noted that, for all the results presented, the track numbers are the total number of unique tracks that occurred. In other words, these numbers will most always be larger than the true physical number of targets, due to the fact that a track lost and regained will be counted twice. In addition, for multiple hypotheses, any time an hypothesis switches to a more likely track, a new track is counted, which can further inflate the numbers. Nonetheless, such a simple track count provides a crude indication of track quality.

The MHT progressively assigns more confidence to tracks as more evidence accumulates that there is indeed a target there:

- All Plots from the MRI are considered to be Potential Tracks
- Potential Tracks are stored for 1, 2 or more successive scans (the Potential Track Life)
- Potential Tracks that are associated within their Life become Tentative Tracks
- Tentative Tracks are coasted for several scans (the Tentative Track Coast) before deletion
- Tentative Tracks that are associated with sufficient updates become Confirmed Tracks
- Confirmed Tracks are coasted for several scans (the Confirmed Track Coast) before deletion

For these tests, the two alternative track promotion criteria of Table 7-1 were used.

Table 7-1 Alternative track promotion criteria

Potential & Tentative Track-Life	Potential Track Life (scans)	Tentative Track Coast (scans)	Tentative Updates to Confirm (scans)	Confirmed Track Coast (scans)
Long	3	3	5	6
Short	1	2	5	6

7.1 SAR Targets in 1.5 m Waves with Radar at 10 m Elevation

The case of a small patrol boat is particularly challenging because the low radar height significantly limits target visibility at long ranges. The following six figures illustrate the problems with SARAIT tracks from a simulation of the Pathfinder/ST II radar 10 m above the sea surface detecting SAR targets in 1.5 m waves. The clutter false alarm probability (Pfa) is 2.5e-4, and the nominal probability of detection (Pd) is 0.8. Average target visibility varied from 10% for distant up- or down-wind targets to 95% for nearby cross wind targets. Note that this trial is to demonstrate parameter tradeoffs rather than performance limits.

The top row of Figure 7-1 displays the correct tracks and overall track statistics for a 3/24 MND processed by the MHT tracker in 24 scan (12s) blocks using the long and short track-life criteria respectively. Only the most likely hypothesis is used to update each track.

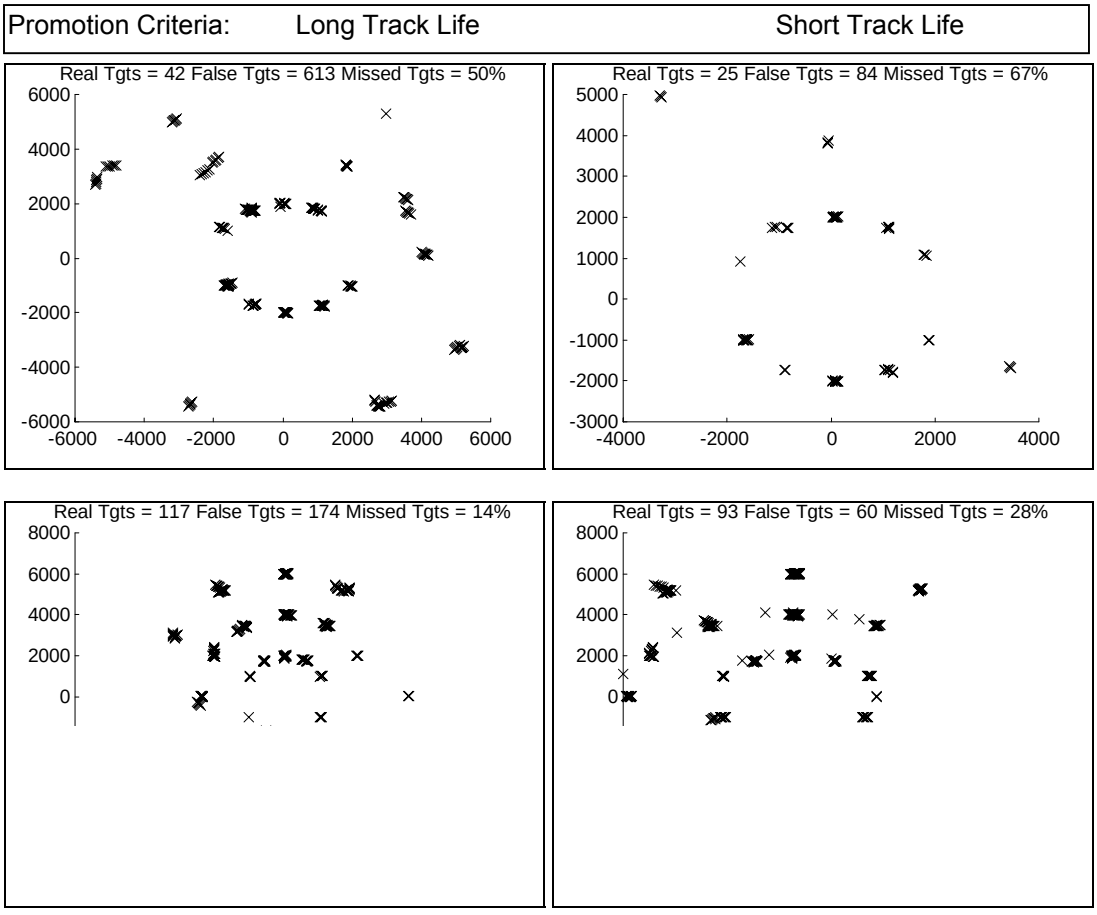

Figure 7-1 3/24 (top) and 3/24-2/4 (bottom) MND Long and Short Track-Life & 1 Hypothesis

Using the single block 3/24 MND, the long track-life promotion criterion detects 613 false target tracks versus only 42 real target tracks and these miss half the real targets. The excessive false tracks arise because the potential tracks are stored for three scans waiting for further associations. There are so many false alarms that random associations commonly promote them to tentative tracks. Moreover, tentative tracks need three successive missed associations to be deleted. Again, the large number of false alarms make further promotion likely. The short track-life criterion tightens up the promotion criteria and reduces the false tracks seven-fold but further reduces the real target tracks and misses even more of the real targets.

The low number of real tracks versus the flood of false tracks highlights the phenomenon of track seduction. By this, we mean that many real target tracks, properly initiated, are being seduced away from their true locations by the much more numerous false alarms, resulting, for all intents and purposes, in false tracks. Track seduction occurs in both the MND (for quasi-stationary tracks) and the MHT.

In an attempt to improve performance in the presence of numerous false alarms, the MND is extended by an outer 2/4 block to reject more of them prior to the MHT. The bottom row of Figure 7-1 shows that this significantly reduces the number of missed targets and (where numerous) the false targets. This indicates that track seduction was indeed a factor in the single block 3 of 24 MND cases. Depending on the overall system requirements, either track-life criterion could be used with the 3/24-2/4 MND. Using the long track-life criterion detects 31 (of 36) targets with 117 target tracks versus 174 false tracks. In contrast, the short track-life criterion correlates detects only 25 targets with 93 target tracks but with proportionally fewer (60) false tracks. Yet longer constraints would be expected to further improve the performance.

The short track-life criterion has one third the false targets but detects somewhat fewer targets. To increase its performance, the number of hypotheses is boosted from 1 to 3. Figure 7-2 shows that increased hypothesis MHT filtered by the more restrictive 3/24-2/4 MND does increase the number of targets from 22 to 28 (and double the number of target tracks from 93 to 152) while slightly increasing the false target tracks from 60 to 84. Note, however, the catastrophic tenfold increase in false target tracks from 84 to 814 that arises with the looser 3/24 MND; that the number of real targets doubles from 12 to 24 (and the number of target tracks quadruples from 25 to 91) cannot outweigh the danger of swamping of the MHT with false tracks.

False confirmed tracks, having a stored history and long life, consume more processing and memory than shorter lived tentative tracks which are in turn more burdensome than potential tracks and false alarm plots. It is therefore important to control the number of tracks by careful parameter setting.

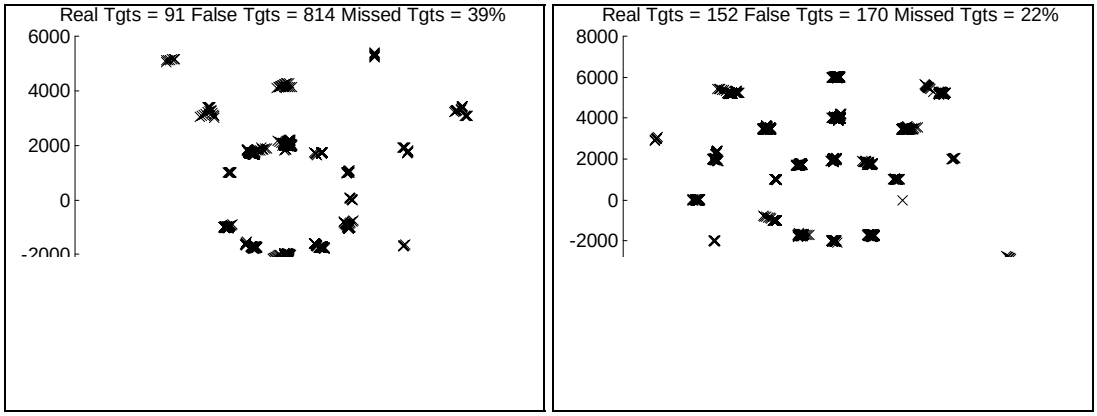

Figure 7-2 3/24 and 3/24-2/4 MND with Short Track-Life & 3 Hypotheses

7.2 SAR Targets in 1.5 m Waves with Radar at 25 m Elevation

Figure 7-3 and Figure 7-4 provide a similar progression of results, simulated for a 25 m high radar similar to the field installation on the CCGS *Bernier*. As before, the nominal Pd is 0.8 but the input Pfa is doubled to 5e-4. The higher radar improves target visibility so the MND is therefore tightened to 4/24 to limit the more numerous MRI false alarms. Track seduction is thereby controlled.

Figure 7-3 illustrates how the short track-life criterion has slightly fewer false targets in both cases. The 4/24 MND with long track-life detects 34 of the targets with 160 tracks versus 55 false tracks and is closely followed by the 4/24-2/4 MND with short track-life which detects 33 targets with 140 tracks versus 42 false tracks. All four trials are very close, detecting between 32 and 34 of the 36 targets with 128 to 160 real target tracks versus 42 to 55 false target initiations.

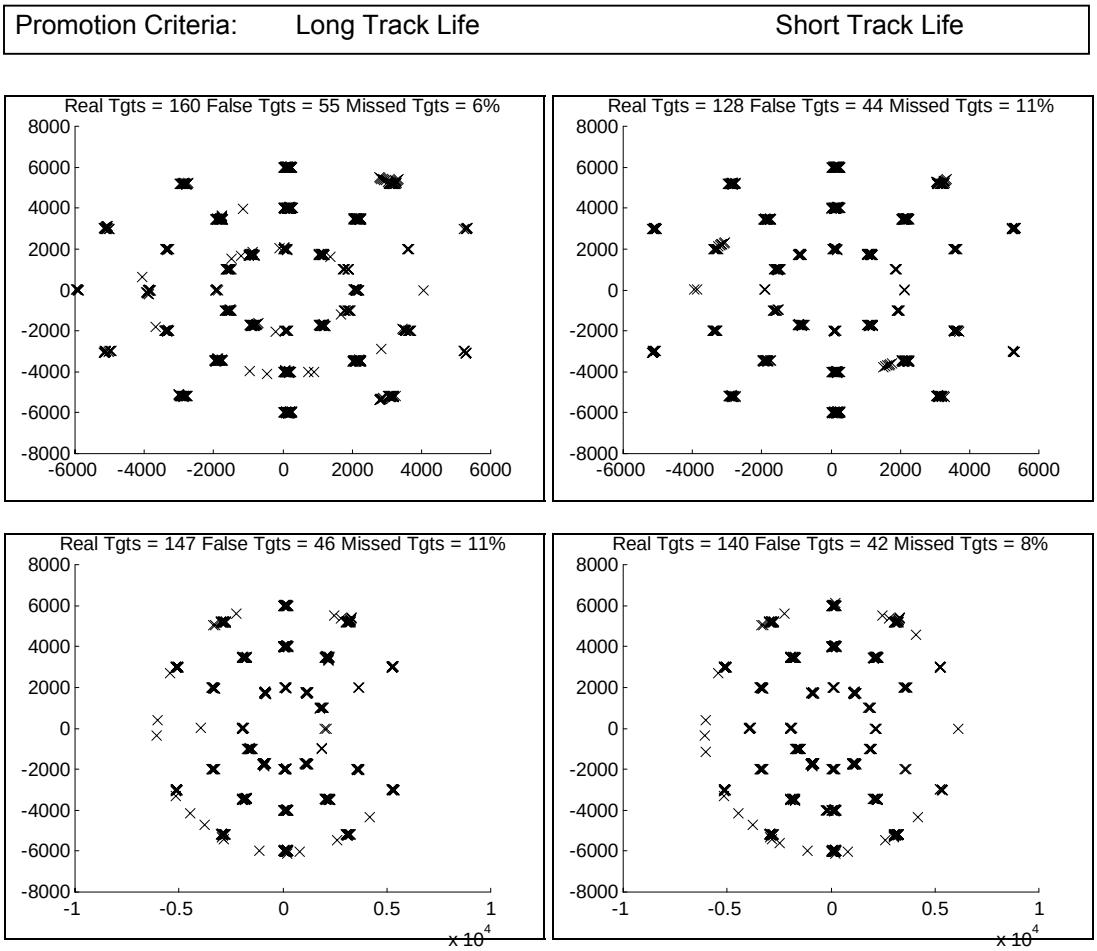

Figure 7-3 4/24 (top) and 4/24-2/4 (bottom) MND with Long and Short Track-Life & 1 Hypothesis

With little target seduction present, increasing the number of hypotheses allows us to "buy back" much of the performance lost when the short track-life criterion case is used. This is a significant result, since recent algorithmic improvements ensure that increasing the hypothesis numbers by moderate amounts will not significantly slow performance. In practice, the total number of tracks and cluster sizes are the time-critical components of the tracker.

Figure 7-4 demonstrates that while the number of real targets does increase (most were detected anyway with 1 hypothesis), the number of real tracks doubles with almost no increase in false tracks. The best result of all the trials is achieved with a 4/24-2/4 MND and MHT which detects all but one target.

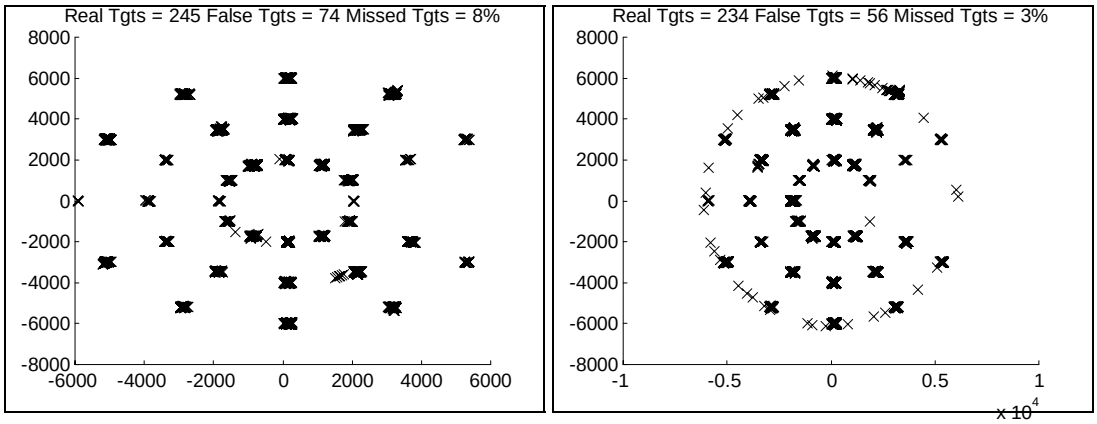

Figure 7-4 4/24 and 4.24-2/4 MND with Short Track Life & 3 Hypotheses

7.3 Summary and Preliminary Conclusions

Individual targets may be tracked several times during testing particularly when multiple hypothesis tracking changes hypothesis. The number of tracks is less important than how many discrete real and false targets are detected. A useful measure of SAR tracking success is therefore the ratio of real to false targets shown in Figure 7-5. The tendency of the tracker to assign multiple tracks to a single target is the measured real track to real target ratio which is then used to deflate the number of false tracks to an estimate of the false target count. The proportion of the 36 real targets that are detected is shown in Figure 7-6 for the same cases.

The top two curves of Figure 7-5 and Figure 7-6 illustrate that high radars on larger vessels (or aircraft) with better visibility can use a tighter MND criterion (4/24 in this case) to reject higher input false alarms yet still detect many more targets and fewer false ones. The benefit of loosening the MND ratio from 4/24 to 6/96 is clear, particularly when the number of hypotheses is increased from 1 to 3. Note that the longer correlation period is essential to limit the false alarms and pick out sustained recurring targets.

The bottom two curves show how radars on smaller boats with less visibility must balance the need for looser MND filtering (3/24) against the increased number of MND false alarm detections. Shortening the MHT track-life criterion is seen to increase the proportion of real targets but reduce their absolute number. This indicates that track seduction in the MND and the MHT is the problem. Loosening the MND ratio from 3/24 to 6/96 is seen to reduce but not solve the problem. Increasing the number of hypotheses degrades performance because the numerous false alarms have more opportunities to seduce otherwise accurate tracks.

The summary conclusions of this preliminary trade-off study are:

- high MRI false alarms must be rejected by the MND to limit MHT track seduction
- low MND ratios (M/N) increase track seduction within the MND which passes more false detections
- longer MND correlation periods are beneficial
    - to detect recurring target detections
    - to reduce the MND ratio for lower visibility targets
    - as long as target drift does not decorrelate detections
- shorter MHT track-life increases the proportion of real targets to false ones but reduces their absolute number in proportion to the degree of track seduction (if any)
- increased MHT hypotheses only improve performance if false MND detections are few and track seduction modest

Monte-Carlo trials are required to define the optimal processing but the above results demonstrate that the SARAIT is robust and can perform well with considerable latitude in the processing parameters. Real data with awash targets and severe competing false alarms will be processed by the SARAIT in Section 9 to further elaborate the many parameter trade-offs.

