

AD-A048 190

12

6

INFORMATION PROCESSING TO MAINTAIN LOCALIZATION IN ASW SURVEILLANCE.

Report to

Naval Analysis Programs Office of Naval Research

11

29 July 1977

12 162p.

15

Under Contract No. N00014-76-C-0676

by

DANIEL H. WAGNER, ASSOCIATES Paoli, Pennsylvania

DDC REPRODUCIBLE DEC 28 1977 B

9

Final rept.

10

By:

William H. Barker Associate

Barry Belkin Barry/Belkin Vice President

Approved:

Lawrence D. Stone Lawrence D. Stone Vice President, Reviewer Daniel H. Wagner Daniel H. Wagner President

DISTRIBUTION STATEMENT A Approved for public release Distribution Unlimited

1473 367 675

LB

ABSTRACT

Problems of optimal surveillance and contact information processing are investigated with ASW applications in view. The optimal allocation of ASW surveillance resources is investigated, by way of Bayesian statistical analysis, in a somewhat idealized surveillance situation involving a moving target and the possibility of false contacts. In the problem considered, the optimal allocation of surveillance effort is shown to be the solution of a certain dynamic programming problem. The optimal allocation is determined numerically in a number of special cases and compared to several simple allocation policies. The properties of these various policies are investigated through analysis as well as through a number of numerical examples. Particular attention is paid to asymptotic behavior of long term surveillance policies. One suboptimal policy, the maximum-information-gain policy, is shown to have a number of very desirable properties.

Procedures for processing ASW contact information are developed with two distinct applications in mind. The first application is directed toward estimating the track of a specified target as well as toward drawing inferences about overall target behavior patterns from contact data on a number of different targets. The approach is to combine, using Bayesian methods, contact data with a scenario-based parametric model for target motion. Statistical estimation procedures are given for estimating the track of a specified target from contact data. Additionally, methods are given for estimating the parameters of the motion model from contact data.

The second application is processing ASW contact information in a multi-target environment where there is ambiguity in assigning contacts to targets. Several Bayesian statistical methods for the systematic generation and updating of target location predictions in a multi-target environment are developed. Of the ones considered, Extended Memory processing is shown to be the only computationally practical method that makes accurate use of the observational data.

ACCESSION for		
NTIS	White Section	<input checked="" type="checkbox"/>
DDC	Buff Section	<input type="checkbox"/>
UNANNOUNCED		<input type="checkbox"/>
JUSTIFICATION		
BY		
DISTRIBUTION/AVAILABILITY CODES		
Dist.	AvAIL	and/or SPECIAL
A		

TABLE OF CONTENTS

ABSTRACT		
PREFACE		iii
SUMMARY		v
CHAPTER I	INTRODUCTION	1
CHAPTER II	OPTIMAL SURVEILLANCE AGAINST A MOVING TARGET IN A FALSE CONTACT ENVIRONMENT	5
	The Tactical Situation	6
	Sensor Response Assumptions	7
	Target Motion Assumptions	9
	Mathematical Structure of K-Stage Surveillance Operations	9
	Evaluating Surveillance Plans	13
	Optimal Surveillance Plans	16
	Evaluating Optimal Surveillance Plans Numerically	19
	Stationary Surveillance Plans	24
	Examples Comparing Optimal and Stationary Surveillance Plans	25
	Asymptotic Behavior of the K-Stage Optimal Surveillance Plan	44
	An Important Special Case: The Homogeneous Sensor	57
	The K-Stage Optimal Surveillance Plan for a Homogeneous Sensor	60
	Examples Involving a Homogeneous Surveillance Sensor	62
	Conclusion	65
CHAPTER III	A STATISTICAL MODEL FOR PROCESSING ASW CONTACT INFORMATION TO ESTIMATE TARGET PATTERNS OF OPERATION	69
	A Parametric Model for Target Motion	70
	The Parametric Target Motion Model in the Face of Uncertainty	72
	Estimating Target Motion Parameters When the Track and Motion Scenario are Known	76

