Intercepting Beam-Agile Radar

I. Vaughan L. Clarkson
P.O. Box 920
Samford Village, Qld., 4520, AUSTRALIA
v.clarkson@ieee.org

***Abstract*—Modern electronically-scanned-array (ESA) radars can switch between beams in a matter of microseconds. This allows great latitude in constructing the beam schedule. From the point of view of an electronic support (ES) receiver, what are the implications? Repeated search tends to make the radar's beam schedule more regular and predictable. Can this be exploited by the receiver to intercept the radar more quickly? What can be done about strongly aperiodic beam schedules? In this paper, we propose two opposing models for a beam-agile radar. One is a periodic model, accounting for radars with a predictable, repetitive search operation. The other is a stochastic model, in which the beam schedule is modelled by a continuous-time Markov chain, accounting for strongly aperiodic radars. Taking the published beam schedule of a hypothetical, agile radar, we conduct numerical experiments under two different assumptions about the ordering of beam dwells in each radar scan. Under the assumption that the dwells are sequentially ordered, we find that the beam schedule appears periodic from the receiver's point of view, to a good approximation. On the other hand, when the radar dwells are randomly shuffled on each scan, the Markov-chain model proves useful in guiding receiver settings.**

I. INTRODUCTION

Electronically scanned (or steered) array (ESA) antennas offer enormous versatility for radar design compared with the older mechanically scanned antenna (MSA) technology. Freed from mechanical inertia, the radar beam may be switched at will from one direction to any other within microseconds [1]. This allows the radar to commingle searching and tracking and other tasks without regard to the current mechanical state of the antenna.

The question to be considered in this paper is how a beam-agile radar's freedom to choose revisit times affects intercept time in an ES receiver that is searching for it. It tests the author's hypothesis that, wherever a modern agile radar employs revisit versatility in a way that makes the beam schedule less periodic, it results in the radar being, in a sense, "easier" to intercept. Stated a different way, the hypothesis is that synchronisation is the most serious impediment to interception, all other parameters being equal. Where the radar itself chooses to break synchronisation, by employing a beam schedule that is manifestly aperiodic, the receiver is a beneficiary. It benefits, according to the hypothesis, because the receiver's probability of intercept (PoI) becomes less sensitive to the receiver's own RF dwell schedule. There exists neither the possibility to intercept the radar especially quickly, relative to other feasible strategies, nor the possibility that the radar will evade interception indefinitely, through synchronisation. The benefit, then, is that the optimisation of the receiver search strategy can be carried out with little reference to such synchronisation-breaking radars. Optimisation instead focuses on the radars with which the possibility of synchronisation exists.

With so much versatility available, what do modern radars actually do? Is there a useful model for testing the hypothesis?

This research is supported by the Commonwealth of Australia as represented by the Defence Science and Technology Group of the Department of Defence.

First, ESA radars need not employ the versatility available to them to be considered desirable by a prospective operator. Stimson *et al.* [1] cite three main advantages of ESA for airborne radar, of which one is beam agility, the other two being reduced radar cross-section and increased reliability. For some applications, reliability is the chief advantage. Exploitation of beam agility may amount to no more than a faster scan. Data sheets for modern radar tend to emphasise fast scan and high reliability as key selling points; see, *e.g.*, [2].

Second, the reputation that modern ESA radars have for low probability of intercept does not stem from their beam agility but from their frequency and waveform agility. We are concerned here only with the problem of having the receiver be tuned to the right band at the right time to intercept energy from the radar. Whether the receiver can then detect the waveform from the intercepted energy is a separate problem. Yet it cannot detect the radar waveform if it does not first intercept its directed energy.

Detailed information about beam scheduling for multifunction radar is hard to come by. The scarcity of publicly available information can be explained by national security and commercial sensitivities. Perhaps for this reason, as observed by Apfeld *et al.* [3], most of the open literature pertaining to ES receiver scheduling either don't mention modern multifunction radar or explicitly state that it is not considered. They note one exception. Glaude *et al.* [4], acknowledging that ESA radars "can produce more or less random patterns", remark that, nevertheless, "due to their scanning function [they] produce roughly periodic illuminations."

The scanning or search function causes the radar to visit and revisit those parts of the sensing volume in which new targets can be expected to be first seen. Each beam position must be revisited regularly. Arguably the simplest scheme that allows for regular revisits is a periodic schedule. Except perhaps in target-dense environments, it is search that makes up the bulk of the radar's schedule. In very target-dense environments, even those radars that schedule tracking separately may fold the tracking back into the search, reverting to "track-while-scan".

These introductory remarks have been in two directions. First, that ESA radars, in their normal operating modes, are likely to use periodic or nearly periodic beam schedules. Second, that, in modes where they produce distinctly aperiodic schedules, the author postulates that intercept time becomes less sensitive to the particular receiver scan strategy. The implication of the first point is that a receiver search strategy which minimises maximum intercept time against periodic emitters—the "minimax" strategy [5]—ought to be effective. The implication of the second is that it is almost solely the receiver's duty cycle in the radar's band that affects intercept time against strongly aperiodic emitters. This suggests that a simple modification to the established minimax optimisation criterion may allow accommodation of strongly aperiodic modes.

