

Clutter-based Test Statistics for Automatic Track Initiation

KENNEDY Hugh Lachlan¹

Abstract Two test statistics based on clutter characteristics are derived. A tentative track is confirmed when the track-is-on-clutter hypothesis is rejected. A constant and known false track rate results when the assumptions of the null hypothesis are true. The first test statistic is based on the clutter density. A high probability of target detection is resulted when the expected distance to the nearest target peak is less than the expected distance to the nearest clutter peak. The second test statistic is based on the clutter amplitude. A high probability of target detection is resulted when the expected amplitude of the target peak is greater than the expected amplitude of clutter peaks. The behavior of the clutter-based test statistics is compared with the target visibility method, using simulated data. All track initiation methods are applied using a track updater based on probabilistic data association (PDA), extended to incorporate peak amplitude information, which is available.

Key words Tracking filters, target detection, automatic track initiation, probabilistic data association (PDA)

In operational wide-area surveillance systems — for example, those described in [1–4] — automatic track initiation is required to reduce operator workload and increase overall system effectiveness. Automatic track initiation is the process whereby tracks on targets are promptly presented to the operator for assessment, without manual intervention. A tentative track is started on all target-like peaks and automatically maintained by the tracking subsystem (the tracker). These tracks are only confirmed and displayed when the track confidence exceeds a nominated upper threshold; tracks with a confidence below a nominated lower threshold are deleted. The ability of the tracker to discriminate between true and false tracks determines the average false track rate, the average track initiation (or establishment) delay, and the average number of tracks held in the tracking database (thus the computational load).

Probabilistic data association (PDA) — which was initially formulated in [5], developed and customized in papers such as [6–7], and recently reviewed in [8–9] — is well-suited to wide-area surveillance systems where the clutter (or false peak) density is typically high, the probability of target detection low, the scan (or revisit) rate low, the surveillance volume large, and the measurement noise high. In these environments, PDA-based trackers provide robustness at a low computational cost, along with reduced development and maintenance costs due to their relative simplicity.

In [10], PDA was extended to incorporate automatic track initiation (and termination), by considering the target-is-not-observable event, in addition to the all-associated-peaks-are-clutter event of the original PDA formulation. The idea has since been developed further, with observability replaced by visibility^[7], perceivability^[11], and existence^[12]. Collectively, these approaches are sometimes referred to as integrated PDA (IPDA). In [13], a fixed-lag smoother is used to refine the estimate of existence; while in [14], the approach is reformulated using an interacting multiple model. State estimation and existence estimation are tightly coupled in IPDA — the gain of the tracking filter is influenced by the existence probability, whereas the existence probability is determined by the target probability density relative to the clutter density. Other PDA-based and non-PDA-based track initiation schemes are described in [15–16] and [17–23], respectively.

In this paper, track initiation schemes based purely on clutter statistics are presented. The problem is formulated as a hypothesis test, with the null hypothesis being that the track is on clutter. A track is confirmed when the null hypothesis is rejected. When the assumptions underlying the null hypothesis are true, a constant and known false track confirmation probability results — the size of the test; furthermore, approaching the track initiation problem in this way has the potential to reduce track initiation delays on targets when the prior assumed measurement noise and process noise distributions (as defined in the R and Q matrices) are inappropriate or when the state estimates are poor. The methods are also suitable for use in legacy systems employing non-Bayesian trackers^[24–25].

Three filters for the automatic initiation of tracks in clutter are described in Section 1. After a brief overview of the target visibility filter, two clutter-only methods are presented. The first is based on the spatial distribution (density) of the clutter peaks; the second is based on the signal strength (amplitude) of the clutter peaks. The parameters of the clutter population are estimated and any track on peak sequences that are unlikely to be due to clutter are confirmed. Test statistics are evaluated using the nearest peak to associate with the track on each update, over the recent history of the track. The power of the tests, i.e., the achieved true track confirmation probability, depends on the target peak characteristics. The amplitude test is powerful when the mean signal strength of target detections is significantly greater than the mean signal strength of the false detections; the density test is powerful when the expected distance to the nearest target peak is less than the expected distance to the nearest clutter peak. The track initiation filters are applied to tracks managed by PDA-based update algorithms, which are outlined in Section 2. It is important to note that, unlike the target visibility track initiation method, the clutter-based track initiation methods do not influence the behavior of the underlying track update algorithm in any way. This decoupling has a number of desirable effects, which are discussed in Sections 4 and 5.

1 Track initiation filters

1.1 Target visibility filter

The automatic track initiation technique described in [7, 10, 26] is briefly presented here as a reference IPDA implementation. During a given update of a given track, the prior probability of target visibility $p_v(k|k-1)$, estimated

Received September 13, 2007; in revised form January 7, 2008
Supported by Technical Knockout Systems Pty. Ltd.
1. Taranaki Road, Edinburgh Parks, Edinburgh SA 5111, Australia
DOI: 10.3724/SP.J.1004.2008.00266

recursively from previous scans, is used to weight the peaks-due-to-target events of the current scan k . The posterior value of $p_v(k|k)$ is then updated using Bayes' rule. The value of p_v is propagated forward in time, from one scan to the next, using a Markov chain,

$$p_v(k|k-1) = \Delta p_v(k-1|k-1) + \bar{\Delta}(1 - p_v(k-1|k-1)) \quad (1)$$

with the elements Δ and $\bar{\Delta}$ corresponding (respectively) to the probability that the target remains, and becomes, visible. The Markov chain effectively prevents p_v from becoming locked at zero or unity during the recursion, thus ensuring that it remains responsive to new data. Track confidence $p_{c_v}^{[26]}$ is used for all track confirmation and deletion decisions; it is a smoothed representation of p_v ,

$$p_{c_v}(k) = (1 - \omega_c) p_v(k|k) + \omega_c p_{c_v}(k-1) \quad (2)$$

where ω_c is an arbitrarily or empirically chosen smoothing parameter, with no statistical or probabilistic basis. A tentative track is confirmed when $p_{c_v} \geq p_{con}$.

