Adaptive Processing Performance Assessment for Over-the-Horizon Radar

David A. Holdsworth

Intelligence, Surveillance and Reconnaissance Division, Defence Science and Technology
PO Box 1500, Edinburgh, 5111, SA, Australia
david.holdsworth@dsto.defence.gov.au

***Abstract*—Performance assessment (PA) provides an important tool for radar system evaluation. In addition to allowing comparison of different signal processing algorithms and sequences on a near real-time or synoptic basis, it can also allow determination of the performance envelope of the radar system by defining the surveillance area over which the system is routinely expected to detect targets. The use of operational PA metrics for over-the-horizon (OTH) radar is typically limited by the availability of truth data, real targets, and real-time signal processing capacity. This paper presents a computationally efficient PA module suitable for real-time operational use. The PA module uses synthetic targets to calculate detection and false alarm probabilities, which are used to estimate a novel "probability difference" (PD) metric. The PD metric provides several advantages over receiver operating characteristic (ROC) derived metrics, and allows the ability to explicitly incorporate target location accuracy. The amplitudes of the synthetic targets are selected to obtain reliable detection statistics while reducing their effect on the signal processing. The PA module is demonstrated in the context of near real-time selection of adaptive processing algorithm appropriate for the prevailing radio frequency interference conditions.**

***Index Terms*—ignore**

I. Introduction

High frequency (HF) skywave over-the-horizon (OTH) radar uses the ionosphere as a propagation medium, allowing the detection and tracking of targets at ground ranges up to 3000 km [1]. The raw data recorded by a uniform linear array of receiving elements is matched filter (range) processed, beamformed and Doppler processed to produce conventionally beamformed (CBF) azimuth-range-Doppler (ARD) data. Constant false alarm rate (CFAR) [2] processing is applied to whiten the CBF data to provide input to the target detector, and the resulting detections provide input to the tracking system to produce target tracks. The performance of OTH radar may be degraded by the external HF interference and noise field, which is often spatially anisotropic and temporally non-stationary. While every attempt is made to select clear frequency channels, it is impossible to completely avoid radio frequency interference (RFI) since natural RF sources (e.g. lightning) and transmissions from other HF users may overlap the receiver bandwidth intermittently and unpredictably.

The HF interference and noise environment motivates the use of adaptive processing in OTH radar. The adaptive processing algorithm yielding best target detection and tracking performance is dependent on the prevailing noise and interference conditions. Spatial adaptive processing (SAP) with spatial degrees of freedom is used to cancel sidelobe RFI [3], while time-varying SAP (TVSAP) is used to cancel non-stationary sidelobe RFI [4]. Range adaptive processing (RAP) with range (fast-time) degrees of freedom is used to reject range correlated RFI incident through the receive antenna beam ("main beam" RFI) [5]. Space-time adaptive processing (STAP) [6] with spatial and range degrees of freedom is used to simultaneously cancel sidelobe and main beam RFI.

Although adaptive processing effectively mitigates interference and noise, it can impact upon target detection, primarily in three ways. First, since adaptive processing training data is typically selected using amplitude thresholding to preclude targets and clutter, low SNR targets can be included in the training data, and hence become attenuated with the interference and noise. Second, adaptive processing can produce "target copy" [3], where targets and/or ionospheric clutter (e.g. meteors, spread Doppler clutter) are copied into other surveillance beams (or ranges) from sidelobes in the adaptive beam (range) pattern. Third, since OTH radar targets are typically oversampled in Doppler, range and beam due to the combination of radar parameter selection and the use of data windows to control sidelobe leakage in each dimension, different target ARD cells may be processed using different adaptive processing weights, potentially leading to target position estimation errors. It is therefore critical that adaptive processing performance assessment techniques explicitly incorporate measurements of target attenuation, target location accuracy, and target detection performance (i.e. ability to discriminate between false alarms and detections).

This paper presents a near real-time air-mode operational adaptive processing performance assessment and selection (PAS) module for OTH radar. Section 2 describes OTH radar PA metric limitations, while section 3 describes existing OTH radar PA metrics and their applicability for adaptive

processing PA. Section 4 introduces the probability difference (PD) metric, while section 5 describes the use of the PD metric in the PAS module. Section 6 uses JORN data to compare the results obtained using PD and existing PA metrics.

II. Over-the-horizon radar performance assessment

There are a number of limitations in developing operational performance assessment (PA) metrics for OTH radar. The major limitation is the lack of truth data available for verifying that radar detections correspond to real targets, rather than ionospheric clutter, signal processing artifacts (e.g. sidelobe target, adaptive processing target copy), or noise and interference artifacts. This limitation may reduce in future as more air and ship targets are instrumented with automatic dependent surveillance-broadcast (ADS-B) and automatic identification systems (AIS), respectively. However, even with this instrumentation, the number of available real targets may be insufficient to produce statistically reliable near real-time PA metrics.

There are a number of alternative approaches that have been applied for OTH radar PA:

1) target tracking: tracking results provide a reliable basis for comparing signal processing algorithms [7]. However, in addition to uncertainty regarding the practicality of some tracking performance metrics [8], tracking requires a pre-defined number of radar dwells for track initiation [8], and often requires user intervention [8,9], thereby making tracking systems unsuitable for near real-time PA.

2) ground truth targets: measurement of the performance of the radar detection system using observations of a live test target whose ground position is known under operationally significant conditions [7,10]. This requires target observation over a sufficient time period, during which the external environmental conditions, target dynamics, and radar system must remain approximately stable, thereby making ground truth targets unsuitable for near-real-time PA. There is also the need to discriminate the target from other targets in the surveillance region, and to exclude these other targets from the analysis as they may corrupt the measurement.