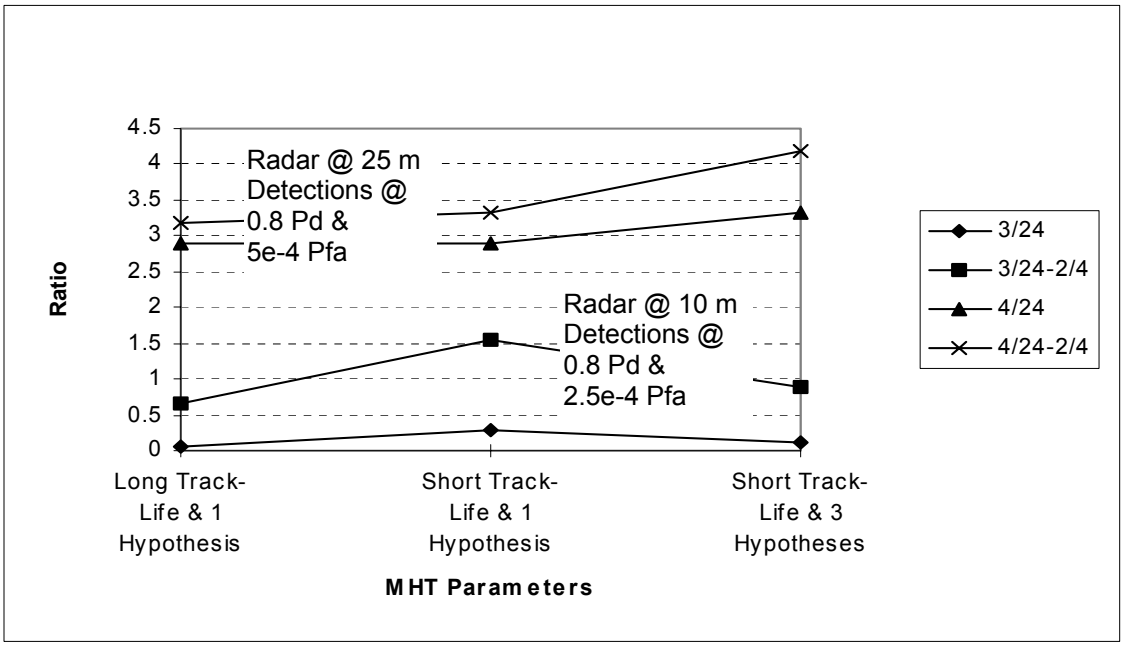

Figure 7-5 Ratio of real to false targets

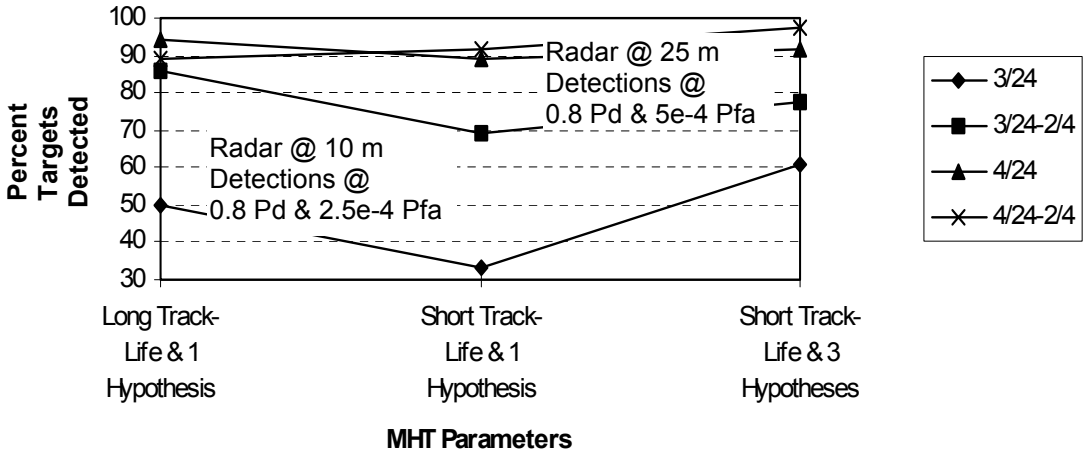

Figure 7-6 Percentage targets tracked

8. SAR Data Recording

Oceans Limited of St. John's was separately contracted by TDC to prepare for and carry out a field trial off the east coast of Newfoundland with the support of the CCGS *J.E.Bernier*. An extensive set of radar data was recorded of calibrated targets in both tethered and drifting configurations. Staff from Sigma Engineering were on board to operate the MRI and from Raytheon Canada to test the SARAIT with live and recorded MRI data.

The performance of any radar processing technique is initially determined by the radar data and the experimental conditions. To accurately assess performance, data was gathered that spanned all the acquisition variables summarized in Table 8-1.

Table 8-1 Data acquisition variables

Variable	Expected Range	Notes
Sea Characteristics		
Sea State	0 to 4+	Depends on weather during trial
Wave Shape	Uniform to Steep	Depends on wind/wave/current geometry
Target Characteristics		
Radar Cross section	0.03 to 0.15 sm 0.2 to 0.6 sm	Typical unaugmented PIW Typical liferafts and augmented PIWs
Extent	1 range/azimuth cell	Point targets
Drift Velocity	0 to 1 kn	Tethered and drifting targets

Dispersion	about 0.4 nmi apart typically 100 m apart	All tethered targets are resolved Multiple survivors & debris fields
Search Vessel Characteristics		
Speed	6 to 10 kn	Typical search speeds
Bearing to Target versus Wave Direction	360 degrees	Test all geometries (large difference from upwind to downwind to crosswind)
Radar Mast Blockage	Ahead vs Astern Port vs Starboard	Test targets @ both to quantify degradation No difference expected
Radar Mast Height	25 m (typical)	One vessel available
Radar Processing		
STC	On or Off	Test both
Pulse Filtering	On or Off	Test both

8.1 Radar Target Characteristics and Ground Truth Data

Twenty-five calibrated radar targets were tethered 0.4 nmi apart in a 5 x 5 square array with a wave-rider buoy (of modest but uncalibrated RCS) tethered 1.5 nmi to the North. Six sizes of target were used:

- A6 Target 0.47 sm Qty 1
- A5 Target 0.31 sm Qty 6
- A4 Target 0.19 sm Qty 6
- A2 Target 0.09 sm Qty 6
- A0 Target 0.03 sm Qty 6
- Wave-Rider 0.5 sm estimated Qty 1

Additional A0, A2 and A4 targets were deployed in various drifting constellations throughout the trial. Flashing lights and a radio beacon were used on the A4 and larger targets to facilitate recovery.

Tests at Cape Spear indicated that a diver's head had an RCS between 0.03 and 0.09 sm and a diver in a lifevest an RCS of about 0.14 sm. The A0, A2 and A4 targets should therefore span the expected range of worst-case SAR targets.

All targets were moored with 191 m anchor cables in about 145 m of water. Allowing for water depth variations, and with no a priori knowledge of the direction of target drift, Oceans Limited calculated the maximum expected "watch circle" diameter to be 260 m. The centre of this watch circle has a deployment error of about one ship length, or 70 m, which must be added in when estimating the potential error in the ground truth array positions.

8.2 Tethered Array Data Gathering

Target detectability with marine radars varies greatly depending on the viewing angle relative to the waves and the wind; visibility is lowest up or downwave and highest crosswind. Performance is also expected to vary with ship speed and blockage effects. A scripted search pattern was therefore sailed that examined all targets in the square array from all possible viewing geometries and ranges in two stages:

1. A perimeter search offset 0.25 nmi from the array with the pulse length changed at the mid-point. The objective was to test how targets are acquired at different ranges and viewing geometries relative to the waves. Target range varies much more slowly than does the viewing geometry. On any leg of the perimeter search, the furthest targets will be 2 to 2.5 nmi distant and therefore around the expected limit of detectability. The closest targets will be 0.25 to 0.5 nmi distant and therefore highly likely to be detected

2. A stepped search with 3 nmi extensions at each corner as shown in Figure 8-1. The objective was to test how quickly targets are acquired and lost as the range (but not the geometry) changes rapidly. The pulse length is changed at the end of each radial to compare short and long pulse performance under comparable geometries and sea states.

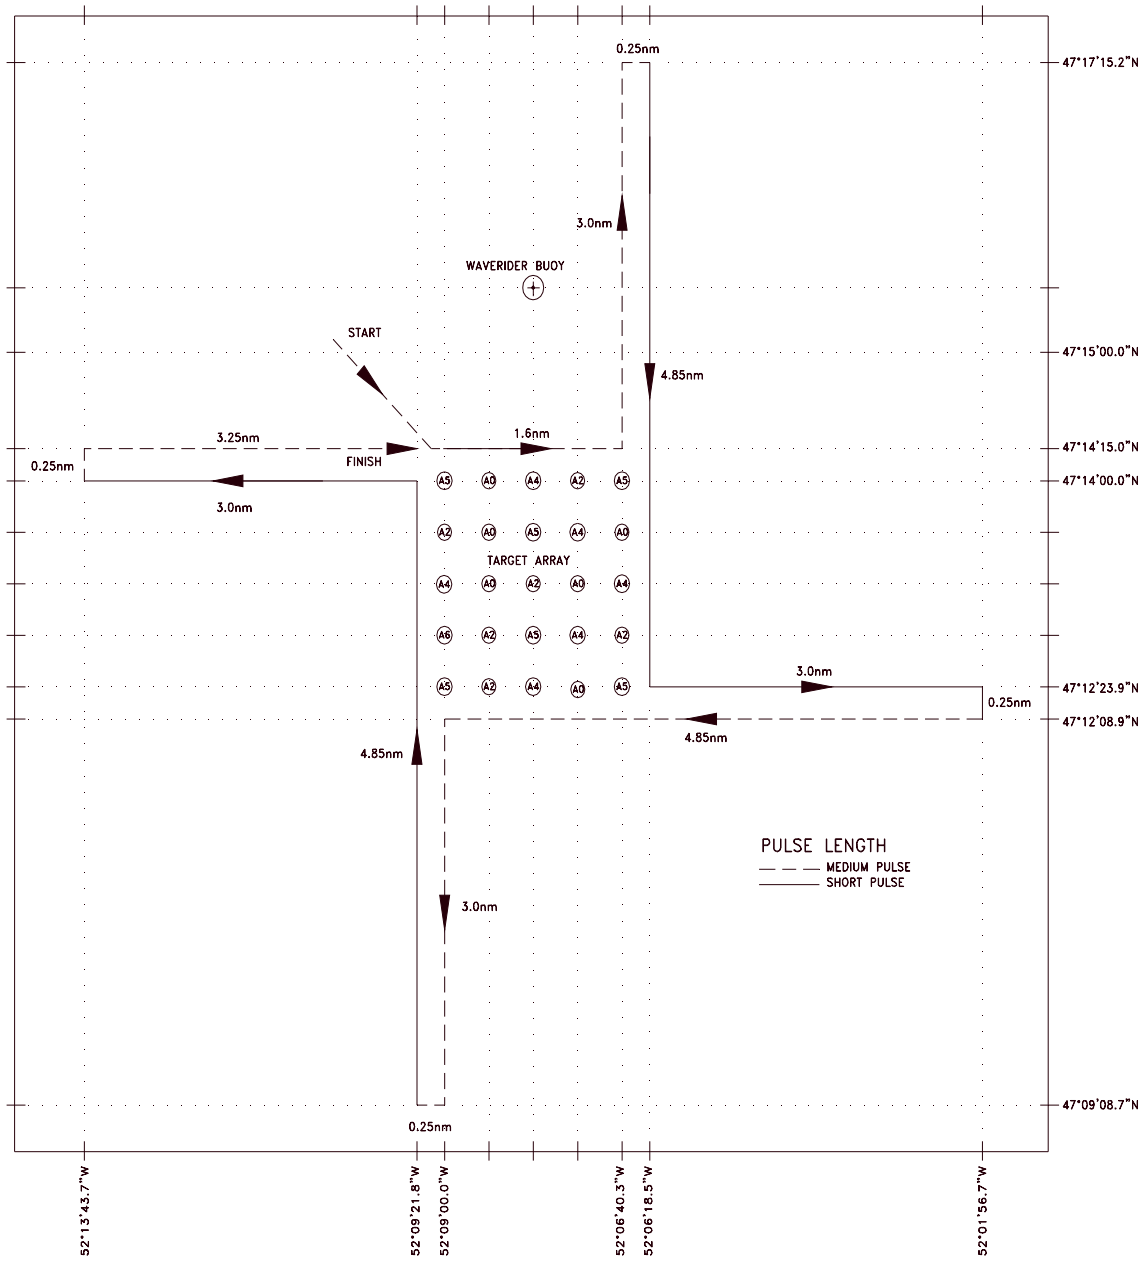

Figure 8-1 Stepped search data gathering sailing pattern

8.3 Drifting Targets Data Collection

Disposable targets were cast adrift en route to (and from) the tethered array and monitored as they receded astern and as they were approached on recovery. DGPS deployment and recovery positions were accurately recorded in a log along with the estimated speed and direction of drift (from experience or charts).

The objective was to measure the performance against drifting targets, both individually and in clusters (both resolved and, seen from differing geometries relative to the wind and swell. To this end, the targets were approached on recovery, and left after deployment, from different directions and in a wide variety of sea states, winds and currents. The target displacement from deployment to recovery will be used to estimate the average drift rate and interpolated position fixes.

8.4 Environmental and Radar Data Log

The following environmental measurements were logged every half hour throughout the trial:

- Local Time;
- Visibility (nmi);
- Cloud (tenths);
- Weather;
- Combined significant wave height and period;
- Significant wind wave height and period;
- Predominant swell height and period;
- Wind speed and direction;
- Air temperature;
- Dew point temperature;
- Sea temperature; and
- Ship speed (kn) and heading.

Throughout the sailing pattern, scrupulous records were kept by Oceans Limited of the timing and nature of all changes to the radar settings, MRI processing and MRI taping.

9. Preliminary Field Testing

The MRI and SARAIT processing is completely parameterized to better explore how best to detect SAR targets in a diversity of sea and weather conditions. The numerous inter-related processing parameters shown in Table 9-1 must be rigorously tested and the results analysed to properly quantify the achievable performance of the integrated 120 rpm Pathfinder/STII - MRI - SARAIT system under the varied conditions of Table 8-1.

An exhaustive combinatorial testing of all the parameters of Table 9-1 was prohibitive but a series of increasingly focussed tests with the complete MRI-SARAIT system usefully quantified the relationships between parameters. Due to the limited time and resources, attention was immediately focussed on long-range detection in high sea states under both low and high wind (and clutter) conditions. Two sets of medium pulse (3 nmi range) data in 3.8 and 3.3 m waves were analysed to explore the performance trade-offs and the results are presented below.

The sequence of tests progressed from the most basic parameters to the most subtle:

1. MRI Scan Average processing, which exploits short-term correlation to increase target strength at wave crests while reducing false alarms,
2. MRI CFAR processing, since with no detections, nothing is possible,
3. MND processing, since long-term correlation reduces the false detection rate and enables practical tracking, and
4. MHT processing, since multiple hypotheses can reduce the false tracking rate while reliably detecting SAR targets.

Table 9-1 Processing Parameters

Variable Processing Parameter	Expected Range	Notes
MRI PARAMETERS		
Plot Extraction Parameters		
Threshold		
MinBlobSize (Resolution Cells)	2 or 3	All targets are less than radar resolution
MaxBlobSize	2 to 20	Larger clusters may add clutter & error
ScanAveraging Enabled	"On" or "Off"	
ScansToAverage	if "on", 2 to 8	Increased integration planned for the MRI
CFARWindowLength	2, 4, 8 or 16	Determined by wave period and clutter
CFAROffset	20	Fixed
CFARRank	20 to 80	Determines false alarm density vs range
SAR AI TRACKER PARAM		
Max. SAR Target Velocity	0 to 1.5 m/sec	Limits expected drift (0 to 3 kn) and spurious correlation in high false alarms
M1, N1	2 of 10 to 15 of 40	10 to 30 % visibility over 1 or 2 waves
M2, N2	1 of 1 to 4 of 4	Promote multi-wave detection sequences
M3, N3	1 of 1 to 4 of 4	Outer loop to extend M of N window to 1 minute or more
Number of Hypotheses	1 to 4	More hypotheses for high FA tracking

9.1 Testing Methodology

For testing, the SARAIT is operated in stand-alone mode with the MRI playback software resident on the MND processor. Test runs are scripted with a text file as shown in Table 9-2.

Table 9-2 SARAIT test script

M1	N1	M2	N2	M3	N3	R min	Az min	R max	Az max	No. Scans	Yr	Mo	Dy	Hr	Mn	Sec
3	20	2	2	2	2	1500.0	225.0	10000.0	45.0	2500	97	11	17	03	25	00
7	40	1	1	2	2	1500.0	225.0	10000.0	45.0	2500	97	11	17	03	25	00
6	40	1	1	2	2	1500.0	225.0	10000.0	45.0	2500	97	11	17	03	25	00

Accurate log records are kept during the data gathering so that parameter changes can be reliably correlated with the processed record.

The following data are output to the screen (and optionally logged) by the SARAIT MND processor once per scan:

- number of input MRI plots,
- M of N processing time (for diagnosis),
- ship position,
- time (UTC), and every N1*N2 scans, and
- number of output MND detections.

The MHT processor receives the accumulated M of N detections every N1*N2 scans and implements MHT for which the following data are output to the screen (and optionally logged):

- location of current and past MND detections, confirmed tracks and current ship position on a map,
- number of input MND detections,
- number of MHT deleted, potential, tentative and confirmed tracks,
- MHT image size and processing time (for diagnosis), and

- elapsed time to process and time stamp interval.