1.2 Clutter density filter

It is assumed that clutter detections are produced by a Poisson process, so that the probability that M clutter peaks are contained within a given volume V is given by

$$p_D(M; \rho, V) = \frac{(\rho V)^M \exp(-\rho V)}{M!} \quad (3)$$

where the parameter ρ is the unit occurrence rate. Distances between consecutive Poisson events are distributed as an exponential variable. This relationship is used to solve queuing problems arising in information systems^[27]. When applied to the problem at hand and used to model the clutter-free volume V_g , with a clutter density of ρ , the exponential distribution takes the following form:

$$f_D(V_g; \rho) = \rho \exp(-\rho V_g) \quad (4)$$

Here, V_g is the (ellipsoidal) clutter-free space around the predicted measurement location, enclosed by the nearest associated peak. It is computed using

$$V_g = V_c D^m \sqrt{|S|} \quad (5)$$

where S is the innovation covariance matrix; V_c is the volume of the unit sphere in the m -dimensional measurement space, computed using

$$V_c = \frac{\pi^{\frac{m}{2}}}{\Gamma(\frac{m}{2} + 1)} \quad (6)$$

and D is the Mahalanobis distance, computed using

$$D = \sqrt{\mu^t S^{-1} \mu} \quad (7)$$

In the above equation, μ is the innovation vector, of length m , with the superscript t denoting its transpose. It is the difference between the predicted value of the target measurement and the observed value of the nearest target measurement. It is assumed that the parameter ρ is known but not necessarily constant from scan to scan. In practice, it is usually estimated by dividing the number of peaks generated in a given scan N_s , which is assumed here to be large, by the surveillance volume of the sensor V_s (see [28] for an alternative approach). As a result of the relationship

between the gamma and chi-squared distributions^[29], (4) yields the following test statistic:

$$2\rho V_g \sim \chi^2(2) \quad (8)$$

If a tentative track is initiated at $k = 0$, and updated using the nearest detection over the next K consecutive scans, then from the reproductive property of chi-squared variables, the following relationship holds if the track is only updated by clutter peaks:

$$Z_D(k) = 2 \sum_{k=1}^K \rho(k) V_g(k) \sim \chi^2(2K) \quad (9)$$

where the k indices denote the scan dependence of ρ and V_g . The null hypothesis (H_{0D}) is rejected if $Z_D(k)$ exceeds a specified confirmation threshold λ_D , where λ_D is selected to give the desired size α_D , using the inverse chi-squared cumulative density function (CDF). The p -value of the test is used as a measure of track confidence on the k -th scan,

$$p_{c_D}(k) = \int_0^{Z_D(k)} \chi^2(Z; v) dZ \quad (10)$$

where v is the degree of freedom.

1.3 Clutter amplitude filter

The amplitude (in dB), or signal-to-noise ratio (SNR), of clutter peaks above the detection threshold is assumed to be distributed as an exponential variable (for empirical suitability and analytical convenience),

$$f_A(a; \theta_A, a_{\min}) = \frac{1}{\theta_A - a_{\min}} \exp\left\{-\frac{a - a_{\min}}{\theta_A - a_{\min}}\right\} \quad (11)$$

where f_A is the probability density function (PDF) of the distribution; a_{\min} is the SNR of the detection threshold, as applied by the peak detector; θ_A is the average SNR of the portion of the clutter population that is above the detection threshold, and a is the SNR of the clutter peak ($a_{\min} \leq a \leq \infty$).

Consequently, the following statistic is distributed as a chi-squared variable, with 2 degrees of freedom:

$$2 \frac{a - a_{\min}}{\theta_A - a_{\min}} \sim \chi^2(2) \quad (12)$$

If a tentative track is initiated at $k = 0$, then receives K updates, and all associated peaks are due to clutter, then from the reproductive property of chi-squared variables, the following statistic results are obtained:

$$2 \frac{\sum_{k=1}^K \{a(k, n'_k) - a_{\min}\}}{\theta_A - a_{\min}} \sim \chi^2(2K) \quad (13)$$

where k is the scan index and n'_k is the index of the nearest associated peak during the k th scan. Similarly, if the SNR of all other peaks (also assumed to be due to clutter) in each scan are summed then

$$2 \frac{\sum_{k=1}^K \sum_{n=1, n \neq n'_k}^{N_s} \{a(k, n) - a_{\min}\}}{\theta_A - a_{\min}} \sim \chi^2(2K(N_s - 1)) \quad (14)$$

where N_s is the number of detections received in each scan. The a_{\min} parameter is a system configuration parameter, its value is therefore known; the θ_A parameter is unknown

and is difficult to estimate when N_s is small, however, it is assumed to be constant over the K scans. It is eliminated when (13) is divided by (14) and a test statistic that is distributed as Snedecor's F variable results when the two equations are divided by their respective degrees of freedom. The following holds when the null hypothesis is true:

$$Z_A(k) = (N-1) \frac{\sum_{k=1}^K \{a(k, n'_k) - a_{\min}\}}{\sum_{k=1}^K \sum_{n=1, n \neq n'_k}^{N_s} \{a(k, n) - a_{\min}\}}, \text{ with}$$

$$Z_A(k) \sim F(2K, 2K(N_s - 1)) \quad (15)$$

The null hypothesis (H_{0A}) is rejected if $Z_A(k)$ exceeds a specified confirmation threshold λ_A , where λ_A is selected to give the desired size α_A , using the inverse F cumulative CDF. The p -value of the test is used as a measure of track confidence on the k th scan,

$$p_{c_A}(k) = \int_0^{Z_A(k)} F(Z; v_1, v_2) dZ \quad (16)$$

where v_1 and v_2 are the degrees of freedom.