II. Strategies for ES Receiver Sensor Scheduling

A. The Pulse Train Model of Radar Interception

In order to maintain surveillance of radar, a common trade-off employed in ES receiver design is to use a relatively narrow receiver bandwidth whose centre frequency can be retuned at intervals to cover the whole surveillance bandwidth. We call such a receiver a *frequency-swept receiver* (FSR). Faced with a similar surveillance trade-off, radars typically make use of relatively narrow-beamwidth antennas which must be scanned in angle. The receiver can only intercept energy from the radar if it's tuned to the radar's frequency at a time when the radar is directing energy towards the receiver. To maximise the opportunity to intercept radars, the problem of optimising the receiver's *sweep schedule* is considered.

The two conditions that must be simultaneously satisfied for intercept to occur, namely, that the receiver be tuned to the right band and that the radar be pointed in the right direction, can be modelled as two rectangular *pulse trains* or *window functions*, one representing each condition. A pulse occurs when the condition represented by that pulse train is satisfied. *Intercept* occurs when pulses from two pulse trains occur simultaneously, *i.e.*, when they *overlap*, however briefly. A distinction is sometimes made with *detection*, wherein overlap for a prescribed minimum duration must occur.

Suppose the pulse trains are periodic. We define *simply* periodic pulse trains in which, in each period, the pulse train is "on" for a certain time interval and then "off" for the remainder. For the receiver, this is representative of a simple periodic sweep schedule. For the radar, it is typical of a circular or perhaps raster scan. When the pulse trains are simply periodic, only three parameters are needed to represent them: the period T, the pulse width τ and the phase ϕ. The phase represents the time offset of a pulse relative to the time origin.

Given the reduction of radar intercept to a pulse-train coincidence problem, an arithmetic of intercept can be developed to answer relevant questions. If both phases are known, when does intercept occur? The answer can be given either in continuous units of time (seconds, minutes, *etc.*) or in the number of looks required by the radar or receiver. If one or both of the phases are unknown, what is the maximum intercept time? Over a given amount of time (seconds or looks), what is the probability of intercept, assuming a uniform distribution in phase for one or both pulse trains?

B. Minimisation of Maximum Intercept Time

The method developed by the present author [5], [6] allows for variation of receiver sweep period between nominated limits and all individual dwells. Maximum intercept time of two simply periodic pulse trains—where the maximum is taken with respect to the relative phase between them—depends on only two parameters, which we denote α and ϵ. The parameter $\alpha = T_2/T_1$ is the ratio of pulse-train periods. The parameter $\epsilon = (\tau_1 + \tau_2 - 2d)/T_1$ is the sum of the pulse widths less twice the minimum required overlap, normalised with respect to the period of the first pulse train. These parameters can be taken together as coordinates on a plane or, more precisely, its first quadrant. There is a simple, recursive procedure for partitioning this (α, ϵ) plane into triangles, wherein each triangle represents a region in which the intercept time, expressed in looks, remains constant.

The following procedure, adapted from [7], produces a schedule.

1) Set T_2 to the minimum allowable sweep period.
2) For each emitter in the threat-emitter list:
   a) Set T_1 to the emitter's scan period. Calculate α.
   b) Set ϵ to the minimum value possible. That is, set $\tau_2 = d$ in calculating ϵ.
   c) The intercept time with that emitter can now be calculated, for instance, by examination of the partitioning of the (α, ϵ) plane.
3) Unless the sum of the dwell times on each band would exceed the current sweep period, repeatedly:
   a) Identify the emitter with the longest intercept time.
   b) Increase the dwell on that band until the intercept time decreases.
4) Identify the longest remaining intercept time from Step 3 and convert it from discrete time (number of looks) to continuous time by multiplication with the sweep period.
5) If this intercept time is smaller than the smallest yet found at other sweep periods then note the current settings.
6) Increase the sweep period by a small amount.
7) If the sweep period is less than the maximum allowable, return to Step 2. Otherwise the noted settings represent the minimax sensor schedule.

The procedure minimises the maximum intercept time, where the maximum is taken with respect to all emitters in the threat-emitter list and with respect to all possible relative phases. The minimum is found with respect to all possible dwell times and all sweep periods. Because this approach to sensor scheduling is alert to synchronisation, the benefits, when compared to naïve periodic scheduling, can be dramatic. Whereas a naïve schedule may be synchronised with several emitters on a threat-emitter list, leading to a theoretically infinite intercept time, the minimax approach can ensure finite intercept time with all of them. By this approach, within the constraints of the model, a guarantee is offered, not merely a probability, that each emitter on the list, if present, will be intercepted within a certain time.