The results of each playback are an MND data file (Table 9-3) and an MHT output file (Table 9-4).

Table 9-3 MND detection file

		Average			Least Squares Fit			
X (m)	*Y (m)*	*X var*	*Y var*	*XY var*	*X var*	*Y var*	*XY var*	*Scan #*
2343.283	906.539	-16.862	-20.241	689.081	3307.303	-1196.01	635.208	5

Table 9-4 SARAIT track output file

Init #	*Trk #*	*X*	*Y*	*Xdot*	*Ydot*	*Xvar*	*Yvar*	*Track Length*	*Track Update*	*Track Misses*
1	1	3123.81	-3472.67	0.0202	-0.6966	513.034	559.563	4	4	0
1	1	3117.88	-3469.92	-0.4335	-0.0179	652.955	774.335	5	5	0

Mon	*Day*	*Yr*	*Hr*	*Min*	*Sec*	*Start Lat*	*Start Long*	*Ship Lat*	*Ship Long*
11	17	97	2	27	18	47.2019	-52.0649	47.2019	-52.0716
11	17	97	2	27	38	47.2019	-52.0649	47.202	-52.0729

These two test files are then analysed as required by the separate Analysis System. The Analysis System is scripted to automatically process the Test File to compare the separate ground truth data with the SARAIT tracks and, if required, with the M of N detections. The number of targets detected and their range and azimuth are automatically extracted along with the false alarm rates (Rfa in false detections per hour) at the outputs of the MRI, MND, MHT and, after rejecting obvious wake and rain clutter artifacts, at the output of the SARAIT.

These tests were done with very limited resources and the results are therefore preliminary. The indicated SARAIT performance must not be taken as anything more than a lower bound on what is achievable. Refinements in the MRI and SARAIT algorithms and in the selected processing parameters are expected to further improve performance. Much more effort will be dedicated in early 1999 to a separate analysis phase which will rigorously determine the best achievable performance under the full range of field conditions.

9.2 Medium Pulse SAR Detections in 3.85 m Seas and Light 3 kn Winds

Tape 22 (Segment 4) was recorded in low clutter conditions with a heavy swell running. The Wave-Rider buoy measured a 3.8 m 9.8 second period swell from 040 degrees. Light 3 kn winds from 140 degrees superimposed small 0.1 m wind driven waves on top of the swell creating a combined 3.85 m sea. Most of the clutter arose from small rain cells slowly drifting through the test area and from sea birds alighting on the gentle seas. This data set is ideal for exploring how best to detect distant awash SAR targets in heavy seas that largely obscure them but introduce little competing clutter.

A series of tests were run to explore the effects of the key MRI and SARAIT processing parameters, particularly at long ranges. Figure 9-1 shows a typical test run with the SAR detections overlaid on the ground truth array positions. All of the array targets are detected by the time the ship has reached the south-east corner of the array. Note that by Tape 22 the array has been in place for over two weeks and several of the larger targets have dragged their moorings.

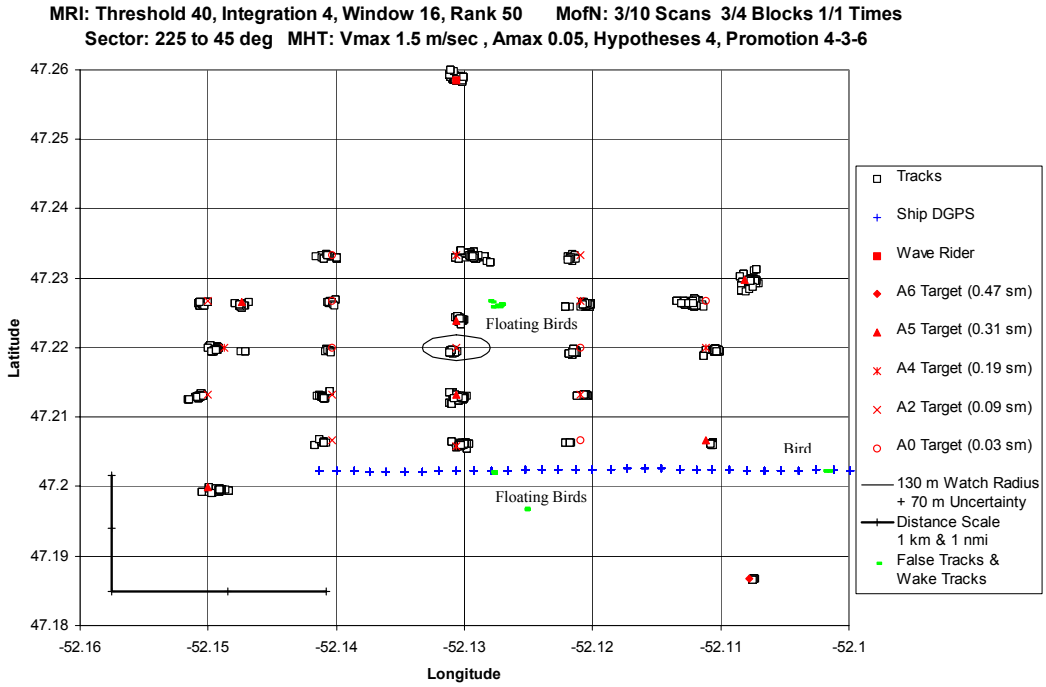

Figure 9-1 SAR target tracks superimposed on the array ground truth (Low Clutter)

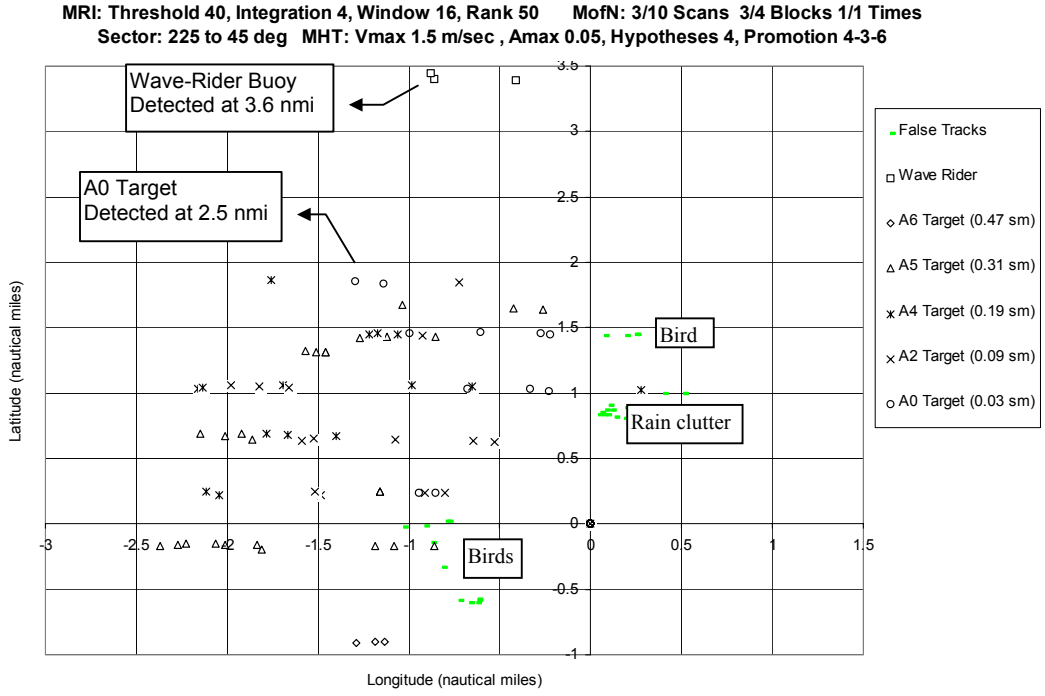

Figure 9-2 SAR target track initiation, range and bearing from the ship (Low Clutter)

Figure 9-2 shows the range and bearing at which the SARAIT initiated tracks on each class of array target. Note the long ranges at which detections occur particularly the Wave-Rider buoy 3.6 nmi to the North and the most distant A0 target 2.5 nmi to the north-west.

Flying birds could be seen as sparse sequences of MND detections because the MND correlation period is only 20 seconds and the required visibility only 9/40. Once the birds were on the water, the MHT initiated tracks on them because of their slow drift speed. Rain cells were also seen as dense clusters of MND detections because of their slow drift in the prevailing 3 kn winds. The MHT eventually initiated several tracks as shown in Figure 9-2. In practice, the AI sub-system in the SARAIT could easily be adapted to reject such tracks based on the extremely high MND detection density.

9.2.1 MRI Multi-Scan Integration

The MRI can linearly integrate up to 8 scans before detection and plot extraction. Linear integration in the absence of clutter effectively improves the signal to-noise ratio for all but the most distant and least visible targets as shown in Figure 9-3 for 1, 2, 4 and 8 scan integration.

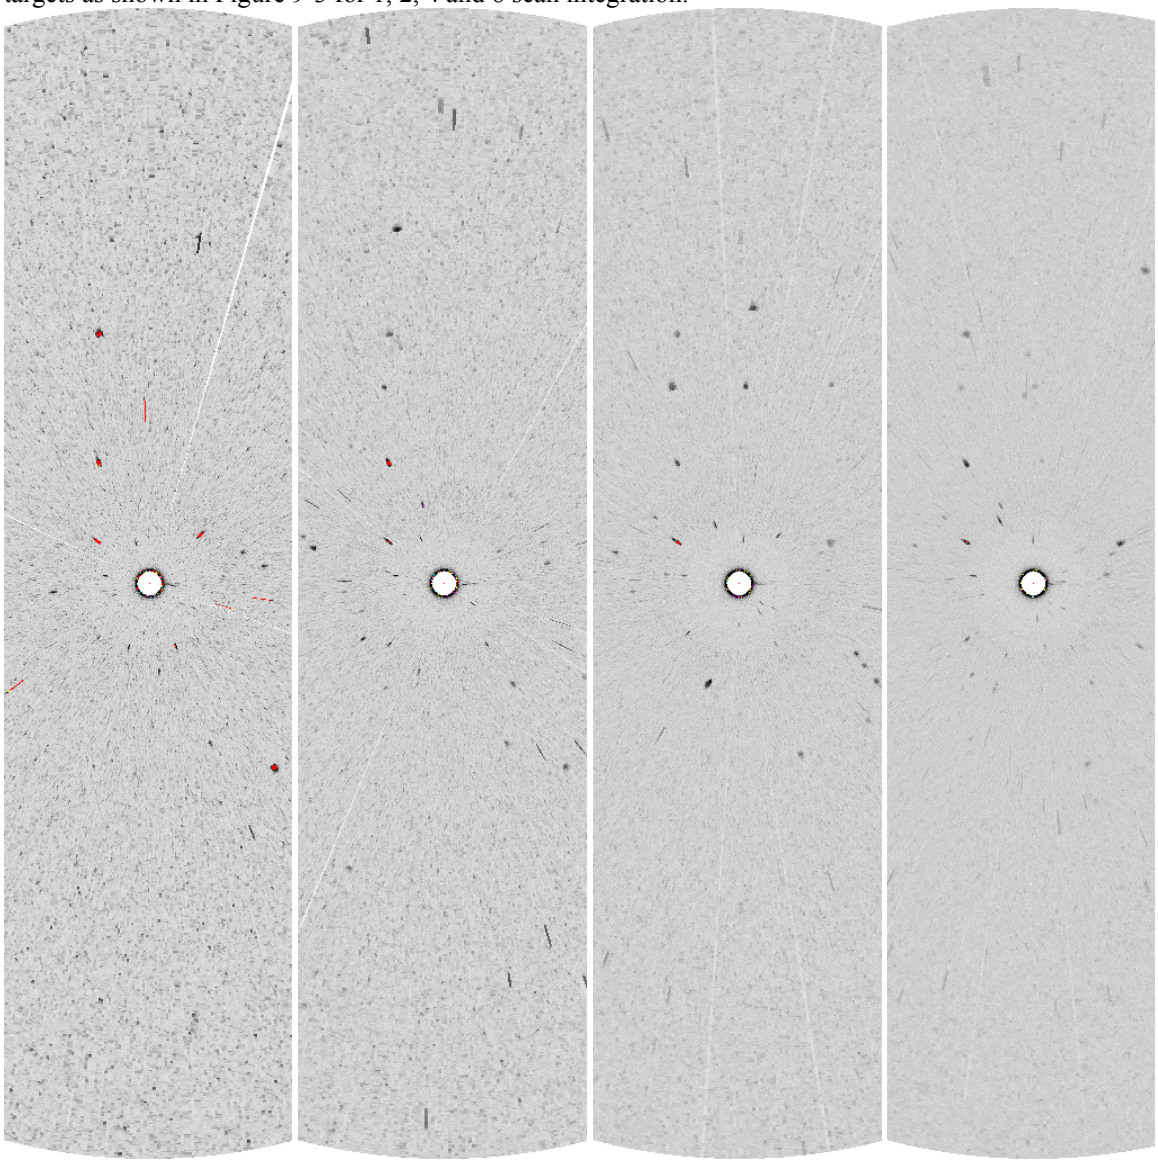

Figure 9-3 MRI integration of 1, 2, 4 and 8 scans (negative image with array to North)

The low clutter conditions ensure that the background noise is white with range out to 3+ nmi and make the constant false alarm ratio (CFAR) detection straightforward. In Figure 9-3, note that the many very small features evident in the single scan are progressively averaged out as more scans are integrated. The array targets themselves have little extent and the smaller ones are indistinguishable from the noise and sea clutter.

Figure 9-4 shows that for the 10 second waves of Tape 22, with identical 3/10,3/4 MND processing of approximately 900 MRI plots per scan, the longest detection range and lowest SARAIT false alarm rate occurs with 4-scan integration. This indicates that most targets at 2 to 3 nmi are detectable for 3 to 4 scans (i.e. 1.5 to 2 seconds) at every wave crest or for an average of 20% every 10 second wave period. Integrating over 2 scans does not average down as much of noise which increases the SARAIT false alarm rate four-fold without maximizing target detection at any range. Integrating over 8 scans does not reduce the false alarm rate any more but limits the detection performance by adding clutter and noise (rather than additional target detections) for half the scans.

The MRI would need to integrate 20 scans to be sure of detecting awash targets at least once on every integration period and to thereby reduce the effects of target occlusion to a minimum. Long MRI integration over one wave period, and preferably several, could therefore significantly improve the performance in low clutter conditions such as these.

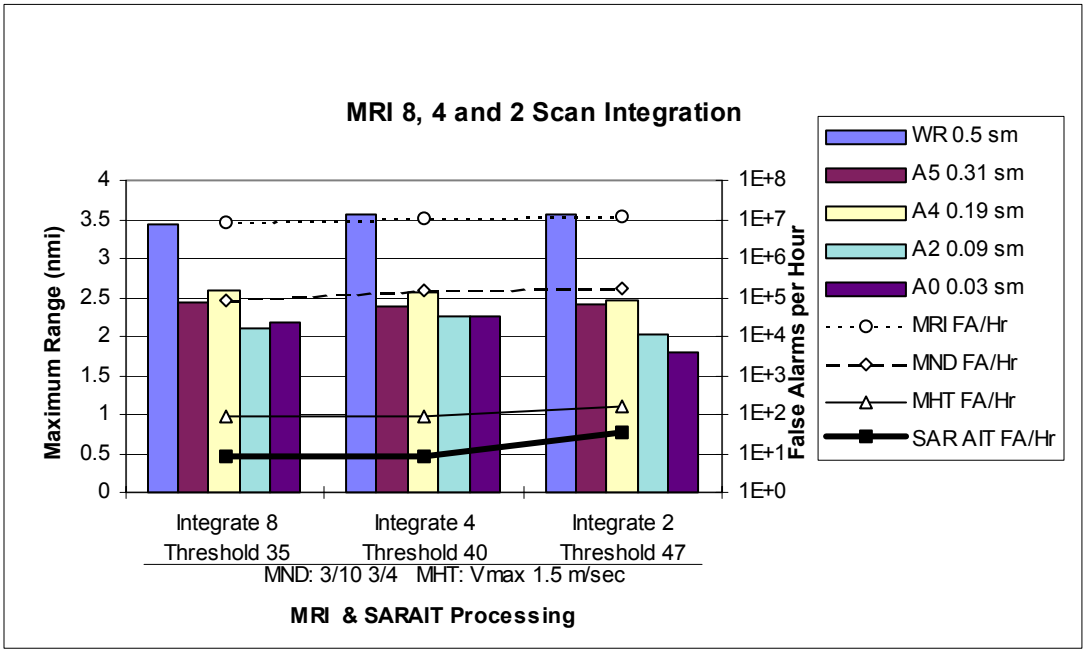

Figure 9-4 MRI integration of 8, 4 and 2 scans in low clutter

9.2.2 MRI Constant False Alarm Rate Detection

In this low clutter case, target detection is limited by receiver noise (which is Gaussian and uniform with range) rather than sea clutter (which has a K or similar long-tailed distribution that varies significantly with range and azimuth). Figure 9-5 illustrates that in this particular case of white background noise, the longest 16 range cell CFAR window is the best statistical estimator in that it delivers a slightly larger effective SNR at all ranges. Gaussian noise is fully described by the average (which equals the median) so varying the CFAR rank had no discernable effect.

Note the presence of numerous birds in the array and to the southeast and also the absence of any obvious detections of the Wave-Rider buoy (at top centre) or the top row of the array. Note also that only the nearest targets are detectable in each of the 3 scans.