TABLE OF CONTENTS (Continued)

	Estimating Target Motion Parameters When the Track is Known but the Motion Scenario is Unknown	80
	Target Contacts	81
	Estimating the Track of a Target from Contact Data	82
	Estimating Target Motion Parameters from Contact Data When the Motion Scenario is Known	86
	Estimating Target Motion Parameters from Contact Data	92
CHAPTER IV	ASW INFORMATION PROCESSING IN A MULTI-TARGET ENVIRONMENT	95
	Background and Introduction	95
	Structure of ASWIPS (ASW Information Processing System)	96
	Discussion of Information Processing Algorithms	97
	Comparison of Processing Algorithms Using ASWIPS	103
REFERENCES		129
APPENDIX A	AUXILIARY FORMULAS FOR EVALUATING GAUSSIAN INTEGRALS	A-1
APPENDIX B	EXTENDED MEMORY PROCESSING METHODOLOGY	B-1

P R E F A C E

This is a report by Daniel H. Wagner, Associates to the Naval Analysis Division of the Office of Naval Research (Code 431) on a research investigation performed under ONR Contract No. N00014-76-C-0676. This report is directed towards developing methodologies for processing ASW surveillance information with the objective of obtaining estimates of target location. It is intended primarily for use by analysts. The methods have been motivated by actual ASW information processing requirements. The goal has been to achieve useful and computationally practical methods which are suitable for real-time assistance to ASW surveillance operations.

We would like to express our appreciation for the excellent cooperation and support that has been given to this work by Mr. J. Randolph Simpson and CDR Ronald James of the Naval Analysis Division of the Office of Naval Research.

Additionally, we would like to recognize the capable efforts of our colleagues B. D. Wenocur who programmed the ASWIPS model discussed in Chapter IV, and R. V. Kohn and S. S. Brown who contributed a section to Appendix A.

SUMMARY

↘ This report is directed towards the problem of providing estimates for a submarine target's location under various conditions. The objective is to develop useful and computationally practical tools for obtaining such estimates. The principal application is the continuous long-term localization of a target or targets through the utilization of surveillance efforts. Much of this current work has been motivated by the methods, results, and conjectures contained in reference [a]. ↙

Chapter I provides a brief introduction, and appendices A and B provide supplementary material. The issue of optimal surveillance in a false contact environment is explored in Chapter II. Processing ASW contact information so as to obtain estimates about the track of a specific target as well as overall long-term behavior patterns is the subject of Chapter III. In Chapter IV, a computationally practical method for processing ASW contact information in a multi-target environment is given. Computational results are included in Chapters II and IV. Unfortunately the work presented in Chapter III has not as yet progressed to that stage. ←

Optimal Surveillance

Chapter II addresses the problem of optimal surveillance against a moving target in a false contact environment. Although the surveillance situation considered is somewhat idealized, our results provide practically useful guidelines for considering adaptive surveillance operations. The goal of this research is to place a practical and useful theory of surveillance on a sound theoretical basis.

The problem of specifying the location of a target as it arises in surveillance problems is substantially different from the problem usually considered in search theory, i. e., the problem of detecting the target with maximum probability. The difference is a consequence of the fact that search theory generally assumes that a detection will also provide the desired localization of the target. Thus, the detection itself is the issue of importance. In surveillance, the issue is localization and its maintenance over time and this goal may be achievable with or without detections or contacts. In addition, surveillance also takes into account the fact that a problem may not end with a contact because of poor localization information.

The surveillance operation considered here involves a moving target located in one of N cells, C_1, C_2, \dots, C_N . The precise cell containing the target is unknown, but at the current time a probability distribution for the target's location has been established. Thus for $l = 1, 2, \dots, N$, let x_l be the probability that the target is currently located in cell C_l . Suppose that our best estimate for the target's location is that cell which contains the target with highest probability.* Indeed if $x_k > x_l$, all $l \neq k$, then the target is most likely to be in cell C_k , and the probability that this estimate is correct is x_k . We seek to apply our surveillance effort so as to maximize, at the end of the surveillance operations, the probability that the target is in the high probability cell.