C. A Randomised Schedule with Near-Linear Intercept Time

A periodic receiver schedule cannot eliminate synchronisation. The minimax method suppresses synchronisation for listed threat emitters but synchronisation can still exist with unlisted emitters. Kelly *et al.* [8] were motivated by the ill effects of synchronisation in proposing jittered scan schedules. However, only simulation studies, and not a theoretical analysis, were ever published.

The present author, with collaborators El-Mahassni and Howard [9], continued the work on stochastic schedules. Like Kelly *et al.*, the objective is to suppress synchronisation universally. The ideal schedule would exhibit an expected intercept time that is 'flat' with respect to the scan period of the emitter.

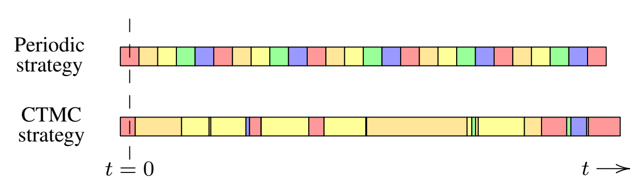

Fig. 1. A periodic schedule compared with a realisation of a CTMC schedule.

The proposed means for controlling the receiver dwells is a continuous-time Markov chain (CTMC). The receiver jumps from band to band (pseudo-)randomly. Figure 1 illustrates a realisation of a CTMC schedule. The authors derive the expected intercept time as a function of an emitter's scan period. The notion of a *maximum expected intercept time*, MEIT, is introduced, where the maximum is taken with respect to the emitter's illumination time. The maximum is attained when the emitter illumination pulse train has a pulse width of 0. It is found that the MEIT is asymptotically linear in scan period.

To optimise the parameters of the CTMC, the authors propose that the CTMC should approach the asymptote as quickly as possible. The MEIT, normalised by scan period, is an infinite sum of damped exponentials. Formally, the criterion for optimisation is to minimise the maximum of the real parts of the exponents, subject to a constraint on the CTMC's *mobility*, *i.e.*, a measure of its 'jumpiness'. This leads to a closed-form solution for the CTMC's defining rate matrix, $\mathbf{Q}$, which is found to have the form $\mathbf{Q} = n(\mathbf{1}\boldsymbol{\pi} - \mathbf{I})/[(n-1)H]$ where n is the number of states/bands, $\mathbf{I}$ represents the identity matrix, $\mathbf{1}$ is a column vector of all ones, H is the mobility and $\boldsymbol{\pi}$ is the stationary distribution of CTMC, *i.e.*, π_i is the probability that, at any given time, the state of the CTMC—the band to which the receiver is tuned—is i.The corresponding MEIT, $f_i(T)$, is found to be

$$f_i(T) = \frac{(1-\pi_i)T}{\pi_i[1-\exp(-nH^{-1}T/(n-1))]}. \tag{1}$$

A useful property of the optimal CTMC is that it is *lumpable* into two states: 'in band' and 'out of band'. That is, if we are interested only in the behaviour of the CTMC in a particular band—to reduce a realisation of the CTMC to a receiver pulse train, for instance—then the n-state CTMC can be replaced by this 2-state CTMC.

III. Further Analysis of the CTMC-Induced Schedule

Since the CTMC strategy is important in subsequent discussion, we extend the analysis of [9]. We begin by reparametrising. Instead of being specified by its mobility H and stationary distribution $\boldsymbol{\pi}$, we prefer the vector of mean sojourn times, $\boldsymbol{\mu}$. The sojourn time of a CTMC state translates to the dwell time on the corresponding band in the context of an ES receiver. In lumping the CTMC into a 2-state model around State i, we have

$$\mu_{\text{in}} = \mu_i = \frac{(n-1)H}{n(1-\pi_i)} \qquad \text{and} \qquad \mu_{\text{out}} = \frac{(1-\pi_i)\mu_i}{\pi_i}. \tag{2}$$

The rate matrix reduces to

$$\mathbf{Q}_{\text{eff}} = \begin{pmatrix} -\mu_{\text{in}}^{-1} & \mu_{\text{in}}^{-1} \\ \mu_{\text{out}}^{-1} & -\mu_{\text{out}}^{-1} \end{pmatrix}$$

where the first state is the 'in' state and the second is 'out'. If the state of the CTMC is sampled every T seconds, we induce a discrete-time Markov chain (DTMC) with transition matrix

$$\mathbf{P}(T) = e^{\mathbf{Q}_{\text{eff}}T} = \frac{1}{\mu_{\text{in}}+\mu_{\text{out}}} \begin{pmatrix} \mu_{\text{in}} + \mu_{\text{out}}e^{-\gamma T} & \mu_{\text{out}}[1-e^{-\gamma T}] \\ \mu_{\text{in}}[1-e^{-\gamma T}] & \mu_{\text{out}} + \mu_{\text{in}}e^{-\gamma T} \end{pmatrix} \tag{3}$$

where

$$\gamma = \frac{1}{\mu_{\text{in}}} + \frac{1}{\mu_{\text{out}}} = \frac{n}{(n-1)H} \tag{4}$$

from (2). The DTMC models whether the receiver is in or out of band on each successive (infinitesimal) illumination at period T.