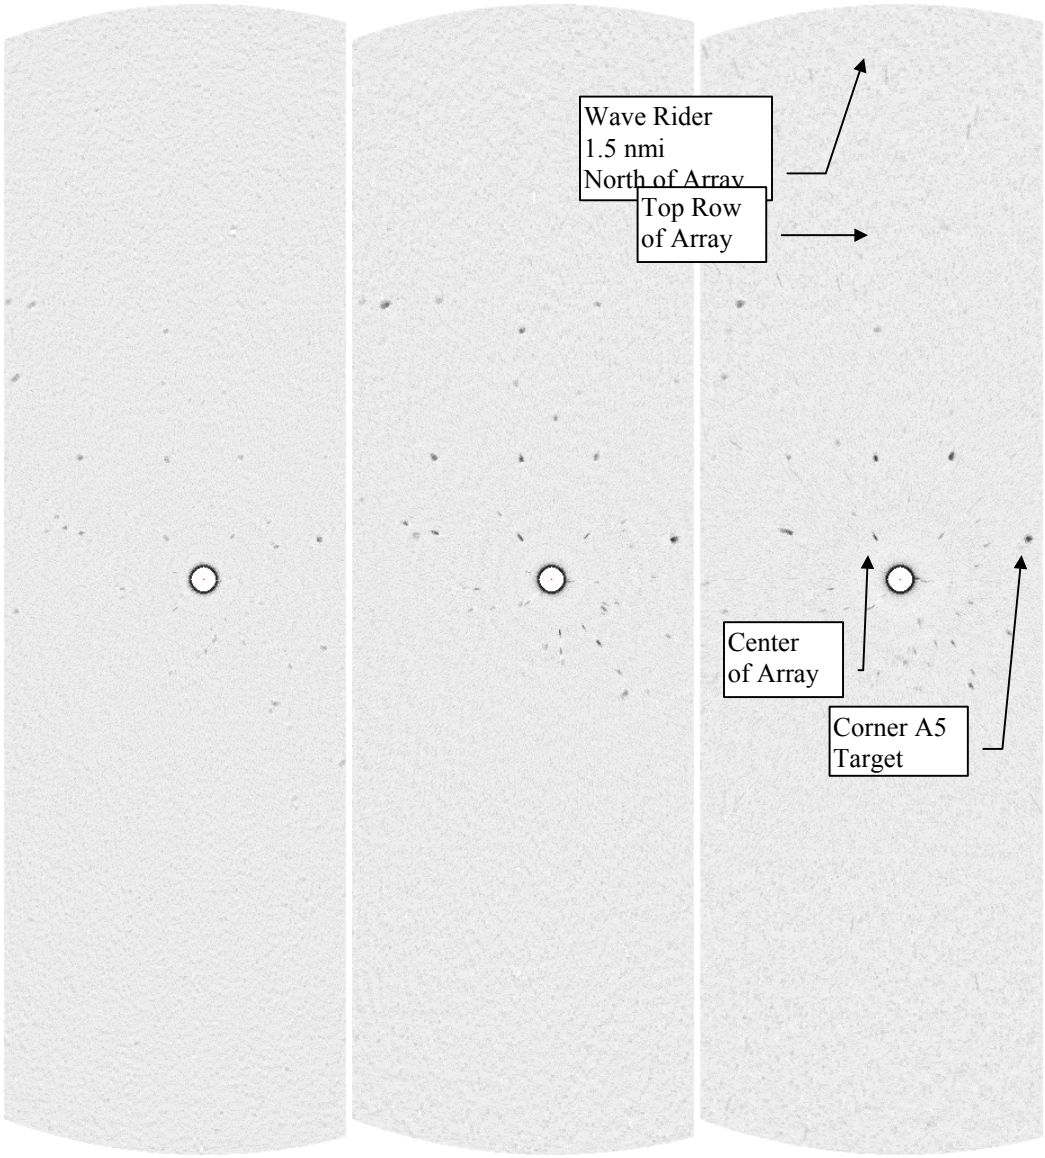

Figure 9-5 MRI CFAR window of 4, 8 and 16 range samples (8 scan integration & negative image)

Lowering the MRI detection threshold increases the number of plots input to the SARAIT. Most of the additional plots are due to clutter but some small fraction will be due to SAR targets which can therefore be used to improve performance. There is typically an optimum MRI threshold that delivers the best balance between the SARAIT false detections and detection range. Figure 9-6 illustrates how lowering the MRI threshold by 2, from 57 to 55, enables the SARAIT to detect the wave-rider buoy 3.5 nmi distant with no additional SARAIT false alarms. There is a modest 50% increase in the MRI, MND and MHT false alarm rate (Rfa) but almost all of these additional MHT false alarms are due to rain clutter which is easily rejected by the SARAIT. A further 2 point reduction to 53 increases the MRI and MND Rfa by another 50% but multiplies the MHT Rtfa by 3 and the SARAIT Rtfa by 10. This threshold behaviour is typical of the SARAIT, as it is of other non-linear systems.

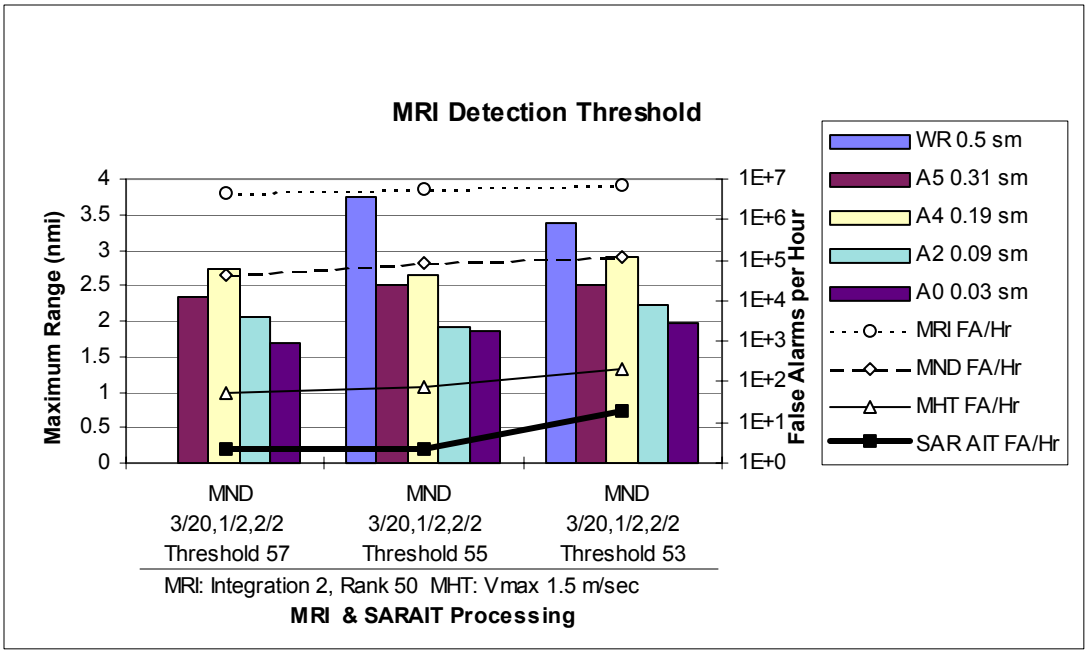

Figure 9-6 Setting the MRI plot threshold to maximize the proportion of small targets

9.2.3 MRI Plot Extraction

The MRI can exclude those plots that are outside a specified range which allows the MRI to be matched to the expected target extent.

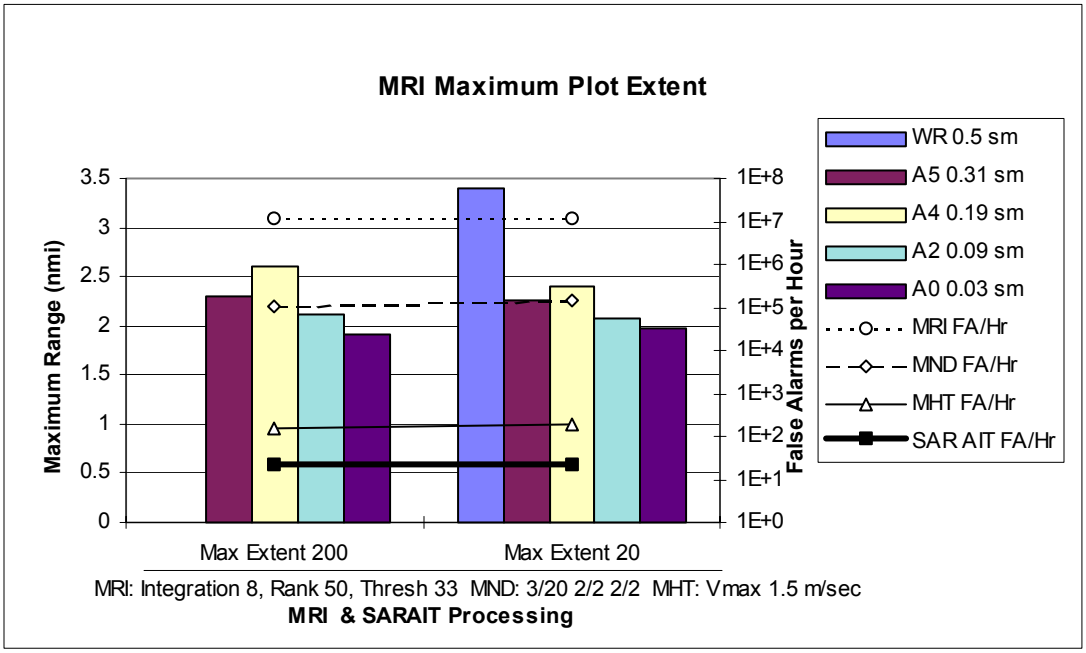

Figure 9-7 Matching the MRI plot extent to the SAR targets

Large clutter features above 20 or 50 cells are usually due to major wave crests, especially at longer ranges. These are considered unlikely to be useful SAR detections even if the target is included because the large number of clutter detections in the centroid will greatly perturb the position. Similarly, very small plots spanning only 1 or 2 cells are most likely due to noise.

Figure 9-7 illustrates how opening the MRI plots to accept unrealistically large features slightly reduces performance against the smallest or the most distant targets while improving it against the nearest large targets.

9.2.4 SARAIT M of N Detection

The SARAIT MND inner sliding window can be configured as a two-stage process to better match SAR plots that are expected to arrive spaced on average by the wave period. For this case of a 9.8 second wave period, Figure 9-8 illustrates how the 10-second correlation period 3/20,2/2,2/2 MND out performs the 20-second 5/40,1/1,2/2 MND. The problem with the 5/40,1/1,2/2 MND is that there are so many clutter detections that two successive 5/40 detections are common. The 3/20,2/2,2/2 MND is more effective because the chances of 3/20 detection coinciding on four successive wave periods is much less. This improvement occurs despite the slightly higher 15% average target visibility of the 3/20,2/2,2/2 MND.

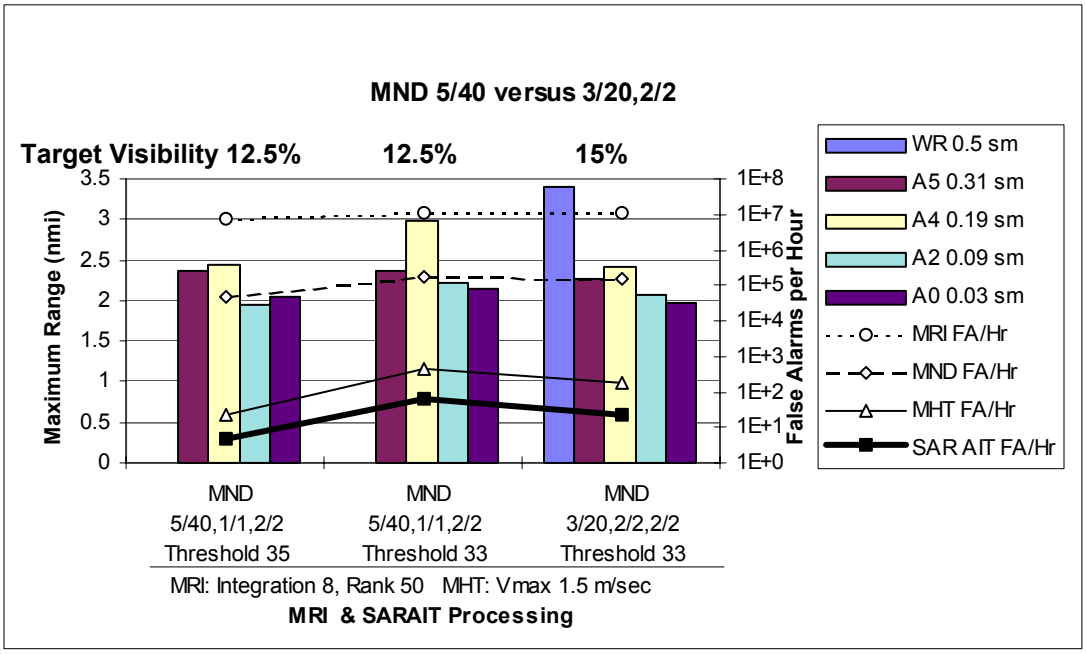

Figure 9-8 Matching the MND inner correlation window to one and two wave periods

The M of N average visibility criterion can be loosened to improve detection of the more distant and therefore less visible targets at some cost in near range detections. Typically, the number and density of MRI plots must be reduced to limit the overall false alarm rate. Comparing the 3/20,2/2,2/2 MND (15% visibility) of Figure 9-8 to the 3/20,1/2,2/2 MND (7.5% visibility) of Figure 9-9, demonstrates range increases of 0.35, 0.25 and 0.23 nmi for the larger targets (wave-rider, A5 and A4) beyond 2.5 nmi where target visibility is typically around 10%. Conversely the maximum detection range for the closer targets where visibility should be higher is reduced by 0.18 and 0.1 nmi for the A2 and A0 targets at 2 nmi.

Reducing the M of N criteria to 2/20,1/2 without reducing the input MRI plots is seen in Figure 9-10 to increase the MND Rfa by 2 and the MHT and SARAIT Rtfa by factors of 5 and 100. Increasing the overall MND correlation time to 120 and 160 scans is seen to have no beneficial effect. In fact, the longer 160-scan window further degrades the performance, probably due to excessive target drift during the MND correlation period.

Maximizing the detection range therefore requires operating the SARAIT at M of N criteria that are comparable to the expected visibility but care is needed to limit the MND Rfa to acceptable levels. For the

somewhat limited test runs of Tape 22, limiting the MND Rfa to 1E5 per hour enables the SARAIT to deliver less than 2 to 5 false alarms per hour.

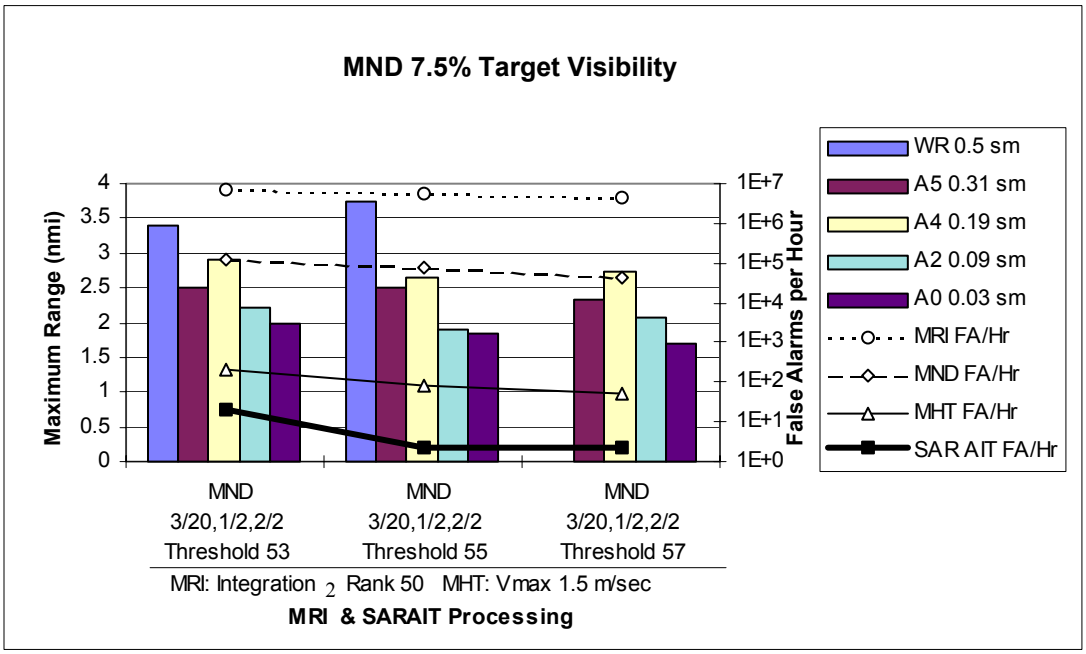

Figure 9-9 Matching the MND to the more distant target visibility

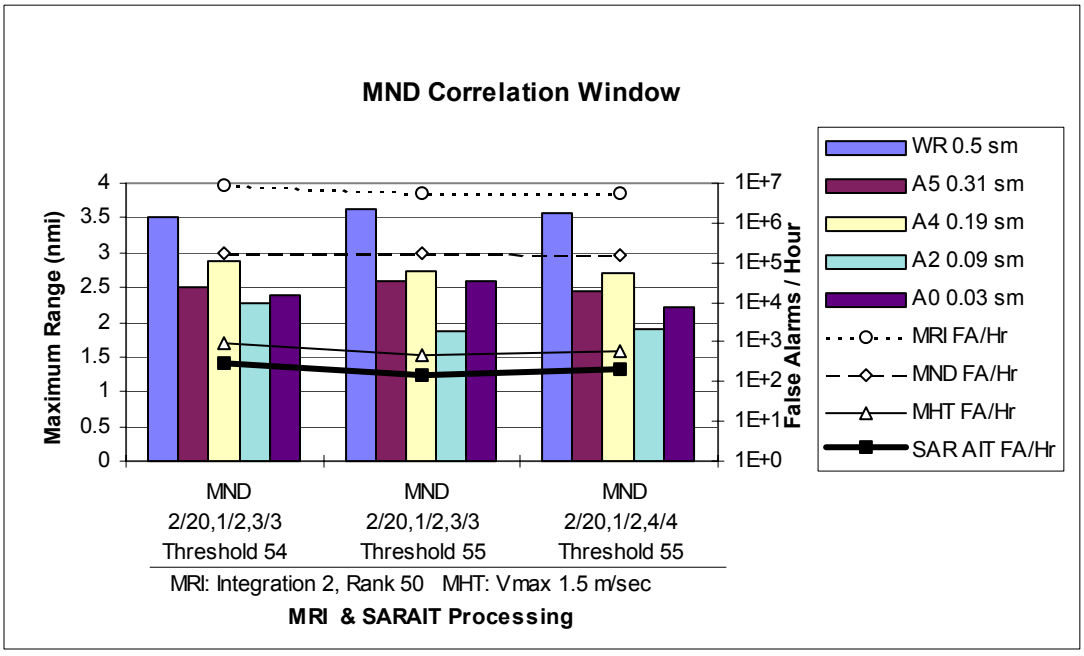

Figure 9-10 Limits to reducing the MND visibility

9.2.5 SARAIT Multiple Hypothesis Tracking

The SARAIT rejects potential tracks that move too fast to be credible SAR targets. The maximum track velocity (Vmax) is used to set the association radius in the MHT and MND and to limit the MHT track velocity. Figure 9-11 illustrates how setting the maximum velocity too low reduces the detection range for all sizes and ranges of targets.