1.4 Clutter filter selection and application

To simplify the notation in the derivation of Z_A , it has been assumed that the number of detections in each scan is constant, although this need not be the case. In both clutter-based methods, after the creation of a tentative track, the analysis window length is allowed to grow up to a maximum length of K_{\max} ; thereafter, a sliding window is used.

The clutter density filter only yields a low and false track production rate when the spatial clutter distribution is spatially uniform and uncorrelated from scan to scan. An acceptable probability of true track promotion is only obtained if the expected distance to the target peak, from the predicted peak location, on each update, on average, is less than the expected distance to the nearest clutter peak.

The clutter amplitude filter only yields a low false track production rate if the incidence of high-amplitude clutter peaks is spatially and temporally uncorrelated. A satisfactory probability of true track promotion is only realized if the average amplitude of target peaks is greater than the average amplitude of clutter peaks.

Both clutter filters may be applied simultaneously and independently where the aforementioned peak conditions prevail. A tentative track is promoted if both or either test is rejected. Either clutter filter can be used with any of the track update filters described in the next section, as the track update and track confidence methods are independent of each other.

2 Track update filters

2.1 PDA

The un-normalized (parametric) PDA event probabilities are evaluated using

$$\beta'_0 = (1 - p_d p_g) \rho \quad (17)$$

and

$$\beta'_i = p_d g_D(\mathbf{z}_i; \hat{\mathbf{y}}, S) \quad (18)$$

where $\hat{\mathbf{y}}$ is the predicted target peak location, S is the innovation covariance matrix, \mathbf{z}_i is the location of the i th associated peak, and g_D is the target likelihood function (a

multivariate Gaussian) in spatial coordinates^[5]. The normalized PDA event probabilities are then evaluated using

$$\beta_i = \frac{\beta'_i}{\sum_{i'=0}^{N_g} \beta'_{i'}}, \quad i = 0, \dots, N_g \quad (19)$$

The subscript i is used to index the event probabilities, with positive indices referring to the i -th associated peak and $i = 0$ corresponding to the all-associated-peaks-are-clutter event, which is assumed either to be due to an absent target detection or a target peak that is detected but is beyond the limits of the association gate. The parameters p_d and p_g correspond to the probability that the target peak lies above the detection threshold, and the probability that the peak lies within the association gate, respectively. Here, p_d is a known constant and p_g is computed using the limit of the ellipsoidal association gate. Following the approach taken in [30], the N_g nearest peaks are always used (i.e. a variable gate size).

The un-normalized event probabilities are weighted by a factor $f_D(N_c)$ which is formed from the product of uniform probability density functions and a Poisson probability mass function, i.e.,

$$f_D(N_c) = \left(\frac{1}{V_g}\right)^{N_c} \times \left(\frac{(\rho V_g)^{N_c}}{N_c!} \exp(-\rho V_g)\right) \quad (20)$$

The argument N_c is the number of clutter peaks assumed in the event hypothesis, i.e., $N_c = N_g - 1$, for the $i > 0$ hypotheses, where one of the peaks is assumed to be due to a target, and $N_c = N_g$ for the $i = 0$ hypothesis, where all peaks are assumed to be due to clutter. After dividing all hypotheses by $f_D(N_g)$, only a factor of ρ/N_g remains in the $i = 0$ hypothesis in (18).

2.2 PDA with target visibility

As [10], the PDA event space is extended to include track initiation and termination as follows:

$$\beta'_{-1} = (1 - p_v) \rho \quad (21)$$

$$\beta'_0 = p_v (1 - p_d p_g) \rho \quad (22)$$

and

$$\beta'_i = p_v p_d g_D(\mathbf{z}_i; \hat{\mathbf{y}}, S) \quad (23)$$

then

$$\beta_i = \frac{\beta'_i}{\sum_{i'=-1}^{N_g} \beta'_{i'}}, \quad i = -1, \dots, N_g \quad (24)$$

In the above expressions, p_v represents the prior probability of target visibility (from previous recursions), i.e., $p_v(k|k-1)$; the posterior probability of target visibility is then evaluated using $p_v(k|k) = 1 - \beta_{-1}$. PDA with target visibility is equivalent to PDA when p_v is fixed at unity.

2.3 PDA with target visibility and amplitude information

Previous studies have shown that the consideration of amplitude information, SNR or signal strength, in a tracking algorithm has the potential to improve overall performance of PDA-based tracking filters^[9, 31-34]. The extent of the benefits of course depends on the expected excess SNR of the target peaks relative to the clutter peaks. Based

loosely on the approach taken in [30], SNR is incorporated into the PDA-with-target-visibility framework using

$$\beta'_{-1} = (1 - p_v) \rho \prod_{j=1}^{N_g} f_A(a_j; \theta_A, a_{\min}) \quad (25)$$

$$\beta'_0 = p_v (1 - p_d p_g) \rho \prod_{j=1}^{N_g} f_A(a_j; \theta_A, a_{\min}) \quad (26)$$

and

$$\beta'_i = p_v p_d g_D(\mathbf{z}_i; \hat{\mathbf{y}}, S) g_A(a_i; \hat{\mu}_A, s_A^2) \prod_{j=1, j \neq i}^{N_g} f_A(a_j; \theta_A, a_{\min}) \quad (27)$$

then

$$\beta_i = \frac{\beta'_i}{\sum_{i'=-1}^{N_g} \beta'_{i'}}, i = -1, \dots, N_g \quad (28)$$