A. Intercept Time At a Specified PoI

Given that the CTMC is initialised randomly with its stationary distribution, the induced DTMC is initialised with probability of being in-band of $\pi_{\text{in}} = \pi_i = \mu_{\text{in}}/(\mu_{\text{in}}+\mu_{\text{out}})$. Therefore the probability of intercept on the first illumination is π_i. Given that it was out of band on the first illumination, the probability of illumination on the second illumination is found from the lower left element of $\mathbf{P}(T)$, namely $\mu_{\text{in}}[1-e^{-\gamma T}]/(\mu_{\text{in}}+\mu_{\text{out}})$. The probability of (first) intercept on the kth illumination—or kth look—of the radar is in general

$$\Pr(K=k) = \begin{cases} \pi_i & k=1, \\ (1-\pi_i)(1-\rho)^{k-2}\rho & k>1 \end{cases} \tag{5}$$

where

$$\rho = \frac{\mu_{\text{in}}(1-e^{-\gamma T})}{\mu_{\text{in}}+\mu_{\text{out}}} = \pi_i(1-e^{-\gamma T}). \tag{6}$$

From (5), we can infer that

$$\Pr(K>k) = \begin{cases} 1-\pi_i & k=1, \\ (1-\pi_i)(1-\rho)^{k-1} & k>1. \end{cases} \tag{7}$$

To achieve a PoI p, we must find k such that $p \leqslant \Pr(K \leqslant k) = 1 - \Pr(K>k)$. If $p \leqslant 1-\pi_i$ then $k=1$, otherwise $1-p \geqslant (1-\pi_i)(1-\rho)^{k-1}$ which means that $k-1 \geqslant [\log(1-p) - \log(1-\pi_i)]/\log(1-\rho)$. Combining the cases $p \lesseqgtr 1-\pi_i$, we get

$$k = 1 + \max\left\{0, \left\lceil \frac{\log(1-p) - \log(1-\pi_i)}{\log(1-\rho)} \right\rceil \right\}. \tag{8}$$

B. Non-Zero Pulse Width

The foregoing analysis assumes that the illumination has zero pulse width. We now relax this assumption, allowing $\tau > 0$. The details of the derivation are omitted; see [10] for details.

To obtain a general expression for $\Pr(K>k)$, we redefine ρ from (6) and π_i as functions of τ so that

$$\rho(\tau) = 1 - \frac{\mu_{\text{out}} + \mu_{\text{in}}e^{-\gamma(T-\tau)}}{\mu_{\text{in}}+\mu_{\text{out}}} e^{-\tau/\mu_{\text{out}}}, \tag{9}$$

$$\pi_i(\tau) = 1 - (1-\pi_i)e^{-\tau/\mu_{\text{out}}}. \tag{10}$$

We observe that $\rho(0) = \rho$ and $\pi_i(0) = \pi_i$. We find that

$$\Pr(K>k) = \begin{cases} 1-\pi_i(\tau) & k=1, \\ [1-\pi_i(\tau)][1-\rho(\tau)]^{k-1} & k>1, \end{cases} \tag{11}$$

which reduces to (7) when $\tau = 0$.

The analogue of (8) results, which is to say that the number of radar looks required to attain a PoI of at least p with the receiver is

$$k = 1 + \max\left\{0, \left\lceil \frac{\log(1-p) - \log[1-\pi_i(\tau)]}{\log[1-\rho(\tau)]} \right\rceil \right\}. \tag{12}$$

In the case of interest in Section V, $T - \tau \gg \gamma^{-1}$. Then $\rho(\tau) \approx \pi_i(\tau)$ and (12) becomes

$$k \approx \left\lceil \frac{\log(1-p)}{\log[1-\pi_i(\tau)]} \right\rceil. \tag{13}$$

C. Non-Zero Minimum Overlap

To calculate the probability of detection, as distinct from intercept, we need to take account of the requirement of a minimum overlap of duration $d > 0$. The analysis is rather lengthy, involving the solution of a delay differential equation, and so it has been omitted; see [10] for details. The result is that we may use (13) to approximate time to detection using the further approximation

$$\pi_i(\tau) \approx 1 - \frac{\mu_{\text{out}}+d}{\mu_{\text{in}}+\mu_{\text{out}}} e^{-\beta(\tau-d)} \qquad \text{where} \qquad \beta = \frac{\mu_{\text{in}}-d}{\mu_{\text{in}}\mu_{\text{out}}}. \tag{14}$$

IV. A Beam-Agile Radar Strawman

Let us now consider a specific proposal for the beam schedule of a modern, agile radar. In the work of Byrne *et al.* [11], an optimisation technique is developed to calculate an ESA radar beam schedule with the aim of unifying the treatment of search and track functions. Dwells are allocated with reference to a cost function. The cost function has two terms: the cost of not detecting targets and the cost of error in the track estimates. The function is minimised with respect to the dwell times on each beam for the upcoming scheduling period or "epoch".