3) radar-propagation-environment-target system modelling: this approach suffers from the inability to accurately represent the detailed temporal and spatial scales of many aspects of the problem necessary for detailed analysis [7], and typically requires significant computational processing capacity, thereby making such modelling truth targets unsuitable for near-real-time PA.

4) synthetic target injection: injecting artificial targets into real radar clutter and noise data. This suffers from the inability to simulate the properties of the target echoes [11]. For instance, it is difficult to simulate the spatial signature of a real target due to the difficulty in measuring signal wavefront distortions and array manifold errors. This is a problem when assessing spatial adaptive processing performance for main beam RFI, as the difference between the actual and simulated spatial signatures can result in an overestimation of the performance of spatial adaptive beamformers.

The adaptive processing performance assessment and selection (PAS) module is required to select the appropriate adaptive processing technique in near real-time without target ground position truth. Given the above considerations, synthetic target injection based PA appears the only feasible option for near-real-time PA.

The OTH radar PA scheme described by [11] randomly injects synthetic targets into dwells of real radar data. The targets allow the estimation of the probability of detection and false alarm as a function of detection threshold, allowing receiver operating characteristic (ROC) curves [12] to be computed. This PA scheme was demonstrated by assessing the performance of a meteor excision algorithm using 10000 synthetic targets. This large number of targets was deemed necessary to achieve sufficiently low estimated probability variances. Further, each target is inserted one at a time to avoid the targets affecting the processing, requiring each processing sequence under assessment to be reapplied for each inserted synthetic target.

Given the limited computing resources available and the near real-time requirement for the PAS, the OTH radar PA scheme of [11] is impractical. Despite this, a similar scheme has been successfully developed, as described in section V.

III. Performance assessment metrics

The relative performance of different signal processing sequences (hereafter processors) can vary on a beam-by-beam and dwell-by-dwell basis due to the spatially anisotropic and temporally non-stationary nature of the HF RFI environment. Although it is feasible to use the output from the best performing processor for each beam (or beam/range) and dwell, this can prove problematic. The use of different processors for each beam (or beam/range) may exacerbate target position errors as different target ARD cells may be processed using different processors. The use of different processors for each dwell may exacerbate target tracking errors as the detection and false alarm statistics may vary for each processor. As a result, the PAS module should select one processor that yields best performance over all surveillance

Metric	Target attenuation	Target accuracy	Target detection
Noise level	No	No	No
Synthetic target SNR gain	Yes	No	No
Peak SNR gain	Yes	No	No
Miss distance	No	Yes	No
Probability of detection/false alarm	Yes	No	No
ROC curve	Yes	No	Yes
Probability difference	Yes	No	Yes
Extended ROC curve	Yes	Yes	Yes
Extended probability difference	Yes	Yes	Yes

TABLE I
OTH RADAR PERFORMANCE ASSESSMENT METRICS AND THEIR ABILITY TO INDICATE TARGET ATTENUATION, TARGET LOCATION ACCURACY, AND TARGET DETECTION PERFORMANCE. EXTENDED ROC CURVES AND PD REFER TO THE USE OF VARIABLE DETECTION GATE SIZE.

beams and multiple dwells. We therefore require a temporally stable scalar PA metric representing processor performance over all surveillance beams.

The following subsections describe different real-time OTH radar PA metrics. The ability of these metrics to evaluate target attenuation, target location accuracy, and target detection performance is summarised in Table 1.

A. OTH radar performance metrics

OTH radars routinely measure environmental performance metrics such as clutter power, background noise level and sub-clutter visibility (SCV, the clutter power to background noise level ratio) for each surveillance beam [10]. The clutter power and SCV are unsuitable for use as adaptive processor metrics as the clutter power may be affected by adaptive processing. The noise-level represents the most useful environmental metric for adaptive processor PA, and can be averaged over all surveillance beams to yield a scalar performance metric.

This mean noise level does not use synthetic target injection, and consequently does not yield information on target attenuation/accuracy or detection performance.

B. Synthetic target SNR gain

Synthetic targets can be used to estimate the relative SNR gain between two processors k and l for each target i using

$$S_{ikl} = \frac{|y_{wl}(d_i, r_i, b_i)|^2}{|y_{wk}(d_i, r_i, b_i)|^2} \tag{1}$$

where $y_{wm}(d, r, b)$ is the whitened ARD data for processor m and (d_i, r_i, b_i) are the positions of the N_t targets. The SNR gain may be estimated using the mean or median SNR gain for all targets.

This metric yields information on target attenuation, but does not yield information on target location accuracy or detection performance.

C. Peak SNR difference

In OTH radar, peak detection is applied by finding all ARD cells with amplitudes exceeding their immediate neighbours. The peak positions $p_j = (d_j, r_j, b_j)$ and SNRs s_j are estimated using an independent interpolation across each ARD dimension using the peak cell and its immediate neighbours.

The peak positions and SNRs can be used to estimate the SNR difference between two processors. This can be achieved by associating peaks from different processors (peak-peak association), and calculating the SNR difference for all associated peaks. Peak-peak association is complicated by the difficulty in verifying that the peaks correspond to real targets. The peak widths estimated during peak interpolation can assist in this process. For synthetic targets, the SNR difference can be calculated by locating the peak closest to each target (target-peak association), and using the peak SNR as the target SNR estimate. The SNR difference between the two processors can then be estimated by calculating the mean or median SNR difference for all synthetic targets.

This metric yields information on target attenuation, but it does not yield information on target location accuracy or detection performance.