The surveillance operation itself is performed in a sequence of discrete time stages, each of duration Δ time units, by a single surveillance sensor. We are permitted only K stages, and we must use the sensor so as to obtain the best possible estimate for the target's position after these K stages. At the start of the surveillance operation, the fixed time at which we must obtain the best estimate of the target's position, the horizon, is $K\Delta$ units into the future. Our measure of effectiveness is the probability that the target is in the high probability cell at the time of the horizon $K\Delta$.

At the beginning of each stage, a cell is chosen and is then investigated for the amount of time Δ . The choice of cell, in general, depends upon the number of stages remaining in the operation (i.e., the amount of time remaining until the horizon), the surveillance capability of our sensor, our target motion assumptions, and the current target location probability distribution. The sensor response results are then used to update the target location probability distribution. Target motion is assumed to take place at the end of each stage, and results in a new target location distribution. The next stage begins with this new target location probability distribution and is the beginning of a surveillance operation involving one fewer stages. The phasing of these various activities is indicated in Figure S-1.


Observe that the surveillance operation described above is dynamic in that the allocation of effort at each stage depends upon the results of the previous stage. This is comparable to, say, a VP operation where flights are flown daily. The results of each day's effort together with the target motion assumptions are then used to decide the allocation of effort for the next day's flights.

In order to illustrate the concepts involved, consider a surveillance operation performed on three cells. Figures S-2 and S-3 compare the expected degree of target localization provided by three surveillance plans against a moving target in a false contact environment. Figures S-2 and S-3 represent the same target motion assumptions and the same surveillance capability. The surveillance capability

* It is possible to consider other measures of target localization, for example, the probability that the target is located in the two highest probability cells. Such generalizations are discussed in detail in Chapter II.

FIGURE 5

PHASING OF VARIOUS ACTIVITIES IN A SURVEILLANCE OPERATION



is defined by the sensor response matrix $R = [r_{ij}]$, where r_{ij} is the probability of a response from investigating cell i if the target is located in cell j . Target motion is assumed to be a Markov chain with the transition matrix M . In both figures the target is restricted to one of three cells. In Figure S-2, however, the prior target location distribution is assumed to be $(1, 0, 0)$ whereas in Figure S-3 it is $(.5, .1, .4)$. The K -stage-optimal surveillance plan is the optimal policy. The 1-stage look-ahead maximum-information-gain policy maximizes at each stage the expected information content of the posterior probability distribution at the end of the stage*. The highest-probability-cell policy investigates at each stage the current highest probability cell. Both the 1-stage look-ahead maximum-information-gain policy and the highest-probability-cell policy are very easy to determine computationally.

In order to understand these figures, consider the situation presented in Figure S-2. The target is assumed to be localized to a single cell at the start of the surveillance operation (the initial target location distribution is $(1, 0, 0)$). Suppose first that no surveillance effort is applied. Because of target motion the degree of target localization will decrease with time. The no surveillance curve (the bottom curve in Figure S-2) represents the expected degree of target localization one can attain after a specified number of stages. The other curves similarly represent the expected amount of target localization one can expect after a specified number of stages using the various surveillance plans. The upper curve is the theoretical maximum expected probability that the target will be localized to a single cell.

Observe that the 1-stage look-ahead maximum-information-gain policy performs almost as well as does the optimal policy. This behavior has been observed in most of the surveillance situations studied to date (see for example Figures II-2 through II-9), and so this policy appears to be a reasonably good suboptimal surveillance policy.

Figures S-2 and S-3 also illustrate, quite dramatically, the conflict between target motion and the application of surveillance effort. In Figure S-2, for example, it is assumed that the target is completely localized to a single cell at the start of the surveillance operation (i. e., the initial target location distribution is $(1, 0, 0)$). Because of this complete localization, a surveillance operation lasting only a few stages will have a high probability of success, even if no surveillance effort is applied. Because of target motion, however, the degree of expected localization decreases rapidly to the limiting values indicated in Table S-1. A surveillance operation against a moving target gains little after a large number of stages from the knowledge that at the beginning of the operation the target was perfectly localized.