A rather detailed simulation is presented in [11] in which a hypothetical radar, equipped with such a beam scheduling algorithm, reacts to a target-rich environment. A record of thirty scheduling epochs is analysed. Significant reduction in the cost function is achieved with respect to a uniform schedule and that demonstrably fewer undetected targets remain in known target-entry regions.

The Byrne *et al.* schedule—or *BWW* schedule for the authors, Byrne, White & Williams—is chosen for three practical reasons:

1) in unifying search, detection and track, it is one of the more advanced beam schedules to appear in the open literature;
2) a detailed account exists of the hypothetical ESA radar, the simulated, target-rich environment and the beam schedule;
3) Marion Byrne, the first-named author of [11], kindly agreed to provide the data in electronic format and indeed even to extend the simulation from 30 epochs to 300 [12].

The hypothetical radar has a 360° azimuth coverage and a $200\,\mathrm{km}$ range. It has 30 equally spaced azimuth beams, numbered 1–30, each with 36° beam width. Each beam partially overlaps with its two neighbours. The epoch is fixed at $1\,\mathrm{s}$. For the required range to be unambiguous, a pulse repetition interval greater than $1.33\,\mathrm{ms}$ is required. We will assume the minimum allowable value.

The simulated environment assumes a certain background arrival rate of targets throughout the surveillance volume, with elevated arrival rates in designated sub-volumes. Over the first 30 epochs ($30\,\mathrm{s}$), five targets arrive and are, soon thereafter, detected and tracked.

In Figure 2, the schedule for the extended simulation scenario is presented. It incorporates the first 30 epochs as presented in [11]. There is a strong tendency, when a target is detected, to devote a greater-than-average dwell time to tracking it in each subsequent epoch.

V. Numerical Experiments

In this Section, the beam schedule of Byrne *et al.* [11] is used for numerical experiments. The aim is to determine the extent to which such a schedule fits either the periodic or Markov-chain models.

Of the beam schedules, Byrne says, "There is no hierarchy of actions within the one-second scan." Specifically, "the beams may be scheduled to use their allocated time in any order" [12]. We will consider two scenarios for ordering the dwells within each epoch.

In the first scenario, we will assume that the beams are scheduled in ascending order of the beam numbers so that the beam, in effect, rotates. The hypothesis to be tested here is that the beam schedule behaves like a periodic emitter with $1\,\mathrm{s}$ period. In that case, the minimax strategy could be used effectively to ensure intercept.

In the second scenario, we will assume that the ordering is randomly shuffled from one epoch to the next. The hypothesis is that the radar behaves, in terms of its intercept properties, as if it were controlled by a Markov-chain strategy. If true, intercept time with the receiver is much less sensitive to the receiver's sweep period.

A. Sequential Beam Schedule

Suppose the beams of the BWW radar are illuminated in a fixed sequence. We examine the extent to which the beam schedule can be modelled as a periodic schedule. Clearly, the beam schedule is not precisely periodic. From Figure 2, we can see that dwell on individual beams can vary markedly. On the other hand, the epoch is unvarying at $1\,\mathrm{s}$. We test the hypothesis that, when the beams are illuminated sequentially, the BWW radar behaves like a circularly scanned radar with a scan period of $1\,\mathrm{s}$ and a beam width of 36°.

In Figure 3, the results of a numerical experiment are plotted. In this experiment, a periodic receiver with similar parameters to the radar attempts to intercept. The receiver monitors ten bands with equal dwell on each band. The sweep period of the receiver is varied from $0.5\,\mathrm{s}$ to $1.5\,\mathrm{s}$. This is the independent variable on the horizontal axis.

The intercept time for the equivalent circular-scanned radar is calculated according to the theory in Section II-B. For the sequentially-scheduled BWW radar, the intercept time is approximated by the time to achieve a 99% PoI. The time is derived from the cumulative distribution function of intercept time, calculated by a method similar to that described by Winsor and Hughes [13]; see [10]. The intercept time is averaged over Azimuth Bins $3, 6, 9, \ldots, 30$.

Agreement is remarkably close. The same synchronisation ratios between radar scan and receiver sweep periods, *i.e.*, $1/2$, $3/5$, $2/3$, $3/4$, $4/5$, 1, $5/4$, $4/3$ and $3/2$, cause intercept-time "spikes" in both. The results support the hypothesis that intercept time with the sequentially-scheduled BWW radar is closely approximated by an equivalent circular-scanned radar, despite the BWW radar's dwell-time variations in response to target stimuli.

TABLE I. Threat-emitter list example.