or upon rearranging and dividing all events by the product of all f_A functions:

$$\beta''_{-1} = \frac{(1 - p_v)}{p_v p_d} \quad (29)$$

$$\beta''_0 = \frac{(1 - p_d p_g)}{p_d} \quad (30)$$

and

$$\beta''_i = \left(\frac{g_D(\mathbf{z}_i; \hat{\mathbf{y}}, S)}{\rho} \right) \cdot \left(\frac{g_A(a_i; \hat{\mu}_A, s_A^2)}{f_A(a_i; \theta_A, a_{\min})} \right) \quad (31)$$

then

$$\beta_i = \frac{\beta''_i}{\sum_{i'=-1}^{N_g} \beta''_{i'}}, i = -1, \dots, N_g \quad (32)$$

where f_A is the clutter likelihood function (an exponential) in the amplitude coordinate – see (11) – and g_A is the target likelihood function (a Gaussian) in the amplitude coordinate^[30], with the parameters $\hat{\mu}_A$ and s_A^2 being the predicted mean and the innovation variance, respectively. The un-normalized and normalized event probabilities are dimensionless. For notational, mathematical, and computational simplicity, amplitude is separated from the spatial variables and is updated in parallel; furthermore, amplitude information is not used during the peak association process; therefore, the dimension of S (thus \mathbf{z} and $\hat{\mathbf{y}}$) is the same for all PDA implementations described here. The filter is factored in this way because it allows the PDA event structure to be reused without modification. The spatial and amplitude dependence of the clutter and target likelihood functions (f and g) are denoted using the D and A subscripts, respectively. Two likelihood functions are used to exploit information in the independent spatial and amplitude coordinates. Use of the Poisson probability mass function and the uniform probability density function as prior clutter distributions, which is central to the PDA formulation – see (1) and (2) in [16] – are only valid over the spatial coordinates; introduction of the amplitude coordinate would violate the assumption of clutter homogeneity.

3 Simulations

A series of Monte Carlo simulations were performed to investigate the track initiation performance of the clutter density and clutter amplitude methods, relative to the target visibility method. The simulations were performed in two groups (I and II). Peak SNR was disregarded in Group I simulations and considered in Group II simulations.

Clutter peaks were (pseudo-) randomly generated using a uniform PDF in two continuous spatial (x, y) coordinates, which results in a Poisson-distributed clutter field (not proven here) when N_s is large. As per the assumed model, the SNR of clutter peaks above the detection threshold was generated using an exponential distribution. The limits of the surveillance area are defined by the parameters x_{\max} , x_{\min} , y_{\max} , and y_{\min} . A single target is moved, at a (nominally) constant velocity, from the point $(x_{\text{beg}}, y_{\text{beg}})$ to $(x_{\text{end}}, y_{\text{end}})$ over the K_{run} scans of each run, with a scan rate of ΔT . The SNR of target peaks was generated according to a Gaussian distribution $N(\mu_A, \sigma_A)$. For a given mean target SNR (μ_A), the variance (σ_A^2) was set to give the specified probability of detection (p_d), which remained fixed for all scans in all simulations. Zero target amplitude process noise (σ'_A) was used in the simulations and filters. Zero-mean Gaussian measurement noise (with a standard deviation of σ_x and σ_y) was added to all target peaks above the detection threshold. On each update, the true target velocity was randomly perturbed by zero-mean Gaussian process noise (with a standard deviation of σ'_x and σ'_y). The initial covariance of the 1st derivative estimates was set equal to the square of the initial target velocity in each dimension. The parameters of a four dimensional state space $\mathbf{x} = [x, \dot{x}, y, \dot{y}]^T$ were estimated using a two dimensional measurement space $\mathbf{z} = [x, y]^T$.

The simulation parameters were as follows: $K_{\text{run}} = 50$, $\Delta T = 1$ s, $p_d = 0.8$, $x_{\max} = 20$ km, $x_{\min} = 10$ km, $y_{\max} = 200$ km, $y_{\min} = 100$ km, $\sigma_x = b(x_{\max} - x_{\min})$, $\sigma_y = b(y_{\max} - y_{\min})$, $\sigma'_x = b'(x_{\max} - x_{\min})$, $\sigma'_y = b'(y_{\max} - y_{\min})$, $b' = 1/1000$, $a_{\min} = 5$ dB, and $\theta_A = 10$ dB. Track initiation performance in a range of different conditions was examined with the number of clutter peaks, measurement noise factor, and mean target SNR selected from $N_s = \{25, 50, 100\}$, $b = \{1/100, 1/50, 1/25\}$, and $\mu_A = \{10$ dB, 15 dB, 25 dB}. In Group I simulations, μ_A was fixed at 10 while N_s and b varied; in Group II simulations, N_s was fixed at 50 while b and μ_A varied.

Three target scenarios (Targ. Type) were examined in each simulation group: absent target (None), slow target present (Slow), and fast target present (Fast). The target-present simulations were used to characterize track initiation delays, while the target-absent simulations were used to characterize false track rates. On the initial scan ($k = 1$), where the target was present, a single tentative track was started on the target peak (the target was always detected on the first scan); when the target was absent, a single tentative track was started on a randomly selected peak. In the latter case, the state covariance matrix P was initialized using the average velocity of the slow and fast targets. The slow- and fast-target scenarios were used to explore the affect of target speed on track initiation delay. In both cases the target started (in km) at (17, 170), then nominally, i.e., without added process noise, moved to (11, 110) in the fast-target case (approx. Mach 3) and (15, 115) in the slow-target case (approx. Mach 1), over the duration of the run.

In an attempt to understand the variables affecting track

initiation performance, a variety of simulation parameter combinations were investigated. Each scenario was repeated 1000 times using an incremented random number seed.