* The information content of the probability distribution (x_1, \dots, x_N) defined in N -cells is the nonnegative number $\sum_i x_i \ln x_i + \ln N$.

FIGURE S-2

COMPARISON OF SURVEILLANCE POLICIES

NOTES: (1) Sensor Response Matrix: (2) Target Motion Matrix:

$$R = \begin{bmatrix} .1 & .01 & .01 \\ .01 & .1 & .01 \\ .01 & .01 & .1 \end{bmatrix} \quad M = \begin{bmatrix} .9 & .05 & .05 \\ .05 & .9 & .05 \\ .05 & .05 & .9 \end{bmatrix}$$

- (3) Initial target location probability distribution (1, 0, 0).
- (4) Cell C_1 is the initial high probability cell.




FIGURE S-3

COMPARISON OF SURVEILLANCE POLICIES

NOTES: (1) Sensor Response Matrix: (2) Target Motion Matrix


$$R = \begin{bmatrix} .1 & .01 & .01 \\ .01 & .1 & .01 \\ .01 & .01 & .1 \end{bmatrix}$$

$$M = \begin{bmatrix} .9 & .05 \\ .05 & .9 \\ .05 & .05 \end{bmatrix}$$

(3) Initial target location probability distribution (.5, .1, .4).

(4) Cell C_1 is the initial high probability cell.

Probability of Correctly Specifying Target's Position



K = Number of Stages in Surveillance Operat

Observe that if no surveillance effort is applied, the initial complete localization to a single cell will decay to a probability of 1/3 that the target is located in any given cell. The interaction between surveillance and target motion is reflected in the difference between the probability of correctly specifying the target's position in the no surveillance case and that provided by the various surveillance plans.

In Figure S-3 the target location distribution at the start of the surveillance plan is (.5, .1, .4), so that the initial probability of correctly specifying the target's location is .5. For surveillance operations lasting less than five stages, the optimal surveillance policy and the maximum-expected-information-gain policy are able to overcome the effects of target motion and slightly improve the extent of target localization. For surveillance plans lasting more than 10 stages, however, the optimal surveillance plan (and therefore the other plans as well) are unable to overcome the effects of target motion, and so the initial localization is better than the expected extent of target localization at the end of the operation.

An extremely interesting aspect of Figures S-2 and S-3 is the rapidity with which the expected degree of target localization converges to a fixed value. Moreover, as indicated in Table S-1, for a given surveillance plan the asymptotic expected target localization depends only on the surveillance plan and not on the initial target location distribution. This indicates that in the case at hand precise knowledge of the initial target location distribution is unimportant to the long term ability of a surveillance system to localize a target.

Because of the speed with which the payoffs for surveillance plans converge to their limiting value, this value represents a good but simple way of comparing surveillance plans without introducing a time horizon. The existence of these limits for the K-stage-optimal surveillance plans are established in Theorems II-2 through II-5.

The surveillance system described and illustrated in Figures S-2 and S-3 is an example of a homogeneous surveillance system. A homogeneous surveillance system is one in which the sensor response matrix is of the form $R = [r_{ij}]$ where

$$r_{ij} = \begin{cases} \mu & \text{if } i \neq j \\ \lambda & \text{if } i = j \end{cases} .$$

Additionally, observe that the target motion transition matrix has the special form $M = [d_{ij}]$ where

TABLE S-1

LIMITING VALUES FOR EXPECTED TARGET
LOCALIZATION PROVIDED BY VARIOUS SURVEILLANCE PLANS

	<u>CASE I</u>	<u>CASE II</u>
Optimal Surveillance Policy	.496	.496
Maximum-Expected- Information-Gain Policy	.464	.464
Highest-Probability- Cell Policy	.440	.440
No Surveillance	.333	.333

$$d_{ij} = \begin{cases} \frac{\delta}{N} & \text{if } i \neq j \\ 1 - \delta \frac{(N-1)}{N} & \text{if } i = j. \end{cases}$$

Setting $\delta=0$ yields the special case of a stationary target. Note that the example discussed in Figures S-2 and S-3 is a homogeneous surveillance system with $\lambda = .1$, $\mu = .01$, and $\delta = .15$.

