OPTIMIZING A SLIDING M-OF-N TRACK INITIALIZER IN CLUTTER

Douglas A. Abraham[a], Douglas J. Grimmett[b] and Jonathan J. Itschner[b]

[a]CausaSci LLC, P.O. Box 627, Ellicott City, MD 21041, USA, abrahad@ieee.org
[b]SPAWAR Systems Center Pacific, 53560 Hull Street, San Diego, CA 92152, USA, grimmett@spawar.navy.mil

Abstract: ***In active sonar applications with moving targets, objects are tracked by associating and combining single-ping measurements (clusters) over short segments of time and space. A common track initialization algorithm (TIA) requires observing M single-ping clusters across any N consecutive pings within a local search region. Design of this TIA is straightforward when the background consists of false alarms (clusters) occurring randomly in the search space, but less tractable in the more realistic scenario of clutter tracks formed on echoes of physical objects (e.g., rock outcrops, fish schools, etc.). In this paper, a process is described for optimizing the sliding M-of-N TIA by choosing M and N to optimize target detection performance while using the cluster decision threshold to maintain a desired false-track rate in the presence of clutter events having a random duration. An example design is shown for a high-duty-cycle system with different coherent processing intervals. Although these result in different values of N, simulation testing illustrates how the design produces the desired false-track rate.***

Keywords: *tracking, track initialization, false-track rate, sliding M-of-N detector*

1. INTRODUCTION

The first step in many target tracking algorithms is to initialize tracks by associating and combining measurements over small regions of space and time. These track initialization algorithms (TIAs) typically use simple motion models and start tracks when the associated data satisfy a detection criteria. For example, it can be required that a measurement be observed within the tracker search gate in M of any consecutive N temporal updates [1, Sect. 3.6.2], which is known as a *sliding M-of-N* detection criterion. Design of the sliding M-of-N TIA entails choosing M, N, and the decision threshold (h) applied to form measurements at each update. Ideally, M and N are chosen to minimize the signal-to-noise power ratio (SNR) required to achieve a desired detection performance and h is chosen to achieve a desired false-

track rate (FTR) given M and N. In practice, M is often the only parameter chosen to optimize performance. When false alarms are dominated by clutter rather than diffuse background, the decision threshold h is typically chosen through trial-and-error to manage the quantity of clusters in each update without regard to M and N. If an average target track duration is known, it can be used to obtain N. When it is not, N is usually chosen heuristically to balance between robustness to signal fading (large N) or target maneuvers (small N). Although it is practical, this design process generally does not result in a constant false-track rate and is more complicated for high duty cycle (HDC) systems with sub-pulse processing [2,3] where the update rate (and therefore N) varies with the coherent processing interval (CPI).

In many scenarios, false tracks arise from physical objects within the scene (e.g., rock outcrops or fish schools in active sonar) rather than randomly occurring false detections arising from diffuse reverberation and noise. To handle the more realistic scenario, a design process is presented in this paper for the sliding M-of-N TIA by representing the occurrence of clutter events with a two-state Markov random process. The clutter-event model is presented in Sect. 2, along with the probability of the sliding M-of-N TIA detecting a clutter event, the average time between false tracks, and the proposed design process. A key requirement in the design is accurate modeling of the measurement decision statistic (typically the cluster peak normalized intensity). Although not presented, the generalized Pareto distribution (GPD) is used to characterize the thresholded cluster peak normalized intensity as was done in [4, 5]. In Sect. 3 optimization of the sliding M-of-N TIA is illustrated with an HDC application entailing different values of N and a simulation analysis used to demonstrate how the design process controls the rate of false tracks.

2. MODELING AND OPTIMIZING TRACK INITIALIZATION

Standard active sonar signal processing includes matched filtering the time series measured in each beam, normalization of the instantaneous intensity by the average background power, and application of a decision threshold to identify data that differs in power level from the background. The set of range/bearing points exceeding the decision threshold are then clustered prior to track initialization. The track initialization algorithm (TIA) associates clusters over multiple consecutive updates when they are consistent with an initial motion assumption (e.g., nearly constant velocity). With respect to cluster detection, this is similar to dividing the range/bearing space into a set of fixed regions and thresholding the peak normalized intensity within each. This suggests modeling track initialization as occurring over a number of independent *channels*,

$$N_{r\theta} \approx \frac{\text{\# beams} \times \text{range window}}{\text{cluster beam extent} \times \text{range extent}}, \tag{1}$$

simultaneously being tested for the beginning of a track. The cluster extents in range and bearing may be measured from real data or derived from the clustering algorithm parameters. Under this modeling assumption where the TIA runs separately on each individual channel simultaneously, the FTR required of the TIA is the system-level FTR (number of false tracks per unit time) divided by $N_{r\theta}$.