Track deletion logic was not implemented; therefore the track either remained in the tentative state or was confirmed. If on the last scan of the run, the target was not located within the 99% probability limits of an elliptical association gate, centered on the track and defined using the measurement noise covariance matrix (R), then the (tentative or confirmed) track was labeled as divergent. Use of the innovation covariance matrix (S) is inadequate for this purpose because it is typically large for divergent tracks (due to a large P matrix). The average track confirmation scan (Avg. Conf. Scan), which is an indication of the track initiation delay (Trk. Init. Del. = Conf. Scan-1), the number of confirmed tracks (Num. Conf.) and the number of divergent tracks (Num. Div.) were determined for each filter, for each run (see Tables 1 and 2). In the target-present scenarios, only confirmed tracks that were not divergent were included in the number of confirmed tracks count and the average track confirmation scan metric.

In Group I simulations, the peak set of a given scenario was processed, in turn, using the target visibility filter (Vis.) then the clutter density filter (Den.). These track initiation filters were implemented using the track update filters described in Sections 2.2 and 2.1, respectively. In Group II simulations, the peak set of a given scenario was processed, in turn, using: 1) the target visibility filter (Vis.); 2) the clutter density filter (Den.) and the clutter amplitude filter (Amp.) with track confirmation declared when either hypothesis is rejected (Den./Amp.); 3) the clutter density filter (Den.) and the clutter amplitude filter (Amp.) with track confirmation declared when both hypotheses are rejected on a given scan (Den. and Amp.). The track initiation filters were supported by the track update filters described in Section 2.3 with p_v held at unity for the clutter filters. The clutter filters were applied in parallel and independent of each other. A "loose" and "tight" test is applied when the clutter filters are combined using 2) and 3), respectively. The target visibility filter parameters were: $p_v(1|1) = 0.5$, $p_{c_v}(1) = 0.5$, $p_{con} = 0.7$, $\omega_c = 0.75$, $\Delta = 0.99$, and $\bar{\Delta} = 0.01$. The clutter filter parameters were: $\alpha_D = \alpha_A = 0.01$ and $K_{max} = 25$. The track update filter parameters used were: $N_g = 4$, $R = \text{diag}([\sigma_x^2, \sigma_y^2])$,

$$H = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix}, \quad F = \begin{pmatrix} 1 & \Delta T & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & \Delta T \\ 0 & 0 & 0 & 1 \end{pmatrix},$$

and

$$Q =$$

$$\text{diag}\left(\left[\left(\frac{1}{2}\sigma'_x\Delta T^2\right)^2, (\sigma'_x\Delta T)^2, \left(\frac{1}{2}\sigma'_y\Delta T^2\right)^2, (\sigma'_y\Delta T)^2\right]\right);$$

$R_A = \sigma_A^2$, $H_A = 1$, $F_A = 1$, and $Q_A = \sigma_A'^2 = 0$. Modeling error sensitivity was not investigated.

The process of filter parameter selection/tuning was as follows: Using Group I data (without SNR), the parameters of the clutter density filter were arbitrarily set and the total number of false tracks generated noted over 1000 repetitions, of the target absent scenario with $N_s = 50$ and $b = 1/50$. The parameters of the target visibility filter were then adjusted to minimize the average track initiation de-

lay, when processing the slow target scenario with $N_s = 50$ and $b = 1/50$, constrained so that number of false tracks generated in the target absent scenario with $N_s = 50$ and $b = 1/50$ approximately equaled the number of false tracks generated by the clutter density filter. The track initiation filter parameters established in this way were then used for the remaining scenarios. The Group I parameters were used, without modification, to process the Group II data.

Table 1 Group I simulation results¹

Targ. type	N_s	b	Filter	Avg. conf. scan	Num. conf.	Num. div.
None	25	1/100	Vis.	14	32	1000
None	25	1/100	Den.	16	95	1000
None	25	1/50	Vis.	14	72	1000
None	25	1/50	Den.	16	94	1000
None	25	1/25	Vis.	14	133	1000
None	25	1/25	Den.	21	103	1000
None	50	1/100	Vis.	13	66	1000
None	50	1/100	Den.	14	92	1000
None	50	1/50	Vis.	14	101	1000
None	50	1/50	Den.	20	107	1000
None	50	1/25	Vis.	16	149	1000
None	50	1/25	Den.	18	110	1000
None	100	1/100	Vis.	11	86	1000
None	100	1/100	Den.	19	103	1000
None	100	1/50	Vis.	13	157	1000
None	100	1/50	Den.	17	126	1000
None	100	1/25	Vis.	19	161	1000
None	100	1/25	Den.	21	115	1000
Slow	25	1/100	Vis.	4.5	998	2
Slow	25	1/100	Den.	4.2	997	3
Slow	25	1/50	Vis.	6.1	989	11
Slow	25	1/50	Den.	5.9	991	9
Slow	25	1/25	Vis.	9	970	23
Slow	25	1/25	Den.	11	977	21
Slow	50	1/100	Vis.	5.3	985	14
Slow	50	1/100	Den.	5	988	12
Slow	50	1/50	Vis.	7.6	964	34
Slow	50	1/50	Den.	7.8	973	26
Slow	50	1/25	Vis.	13	781	206
Slow	50	1/25	Den.	19	765	154
Slow	100	1/100	Vis.	6.6	965	35
Slow	100	1/100	Den.	6.2	969	31
Slow	100	1/50	Vis.	9.6	839	155
Slow	100	1/50	Den.	13	884	107
Slow	100	1/25	Vis.	14	308	676
Slow	100	1/25	Den.	22	306	482
Fast	25	1/100	Vis.	5.3	974	26
Fast	25	1/100	Den.	4.7	984	16
Fast	25	1/50	Vis.	6.8	932	68
Fast	25	1/50	Den.	6.9	970	30
Fast	25	1/25	Vis.	9.6	700	297
Fast	25	1/25	Den.	14	836	148
Fast	50	1/100	Vis.	6.2	904	96
Fast	50	1/100	Den.	6.1	946	54
Fast	50	1/50	Vis.	8.4	729	269
Fast	50	1/50	Den.	11	848	150
Fast	50	1/25	Vis.	10	158	840
Fast	50	1/25	Den.	23	356	575
Fast	100	1/100	Vis.	7.3	695	304
Fast	100	1/100	Den.	9.3	851	145
Fast	100	1/50	Vis.	9.6	246	753
Fast	100	1/50	Den.	19	524	461
Fast	100	1/25	Vis.	8	1	999
Fast	100	1/25	Den.	27	12	969