The optimal whereabouts search, introduced by Kadane in reference [b], is the special case where $\mu=0$ and $\delta=0$. In this situation the target is stationary and a sensor response can occur only if the target is in the cell being investigated. Thus a sensor response completely localizes the target to a single cell. The allocation of surveillance effort in an optimal whereabouts search with discrete effort is always to deploy the sensor to the second highest probability cell. See, for example, section 4.4 of reference [c]. Note in particular that the optimal allocation depends only on the current target location probability distribution and not on the horizon. The resulting surveillance plan thus yields uniformly optimal probabilities of localizing the target for any possible horizon.

We now view the K-stage optimal surveillance plan for the homogeneous sensor as a generalization of the optimal whereabouts search. Remarkably, as shown in Chapter II, our preliminary theoretical analysis, together with our numerical results, indicate that the optimal allocation of surveillance effort when $\mu > 0$ is the same as for the optimal whereabouts search, i.e., when $\mu=0$. Thus for any homogeneous surveillance sensor and any target motion matrix of the specified form, we conjecture that the K-stage optimal surveillance plan requires that we allocate all our effort to the second highest probability cell.¹

If this conjecture can be proven, it will have a number of important consequences. First note that such a surveillance plan depends only on the current target location probability distribution and not on the number of stages in the operation and is thus a 1-stage surveillance plan. This plan results in uniformly optimal probabilities of localizing the target for each possible choice of horizon. Moreover, we feel that it is reasonable to model many operational situations with a homogeneous sensor. Since in this case the K-stage optimal surveillance plan would give optimal results for every horizon, it has potential for widespread applications. Finally observe that if this conjecture can be established, then the optimal surveillance plan in the case at hand can be specified for any number of cells without resorting to complicated computational optimization methods.

¹Since this was written, a counterexample has been found by J. R. Weisinger.

A Statistical Model for Processing ASW Contact Information to Estimate Target Patterns of Operation

Two major problems which an ASW planner must frequently face in the presence of sparse contact data of various types and quality are (1) to obtain an estimate for the track of a specified target, and (2) to make inferences about overall target behavior patterns on the basis of contact data on several targets. The objective of Chapter III is to outline a Bayesian method for obtaining these estimates. Unfortunately, this work has not yet been developed to the point where numerical results can be computed.

The goal in undertaking such a study is to provide ASW information processing systems with the capability of combining historical data (in the form of prior estimates on target patterns of operation) with contact data on targets of current interest. It is hoped that this marriage will provide improved target location estimates in that it will methodically exploit contact data on all previous targets.

Our approach is based on a parametric model for target motion. The object is to use the available contact data to obtain posterior estimates for the parameters which describe target motion. A major consideration here is the development of a parametric model for target motion which is rich enough to model real world situations but which is also computationally tractable.

The approach considered here is most applicable in the case of transiting targets. Since the approach is Bayesian, it requires a general form for patterns of motion characterized by parameters for which there are reasonable prior estimates. These prior estimates may be based, for example, on past experience or on certain operational or geographical constraints. The Bayesian approach, however, enables one to obtain estimates of target operation patterns in the presence of sparse data.


The parametric model for target motion is based on the notion of weighted target scenarios. We postulate the existence of a finite number of such scenarios, or patrol track plans, which a target might follow during a specified phase of its mission. Each scenario may be thought of as corresponding to a basic geometric pattern of target motion.