The best balance between long range SARAIT detection and false alarms occurs with 1.5 m/sec (3 kn) Vmax. This is consistent with direct observation of the raw radar and MND detections and is probably due to the high 3.5 m waves pushing the targets back and forth around the circumference of the watch circle like guard dogs on a long chain. Figure 9-12 shows a typical tethered target, as seen by the radar in short pulse mode, moving through a 50 m arc to the southeast over a 284-second observation period. Although the *average* velocity is only 0.2 m/sec, the measured target position frequently jumps 15 to 25 m along the north-south axis in the 10 seconds between MND detections. This short-term velocity often exceeds 1.5 m/sec. The position jumps by more than the system error resulting from the radar accuracy (10 m in short pulse mode for nominal targets) and the much smaller DGPS errors and is therefore likely to reflect real but short-lived target motion. Because the axis of this rapid motion is similar to the swell direction, the commonly observed back and forth action of the high swell is probably indicated.

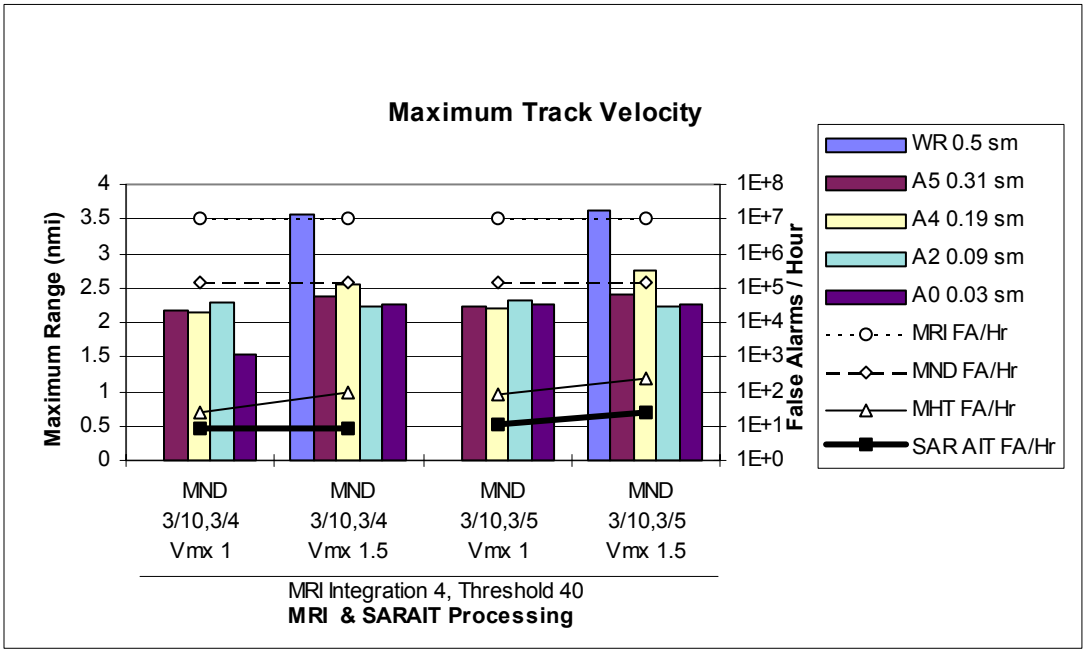

Figure 9-11 Matching the SARAIT maximum track velocity to the target drift in light clutter

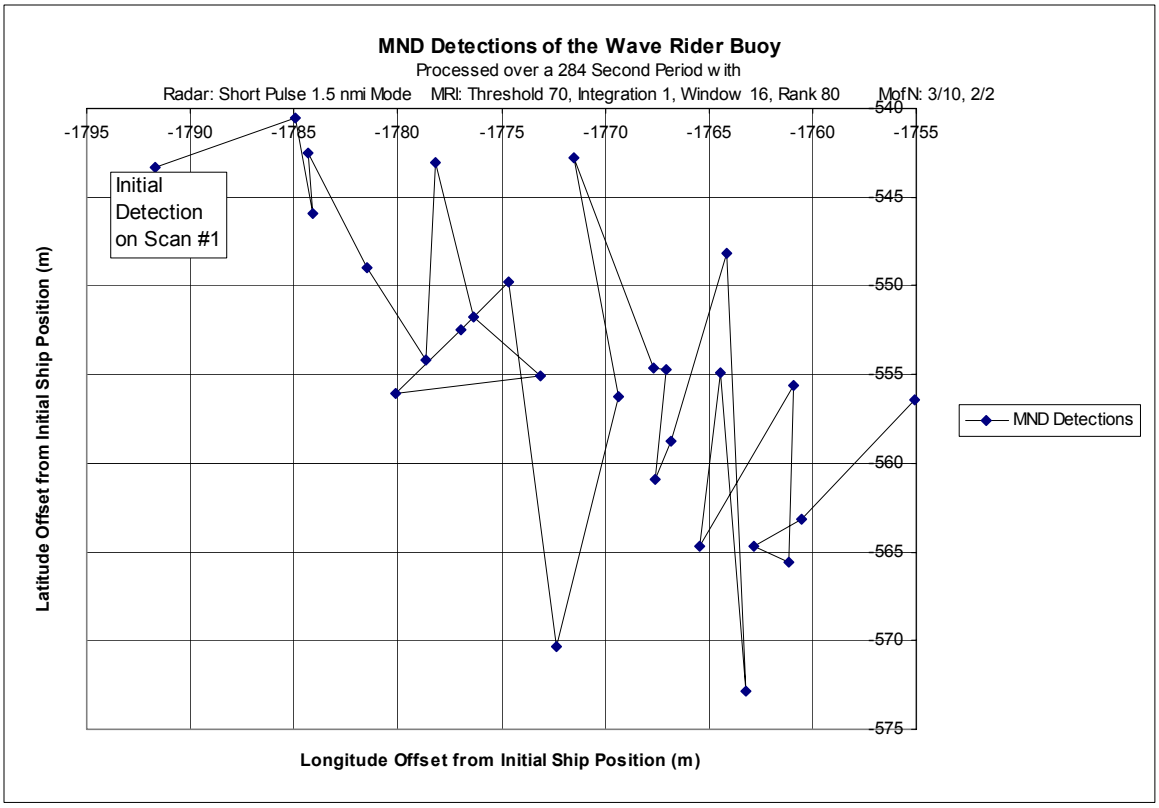

Figure 9-12 M of N Detections of a tethered Wave-Rider buoy showing slow drift in 3.8 m swell

9.2.6 Conclusions from Tests in High Seas with Weak Clutter

Figure 9-13 shows several alternative approaches to achieving a low SARAIT Rtfa of 2 false alarms per hour. The best performance is achieved when the MND matches the wave period and target visibility.

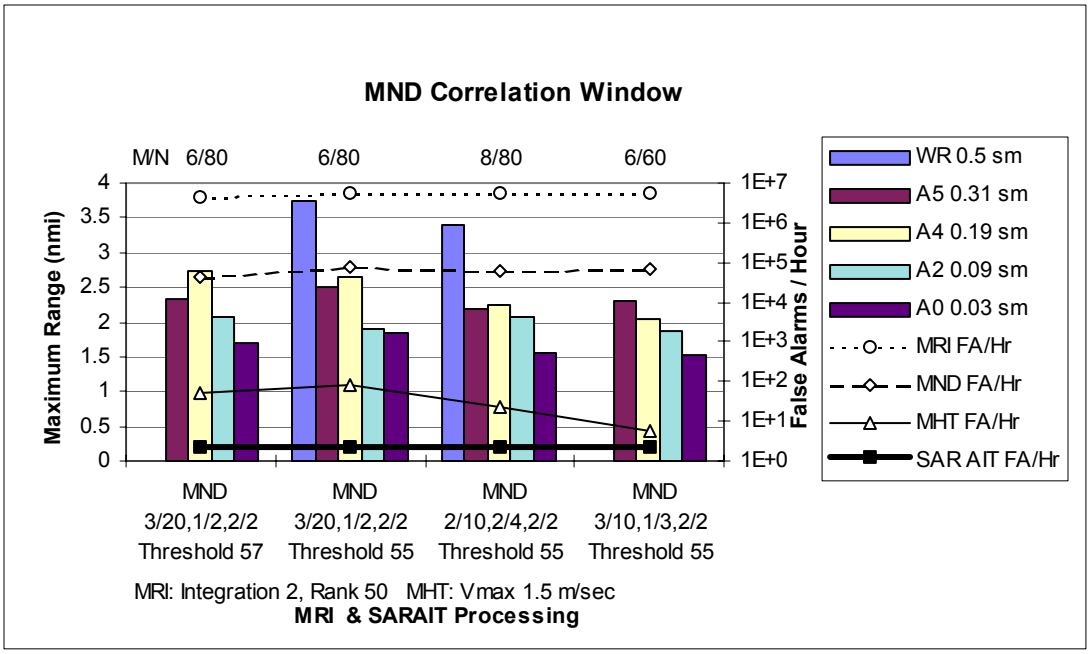

Figure 9-13 Matching the MND to the wave period and target visibility

Figure 9-14 demonstrates the essentially robust nature of the SARAIT. The Wave-Rider buoy is consistently detected at 3.5 nmi whether 2, 4 or 8 scans are integrated in the MRI. The SAR array targets are all detected at around 2 to 2.5 nmi irrespective of their size. This indicates that target visibility is the main limitation under low clutter conditions such as these.

Reasonable detection ranges with 2 to 20 false tracks per hour are therefore achieved in these large waves under low clutter conditions by:

- setting the MRI threshold for 750 to 1500 plots per scan (Rfa = 1e7),
- with an MND ratio from 7.5 to 25% giving 200 to 500 MND detections per interval (Rfa = 1e5/hr),
- with an inner MND window less than or equal to the average wave period (10 seconds or 20 scans),
- with an overall MND window spanning several wave periods, and
- with a maximum track velocity of 1.5 m/sec (3 kn) to allow for sudden swell driven motion.

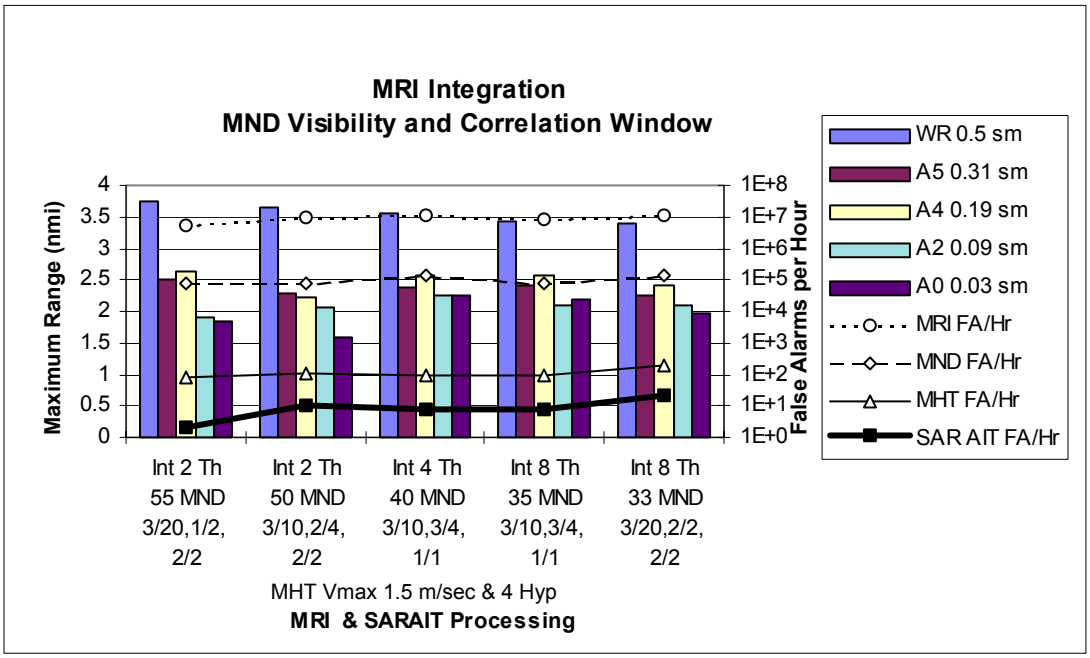

Figure 9-14 Typical results for diverse MRI and SARAIT settings

9.3 Medium Pulse SAR Detections in 3.3 m Seas and Gusty 10 kn Winds

Tape 21 (Segment 1) was recorded three hours earlier, than Tape 22 (Segment 4), when winds were still blowing gustily at 10 kn and veering from around 280 degrees at the start of the record to 0 degrees at the end. The 0.3 m wind-blown seas were therefore somewhat confused and the sea clutter intense. The underlying 3.1 m swell was again from 40 degrees but had a slightly shorter 8.6 second period.

Detecting the SAR targets under these conditions was particularly difficult because the wind was blowing strongly across the swell and therefore along the troughs where target visibility would normally be highest. Figure 9-15 shows the SAR target and false alarm tracks with the ship positions (for each track update) overlaid on the true target positions. The array targets are observed from most bearings as the ship sails 3/4 of a complete circuit in this segment. Figure 9-16 shows how the detection range varies in azimuth, reaching a maximum of 1.7 nmi midway between the swell and wind direction, a minimum looking into the steep-faced (since still building) wind-driven waves and between 1 and 1.4 nmi elsewhere.

Note also, on Figure 9-15, the presence of obvious wake tracks, as distant as 1 nmi, along the path of the ship. Figure 9-16 shows how the (non-wake) false alarms are mostly upwind with the balance downwind.

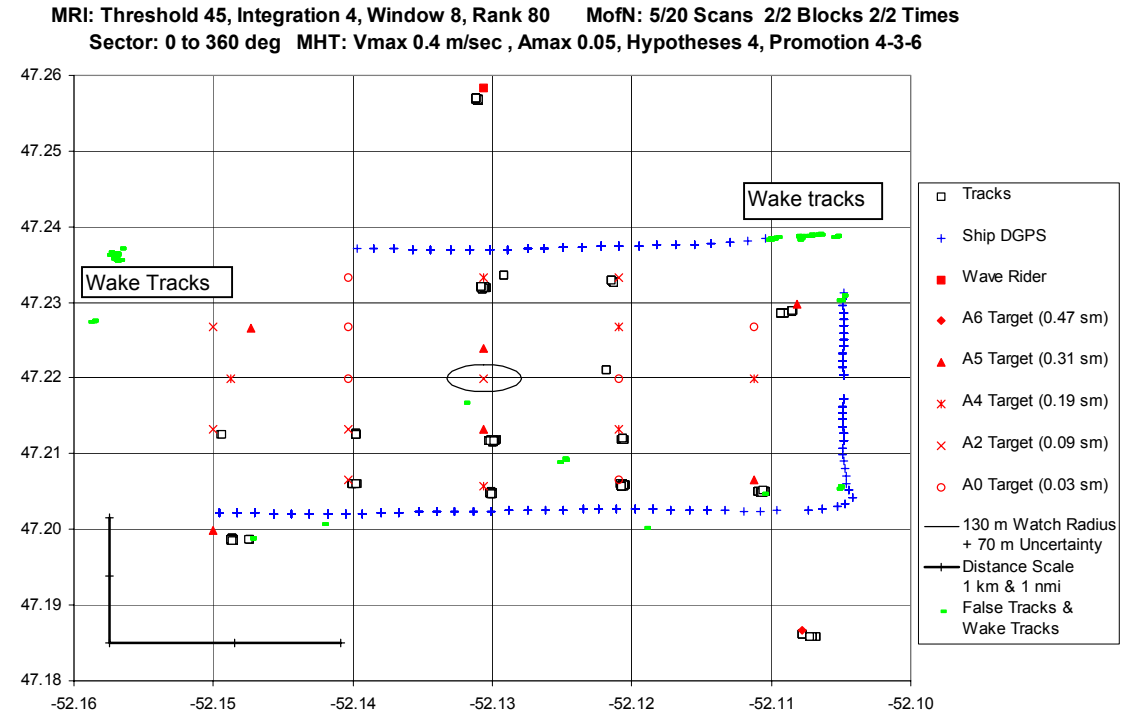

Figure 9-15 SAR target tracks superimposed on the array ground truth (High Clutter)