¹ μ_A fixed at 10 dB, SNR information ignored

Table 2 Group II simulation results²

Targ. Type	μ_A	b	Filter	Avg. Conf. Scan	Num. Conf.	Num. Div.
None	NA	1/100	Vis.	7.3	7	1000
None	NA	1/100	Den. Amp.	20	161	1000
None	NA	1/100	Den.&Amp.	28	1	1000
None	NA	1/50	Vis.	6.2	17	1000
None	NA	1/50	Den. Amp.	17	191	1000
None	NA	1/50	Den.&Amp.	17	4	1000
None	NA	1/25	Vis.	9.8	33	1000
None	NA	1/25	Den. Amp.	20	184	1000
None	NA	1/25	Den.&Amp.	NA	0	1000
Slow	25	1/100	Vis.	4.4	998	2
Slow	25	1/100	Den. Amp.	3	998	2
Slow	25	1/100	Den.&Amp.	5.5	998	2
Slow	25	1/50	Vis.	5	998	2
Slow	25	1/50	Den. Amp.	3.5	998	2
Slow	25	1/50	Den.&Amp.	7.8	998	2
Slow	25	1/25	Vis.	6.1	998	2
Slow	25	1/25	Den. Amp.	5.5	1000	0
Slow	25	1/25	Den.&Amp.	18	978	0
Slow	15	1/100	Vis.	5.4	991	9
Slow	15	1/100	Den. Amp.	4.2	996	4
Slow	15	1/100	Den.&Amp.	8.3	996	4
Slow	15	1/50	Vis.	7.9	994	6
Slow	15	1/50	Den. Amp.	5.6	996	4
Slow	15	1/50	Den.&Amp.	11	996	4
Slow	15	1/25	Vis.	12	897	81
Slow	15	1/25	Den. Amp.	11	974	24
Slow	15	1/25	Den.&Amp.	22	894	24
Slow	10	1/100	Vis.	5.9	976	22
Slow	10	1/100	Den. Amp.	5.1	993	7
Slow	10	1/100	Den.&Amp.	28	449	7
Slow	10	1/50	Vis.	10	903	82
Slow	10	1/50	Den. Amp.	7.8	984	16
Slow	10	1/50	Den.&Amp.	28	362	16
Slow	10	1/25	Vis.	18	515	395
Slow	10	1/25	Den. Amp.	18	837	121
Slow	10	1/25	Den.&Amp.	31	140	121

² N_s fixed at 50

4 Discussion

On the one hand, the IPDA method is desirable because the filter gain is reduced (due to the lower weighting of the peak-is-due-to-target events) as target visibility falls, so that tracks maintain their current heading in clutter during periods of reduced signal strength; on the other hand, the method is undesirable during track initiation in stressful conditions and during maneuvers because the tracking filter does not have the gain required to remove position and/or velocity bias. Use of the target visibility filter reduces the number of confirmed tracks and increases the number of divergent tracks relative to the clutter density filter. This behavior is most apparent in the more stressful scenarios in Table 1 (i.e. fast target, $N_s \geq 50$ and $b \geq 1/50$). In these cases, the target visibility filter results in faster track initiation than the clutter density filter, on average (as indicated

by Avg. Conf. Scan); however, this is because most of the target visibility tracks diverge, whereas clutter density filter tracks struggle for many scans but manage to follow the target and eventually get confirmed. Fig. 1 shows that the target visibility filter is much less likely to confirm tracks after the first 10 scans, in one of the more challenging scenarios. The average track confirmation scan is increased as a result. In the less stressful scenarios the behavior of the filters is similar, although less pronounced.




Fig. 1 Aggregate results (1000 repetitions) for a stressful scenario (Targ. Type = Fast, $N_s = 100$, $b = 1/50$)

The tendency of the target visibility filter to suppress poor tracks on targets also acts to suppress false tracks on weakly correlated sequences of clutter peaks (and delete them quickly). The results in Table 1 show that the number of false tracks produced by the clutter density filter is approximately constant (around 100), in all target-absent scenarios; while the false track rate of the target visibility filter depends on the scenario parameters. Furthermore, use of prior target models by the target visibility filter consistently results in a lower false track rate in the low measurement noise scenarios ($b = 1/100$) and a higher false track rate in the high measurement noise scenarios ($b = 1/25$), for all clutter densities ($N_s = 25, 50$, and 100), when compared with the clutter density filter. Without exception, the number of false tracks generated by the target visibility filter increases with the clutter density, for any given level of measurement noise.

Use of average confirmation delays obscures a significant difference between the clutter filters and the target visibility filter – use of a smoothing factor (ω_c) of 0.75 makes it impossible for a track to be confirmed before third scan; the limit increases as the smoothing factor increases, even if all of the nearest associated peaks coincide with the predicted peak location during the life of the tentative track. The clutter filters permit track confirmation on any scan, except for $k = 1$ (the initiation scan). As a consequence, average track initiation delays for the clutter density filter are less than the target visibility filter, in the benign scenarios with low measurement-noise ($b = 1/100$) and slow targets. The different track confirmation behavior at low scans is evident in Fig. 1.