Each scenario is characterized by a mean target track and corresponding covariance matrix. Once a scenario has been chosen for a given target, the target must move roughly according to the mean track of the scenario. The target, however, is permitted to operate with some deviation from the mean track. For example, the target may move faster or more slowly than the specified mean track, or it may vary its course along the mean track. The extent of these perturbations in target motion and their correlation are determined by the covariance matrix associated with the scenario.

An example of some basic tracks which might be used to define scenarios is given in Figure S-4. Additionally, two sample target tracks drawn from two different

FIGURE S-4

EXAMPLE OF TARGET SCENARIOS AND SAMPLE TARGET TRACKS



target scenarios are indicated in Figure S-4. In our model we deal with discrete time and a track is specified by giving the target's location at times $t = 1, \dots, \tau$.

To formulate more precisely our parametric model for target motion, suppose that there are J possible target scenarios, S_1, S_2, \dots, S_J . Let p_j be the probability that a target will follow scenario S_j . Additionally, assume that the conditional distribution function of the target's track $Z = (z_1, \dots, z_\tau)$ is a multivariate normal distribution with mean μ_j and covariance matrix Σ_j^{-1} . Note that μ_j is an element of $\mathbb{R}^{2\tau}$ and that Σ_j^{-1} is a $2\tau \times 2\tau$ positive definite symmetric matrix.

Now assume that the parameters in the target motion model described above are unknown to us. By this is meant, in particular, that the following quantities are not known:

- i) the vector $p = (p_1, p_2, \dots, p_J)$ which gives the prior probabilities that a target will move according to a given scenario,
- ii) the mean target paths $\mu_1, \mu_2, \dots, \mu_J$ of the scenarios, and
- iii) the covariance matrices $\Sigma_1^{-1}, \Sigma_2^{-1}, \dots, \Sigma_J^{-1}$ of the scenarios.

We do assume, however, that the number of possible target scenarios, J , has already been established.

The object of our investigations is now twofold. First, the contact data obtained on a specified target are used to obtain a Bayesian estimate for the track of the target. Secondly, the contact data are processed to obtain Bayesian estimates for the parameters p, μ_1, \dots, μ_J , and $\Sigma_1^{-1}, \dots, \Sigma_J^{-1}$. Computational methods for obtaining these estimates are given in Chapter III.

ASW Information Processing in a Multi-Target Environment

In Chapter IV we develop Bayesian statistical methods for processing ASW information in a multi-target environment. Specifically, suppose that a force of submarine targets is known to be operating in a particular ocean area and that periodic location estimates are to be generated for each target. Potential sources of information to assist in target force localization might include: ocean characteristics that affect navigability such as water depth; submarine operational characteristics such as speed range and frequency of course changes; and surveillance information such as the direct observation of port arrivals and departures and contacts generated by ASW sensors. This localization problem leads naturally to the consideration of ASW information processing systems, i.e., systems designed to accept as input diverse ASW information of the type described and generate as output periodic estimates of the locations of the target submarines.

The methods developed under this problem are applicable to issues of vulnerability to surveillance as well as to issues of ability to exercise surveillance.

The principal emphasis in Chapter IV is on the development of Bayesian statistical methods for the systematic generation and updating of target force location predictions. It should be noted that the surveillance problem considered is an intrinsically multi-target problem as opposed to a composite of many isolated single-target problems. This results first from the fact that except for port departures and arrivals, the submarine targets are treated as observationally indistinguishable. Secondly, all targets are presumed to make potential use of the entire operating area available to the force so that individual targets are not restricted to operate in disjoint subregions. Under these restrictions it will frequently be the case that there is considerable uncertainty about the exact identity of the target that generated a sensor contact. This uncertainty leads to correlation in estimates of the locations of the targets comprising the force under surveillance even though the underlying target motions may be statistically independent. It is from this intrinsic multivariate nature of the problem that the information processing difficulties considered in Chapter IV arise.

A substantial portion of the discussion in Chapter IV deals with specific multi-target Bayesian information processing algorithms and their comparative evaluation based on numerical examples. A particular such algorithm which we call Extended Memory processing will emerge as the only one of the Bayesian methods we consider that both accurately reflects the localization information inherent in the observational data and is at the same time computationally practical.