2.1. False tracks from clutter events

In many scenarios false tracks are produced by clutter arising from spatially isolated physical objects (e.g., rock outcrops, seaweed, or fish schools) rather than randomly occurring thresh-

old crossings from diffuse reverberation or ambient noise. This implies that clutter event clusters do not occur randomly (i.e., uniformly distributed) over time. However, it is reasonable to assume they are uniformly distributed over the spatial scene and so are equally likely to occur in any tracker channel. Within a given tracker channel, this clutter-event model describes a two-state random process where the data are either from diffuse background or a clutter event.

If the random process is Markovian (i.e., the statistical characterization of a future state can be described by the current state) and stationary, the dwell time in the diffuse-background or clutter-event states is a geometric random variable (using the alternative definition having a minimum of one sample [6, Ch. 23]) with success probability equal to the inverse of the average length. The probability mass function of the length of the ith clutter event (C_i) is then

$$f_{\mathrm{G}}[c; \bar{C}^{-1}] = \frac{1}{\bar{C}}\left(1 - \frac{1}{\bar{C}}\right)^{c-1} \quad \text{for } c = 1, 2, 3, \ldots \tag{2}$$

where $\bar{C}$ is the average length of a clutter event. Evidence supporting the geometric distribution (and therefore the Markov-random-process assumption) is found using data from the NATO Center for Maritime Research and Experimentation's (CMRE's) Littoral Continuous Active Sonar 2015 Experiment (LCAS15; see Sect. 5 for full acknowledgement). The data were processed for detection and through tracking to form a set of clutter tracks as described in [7]. The exceedance distribution function (EDF) of clutter-track length is shown in Fig. 1 as estimated from the LCAS15 data (along with a 90% confidence interval) and using the geometric model while conditioning on the track length being at least 17 updates long (17 was the shortest track length). Except for the shorter track lengths, the fit of the geometric model appears quite good. Note that the distribution of the length of a track produced by a sliding M-of-N detector is approximately geometric at large values (e.g., see the approximation in [8]). To ensure the results shown in Fig. 1 describe the clutter events and are not affected by the tracker, the EDF expected from the track-end logic (which required N consecutive coasts) is also shown using the approximation in [8]. With the length-distribution EDF of the tracker end-logic so close to one, the tracker is expected to produce much longer tracks than that seen in the data if the clutter event were to continue. As such, the clutter tracks are believed to end because the clutter event ends and are not affected by the tracker end logic.

The waiting time B_i before the ith clutter event is also modeled as a random variable. However, because false tracks are assumed to arise during clutter events and not during periods of diffuse background, only its mean $\bar{B}$ is required and no assumptions must be made about its distribution.

2.2. Average time between false tracks

The false-alarm performance metric of interest is the false-track rate (FTR), which is the number of false tracks observed per unit time. The TIA is designed using the inverse of this metric, $F = 1/\mathrm{FTR}$, which is the average time between false tracks. As previously noted, when the TIA operates on a single tracker channel, it is expected to observe a diffuse background interrupted at times by clutter events. In most scenarios false tracks will arise from the clutter events and not the diffuse background. However, the average extents of the background and clutter states should be included when designing the TIA. For example, if the diffuse background is much more prevalent than the clutter (i.e., $\bar{B} \gg \bar{C}$), the TIA can operate aggressively with respect to the desired FTR because it is so often encountering diffuse background and only rarely observes a clutter event.