The use of amplitude information in the Group II simulations results in a five- to ten-fold reduction in the number of false tracks generated by the target visibility filter (comparing Table 1 with Table 2). Similar reductions in the number of divergent true tracks, when the average target SNR is

significantly greater than the clutter ($\mu_A \geq 15$ dB), were realized for all track initiation methods, due to the enhanced operation of the PDA-based track updater; however, the number of divergent tracks produced by the clutter initiation methods (Den.|Amp. and Den. & Amp.) remain less than or equal to the number generated by the target visibility method in all Group II scenarios. Use of the Den.|Amp. track initiation logic nearly doubles the false track rate but reduces the track initiation delays, relative to the corresponding Group I results; whereas the Den. & Amp. test all but eliminates false tracks, at the expense of significantly longer track initiation delays, especially in the low SNR scenarios ($\mu_A = 10$ dB).

Supplementary simulations were performed to verify that the probability of false track confirmation in the target absent scenarios, on any given scan, is approximately equal to the size of the clutter density test. This relationship is not reflected in the results of Tables 1 and 2. The absent target scenario with 100 clutter peaks was modified so that the track confirmation was only applied on the last scan of the run, in an attempt to eliminate scan-dependent correlations. Runs with a maximum of 2 to 50 scans were created; each was repeated 1000 times using a size of 0.01. If the clutter has been correctly modeled, and the filters/simulations correctly implemented, then the total number of confirmed tracks in Fig. 2 should be approximately equal to 0.01×1000 , on every scan. The agreement should improve as the number of clutter peaks and repetitions increases. If there is a slight bias towards zero at high scans, it is due to edge effects, i.e., as tentative tracks wander beyond the limits of sensor coverage, the expected distance to the nearest clutter peak increases, making false track confirmation less likely. The analysis was repeated for the clutter amplitude filter (see Fig. 3).

Two clutter-based test statistics have been defined in this paper. Their suitability in a given application depends on the extent to which the clutter distributions satisfy the assumptions of the null hypothesis and the expected power of the test (as discussed in the opening paragraph of Section 1). In practice, the assumption of clutter uniformity is more likely to be met when the parameters of the clutter statistics are estimated in the local vicinity of each track, rather than over the entire surveillance volume. Where appropriate, both clutter tests are applied separately and independently of each other.




Fig. 2 The number of confirmed (false) tracks produced by the clutter density filter is compared with the theoretically expected values




Fig. 3 The number of confirmed (false) tracks produced by the clutter amplitude filter is compared with the theoretically expected values

The fact that the clutter filters only use the nearest associated peak, whereas the PDA filter uses many nearest peaks, may appear to be suboptimal; however, regardless of the track initiation technique used, a PDA track is unlikely to follow a target when the target peak is not the nearest peak, most of the time; therefore, use of the nearest peak is a justifiable simplification.

5 Conclusions

Due to the coupling between the track initiation filter and the track update filter, the target visibility method tends to suppress both true and false tentative tracks on weakly correlated peak sequences. The suitability of the target visibility filter relative to the clutter density and clutter amplitude filters, therefore, depends on the cost of a type I error (false track confirmed) relative to the cost of a type II error (true track not confirmed), where “cost” is a user defined system requirement. The clutter filters do not influence the behavior of the track update filter and, unlike the target visibility filter, the false track rate does not depend on the clutter density or the parameters of the track update filter; consequently, the clutter filters yield a constant (and known) false alarm rate when the conditions assumed under the null hypotheses are true. The clutter density and clutter amplitude filters, respectively, result in fast true track initiation when the expected distance to the nearest target peak is less than the expected distance to the nearest clutter peak and when the expected target peak SNR is greater than the expected clutter peak SNR. It is shown that the use of peak amplitude information has the potential to enhance the track initiation performance of all methods, when the expected target SNR is significantly greater than the average clutter SNR. Like all model-based techniques, the effectiveness of the clutter-based methods in any given application will be determined by the extent to which the underlying statistical assumptions of the models are satisfied.

References

- 1 Colegrove S B. Project Jindalee: from bare bones to operational OTHR. In: Proceedings of the Record of the IEEE 2000 International Radar Conference. Alexandria, USA: IEEE, 2000. 825–830
- 2 Wang W, Peng Y N, Quan T F, Liu Y T. HF OTHR target detection and estimation subsystem. *IEEE Aerospace and*