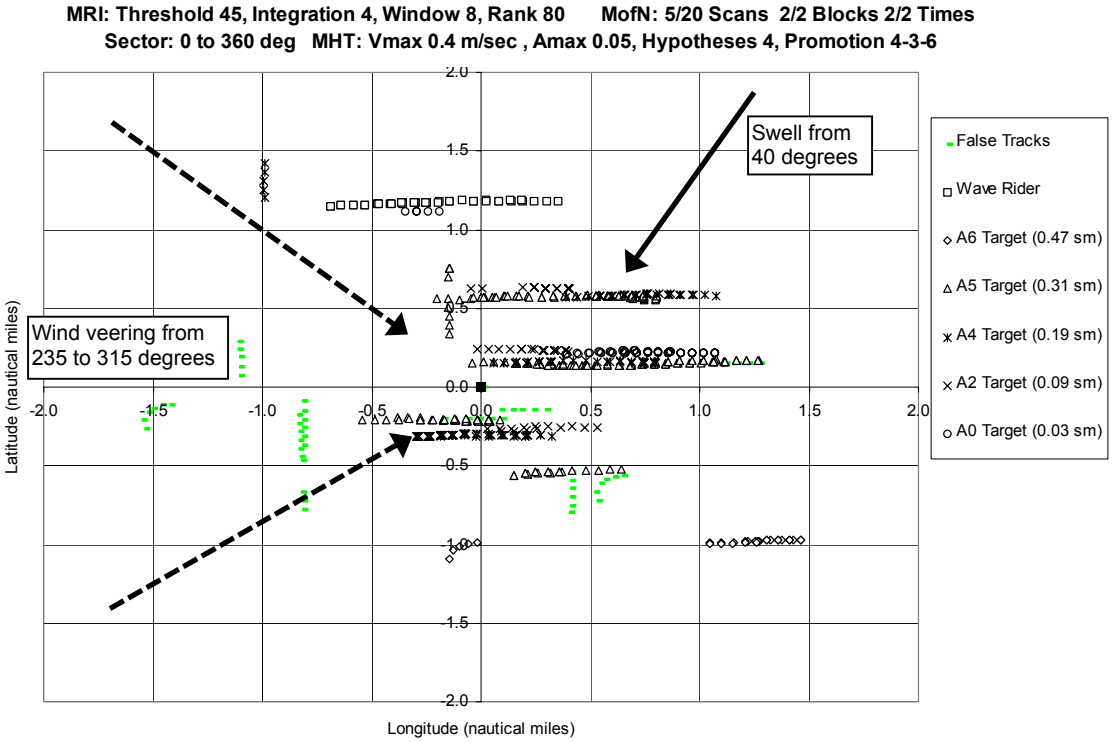

Figure 9-16 SAR target tracks, range and bearing from the ship (High Clutter)

9.3.1 MRI Multi-Scan Integration

Figure 9-17 shows the MRI PPI at the start of the segment when the wind was blowing strongly from 235 degrees. The steep-faced wind-blown waves raise the clutter level everywhere by making the swell more visible but this effect is particularly severe to windward. The smoothing effect of integrating two, four and eight scans in the MRI is evident as is the wake, particularly at the turn. The effect of integration on the targets depends on the local clutter. For example, the Wave-Rider buoy is in a low clutter region where integration can increase the SNR of larger targets but the array targets are in a high clutter region where integration is less effective or even, for the smaller targets, counterproductive.

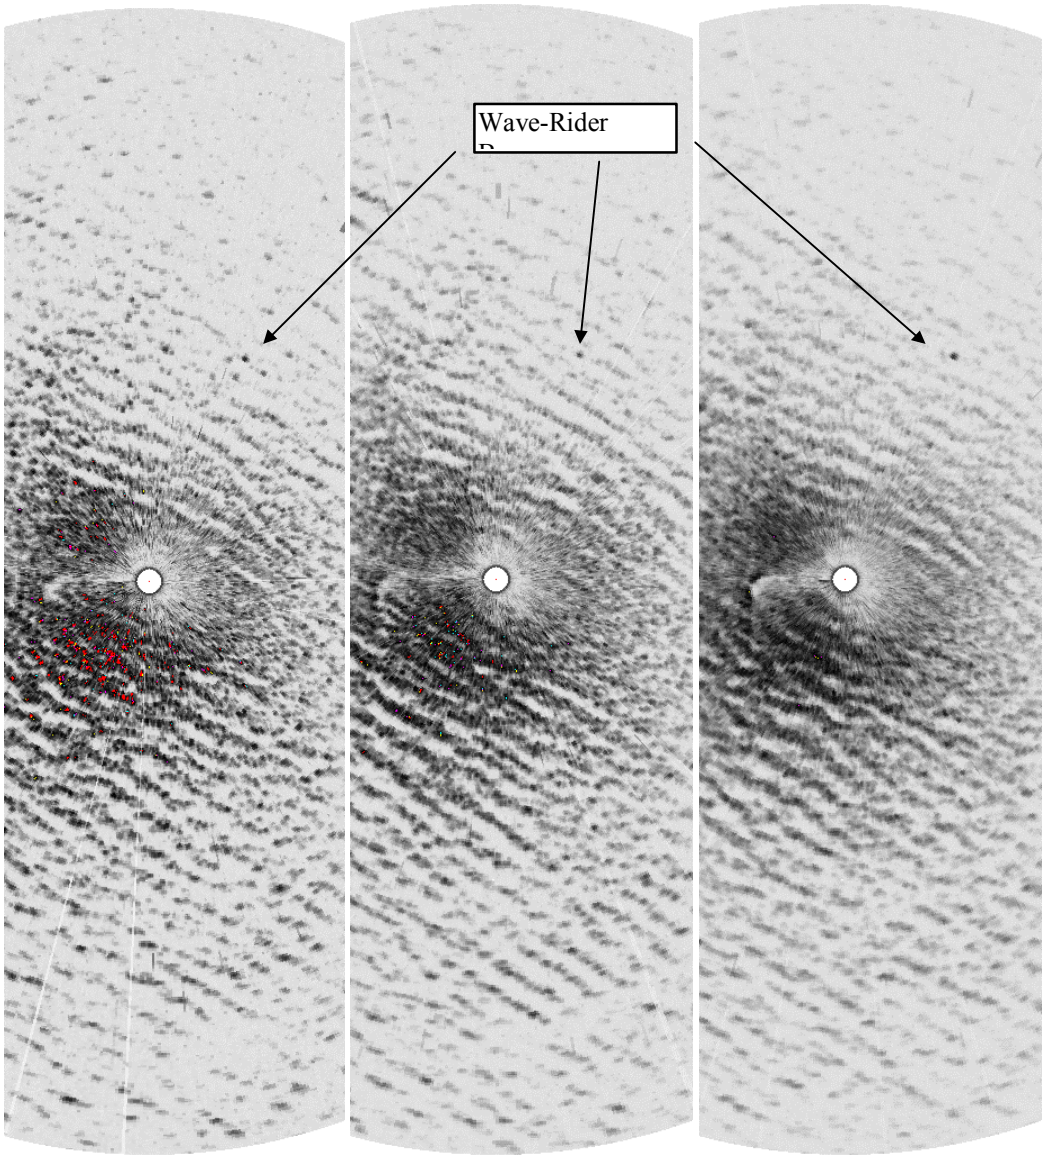

Figure 9-17 MRI Integration of 2, 4 and 8 scans

Figure 9-18 and Figure 9-19 show how the presence of clutter limits the benefit of scan-to-scan integration. Moving from single scan to two scan integration (with constant MRI Rfa) is seen to significantly improve detection of most targets but particularly the smaller A0 and A2 targets with no increase in the SARAIT false alarms. More wake tracks are also detected by the MHT but easily removed from the SARAIT tracks. Further increases from 4 to 8 scan integration are seen to increase the SARAIT false alarm rate four-fold

while only modestly improving detection of the largest Wave-Rider and A6 target (which are least perturbed by the additional clutter being averaged in) and degrading detection of the smallest A0 target.

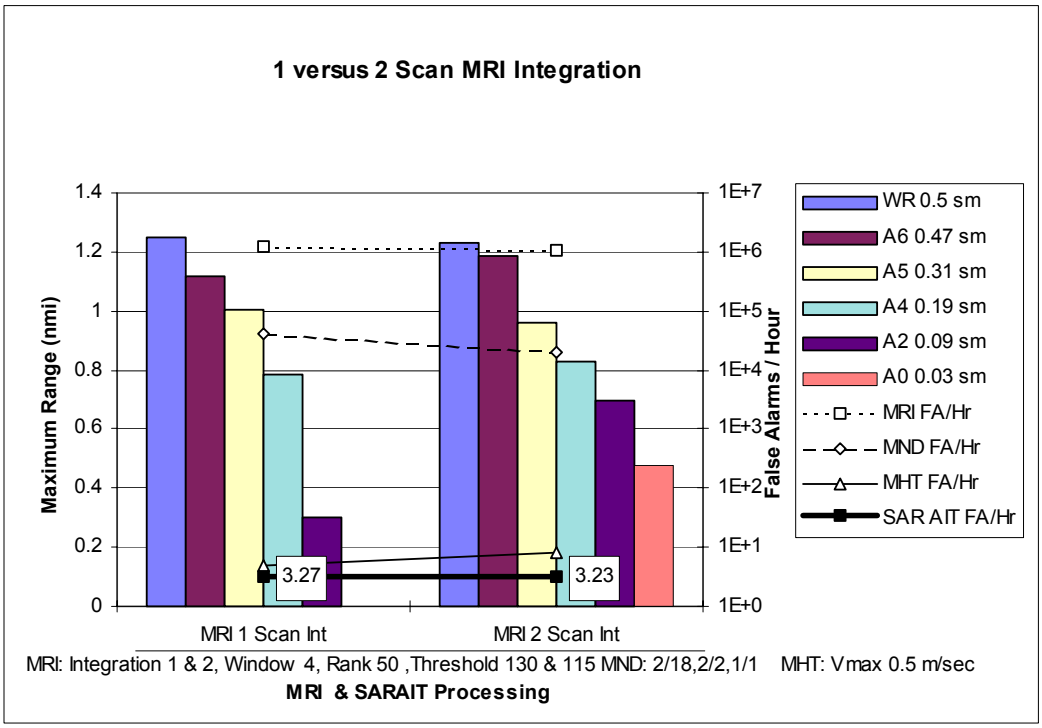

Figure 9-18 MRI integration of 1 and 2 scans in high clutter

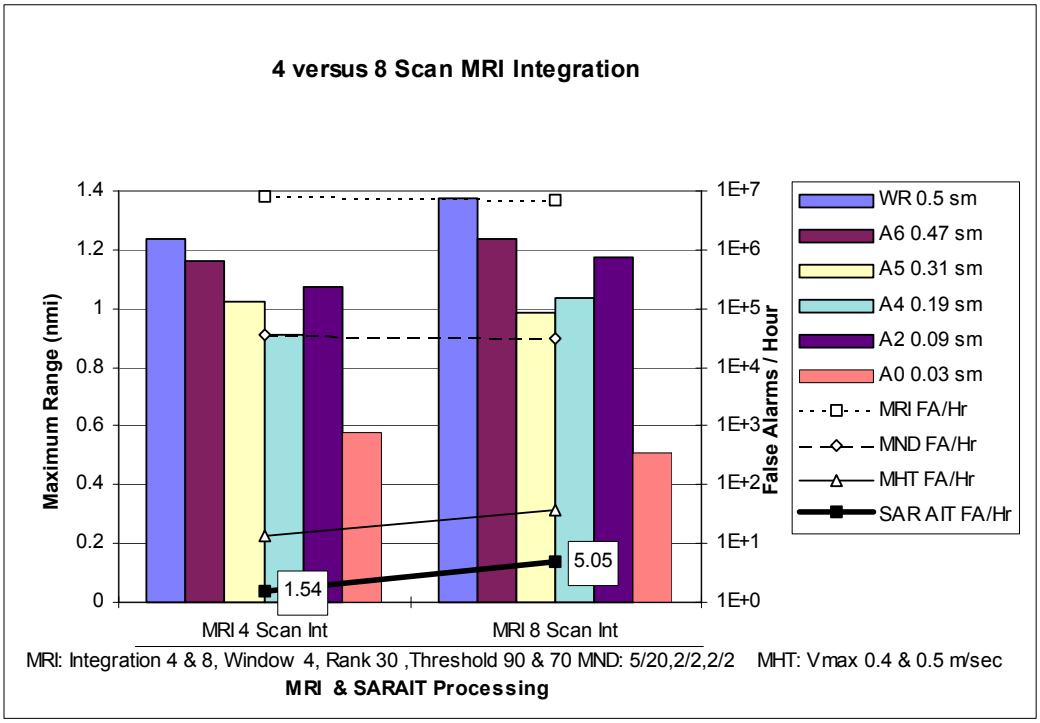

Figure 9-19 MRI integration of 4 and 8 scans in high clutter

9.3.2 MRI Constant False Alarm Rate Detection

Minimizing detections of the underlying swell, sea clutter and sea spikes while maximizing detections of targets requires differing CFAR strategies. As seen in Figure 9-20, shorter CFAR windows track the swell and therefore reduce the number and intensity of the swell detections while increasing the number and detectability of small, close-in, wind-driven features, particularly to windward. Smaller targets are also better detected, particularly near wave crests where neighbouring clutter may be strong. Longer windows better detect the larger wave features and higher SNR targets while rejecting more of the smaller near-in features and lower SNR targets.

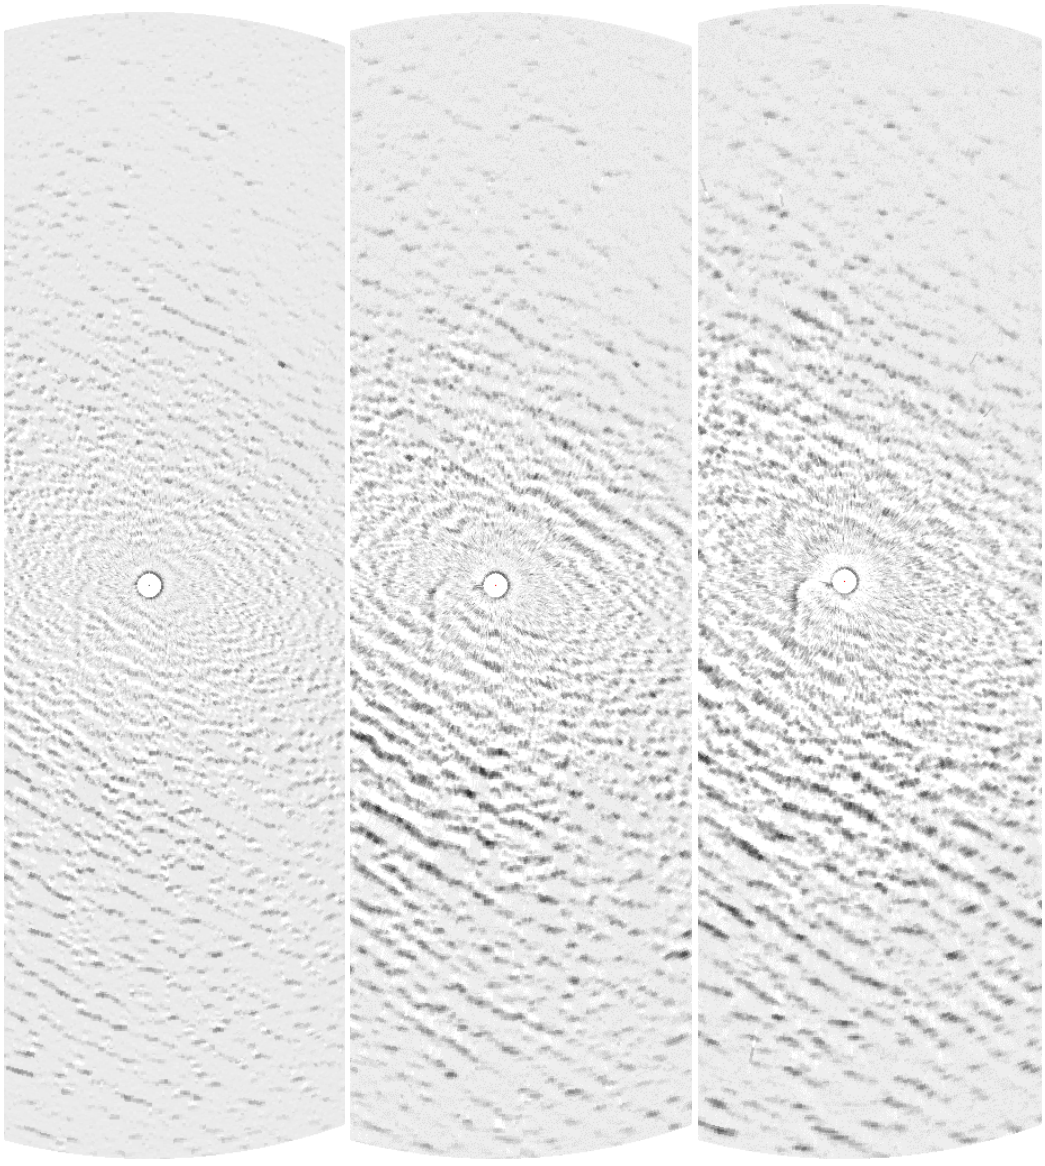

Figure 9-20 MRI CFAR window of 4, 8 and 16 range samples (4 scan integration & negative image)

Clearly, the performance depends on the particular target SNR and the surrounding clutter and is therefore range and azimuth dependent. Short of running parallel detectors or at least adapting the MRI CFAR in range and azimuth, a compromise must be struck based on the desired performance. In this particular case with 4 scan integration and a CFAR rank of 80, increasing the CFAR window from 4 to 8 reduces the number of SARAIT false alarms tenfold yet increases the detection range of all but the smallest A0 targets.

Further analysis comparing the effect of MRI integration, CFAR rank and window length will be needed to identify the best compromise for each target size at each range and azimuth.