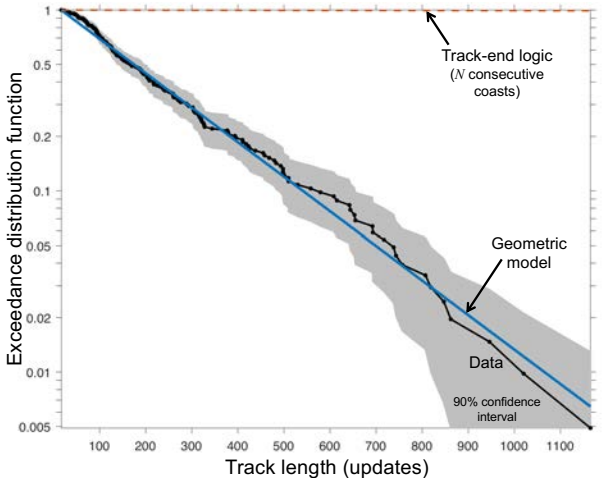

Fig. 1: Estimated and modeled exceedance distribution function of clutter track lengths formed from LCAS15 data with a CPI of $T_p = 5\,s$.

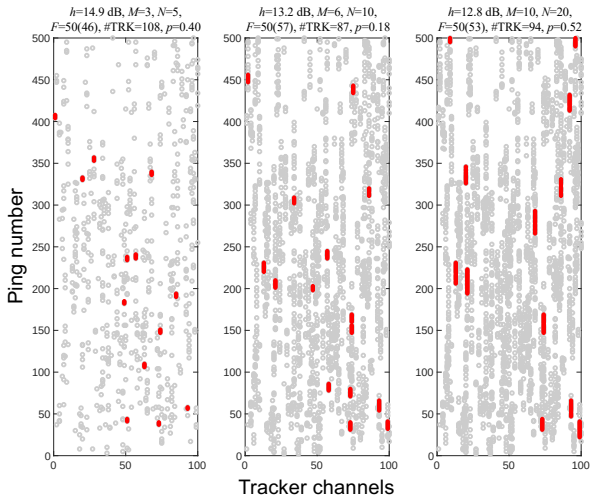

*Fig. 2: Simulated false clusters and tracks for design using $N = 5$, 10, *and* 20 *and an average of* $F = 50$ *samples between tracks.*

Once the ith clutter event occurs, the data feeding the TIA during the clutter event, X_c for $c = 1, \ldots C$, are assumed to come from a distribution $f_c(x)$ so the probability of observing a cluster is

$$p_0 = \int_h^\infty f_c(x)\,\mathrm{d}x, \tag{3}$$

where h is the cluster detection decision threshold. When the sliding M-of-N TIA is applied to these data, the probability of the TIA detecting the clutter event (i.e., producing a false track) within its C_i samples after averaging over the geometrically-distributed C_i can be approximated by

$$P_{\text{ft}} \approx \left\{ \sum_{c=0}^{N} \frac{1}{\bar{C}} \left(1 - \frac{1}{\bar{C}}\right)^{c-1} \left[1 - F_{\text{Bin}}(M-1; c, p_0)\right] \right\} + \left(1 - \frac{1}{\bar{C}}\right)^{N} \left[1 - \frac{\rho \tilde{p}^{N+1}}{\bar{C} - \tilde{p}(\bar{C}-1)}\right] \tag{4}$$

where F_{Bin} is the cumulative distribution function (CDF) of the binomial distribution, $\tilde{p} = (Q'_3/Q'_2)^{1/N}$, and $\rho = [Q'_2]^3/[Q'_3]^2$ with Q'_L depending on M, N, and p_0 as described in [8] (or see [9], [10, eqs. 4.4 & 4.5]). The latter term in (4) exploits Naus' approximation [8] for the probability of a sliding M-of-N test stopping within c samples.

Assuming each clutter event is independent of the others and identically distributed in terms of its length, the probability that the first false-track detection occurs in the ith clutter event is

$$\Pr\{\text{stop in } i\text{th clutter event}\} = (1 - P_{\text{ft}})^{i-1} P_{\text{ft}} \tag{5}$$

for $i = 1, 2, \ldots$, under the assumption that no TIA detections occur during the periods of diffuse background. This implies that I, the random stopping time of the TIA (in units of numbers of clutter events), is a geometric random variable with success probability P_{ft} and PDF as in (2) (replacing $1/\bar{C}$ with P_{ft}). The average number of clutter events before one produces a false track is therefore $1/P_{\text{ft}}$. This can then be used to form the average number of samples between

false tracks in a single tracker channel,

$$F_{\text{1-ch}} = E\left[\sum_{i=1}^{I} B_i + \sum_{i=1}^{I-1} C_i\right] = \bar{B}E[I] + \bar{C}E[I-1] = \frac{\bar{B}+\bar{C}}{P_{\text{ft}}} - \bar{C}, \tag{6}$$

where the beginning of the detected clutter event is taken as the time of the false track. Recall that this is related to the system-level average time between false tracks through $F_{\text{1-ch}} = N_{r\theta}F$.