D. Miss distance

For synthetic targets, the peak positions p_j can be used to provide an estimate of the miss distance between the peaks and synthetic targets using target-peak association. A detection miss distance metric may be estimated by taking the mean, median or root-mean-square (RMS) difference between the peak and simulated target positions. Likewise, a false alarm miss distance metric may be estimated by taking the mean, median or root-mean-square (RMS) difference between the peak and simulated target positions using the original data without synthetic targets added (i.e. target-free data).

The detection miss distance yields information on target location accuracy, but does not yield information on target attenuation or detection performance.

E. Probability of detection and false alarm

For real targets, the probability of detection p_{det} can be determined by forming a finite-width ARD detection gate around known target positions, and dividing the number of

detections within each gate by the number of detection gates used. Likewise, the probability of false alarm $p_{\rm fa}$ can be determined using the same process, but with detection gate placed in positions expected to be free of targets. A detection is registered if the square of the maximum cell amplitude within the detection gate exceeds an SNR threshold t, yielding probability of detection $p_{\rm det}(t)$ and false alarm $p_{\rm fa}(t)$ as a function of t. Due to the aforementioned limitations in obtaining target truth, it is common to use synthetic targets inserted into areas of ARD space where there are expected to be no real targets. $p_{\rm det}$ is found using detection gates around the synthetic target positions for data with synthetic targets (target-added data), while $p_{\rm fa}$ is found using detection gates in the same positions for original target-free data. Note that we hereafter denote $p_{\rm det}(t)$ and $p_{\rm fa}(t)$ as $p_{\rm det}$ and $p_{\rm fa}$ for notational convenience.

The probability of detection yields information on target attenuation. However, it does not necessarily yield information on target location accuracy as biases in target positions of adaptively processed data are generally smaller than the ARD cell size, while detection gate sizes are typically of the order of the ARD cell size. Although $p_{\rm det}$ and $p_{\rm fa}$ can be combined to assess detection performance using Receiver Operating Characteristic (ROC) curves as described in the following subsection, they do not assess detection performance individually.

The peak positions p_j and SNRs s_j can also be used as the basis for $p_{\rm det}$ and $p_{\rm fa}$ estimation in OTH radar, with a detection registered if the SNR of any peak within the detection gate exceeds the SNR threshold. This approach is intuitively sensible since peak detections provide the input for OTH radar target tracking systems. However, the use of peak derived $p_{\rm det}$ and $p_{\rm fa}$ typically results in maximum $p_{\rm det}$ and $p_{\rm fa}$ values less than unity, which obviate some of the practicality of using ROC curves to compare detection performance. The maximum peak derived $p_{\rm fa}$ is typically significantly less than unity as the peak detection process protects against peak detection in noise, while the maximum peak derived $p_{\rm det}$ is occasionally less than unity as ionospheric clutter, noise and interference artifacts, and processor losses can result in simulated targets not being peak detected. Despite these issues, peak derived $p_{\rm det}$ and $p_{\rm fa}$ are used in the PAS.

F. Receiver operating characteristic curves

Receiver operating characteristic (ROC) [12] curves have long history of use in signal processing detection theory, and have recently been extended for analysing the behaviour of diagnostic systems and medical decision making. ROC curves capture the relationship between $p_{\rm det}$ and $p_{\rm fa}$ as the detection threshold is varied, providing a detection performance comparison measure. There are typically two methods of estimating a scalar ROC performance metric: $p_{\rm det}$ at a specified $p_{\rm fa} = p'_{\rm fa}$ (i.e. $p_{\rm det}(p'_{\rm fa})$), and the area under the ROC curve

$$A = \int_{p_{\rm fa}=0}^{1} p_{\rm det}(p_{\rm fa})\, dp_{\rm fa}, \tag{2}$$

also known as the Mann-Whitney test statistic, which estimates the probability that a randomly selected detection will be ranked higher than a randomly selected false alarm. The $p_{\rm det}(p'_{\rm fa})$ metric is often used for assessing the output of processors providing input to tracking systems, since some tracking systems require the inputs to conform to a specified $p_{\rm fa} = p'_{\rm fa}$. For processor comparisons, the area metric A requires $p_{\rm det}$ to be resampled at fixed $p_{\rm fa}$ values so the area estimate for each processor is calculated identically.

Although the $p_{\rm det}(p'_{\rm fa})$ and A metrics yield information on target attenuation and detection performance, they do not yield any information on target location accuracy.

An example of typical peak derived $p_{\rm det}$, $p_{\rm fa}$, and ROC curves obtained by injecting synthetic targets into the JORN data-set commencing 0500 UT, 22 April 2008 (hereafter 20080422-050000) using CBF and four adaptive processors (AP1 to AP4) are shown in Figures 1 and 2. Although AP2 clearly produces the largest $p_{\rm det}$ for all $p_{\rm fa}$, the remaining processors each produce the second largest $p_{\rm det}$ at different values of $p_{\rm fa}$. Determination of the best of these remaining processors using $p_{\rm det}(p'_{\rm fa})$ is therefore sensitive to the choice of $p'_{\rm fa}$. Estimation of A is also problematic as the minimum/maximum $p_{\rm fa}$ values differ for each processor. The maximum $p_{\rm fa}$ for AP2 and AP4 exceed those of the other processors, indicating AP2 and AP4 produce more false alarms. If A is calculated over the available range of $p_{\rm fa}$ values, the AP2 and AP4 estimates will be increased despite their poorer false alarm performance. This problem can be avoided by selecting minimum/maximum $p_{\rm fa}$ values common to all processors, but this can introduce inconsistency in comparing or combining metrics from multiple radar dwells. The area estimation problems do not occur for cell amplitude derived ROC curves as the maximum $p_{\rm fa}$ is unity. A further implication of using peak derived ROC curves is that A no longer yields the Mann-Whitney test statistic, and in fact has no apparent statistical relevance. Although it is tempting to avoid these issues by renormalising the ROC curves so the maximum $p_{\rm det}$ and $p_{\rm fa}$ values are unity, this introduces further undesirable outcomes.