- Electronic Systems Magazine*, 1999, **14**(4): 39–45
- 3 Wise J C. Summary of recent Australian radar developments. *IEEE Aerospace and Electronic Systems Magazine*, 2004, **19**(12): 8–10
 - 4 Shaw S W, Arnold J F. Design and implementation of a fully automated OTH radar tracking system. In: Proceedings of the Record of the IEEE International Radar Conference. Alexandria, USA: IEEE, 1995. 294–298
 - 5 Bar-Shalom Y, Tse E. Tracking in a cluttered environment with probabilistic data association. *Automatica*, 1975, **11**(5): 451–460
 - 6 Fortmann T, Bar-Shalom Y, Scheffe M. Sonar tracking of multiple targets using joint probabilistic data association. *IEEE Journal of Oceanic Engineering*, 1983, **8**(3): 173–184
 - 7 Colegrove S B, Davey S J. PDAF with multiple clutter regions and target models. *IEEE Transactions on Aerospace and Electronic Systems*, 2003, **39**(1): 110–124
 - 8 Kirubarajan T, Bar-Shalom Y. Probabilistic data association techniques for target tracking in clutter. *Proceedings of the IEEE*, 2004, **92**(3): 536–557
 - 9 Bar-Shalom Y, Kirubarajan T, Lin X. Probabilistic data association techniques for target tracking with applications to sonar, radar and EO sensors. *IEEE Aerospace and Electronic Systems Magazine*, 2005, **20**(8): 37–56
 - 10 Colegrove S B, Davis A W, Ayliffe J K. Track initiation and nearest neighbours incorporated into probabilistic data association. *Journal of Electrical and Electronics Engineering*, 1986, **6**(3): 191–198
 - 11 Li N, Li X R. Tracker design based on target perceivability. *IEEE Transactions on Aerospace and Electronic Systems*, 2001, **37**(1): 214–225
 - 12 Musicki D, Evans R, Stankovic S. Integrated probabilistic data association. *IEEE Transactions on Automatic Control*, 1994, **39**(6): 1237–1241
 - 13 Chakravorty R, Challa S. Augmented state integrated probabilistic data association smoothing for automatic track initiation in clutter. *Journal of Advances in Information Fusion*, 2006, **1**(1): 63–74
 - 14 Bar-Shalom Y, Chang K C, Blom H A P. Automatic track formation in clutter with a recursive algorithm. In: Proceedings of the 28th IEEE Conference on Decision and Control. Tampa, USA: IEEE, 1989. 1402–1408
 - 15 Jing Z L, Zhou H R, Wang P D. Tracking initiation and termination of multiple maneuvering targets in a dense multi-return environment. In: Proceedings of the 29th IEEE Conference on Decision and Control. Honolulu, USA: IEEE, 1990. 2270–2275
 - 16 Chen H M, Li X R, Bar-Shalom Y. On joint track initiation and parameter estimation under measurement origin uncertainty. *IEEE Transactions on Aerospace and Electronic Systems*, 2004, **40**(2): 675–694
 - 17 Caprari R S, Goh A S. Detect-track-confirm filter with minimal constraints. *IEEE Transactions on Aerospace and Electronic Systems*, 2004, **40**(1): 336–345
 - 18 Zhu H Y, Han C Z, Han H, Zuo D G, Wen R. A track initiation and termination approach based on the combination of SPRT and EM algorithm. In: Proceedings of the 6th International Conference of Information Fusion. Queensland, Australia: IEEE, 2003. 1282–1286
 - 19 Davey S J, Rutten M G. A comparison of three algorithms for tracking dim targets. In: Proceedings of IEEE Conference on Information, Decision and Control. Adelaide, Australia: IEEE, 2007. 342–347
 - 20 Sun H M, Chiang S M. Tracking multitarget in cluttered environment. *IEEE Transactions on Aerospace and Electronic Systems*, 1992, **28**(2): 546–559
 - 21 Van K G. Sequential track extraction. *IEEE Transactions on Aerospace and Electronic Systems*, 1998, **34**(4): 1135–1148
 - 22 Hu Z J, Leung H, Blanchette M. Statistical performance analysis of track initiation techniques. *IEEE Transactions on Signal Processing*, 1997, **45**(2): 445–456
 - 23 Davey S J, Gray D A. Integrated track maintenance for the PMHT via the hysteresis model. *IEEE Transactions on Aerospace and Electronic Systems*, 2007, **43**(1): 93–111
 - 24 Benedict T, Bordner G. Synthesis of an optimal set of radar track-while-scan smoothing equations. *IEEE Transactions on Automatic Control*, 1962, **7**(4): 27–32
 - 25 Castella F R. Sliding window detection probabilities. *IEEE Transactions on Aerospace and Electronic Systems*, 1976, **12**(6): 815–819
 - 26 Colegrove S B. Advanced Jindalee Tracker: Probabilistic Data Association Multiple Model Initiation Filter, Technical Report No. DSTO-TR-0659, Defence Science and Technology Organization, 1999
 - 27 Kleinrock L. *Queueing Systems – Volume 1: Theory*. New York: John Wiley and Sons, 1975
 - 28 Li X R, Li N. Integrated real-time estimation of clutter density for tracking. *IEEE Transactions on Signal Processing*, 2000, **48**(10): 2797–2805
 - 29 Bain L J, Engelhardt M. *Introduction to Probability and Mathematical Statistics (Second Edition)*. California, USA: Duxbury Press, 1992
 - 30 Colegrove S B, Davey S J. On using nearest neighbours with the probabilistic data association filter. In: Proceedings of IEEE International Radar Conference. Alexandria, USA: IEEE, 2000. 53–58
 - 31 Lerro D, Bar-Shalom Y. Interacting multiple model tracking with target amplitude feature. *IEEE Transactions on Aerospace and Electronic Systems*, 1993, **29**(2): 494–509
 - 32 Askar H, Chen Y, Li Z M. Tracking dim moving point target using modified probabilistic data association filter. In: Proceedings of IEEE International Conference on Communications, Circuits and Systems and West Sino Expositions. IEEE, 2002. 620–624
 - 33 Slocumb B J. Comparison of two augmented PDA filters. In: Proceedings of the 1997 American Control Conference. Albuquerque, USA: IEEE, 1997. 3688–3692
 - 34 Li X R. Tracking in clutter with strongest neighbor measurements—part I: theoretical analysis. *IEEE Transactions on Automatic Control*, 1998, **43**(11): 1560–1578



KENNEDY Hugh Lachlan Senior algorithm specialist at BAE Systems, Australia. He received his bachelor and Ph.D. degrees from the University of New South Wales, Sydney, Australia, in 1993 and 2000, respectively. His research interest covers tracking, data fusion, DSP, image processing, radar, sonar, acoustics, and electronic warfare.
E-mail: hugh.kennedy@baesystems.com, hkennedy@tkosystems.com.au