2.3. Optimizing the track initializer

The objective in optimizing the TIA is to minimize the SNR required to achieve a desired detection performance specification while maintaining a desired false-track rate. If the false-track rate FTR_{h} is in units of false tracks per hour, then the desired average number of updates (e.g., pings or CPIs) between false tracks is

$$F = \frac{3600}{\text{FTR}_{\text{h}} \times T_{\text{up}}} \tag{7}$$

where T_{up} is the time between updates in units of seconds.

In most applications, the temporal extent over which the TIA is operated ($\sim NT_{\text{up}}$) is chosen to balance between robustness to signal fading and target maneuvers and produces a fixed value of N. Systems that vary the update rate yield a range of N over which the optimization must occur. For each value of N and for $M = 1, \ldots, N$, target detection performance must be assessed for a given SNR as follows.

1. Given M and N, find p_0 achieving

$$P_{\text{ft}} = \frac{\bar{B}+\bar{C}}{N_{r\theta}F + \bar{C}} \tag{8}$$

   as a functional inversion of (4).

2. Obtain the decision threshold h producing the desired p_0 as a functional inversion of (3). This requires modeling the clutter-cluster peak normalized intensity, which can be done using the GPD model as described in [4, 5].

3. Evaluate the probability of observing a target cluster as $p_1 = \int_h^\infty f_{\text{t}}(x; S)\,\mathrm{d}x$ for SNR S where $f_{\text{t}}(x; S)$ is the target-cluster peak normalized intensity, which can be approximated through the single-resolution-cell value when SNR is high enough.

4. Assess TIA target detection performance using p_1. Detection performance is quantified by the probability of detection for finite-duration signals or latency for infinite-duration signals.

When optimizing for a *design SNR*, the (M, N) combination maximizing the desired detection-performance measure is chosen. When optimizing to minimize the SNR required to achieve a performance specification, steps 3 and 4 are iterated for the given (M, N) combination to search over SNR until the desired performance is obtained. The (M, N) combination minimizing the SNR required to achieve the performance specification is then chosen.

3. DESIGN EXAMPLES

3.1. TIA optimization

The following parameters have been used to illustrate the design process for a high-duty-cycle active sonar system having a variable coherent processing interval (CPI):

Processing parameters:
CPI and spacing: $T_p = 5$ s or 1.25 s with 50% overlap
Track initialization window: $T_a = 60$ s yielding $N = 23$ or 95
Number of tracker channels: $N_{r\theta} = 1000$
Performance specification:
False-track rate: $\text{FTR}_\text{h} = 10$ per hour
Detection: static $P_d = 0.5$ or 0.9, latency equal to T_a
Signal and clutter modeling:
Signal model: deterministic or Gaussian-fluctuating point target
Clutter-event average length and spacing: $\bar{C} = 120$ s and $\bar{B} = 600$ s
Cluster intensity distribution: GPD parameters as estimated from the data with $h_0 = 12$ dB and coasting probability $q_c(h_0) = 0.6$

The SNR required to achieve each detection performance specification is shown in Fig. 3 for $T_p = 5$ and 1.25 s as a function of M for both signal models. For $T_p = 5$ s (left figure) it is seen that using $M = 10$ for the Gaussian signal required the least SNR to obtain each of the desired operating points (∇'s, solid black for P_d & colored for latency), whereas for the deterministic signal $M = 19$ was best for achieving $P_d = 0.5$ and $M = 18$ for $P_d = 0.9$ or latency equal to T_a. Also shown are the optimal values of M obtained using Shnidman's [11] method (■) and the minimax-optimal value for K-distributed clutter (▲) from [9] using $F = (\bar{B} + \bar{C})/P_\text{ft}$. The SNR loss in using the approximately optimal values of M appears to be small enough (generally < 1 dB, less when N is small) to be useful in practice (e.g., as they are or to initialize a search for the optimal, which may be preferred when N is large). In both design examples, the lower threshold (h_0) caused a limitation in the design at some point as M increased (i.e., the desired p_0 required a threshold $h < h_0$). In both cases, however, the optimal value of M was below the point of limitation. Interestingly, the SNR required to achieve latency equal to T_a (the dots in each figure) was nearly identical to that for achieving a static $P_d = 0.5$.