Although we would prefer to use ROC derived metrics due to their widespread use throughout the signal processing community, they appear incompatible with our preference for using peak derived $p_{\rm det}$ and $p_{\rm fa}$.

IV. The probability difference metric

A number of alternative methods for incorporating $p_{\rm det}$ and $p_{\rm fa}$ into a scalar detection metric have been investigated. One

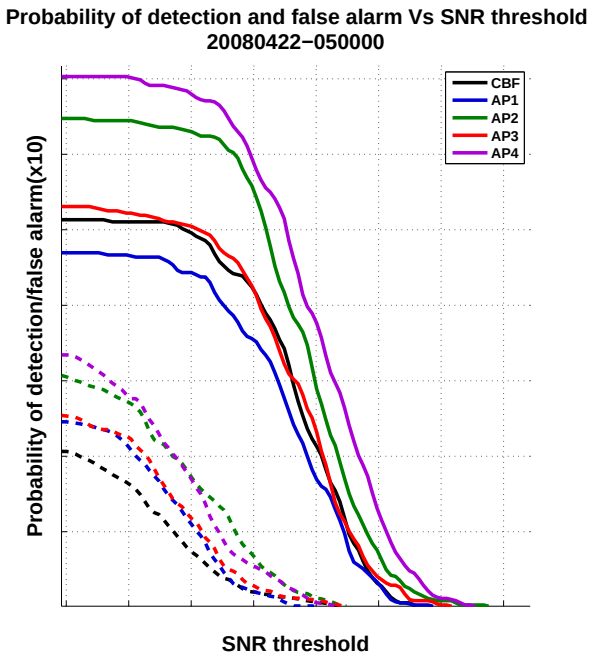

Fig. 1. Probability of detection (solid line) and false alarm (dashed lines) for the application of different processors to JORN data-set 20080422-050000: CBF (black), AP1 (blue), AP2 (green), AP3 (red), AP4 (purple). The probability of false alarm is multiplied by a factor of ten for clarity.

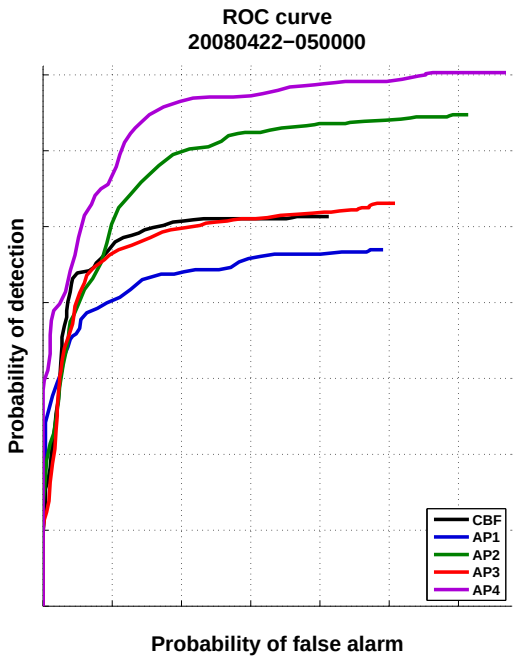

Fig. 2. ROC curves for the application of different processors to JORN data-set 20080422-050000: CBF (black), AP1 (blue), AP2 (green), AP3 (red), AP4 (purple).

method which may appear desirable is using the likelihood ratio

$$L(t) = \frac{p_{\rm det}}{p_{\rm fa}}. \tag{3}$$

This can be used to produce a scalar detection metric by evaluation at a specified $t = t_0$, (i.e. $L_0 = L(t_0)$), or by integration over a specified range of values $t = t_{\min}, ..., t_{\max}$

$$L_A = \int_{t_{\min}}^{t_{\max}} L(t)\, dt. \tag{4}$$

However, L_0 suffers from similar problems as described above $p_{\rm det}(p'_{\rm fa})$, while L_A is unstable since $L(t) \rightarrow \infty$ as $p_{\rm fa} \rightarrow 0$. This instability can be eliminated using the integrated log likelihood ratio. However, the SNR threshold integrated log likelihood ratio is an unbounded estimator, and is therefore deemed unsuitable for PAS use.

The statistic selected for incorporating $p_{\rm det}$ and $p_{\rm fa}$ into a scalar detection metric for the PAS is the probability difference (PD) metric

$$p_\Delta(t) = p_{\rm det}(t) - p_{\rm fa}(t), \tag{5}$$

which is used to produced a scalar metric by integrating over SNR threshold

$$p_A = \int_{t_{\min}}^{t_{\max}} p_\Delta(t)\, dt. \tag{6}$$

The SNR threshold integrated PD metric is simpler to calculate than the ROC area metric as it does not require $p_{\rm det}$ resampling or $p_{\rm det}$ and $p_{\rm fa}$ normalisation. And unlike the SNR threshold integrated log likelihood ratio, the integrated PD metric is bounded within range the [0, 1].

The SNR threshold integrated PD yields information on target attenuation and detection performance, but no information on target location accuracy.