Emitter number	Band	Scan period (µs)	PRI (µs)	Beamwidth ($^\circ$)
1	A	8.4×10^6	2.38633×10^3	1.3
2	B	2.97×10^6	1.37792×10^3	2.6
3	C	10.5×10^6	9.38	2.1
4	D	1×10^6	1.33333×10^3	36

As a second test of the accuracy with which the intercept time with the sequentially-scheduled BWW radar is modelled by a circular-scanned radar, let us now examine the effectiveness of the minimax strategy in choosing a receiver dwell schedule that takes account of it. In Table I, we set out an example threat-emitter list. The first three entries on the list constitute an example that the present author used to illustrate the effectiveness of the minimax strategy in [5]. With just those three entries, optimised over sweep periods between $0.5\,\mathrm{s}$ and $1.5\,\mathrm{s}$, allowing a minimum of 5 consecutive radar pulses for detection, a best sweep period of $0.6769\,\mathrm{s}$ is found for the receiver. A maximum intercept time of $25.04\,\mathrm{s}$ is indicated with Emitter 1.

With the fourth entry added to the threat-emitter list, representing the BWW radar, the minimax schedule barely changes. The best sweep period becomes $0.6770\,\mathrm{s}$. The maximum intercept time of $25.05\,\mathrm{s}$ is hardly altered for Emitter 1. The intercept time for

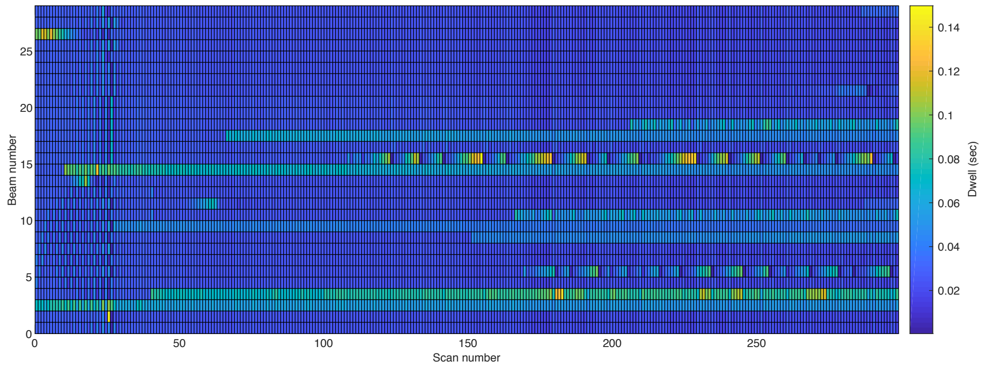

Fig. 2. Sample beam schedule for 300 epochs [12].

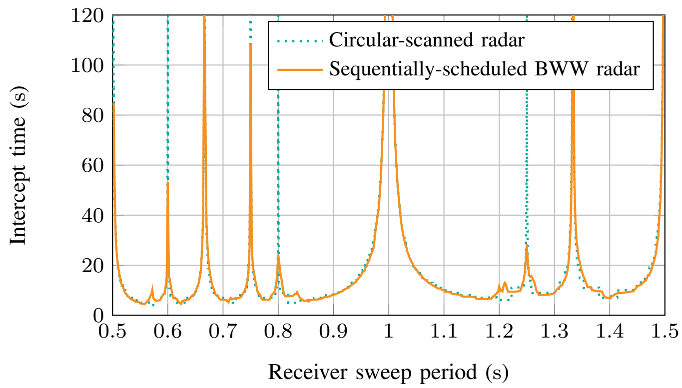

Fig. 3. Intercept time for a periodic receiver with the sequentially-scheduled BWW radar and a circular-scanned radar.

Emitter 4, the BWW radar, is $18.96\,\mathrm{s}$. Just $10.77\,\mathrm{ms}$ of receiver dwell is needed on the BWW radar's band (Band D) in each sweep.

How well does this represent intercept times with the true BWW beam schedule? When the time required for 99% PoI is averaged over Azimuth Bins $3, 6, \dots, 30$, the value is found to be $21.23\,\mathrm{s}$, just over two seconds longer than indicated by the minimax solution.

B. Shuffled Beam Schedule

Suppose we randomly shuffle the order in which each beam is scheduled within an epoch, while retaining the indicated dwell time. Our hypothesis is that the shuffle-scheduled BWW radar behaves like a Markov-chain-controlled radar from the viewpoint of intercept time.

In substituting a Markov-chain-controlled radar, we must equip it with the right parameters. Each beam of the BWW radar partially overlaps its two neighbours. With the receiver situated in Azimuth Bin b, it is illuminated by Beams $b-1$, b and $b+1$, modulo 30, which is the number of defined azimuth bins and beams. The receiver is illuminated on as many as three separate occasions in any $1\,\mathrm{s}$ epoch. Therefore, our substitute Markov-chain-controlled radar will have thirty states with a mean dwell time in each state of $1/30$th of a second. Three of these thirty states illuminate the receiver. We can lump this CTMC into a smaller CTMC having only ten states, each again with a mean dwell time of $1/30$th of a second, in which only one of the ten states illuminates the receiver. This CTMC can be further lumped, as described in Section II-C, into two states, illuminating and non-illuminating or "in" and "out". Note that the role of radar and receiver is reversed from Section II-C, with the CTMC now controlling the radar. We have $\mu_{\text{in}} = 33.33\,\mathrm{ms}$ and $\mu_{\text{out}} = 0.3\,\mathrm{s}$.