In comparing the two design cases it can be seen that a lower SNR is required when $T_p = 1.25$ s (~ 1.1 dB lower for the deterministic signal and ~ 3.2 dB lower for the Gaussian-fluctuating signal to achieve $P_d = 0.5$). Countering this is the increase in SNR achieved by the longer CPI, which is only expected to be proportionate (6 dB increase in SNR when going from $T_p = 1.25$ s to 5 s) when there is no coherence loss or energy spreading loss (e.g., see [2]). When coupled with a spreading-loss model [12], the procedure presented here thus allows optimization over N while controlling false-track rate.

3.2. Controlling false-track rate

To illustrate how the tuning process controls the false-track rate, contact-level data were simulated using the GPD with parameters $h_0 = 12$ dB, $\gamma = 0.5$, $\lambda_0 = 5$, and $q_c(h_0) = 0.6$ over $N_{r\theta} = 100$ tracker channels for $N_p = 5000$ pings. The average clutter-event length was set to $\bar{C} = 48$ samples with an average of $\bar{B} = 240$ samples separating them. Three different TIAs were then designed for $N = 5$, 10, and 20 with M chosen according to Shnidman's rule

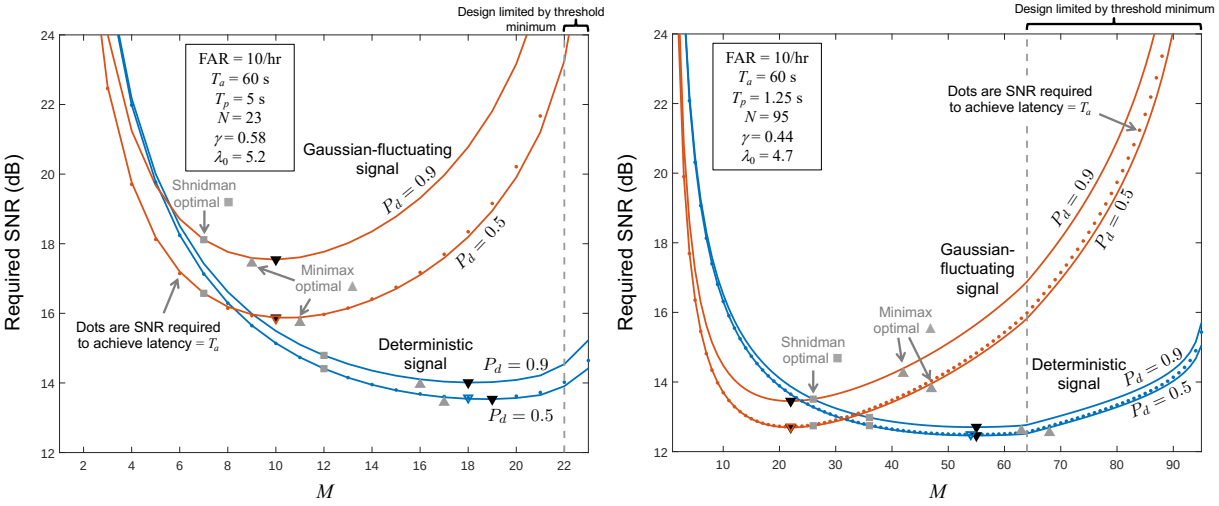

Fig. 3: SNR required to achieve a desired performance specification as a function of M for $T_p = 5\,s$ yielding $N = 23$ (left) and $T_p = 1.25\,s$ yielding $N = 95$/ (right).

for a deterministic signal. In each case, the desired average number of updates between false tracks was $F = 50$. Ten percent of the simulated data are shown in Fig. 2 where the gray circles represent clusters passed to the TIA and the connected red dots represent tracks that are formed. The TIAs with larger values of M (generally corresponding to designs with a larger value of N) have lower decision thresholds [12.8 dB for the $(10, 20)$ TIA vice 14.9 dB for $(3, 5)$] and so produce longer tracks.