A. Extended ROC and probability difference metrics

The metrics presented so far in this paper do not fully encapsulate the characterises required by the PAS: target attenuation, target location accuracy, and target detection performance. While the detection miss distance yields information on target location accuracy, it does not explicitly measure target attenuation, and it is not clear or straightforward to determine how to incorporate false alarm information. Similarly, although the ROC/PD metrics combine detection and false alarm information, they do indicate how accurately targets are detected within a detection gate. It was therefore decided to develop new metrics combining the benefits of the miss distance and ROC/PD metrics.

The miss-distance d can be incorporated into a detection metric by varying the cell width of the detection gate g used, thereby producing a metric dependent on $d(= g)$. A single integrated metric B_{int} can then be obtained by calculating the weighted sum of a selected gate-width dependent metric $B(g)$

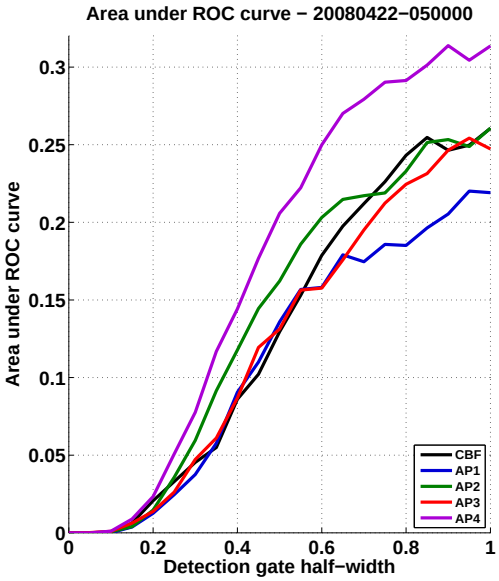

Fig. 3. Unnormalised ROC curve area metric $A(g)$ as a function of detection gate-width g for the data-set 20080422-050000: CBF (black), AP1 (blue), AP2 (green), AP3 (red), AP4 (purple).

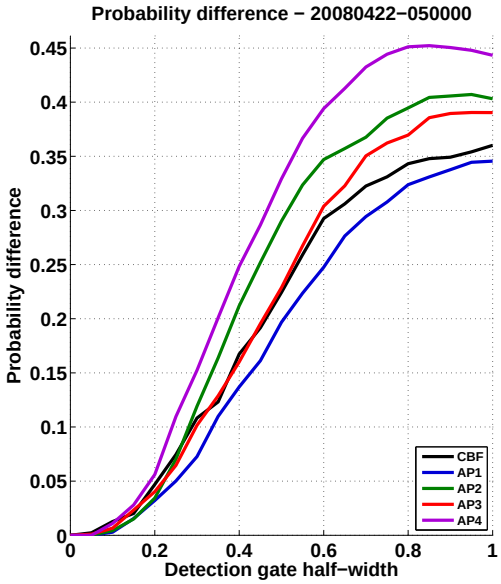

Fig. 4. Integrated PD metric $p_A(g)$ obtained as a function of detection gate-width g for the data-set 20080422-050000: CBF (black), AP1 (blue), AP2 (green), AP3 (red), AP4 (purple).

using

$$B_{int} = \int_{g_{\min}}^{g_{\max}} B(g)\, w(g)\, dg, \tag{7}$$

where $w(g)$ is a weighting function, and $g_{\min}$ and $g_{\max}$ are the minimum and maximum detection gates sizes. The weighting function $w(d)$ can be selected to reflect the desired miss-distance d distribution.

Figures 3 and 4 shows the unnormalised ROC curve area metric $A(g)$ and the integrated PD metric $p_A(g)$, respectively, as a function of detection gate-width g for the data-set 20080422-050000. The integrated PD metric $p_A(g)$ shows significantly smoother variation than $A(g)$. The larger variation of $A(g)$ is a consequence of different $p_{\rm det}$ resampling at each g, and due to differences in maximum $p_{\rm fa}$ at each g, which as noted above also introduces an undesirable bias to larger values for larger maximum $p_{\rm fa}$. Although using the normalised ROC curves eliminates the bias to smaller values for larger maximum $p_{\rm fa}$, the different normalisation required for each g introduces significant extra variation into $A(g)$. These factors further justify the use of integrated PD metric in the PAS.

The SNR threshold and gate-width integrated probability difference

$$\alpha = \int_{g_{\min}}^{g_{\max}} \int_{t_{\min}}^{t_{\max}} p_{\Delta}(t,g)\, w(g)\, dg\, dt \tag{8}$$

is hereafter referred to as the probability difference (PD) metric. We impose the constraint

$$\int_{g_{\min}}^{g_{\max}} w(g)\, dg = 1 \tag{9}$$

to limit the PD metric to be bounded within [0, 1]. However, the unity value for a "perfect" detector will never be attained in practice.

B. Use of the probability difference metric

Although the PD metric fulfils the requirements of the PAS, it can be difficult for OTH radar users to meaningfully interpret. We therefore propose two PD derived statistics for practical use. First, the PD metric α can be expressed as a detection percentage D relative to a reference processor

$$D = \frac{\alpha - \alpha_{\rm ref}}{\alpha + \alpha_{\rm ref}} \times 50, \tag{10}$$

where $\alpha_{\rm ref}$ is the PD metric of the reference processor. Second, a relative SNR gain can obtained by determining the mean additive factor (in dB) required to be applied to the reference processor peak SNRs s_j to yield the same PD metric as the processor under assessment. This SNR gain can be approximated using

$$S_{\rm pd} = \frac{\alpha - \alpha_{\rm ref}}{\alpha'_{\rm ref} - \alpha_{\rm ref}} 10^{(T/10)}, \tag{11}$$

where $\alpha_{\rm ref}$ is the PD metric obtained by increasing the reference processor peak SNRs by T dB. CBF data provides a convenient reference processor for implementing these metrics for OTH radar adaptive processor comparisons.