As a foil to the radars, we propose an FSR with the same parameters as in Section V-A. The receiver sweeps ten bands, with equal dwell in each band and a sweep period between $0.5\,\mathrm{s}$ and $1.5\,\mathrm{s}$.

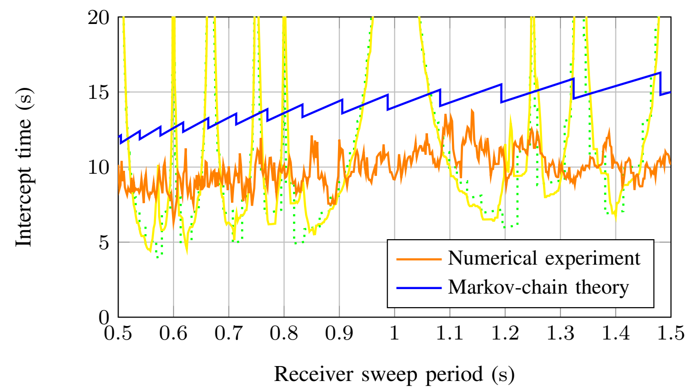

Fig. 4. Intercept time for a periodic receiver with the shuffle-scheduled BWW radar and a CTMC-controlled radar. (Light yellow and green plots are repeated from Figure 3.)

In Figure 4, we present the results of the numerical experiment. In light yellow and green, to provide some reference for the new results, we re-plot the results from Figure 3. The results of the numerical experiment are in orange. The receiver with a periodic schedule attempts to detect the BWW radar with a shuffled beam schedule. The detection time is determined from 99% probability of detection and averaged across 10 azimuth bins, as in Section V-A. In blue is plotted the time to detection, again at 99% probability of detection, calculated from (13) and (14), under the assumption that the radar is controlled by a CTMC.

We observe that the intercept time is much less variable as a function receiver sweep period for both the shuffle-scheduled BWW radar and the CTMC-controlled radar than it was in Section V-A for the sequentially-scheduled BWW radar. There is no evidence of synchronisation causing large excursions in the intercept time. We also observe that the CTMC-controlled radar is a conservative approximation to the shuffle-scheduled BWW radar in terms of intercept time. They exhibit the same general behaviour as a function of receiver sweep period. Both slowly increase with the sweep period.

Although the agreement between shuffled-scheduled BWW radar and the CTMC-controlled radar is not as close as we found between the sequential-scheduled BWW radar and the circular-scanned radar in Section V-A, the agreement is still adequate. We will use the CTMC-controlled radar as a conservative "doppelganger" for the shuffle-scheduled BWW radar in computing intercept and detection times for receiver schedule optimisation.

TABLE II. THREAT-EMITTER LIST ADDENDUM.

Emitter number	Band	Mean revisit (μs)	PRI (μs)	Mean illumination (μs)
5	E	333.333×10^3	1.33333×10^3	33.3333×10^3

Let us now perform another test in which we attempt to minimise the maximum intercept time. We retain the four emitters in the threat-emitter list of Table I but add a fifth emitter as listed in Table II. This fifth emitter represents the shuffle-scheduled BWW radar. In effect, we now have the BWW radar appearing twice on the threat-emitter list. One instance of the BWW radar has a sequential beam schedule, the other, in a different band, has a shuffled beam schedule.

The tabulated parameters in Table II are slightly different from those in Table I. Instead of scan period, we have mean revisit time and, instead of beamwidth, we give mean illumination time. Mean illumination time corresponds to the μ_{in} parameter, the expected sojourn time of the lumped CTMC in the illuminating state. Mean revisit time is the expected elapsed time between consecutive visits to the illuminating state, and so is equal to $\mu_{\text{in}} + \mu_{\text{out}}$.

The minimax optimisation technique as described in [5] is not designed to compute a schedule when the threat-emitter list contains such an emitter. Strictly speaking, the maximum intercept time with a stochastic emitter must be infinite. Instead, we augment the minimax process to minimise the detection time with aperiodic emitters of this type using the approximation (13).

With the procedure adapted in this way, we optimise (in the newly approximate sense) the receiver dwell schedule over the five bands. The sweep period is optimised, as before, over the range from $0.5\,\text{s}$ to $1.5\,\text{s}$, with five consecutive emitter pulses required for successful detection. The approximately optimal settings are found to be not much different from those of Section V-A. The best sweep period is now $677.3\,\text{s}$, just $0.3\,\text{s}$ longer than the previous case. The maximum intercept time over all five emitters is $25.06\,\text{s}$, just $0.01\,\text{s}$ longer. The dwell allocated to the sequentially scheduled BWW radar, Band D, is reduced even further to $7.74\,\text{ms}$. The dwell allocated to the shuffle-scheduled BWW radar, Band E, is a little longer, at $23.15\,\text{ms}$, but this still represents only $3.42\,\%$ of the sweep period. Intercept time with this radar is equal longest, with Emitter 1, *i.e.*, $25.06\,\text{s}$.