The false-track counts over the complete simulation (108, 87, and 94, respectively, for $N = 5$, 10, and 20) yield p-values of 0.40, 0.18, and 0.52 probability of observing a count farther from the expected number (100) based on the design FTR and assuming a Poisson-distributed count. This suggests the design process successfully controls the false-track rate. Although not shown, similar results (p-values of 0.20, 0.18, and 0.74) were seen when the $(6, 10)$ TIA was designed for three different false-track rates ($F = 10$, 50, and 250 pings between false tracks).

4. CONCLUSIONS

An approach for the design of a sliding M-of-N track initialization algorithm (TIA) was presented where M and N are chosen to minimize the SNR required to achieve a detection performance specification and the cluster-level detection decision threshold is chosen to control the false-track rate. A two-state Markov random process was used to represent clutter-event occurrences, which were assumed to be the dominant source of false tracks. This model implies the clutter events have a geometrically distributed length, support for which was obtained by analyzing data from CMRE's LCAS15 experiment. The peak normalized intensity of clutter clusters was modeled using a generalized Pareto distribution (GPD), which was also supported by the LCAS15 data although not shown here. Design examples illustrated how to optimize the TIA and how it successfully controls false-track rate.

5. ACKNOWLEDGEMENTS

This work was sponsored by the Office of Naval Research; Code 321 Undersea Signal Processing.

Parts of this work were made possible by the LCAS Multi-National Joint Research Project (MN-JRP), including as Participants the NATO Centre for Maritime Research and Experimentation, the Defence Science and Technology Organisation (AUS), the Department of National Defence of Canada Defence Research and Development Canada (CAN), the Defence Science and Technology Laboratory (GBR), Centro di Supporto e Sperimentazione Navale-Italian Navy (ITA), the Norwegian Defence Research Establishment (NOR), the Defence Technology Agency (NZL), and the Office of Naval Research (USA).

REFERENCES

[1] A. Farina and F. A. Studer. *Radar Data Processing: Volume 1—Introduction and Tracking*. John Wiley & Sons, 1985.

[2] S. M. Murphy and P. C. Hines. Sub-band processing of continuous active sonar signals in shallow water. In *Proceedings of 2015 MTS/IEEE Oceans Conference*, pages 1–4, Genova, Italy, May 2015.

[3] R. Plate and D. Grimmett. High duty cycle (HDC) sonar processing interval and bandwidth effects for the TREX'13 dataset. In *Proceedings of 2015 MTS/IEEE Oceans Conference*, pages 1–10, Genova, Italy, May 2015.

[4] B. R. La Cour. Statistical characterization of active sonar reverberation using extreme value theory. *IEEE Journal of Oceanic Engineering*, 29(2):310–316, 2004.

[5] J. M. Gelb, R. E. Heath, and G. L. Tipple. Statistics of distinct clutter classes in mid-frequency active sonar. *IEEE Journal of Oceanic Engineering*, 35(2):220–229, April 2010.

[6] C. Forbes, M. Evans, N. Hastings, and B. Peacock. *Statistical Distributions*. Wiley, Hoboken, New Jersey, fourth edition, 2011.

[7] D. Grimmett and R. Plate. High duty cycle sonar performance as a function of processing time-bandwidth for LCAS'15 data. In *Proceedings of the Underwater Acoustics Conference & Exhibition*, pages 185–192, Skiathos, Greece, September 2017.

[8] J. I. Naus. Approximations for distributions of scan statistics. *Journal of the American Statistical Association*, 77(377):177–183, 1982.

[9] D. A. Abraham. Optimization of M-of-N detectors in heavy-tailed noise. In *Proceedings of 2018 IEEE/MTS Oceans Conference*, Charleston, SC, 2018.

[10] J. Glaz, J. Naus, and S. Wallenstein. *Scan Statistics*. Springer-Verlag, New York, 2001.

[11] D. A. Shnidman. Binary integration for Swerling target fluctuations. *IEEE Transactions on Aerospace and Electronic Systems*, 34(3):1043–1053, 1998.

[12] D. E. Weston. Correlation loss in echo ranging. *Journal of the Acoustical Society of America*, 37(1):119–124, January 1965.