Examples IV-1 and IV-2 of Chapter IV will show that even when there are as few as two or three targets involved, Bayesian methods that fail to account adequately for correlation in target location estimates can break down badly. In particular, two such methods that we call sequential processing and parallel processing and considered in Chapter IV may lead to complete permutation of target identities and other anomalous behavior. These methods also lead to the so-called quiet target problem that occurs when one of two targets involved in an identity confusion is substantially more quiet than the other. In such a situation Bayesian processing methods that treat target locations as independent tend to react to a stream of contacts on the noisier target by concentrating the target location distributions of both targets in essentially the same location. As a result, valid representation of the position distribution of the quieter target is completely lost. Examples IV-1 and IV-2 show that the Extended Memory method solves the quiet target problem. The theoretical development of Extended Memory is given in Appendix B.

In order to provide the Bayesian information processing methods with a framework to be implemented, tested, and evaluated, we have developed a small-scale information processing system described in Chapter IV called ASWIPS (ASW Information Processing System) on a Prime 400 mini-computer.

The processing system ASWIPS is currently configured to handle up to three targets in a discrete 10 x 10 cellular grid. Time is discretized into information processing update stages. Target motion is taken to be a symmetric random walk with a transition from the current cell to any one of the adjacent cells equally likely. The output of ASWIPS at the end of each update stage consists of estimates of the current locations of all targets in terms of probability distributions. These target location distributions can then be compared with the actual target tracks used to generate the simulated contact data input to ASWIPS. Such comparisons then provide the basis for the evaluation of the predictive capability of various processing approaches. In particular, the calculations involved in Examples IV-1 and IV-2 and the conclusions these examples support are based on ASWIPS.

INFORMATION PROCESSING TO MAINTAIN
LOCALIZATION IN ASW SURVEILLANCE

CHAPTER I

INTRODUCTION

The purpose of this report is to develop methods for processing ASW surveillance information so as to obtain localization estimates on submarine targets on a continuing basis. Such methods need be obtained as useful and computationally practical tools. Potential applications include ASW information processing of surveillance and intelligence data, with an objective of providing improved tactical ASW mission planning. This report is intended primarily for use by analysts.

A predecessor report, reference [a], established much of the framework for subsequent investigations in ASW surveillance information processing. This earlier report attacked two issues at the core of ASW information processing. The first issue concerns a collection of *fixed sensors*, in a false contact environment, with known detection capabilities. The problem was to obtain an estimate of target position on the basis of sensor contact information (both positive and negative). The approach taken was to determine weighted scenarios for target motion and to compute posterior target location distributions and scenario weights on the basis of the contact information. Much of the work contained in Chapter II of this present report was motivated by this earlier reference.

The second issue concerned a movable sensor with known detection capability operating in a false contact environment. The problem was to allocate the sensor so as to serve certain tactically useful purposes, e.g., to localize the target to a specified number of cells. This issue was examined in an exploratory way using Monte Carlo techniques. On the basis of this analysis a number of important conjectures were formulated.

The first such conjecture concerns a concrete connection between information theory and search theory, i.e., that the optimal detection search extracts information from the target location probability distribution at the maximum possible rate. Such an allocation maximizes the entropy of the posterior distribution given failure to detect. This conjecture was established in reference [d].

In contrast it was also conjectured in reference [a] that the opposite connection exists between information theory and surveillance theory, i.e., that a reasonable surveillance plan is to allocate surveillance effort so as to place information into the target location probability distribution at the maximum possible rate. Such an allocation minimizes the expected entropy of the posterior distribution. (Much additional numerical evidence to support this conjecture is given in Figures II-2 through II-9.)