Given that the BWW radar with a shuffle-scheduled beam schedule is not, in truth, a CTMC-controlled radar, how closely do minimax's predicted intercept times match the observed ones? For a single realisation, the 99% detection time, averaged over the ten azimuth bins, is found to be $21.39\,\text{s}$. Again, and as expected, the minimax calculation is tolerably accurate but conservative.

VI. CONCLUSIONS

Modern ESA radar, freed from mechanical constraints, has great flexibility in constructing its beam schedule. Yet the demands of the radar's search function tend to ensure that the beam schedule exhibits regularity from one epoch to the next. At one extreme, the radar could be modelled as a traditional periodic-scan radar. At the other, it could be modelled as a highly unpredictable Markov chain of beam dwells.

To test how well these models match a modern radar, we conducted numerical experiments with the beam schedule of the BWW radar [11], [12]. The published BWW schedules do not specify the ordering of the beam dwells within a scan, so we experimented with a regular order and a randomly shuffled order. We found that the regularly ordered beam dwells can exhibit synchronisation with an ES receiver. In contrast, intercept time is relatively insensitive to the sweep period of the receiver when the schedule is randomised.

We revisited the minimax technique [5] for calculation of optimal periodic sweep schedules. We described a simple extension to the algorithm that allows periodic and aperiodic radar scans to be accounted for in deriving the schedule. We demonstrated the algorithm with numerical examples.

VII. ACKNOWLEDGEMENTS

The author would like to acknowledge the help and advice of collaborators at the Defence Science & Technology Group, namely Claire Rivett, Marion Byrne, Jason Williams and Kruger White.

REFERENCES

[1] G. W. Stimson, H. Griffiths, C. Baker, and D. Adamy, *Stimson's Introduction to Airborne Radar*, 3rd ed. Edison, NJ: SciTech Publishing, 2014.

[2] Israel Aerospace Industries ELTA Systems, "Very short range air defence radar (VSHORAD): ELM-2026B," online product brochure, retrieved 1st May, 2017. [Online]. Available: http://www.iai.co.il/Sip_Storage//FILES/6/41416.pdf

[3] S. Apfeld, A. Charlish, and W. Koch, "An adaptive receiver search strategy for Electronic Support," in *Proc. Sensor Signal Process. Defence*, Edinburgh, Sep. 2016, pp. 1–5.

[4] H. Glaude, C. Enderli, J.-F. Grandin, and O. Pietquin, "Learning of scanning strategies for Electronic Support using predictive state representations," in *Proc. IEEE Internat. Workshop Machine Learning Signal Process.*, Sep. 2015.

[5] I. V. L. Clarkson, "Optimisation of periodic search strategies for electronic support," *IEEE Trans. Aerospace Electron. Syst.*, vol. 47, no. 3, pp. 1770–1784, Jul. 2011.

[6] ——, "Sensor scheduling for electronic support: Recent developments and future directions," Defence Science & Technology Organisation, Tech. Rep., May 2017, final Report for Research Agreement 2014/MyIP 5770, in review.

[7] E. Köksal, "Periodic search strategies for electronic countermeasure receivers with desired probability of intercept for each frequency band," Master's thesis, Middle East Technical University, Jan. 2010. [Online]. Available: http://etd.lib.metu.edu.tr/upload/12611398/index.pdf

[8] S. W. Kelly, G. P. Noone, and J. E. Perkins, "The effects of synchronisation on the probability of pulse train interception," *IEEE Trans. Aerospace Electron. Syst.*, vol. 32, no. 1, pp. 213–220, Jan. 1996.

[9] I. V. L. Clarkson, E. D. El-Mahassni, and S. D. Howard, "Sensor scheduling in electronic support using Markov chains," *IEE Proc. Radar Sonar Navigation*, vol. 153, no. 4, pp. 325–332, Aug. 2006.

[10] I. V. L. Clarkson, "Sensor scheduling for electronic support to intercept modern versatile radar," Defence Science & Technology Organisation, Tech. Rep., Jul. 2017, final Report for Contract ATM 0553, in review.

[11] M. Byrne, K. White, and J. Williams, "Scheduling multifunction radar for search and tracking," in *Proc. Internat. Conf. Inform. Fusion*, Jul. 2015, pp. 945–952.

[12] M. Byrne, Mar. 2017, personal communication.

[13] C. Winsor and E. J. Hughes, "Optimisation and evaluation of receiver search strategies for electronic support," *IET Radar Sonar Navigation*, vol. 6, no. 4, pp. 233–240, Apr. 2012.