V. The performance assessment and selection module

The performance assessment and selection (PAS) module is required to select the appropriate processors for the prevailing noise and interference conditions in near real-time. The processors available for selection include different adaptive processors and CBF (i.e. no adaptive processing), and also incorporate common standard OTH radar processing algorithms transient mitigation (e.g. impulses, meteors) [13] and data extrapolation (Datex) [14].

In order to obtain accurate PD metrics the synthetic targets should be inserted into ARD cells free of real targets, surface clutter and ionospheric clutter. These are the same ARD cells we wish to select for adaptive processing training data selection, and we therefore adopt the adaptive processing strategy of forming a clutter-target mask (CTM) listing the range-Doppler (RD) cells containing clutter and strong targets in any surveillance beam. Synthetic targets are sequentially inserted into each surveillance beam by selecting random RD locations that are not listed in the CTM. These positions are then included in the CTM to avoid further targets being added to these locations. This process is repeated until the specified number of targets have been inserted in each beam, or the CTM is fully occupied. This approach ensures there is only one target at any RD position. This allows $p_{\rm det}$ to be calculated using a detection gate around target positions (d_i, r_i, b_i), and $p_{\rm fa}$ to be calculated using detection gates around non-target positions $(d_i, r_i, b_i), \forall\ b_k$ s.t. $|b_k - b_i| > b_g$, where b_g is a beam "guard" to allow for beam oversampling. This approach provides three advantages over to the standard approach of estimating $p_{\rm fa}$ by forming detection gates about the target positions using the target-free data. First, it obviates the need to apply each processor to the target-free data, effectively halving the processing time. Second, it significantly increases number of detection gates used for $p_{\rm fa}$ estimation, thereby reducing the variance of $p_{\rm fa}$. Third, it also includes a measure for spatial adaptive processing target copy [3], as target copies will be observed at synthetic targets RD locations in other surveillance beams, and therefore increase $p_{\rm fa}$ and reduce α.

The PAS uses six parallel processing streams. Five streams are used for applying the five processors to target-added data to produce $p_{\rm det}$ and $p_{\rm fa}$ curves, with the remaining stream applying the selected "best" processor to target-free data, which is subsequently used for peak detection and target tracking. The five processors consist of CBF (i.e. no adaptive processing) and four adaptive processors. The adaptive processors may incorporate different adaptive processing techniques (e.g. SAP, RAP, etc), or multiple instances of selected adaptive processing techniques using different input parameter sets. This architecture requires the best processor to be selected before processing each dwell. This is achieved by combining $p_{\rm det}$ and $p_{\rm fa}$ from all previous dwells within a time period T_m of the current dwell to obtain a time integrated PD metric for each processor, with the processor yielding the highest PD metric nominated the best processor for mitigating the current noise and interference characteristics. The value of T_m is chosen to be large enough to reduce the variance of the estimated probability variances, yet small enough to be able to respond to temporal characteristics of the noise and environment, which for OTH radar is typically dependent on ionospheric characteristics.

There are a number of factors to consider in selecting the target SNR and number of targets for adaptive processing assessment. The role of adaptive processing is to uncover "weak" targets that cannot be detected in the CBF data. This must be achieved without deteriorating CBF detectable targets. Since the most significant adaptive processing deterioration is expected to occur for targets included in the adaptive training data, it appears sensible to insert the targets at an SNR where they will be excluded from the CTM and hence included in the adaptive processing training data. The PD metric then captures the maximum impact of the adaptive processing upon the weak targets it seeks to uncover. We also need to consider that synthetic targets can reduce adaptive processing effectiveness in two ways: strong targets reduce training data availability, while weak targets not included in the CTM can modify training data covariances. The target SNRs should therefore be small enough to minimise target effects on the processing, yet large enough to yield significant enough differences between $p_{\rm det}$ and $p_{\rm fa}$ to improve the statistical reliability of the PD metrics. Similarly, the number of targets must be small enough to minimise target effects on the processing, yet large enough to maximise the statistical reliability of the PD metric.

The effects of target SNR and number of targets have been investigated by calculating the noise level difference obtained by applying the adaptive processors to the target-added and target-free data for various JORN data-sets collected in different noise and interference environments. Values of 20 targets per beam and target SNRs of 9 dB have been found to produce statistically reliable PD metrics and noise level differences significantly below 1 dB in most cases, and have therefore been adopted as default parameters.

VI. Results

The PAS is demonstrated using JORN Longreach data recorded at a carrier frequency of 15.463 MHz between 0456 and 0502 UT 22 April 2008. During this period an amplitude modulated (AM) RF transmission was received. The range-Doppler plot of the data for beam 1 of the dwell commencing 050000 is shown in Figure 5a. The AM carrier signal produces a "narrow-band" frequency component at Doppler $d = 32$, while the AM sidebands produce a "wide-band" noise component that spreads across all ranges and Dopplers. A meteor echo is observed at range $r = 11$, and several