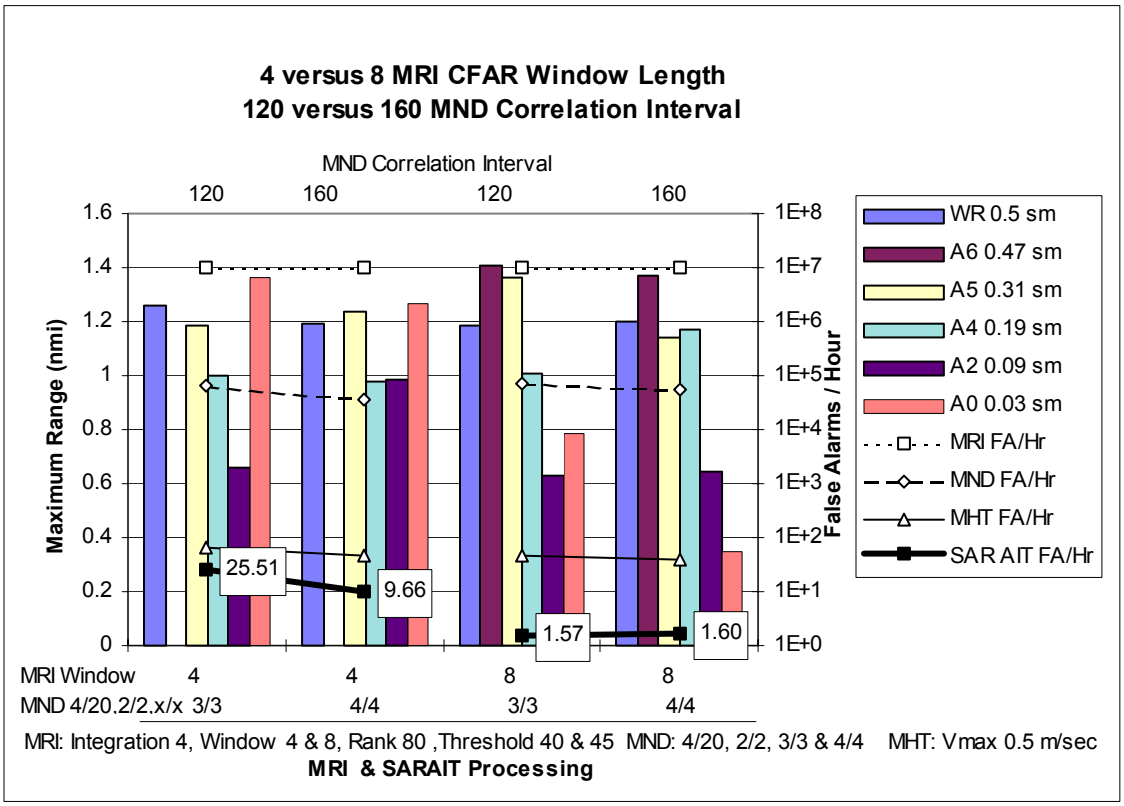

Figure 9-21 MRI CFAR window of 4 and 8

The MRI implements an ordered statistic CFAR detector to better reject non-Gaussian features such as sea spikes. For each range sample, the surrounding range samples within the CFAR window are ordered, the ranked (in percent) largest value is taken as the local background and subtracted from the range sample under consideration and a plot declared if the difference exceeds the threshold. A CFAR rank of 50 is equivalent to the normal cell averaged CFAR **if** the clutter is Gaussian.

Lowering the CFAR rank rejects more of the sea spikes and other intense features (typically having a long-tailed distribution) from the local clutter estimate which is therefore closer to the quiet background clutter level (typically shorter-tailed though not yet Gaussian). This means that numerous weak features including the smaller targets are detected, particularly at close range where the signal strength is highest, along with the usual strong returns from sea spikes. The density of the near-in clutter plots therefore increases. For good SARAIT performance, the number and density of the MRI plots must be limited to avoid excessive spurious correlation in the MND and a resulting loss of performance. In Figure 9-22 for rank 30 CFAR, 100 MRI plots per scan are correlated into 300 MND detections per update but these are very dense to windward with few occurring beyond 1 nmi. This is seen into artificially limit the detection range of the larger targets while ensuring that the smallest A0 targets are at least detected. However, the high MRI and MND detection density increases the SARAIT false alarm rate tenfold when compared to the rank 50 and 70 cases.

Increasing the CFAR rank estimates the local background based on the more intense clutter and therefore detects fewer of the smallest features. The MRI detections are more uniformly dispersed in range and azimuth because they are dominated by sea spikes and larger targets. This is not necessarily a disadvantage because the MND can still detect sparse sequences of detections from even the smallest targets particularly

when the MRI plot density is so reduced. Figure 9-22 shows how increasing the CFAR rank increases the maximum detection range for the larger targets but reduces it for the smaller ones. Raising the CFAR rank to 50 and 70 also reduces the SARAIT false alarm rate to 2 per hour. As usual, a compromise is required to detect both the smallest and the largest targets at both long and short range.

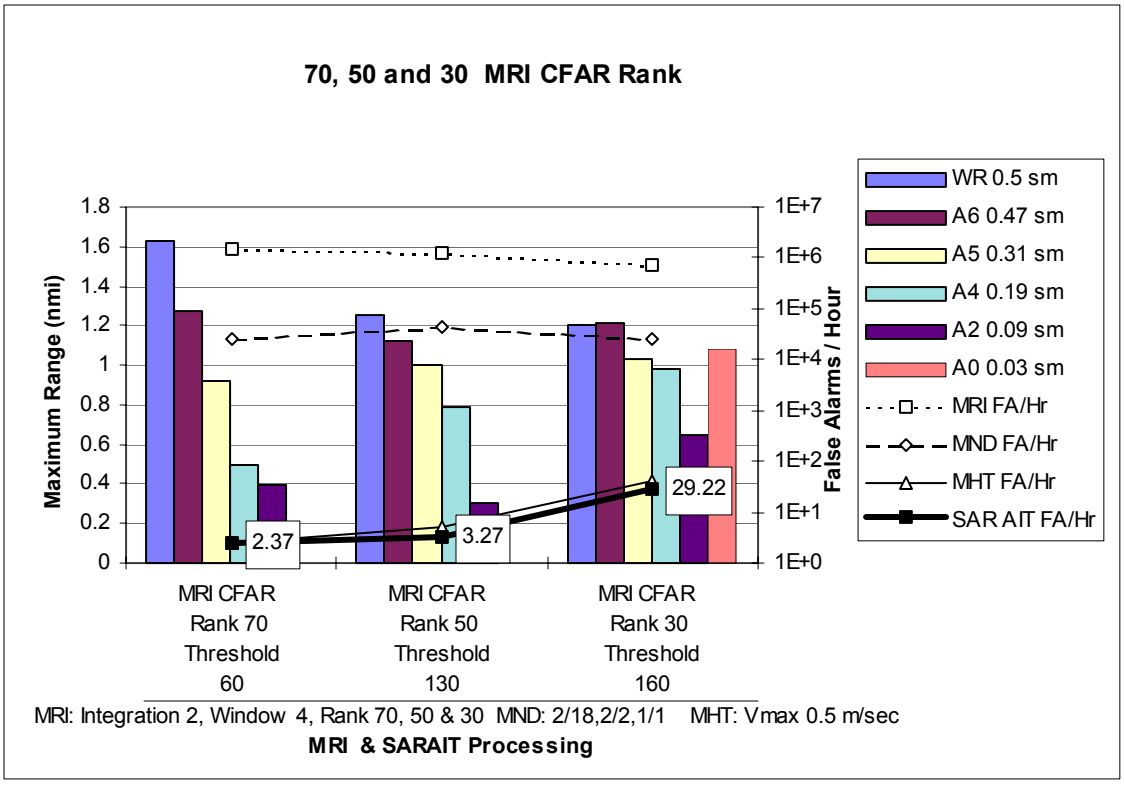

Figure 9-22 MRI CFAR rank of 70, 50 and 30

9.3.3 MRI Plot Extraction

The target extent is known to be very small. The MRI plot extractor extracts only those plots that meet the minimum and maximum extent parameters. As discussed in Section 9.2.3, limiting the plot extent to between 3 and 20 cells is effective. Lowering the minimum extent from 3 to 2 detects much more noise, typically doubling the number of plots and greatly increasing the likelihood of false MND associations. Increasing the extent beyond 20 rarely increases the number of useful target detections since most large features are dominated by large breaking waves.

9.3.4 SARAIT M of N Detection

Matching the inner MND window to the 8.6 second swell period (18 scans) typically reduces the SARAIT false alarm rate and delivers the longest range detection with the shortest MND correlation interval. Figure 9-23 contrasts two sets of MND with correlation windows that span one wave period twice (5/20,2/2) and two wave periods once (10/40,1/1). Both sets of MND are run with 80, 120 and 160 scan correlation intervals and identical 25% average visibility constraints. The most obvious difference is in the SARAIT Rtfa which for the 5/20,1/1 MND is consistently about 25% of the 10/40 Rtfa. The best balance of low Rtfa and long range detection performance is achieved with the shortest (80 scan) 5/20,2/2,2/2 MND. Extending the inner 5/20,2/2 for 3/3 and 4/4 repetitions is seen to be counter-productive.

The 10/40,1/1 MND has difficulty detecting the smallest A0 target due to the more numerous false alarms that arise with this looser constraint. Comparably long ranges are only achieved by the 10/40,1/1 MND on its longest 160-scan correlation interval (i.e. 10/40,1/1,4/4 MND) and with four times as many false alarms.

Note that the reduction between the MHT and the SARAIT Rtfa indicates the number of wake tracks which are rejected by the AI subsystem. Although easily rejected, the wake does consume MND and MHT resources and does not become an irrelevance until the correlation window is 160 scans for the 10/40 MND and 120 scans for the 5/20,2/2 MND.

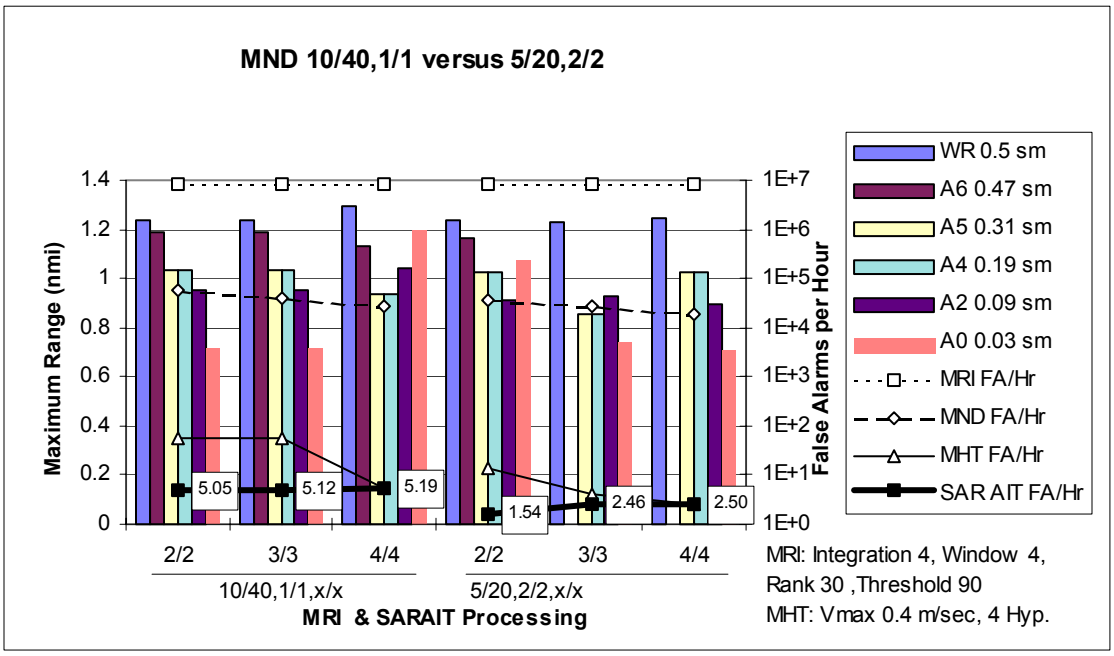

Figure 9-23 Matching the MND inner correlation window to one and two wave periods

9.3.5 SARAIT Multiple Hypothesis Tracking

The maximum track velocity must be reduced when the density of MRI plots increases to limit the potential for spurious correlation in the MND and later in the MHT. Figure 9-24 shows that reducing the maximum velocity from 0.6 to 0.4 m/sec drops the MHT Rtfa tenfold. Proportionally more of the 0.4 m/sec MHT tracks are due to the wake so the SARAIT Rtfa drops about twentyfold.

Figure 9-25 illustrates the significant benefits of using more hypotheses in the MHT to combat the confusion inherent in a high false alarm cluttered environment, particularly when the MND is configured loosely to pass numerous detections. Note how raising the number of hypotheses from 2 to 4 increases the range at which most targets are detected with no appreciable increase in SARAIT false alarms.

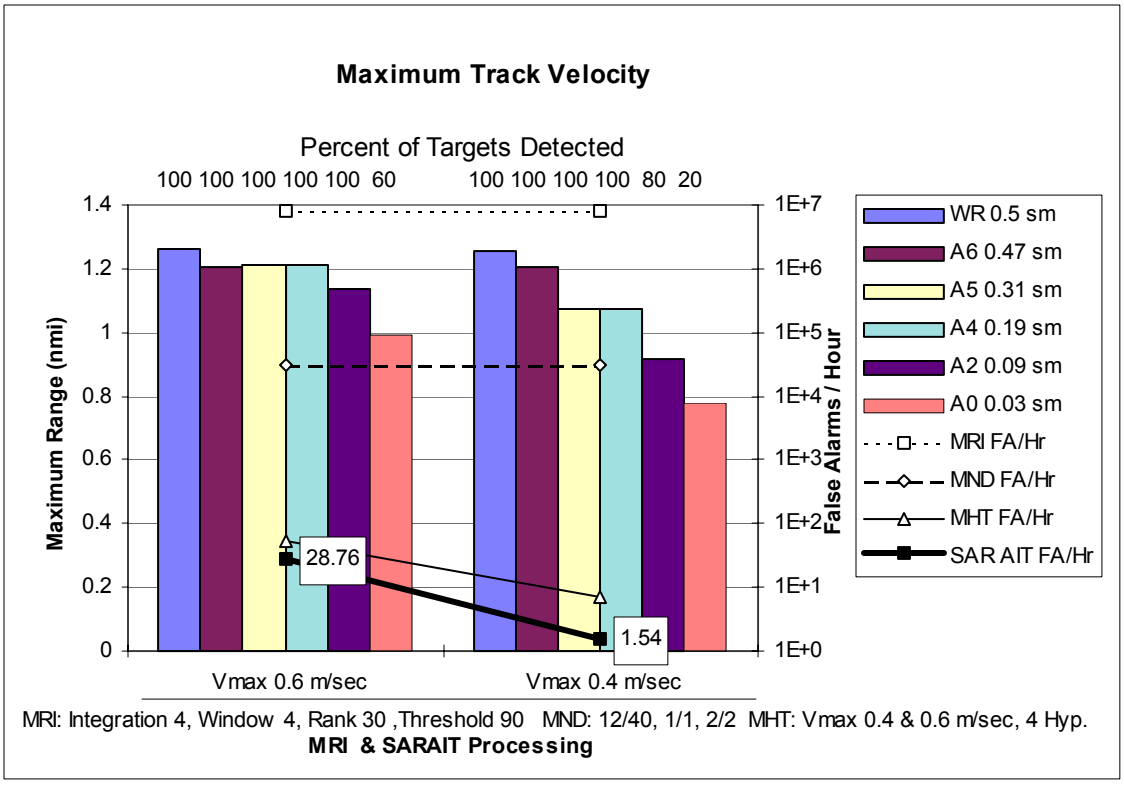

Figure 9-24 Reducing the SARAIT maximum track velocity in heavy clutter

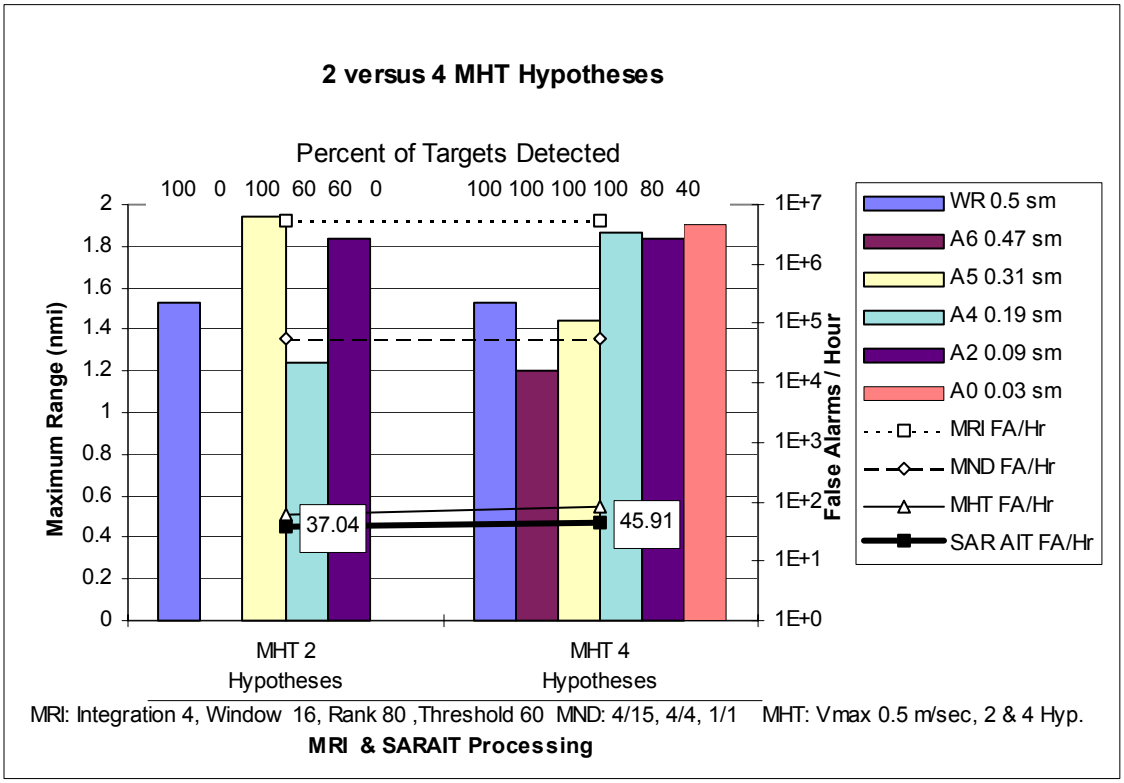

Figure 9-25 Increasing the MHT hypotheses to improve tracking in high false alarm conditions

9.3.6 Conclusions from Tests in High Seas with Severe Clutter

The SARAIT has demonstrated that it can reliably detect awash SAR targets in high 3.3 m waves with severe sea clutter. The testing on this single data set probed the inter-relationships between the many processing parameters but there was insufficient time to explore the performance limits.