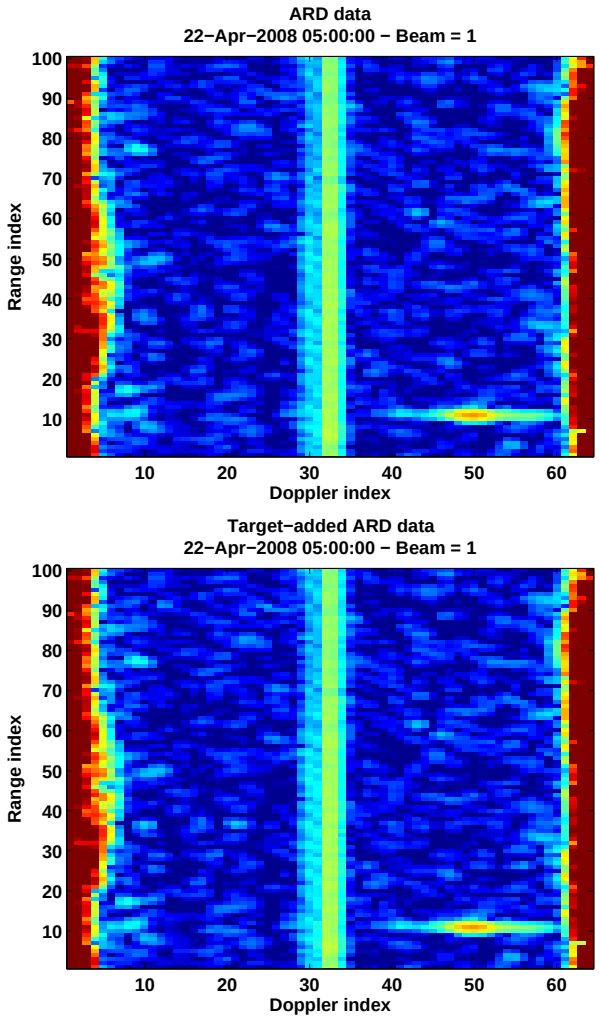

Fig. 5. Range-doppler plots of target-free (top) and target-added data (bottom) for beam 1 of data-set 20080422-050000.

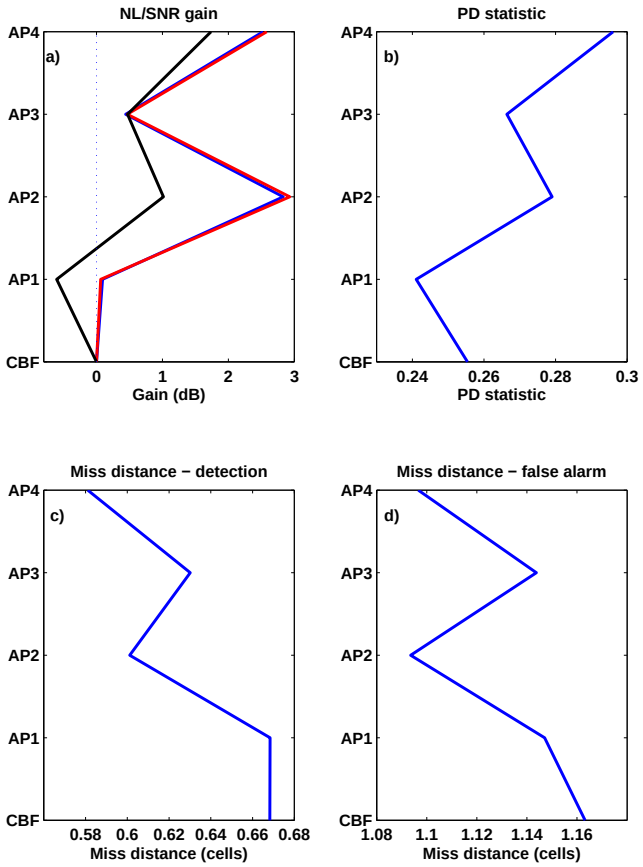

Fig. 6. Mean PA metrics for data-set 20080422-050000: a) Noise level reduction relative to CBF (red), SNR gain relative to CBF (blue), PD-SNR gain relative to CBF (black). The dotted line indicate 0 dB gain; b) PD metric; c) Detection miss distance, d) False alarm miss distance.

possible target echoes are observed across range between $d = 6, .., 13$. Twenty synthetic targets with SNR = 9 dB have been injected into each of the N_b = 22 surveillance beams for PAS demonstration. The target-added range-Doppler plot for beam 1 is shown in Figure 5b. Comparison of the target-free and target-added reveals little obvious evidence of the injected targets. The most prominent targets are those injected at $(d, r) = (10, 37)$ and $(d, r) = (22, 37)$.

The mean values of selected PA metrics for each adaptive processor are shown in Figure 6. The mean noise-level reduction and SNR gain metrics are plotted with respect to the CBF results. The adaptive processors reduce the mean noise level and improve the mean target SNR. However, the mean SNRs gain are slightly lower (< 0.1 dB) than the mean noise-level reduction, indicating that some adaptive processing target attenuation is occurring. The adaptive processors reduce the detection miss distance, indicating they improve target location accuracy relative to CBF. However, the adaptive processors also reduce the false alarm miss distance, indicating they produce more false alarms than CBF. The PD metric indicates that AP4 produces the best target detection performance, yielding an effective SNR gain of 1.73 dB. Although AP2 yields a higher noise level reduction than AP4, it produces increased target attenuation and increased detection miss distance, leading to reduced target detection performance. Although AP1 yields small mean noise level and target SNR gains, the increased false alarm rate (i.e. increased false alarm miss distance) results in a detection performance loss relative to CBF.

The temporal variation of the 8-dwell time-integrated PD metric for each adaptive processor is shown in Figure 7. The large variation in the PD metric for dwells 1 to 7 result from the unavailability of 8 completed dwells, emphasising the need for a time-integrated the PD metric. The results

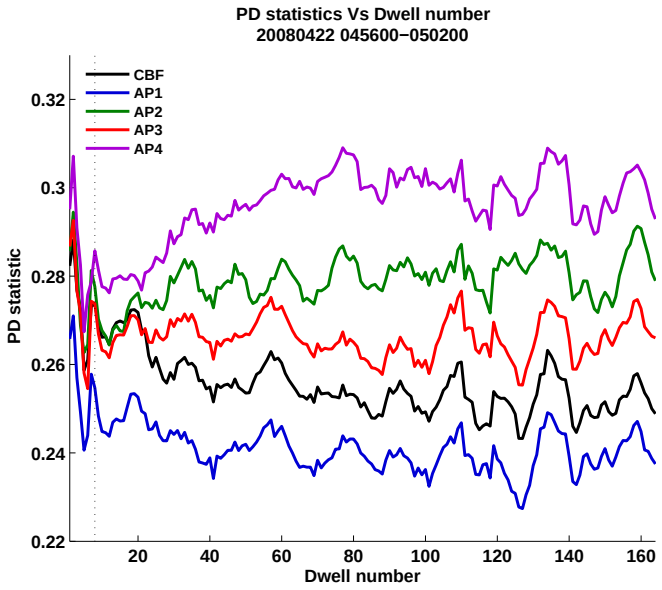

Fig. 7. Time-integrated PD metrics for the PAS adaptive processors: CBF (black), AP1 (blue), AP2 (green), AP3 (red), AP4 (purple). The dotted line indicates dwell number 8.

for dwell 8 onwards indicate the time-integrated PD metric is temporally stable, and can be successfully employed for stable adaptive processor selection. The use of eight dwells corresponds to 3520 synthetic targets (i.e. 8 dwells, 22 beams, 20 targets per beam), a factor of three less than that required by [11] to obtain sufficiently low estimated probability variances. Furthermore, this performance is achieved by applying each processor eight times (i.e. for eight dwells), rather than for 10000 times as required by [11].

VII. Conclusions

This paper presents a computationally efficient PA module suitable for near real-time operational use. The PA module uses synthetic targets to calculate detection and false alarm probabilities, which are used to estimate a novel "probability difference" (PD) metric. The PD metric provides several advantages over receiver operating characteristic (ROC) derived metrics, and allows the ability to explicitly incorporate target location accuracy. The PD metric forms the basis of the performance assessment and selection (PAS) module for near real-time selection of adaptive processing algorithm appropriate for the prevailing radio frequency interference conditions. The PAS module has been successfully demonstrated using operational JORN data.

References

[1] A. Cameron, The Jindalee operational radar network: its architecture and surveillance capabilities, *Proceedings IEEE International Radar conference*, 9-12 May, Arlington, Virginia, USA, 2005.

[2] M. D. E. Turley, Hybrid CFAR techniques for HF radar, *Radar 97*, Edinburgh, United Kingdom, 14 -16 October, 1997.

[3] M. D. E. Turley, Skywave radar spatial adaptive processing with quiescent pattern control, *5th International Symposium on Signal Processing and its Applications*, ISSPA 99, Brisbane, Australia, 22-25 August, 1999.

[4] G. A. Fabrizio, A. B. Gershman, and M. D. Turley, Robust Adaptive Beamforming for HF Surface Wave Over-The-Horizon Radar, *IEEE Transactions on Aerospace and Electronic Systems*, 40(2), 510 - 525, 2004.

[5] M. Turley and M. L. Lees, An adaptive impulsive noise suppressor for FMCW radar, *Proceedings IREECON87, 21st International Electronic Convention and Exhibition*, Sydney, Australia, 665-668, 1987.

[6] G. A. Fabrizio, G. J. Frazer, M. D. Turley, STAP for Clutter and Interference Cancellation in a HF Radar System, in *International conference on Acoustics, Speech and Signal Processing (ICASSP)*, Toulouse, France, May 14-19, 2006.

[7] M. D. E. Turley and M. A. Tyler, A technique for estimating the detection performance of a skywave over-the-horizon radar, *Proceedings IEEE International Radar conference (RADAR 2009)*, Bordeaux, France, October 12-16, 2009.

[8] S. B. Colegrove, S.J. Davey and B. Cheung, A Tracker Assessment Tool for Comparing Tracker Performance, Technical report, DSTO-TR-1694, Defence Science and Technology Organisation, Australia, 2005

[9] S. A. Mabbs, A performance assessment environment for radar signal processing and tracking algorithms, *Proceedings IEEE Pacific Rim Conference on Communications, Computers and Signal Processing*, Victoria, BC, Canada, 19-21 May 1993.

[10] D. J. Kewley and I. W. Dall, Performance assessment criteria for OTH radar, *International Conference on Radar*, Brighton, United Kingdom, 12-13 Oct, 1992.

[11] J. Praschifka, J. Lomen, and M. D. E. Turley, Analysis of detection performance for signals in unknown noise, as applied to HF skywave radar, *Processing International symposium on signal processing and its applications (ISSPA)*, Gold Coast, Australia, 23-30 August, 1996.

[12] H. L. van Trees, Detection, Estimation, and Modulation Theory - Part I: Detection, Estimation, and Linear Modulation Theory, *New York, John Wiley & Sons*, 1968.

[13] M. Turley, Impulsive noise rejection in HF radar using a linear prediction technique, *IEEE conference on radar*, Adelaide, Australia, 3 - 5 September, 2003.

[14] M. D. E. Turley, Signal Processing Techniques for Skywave Radar Maritime Surveillance, in *Proceedings Defence Applications of Signal Processing*, Fraser Island, Australia, 10-14 December, 2006.