One key conclusion is that the MRI-SARAIT system is robust; there are many ways to set up the system to detect the SAR targets with from 2 to 5 false alarms per hour. Higher false alarm rates up to 30 per hour would probably be acceptable in good visibility when lookouts could quickly be cued to assess whether each detection is real or not. In the night or poor visibility with only the radar, no more than 5 false alarms are likely to be acceptable.

Figure 9-26 illustrates several successful approaches that expose the broad trends that will lead to the best possible performance in the next phase:

- MRI integration over the number of scans targets will be at each wave crest (2 to 4 in this case),
- MRI CFAR window set to 8 for lowest Rfa overall and to 4 for maximum small target range,
- MRI CFAR rank of 30 for nearby small targets to 80 for larger distant targets,
- MND inner window matched to the expected target visibility (10 to 25% depending on range, size and azimuth) and the wave period (20 scans in this case),
- total MND correlation interval of 2 to 6 wave periods (40 to 120 scans in this case),
- MHT hypotheses set to 4 or more in high clutter environments,
- MHT and MND maximum velocity reduced to 0.5 m/sec (i.e. 1 kn) to limit SARAIT false alarms, and
- AI rejection of obvious wake tracks.

While the best processing scheme remains to be defined, the main lesson from this testing program is that the MRI and SARAIT processing needs to vary with range and azimuth. Globally setting the processing parameters forces compromises that limit the SARAIT to detecting SAR targets out to 1.5 nmi. This is of great value operationally but falls short of what the MRI-SARAIT system is capable of.

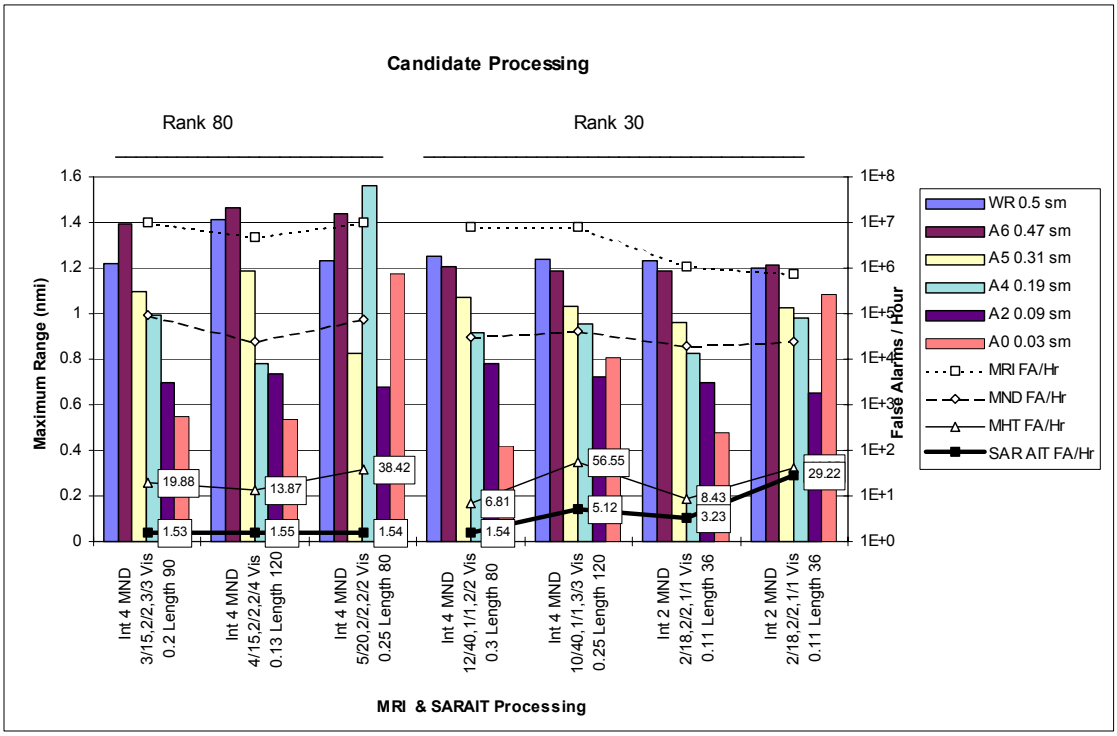

Figure 9-26 Candidate processing schemes

10. Conclusions

The prototype Search and Rescue Artificial Intelligence Tracker (SARAIT) operates in real time to process Modular Radar Interface (MRI)[8] plots, from a re-geared 120 rpm Pathfinder Mark 2 radar[9], into reliable detections of small, awash, slowly drifting targets such as liferafts, persons in water (PIW) and wreckage. The SARAIT has been operated during offshore data gathering trials on the CCGS J.E.*Bernier* and has undergone a limited set of testing on a small subset of the extensive recorded data set. The preliminary tests while sailing at 8 to 10 kn in 3.3 to 3.8 m seas demonstrate that the SARAIT detects very small PIW-sized targets at 1 to 2 nmi (depending on the clutter intensity) and small liferaft-sized targets at 2 to 3.5 nmi, all with less than 5 false detections per hour. These are only the lower bounds and longer detection ranges are expected to result from the more involved testing planned for early 1999. The estimated cost of the SARAIT is only $23,000.

The SARAIT incorporates substantial extensions to the underlying Artificial Intelligence Tracker (AIT) that were designed and implemented in this phase:

- two dual Pentium Pro single board computers in a single industrial PC enclosure to sustain real-time processing of 1000+ MRI plots per scan at 120 rpm,
- new M of N Detector (MND) algorithm to correlate over several wave periods and thereby increase the probability of detecting intermittently visible targets such as awash wreckage and PIWs,

[8] The Modular Radar Interface is a product of Sigma Engineering, St. John's, Nfld., Canada.
[9] The Pathfinder Mark 2 is a product of Raytheon Marine Company, Manchester, N.H., U.S.A.

- MND software module in C that can correlate more than 160 scans (80 seconds) of 1000 plot MRI data in real time,
- extensions to the Multiple Hypothesis Tracking (MHT) software module to initiate and maintain high confidence tracks on only the slow moving SAR targets, and
- new communications software to enable the SARAIT to receive MRI plots and parameters and to send tracks and control messages.

To detect the smallest targets, such as a person in water (PIW), at maximum range, the radar detector must be operated at maximum sensitivity which necessarily entails detecting numerous noise and sea clutter features. The maximum detection range for a given target depends on the target reflectivity, the sea clutter, the combined sea state and the swell. The swell is important because SAR targets are often awash and, however high their radar cross section, are frequently hidden by the wave crests. The longer the range, the lower the target visibility. Targets along the direction of the swell will typically be detectable only 10 to 25% of the time. Across the swell, targets are detectable 25 to 40% of the time. The SARAIT correlates across several wave periods using an M of N Detector (MND) that detects infrequently recurring detections that are more likely to be real targets.

Typically, there are so many false plots that random associations raise the MND false alarm rate. Sea and rain clutter are more important than receiver noise because they are often correlated over a few seconds and therefore, for several scans, appear more like real SAR targets. At maximum sensitivity, the MRI false alarm rate is typically reduced 100-fold by the MND but is still well above the desired rate of 5 to 30 false alarms per hour.

The Multiple Hypothesis Tracker (MHT) is therefore used to correlate sequences of MND detections into high confidence tracks. The MHT algorithm is a particularly effective detector of spatial correlation in high false alarm conditions because it retains several candidate interpretations of both pre- and post-detection tracks. Tracks are confirmed as high confidence detections after the evidence has accumulated over many scans. The MND false alarm rate is typically reduced 1000-fold by the MHT to under 20 false alarms per hour. The Artificial Intelligence subsystem is not yet complete but will identify non-SAR wake and rain cell tracks to further reduce the false alarm rate and automatically adapt the MRI and SARAIT processing parameters to the prevailing conditions. The final result when operating at maximum sensitivity is 2 to 10 SARAIT false alarms per hour.

The SARAIT has been tested on two small sets of data gathered in 3.8 and 3.3 m seas, the first data set has very little clutter and the second a great deal. For these preliminary tests, the measure of performance is the maximum detection range for each class of target balanced by an operationally acceptable SARAIT false alarm rate.

The testing on these two data sets probed the interrelationships between the many processing parameters but there was insufficient time to explore the performance limits. Determining these limits under the full range of sea states and other conditions must wait for the next phase of focussed testing in early 1999 on the extensive data set gathered off Newfoundland. The measured performance is therefore a lower bound and does NOT reflect the achievable SARAIT performance.

Under low clutter conditions most false alarms are due to relatively uncorrelated receiver noise and sea spikes so the MND can detect even the least visible SAR targets by correlating over several wave periods. Because of this, the maximum SARAIT detection range is approximately the same at all bearings. Under these conditions, the SARAIT processing is particularly robust and insensitive to the exact choice of processing. Figure 10-1 demonstrates that the Wave-Rider buoy is consistently detected at 3.5 nmi whether 2, 4 or 8 scans are integrated in the MRI. Similarly, the SAR array targets are all detected at around 2 to 2.5 nmi irrespective of the processing or their size. This confirms that target visibility is the main limitation under low clutter conditions such as these.

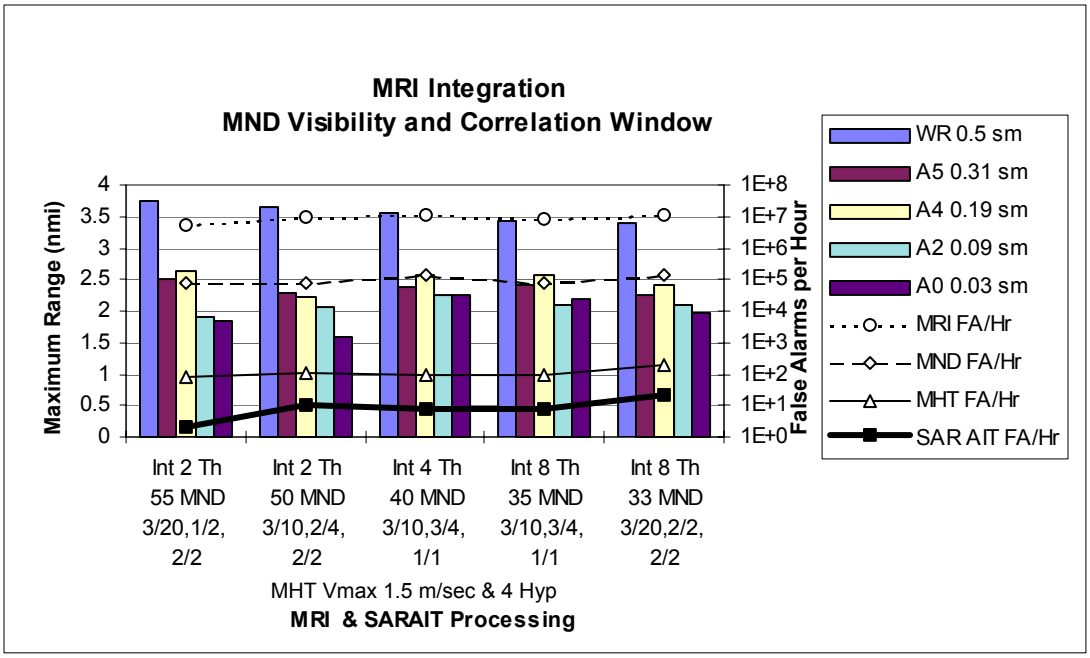

Figure 10-1 Typical SARAIT performance in 3.8 m seas and low clutter

Reasonable detection ranges with 2 to 20 false tracks per hour are therefore achieved under these low clutter conditions by:

- setting the MRI threshold for 750 to 1500 plots per scan (Rfa = 1e7),
- with an MND ratio from 7.5 to 25% giving 200 to 500 MND detections per interval (Rfa = 1e5/hr),
- with an inner MND window less than or equal to the average wave period (10 seconds or 20 scans),
- with an overall MND window spanning several wave periods, and
- with a maximum track velocity of 1.5 m/sec (3 kn) to allow for sudden wave-driven motion.

The SARAIT can also reliably detect awash SAR targets in severe sea clutter and similarly high 3.2 m swells albeit at shorter ranges and not into the face of steep wind-driven 0.3 m waves where the clutter is most intense. In high winds, the sea clutter varies enormously in range and azimuth while the MRI settings are global. The MRI must therefore compromise between dense detections close-in and upwind and sparser detections at longer ranges and down-wind. Nonetheless, the Wave-Rider buoy and the A6 and A5 targets, are typically detected at 1.5 nmi and the A0 targets at 0.5 to 1 nmi. Longer range detections are expected in the next phase once the clutter dependencies of the MRI and SARAIT processing are explored more completely.

There are many ways to set up the system to detect the SAR targets with operationally acceptable false alarm rates. Figure 10-2 illustrates several examples of successful high clutter processing that expose the broad trends that will lead to the best possible performance in the next phase:

- MRI integration over the number of scans targets will be at each wave crest (2 to 4 in this case),
- MRI CFAR window set to 8 for lowest Rfa overall and to 4 for maximum small target range,
- MRI CFAR rank of 30 for nearby small targets to 80 for larger distant targets,
- MND inner window matched to the expected target visibility (10 to 25% depending on range, size and azimuth) and the wave period (20 in this case),
- total MND correlation interval of 2 to 6 wave periods (40 to 120 scans in this case),
- MHT hypotheses set to 4 or more in high clutter environments,
- MHT and MND maximum velocity reduced to 0.5 m/sec (i.e. 1 kn) to limit SARAIT false alarms, and
- AI rejection of obvious wake tracks.

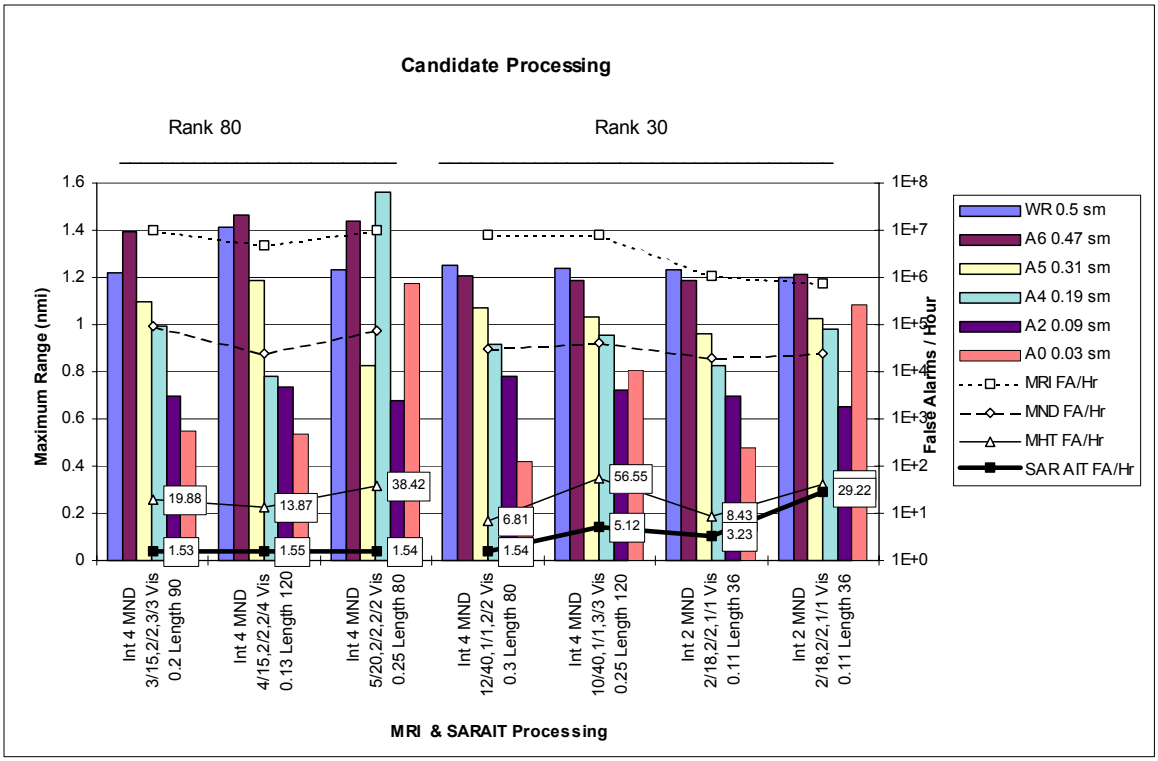

Figure 10-2 Typical SARAIT performance in 3.3 m seas and severe clutter

While the best processing scheme remains to be defined in the next phase, the main lesson from these preliminary tests is that the MRI and SARAIT processing needs to vary with range and azimuth. Globally setting the processing parameters forces compromises that limit the SARAIT performance.

The SARAIT has therefore reliably detected PIW-sized 0.03 m^2 targets in 3.3 to 3.8 m waves at 2 nmi in low clutter and 1 nmi in high clutter with 2 to 10 false tracks per hour. Small life raft-sized 0.5 m^2 targets are detected at 3.5 and 2 nmi depending on the clutter level. These targets were detected from a ship sailing an 8 to 10 kn search pattern as is typical for Canadian marine search and rescue. Since conventional radars cannot detect such small awash targets under comparable conditions, the SARAIT has been demonstrated to deliver an operationally significant advance that should greatly improve the effectiveness of marine search and rescue.