To understand this difference in allocation of effort it is necessary to indicate the relationship between search and surveillance. Search and surveillance are closely related but essentially different activities with different goals and correspondingly different methods of achieving these goals. The principal objective of search is to obtain a target detection. Moreover, search theory generally assumes that the detection also provides the desired target localization. In contrast, surveillance is concerned with localization and its maintenance over time, and this goal can be achieved with or without detections or contacts. In particular, a surveillance problem may not end with a contact because of poor localization information.

Chapter II is an investigation into the optimal allocation of ASW surveillance resources in a somewhat idealized surveillance setting. The analysis is intended to establish a practical and useful theory of surveillance on a sound theoretical basis. The problem considered involves a movable sensor which is deployed in a sequential fashion against a single stochastically moving target in a false contact environment. Among our results, we establish that a good suboptimal allocation of surveillance effort is to maximize incrementally the expected information gained in the posterior target location distribution. This surveillance policy is called the maximum-information-gain plan. The advantage of the maximum-information-gain plan is that it is easily computed in an incremental manner and does not depend on the time horizon. Finding the optimal plan, on the other hand, requires the use of techniques such as dynamic programming and quickly becomes impractical for target distributions with large numbers of cells (say 10 or more).

The objective in Chapter III is to develop a Bayesian method for processing surveillance contact information so as to obtain probability estimates for a single target's track and to combine contact data on a number of targets to estimate general patterns of operation. Our approach is to develop a Bayesian method for combining ASW contact information with a scenario-based parametric model for target motion. We assume that, based on past experience and general operational considerations, we can specify a finite number of general operating plans or basic tracks called scenarios, which could be followed by a target. These scenarios may be thought of as corresponding to the basic geometric patterns of target motion. It may be assumed that these basic patterns are selected for each target in random fashion and then specified for the target in its operations order. Once the basic scenario for the target has been specified,

the target then "chooses" its own particular variation. For example, the target may move locally faster or more slowly than the basic scenario or it may vary its track about the basic scenario.

Bayesian methods have been extensively used in developing real-time computer programs to produce a sequence of updated probability distributions for target location. The target motion model assumed in Chapter III is closely related to models developed by H. R. Richardson and T. L. Corwin. However, the emphasis in this report is on using contact data to revise the motion models or scenarios in a Bayesian fashion. In previous work, Bayesian methods were employed to use contact data to predict target locations but not to revise target motion scenarios.

Observations on a single target are assumed to be in the form of contacts with possibly varying degrees of localization. There are two pieces of information which we wish to obtain from these contact data: First, what is the best estimate of the present target's track, and second, what do contact data on this particular target tell us about general operating patterns. The first question is of interest primarily in situations where contact data are sparse. We answer these two questions in Chapter III by devising methods for computing the posterior distributions for the present target's track and the posterior distribution on scenarios given a series of contact data. The main effort is devoted to developing a class of prior distributions or models for target motion which is rich enough to represent real situations but which is still computationally tractable.

A totally different class of problems in ASW information processing is posed by a multi-target environment. In Chapter IV the concept of Extended Memory processing is introduced with the objective of resolving contact ambiguities on multiple targets. When trying to estimate target location distributions in a multiple target environment, the amount of computer storage required to retain the probability distributions can quickly become excessive unless the target distributions are all mutually independent. However, even when the prior target distributions are assumed to be independent, they lose their independence as soon as one obtains an ambiguous contact, i. e., one which cannot be positively identified as being on a specific target.

In order to resolve this problem the concept of Extended Memory processing considers all reasonable assignments of contacts to targets and computes the probability that each of these assignments is correct. For numerical reasons, assignments with very low probability are excluded from the list. The crucial point is that conditioned on an assignment of contacts, the target distributions are mutually independent. The composite target distributions are then obtained as averages with respect to the assignment probabilities of these independent distributions. The result is that the storage requirements are roughly linear in the number of targets rather than exponential.

In addition, by considering single contacts in the context of an assignment of all the contacts to targets, it often happens that a contact which is ambiguous at the present time will be resolved by future contacts. Thus, the Extended Memory processing has the capability of deferring judgment on a contact until more information is obtained.

Overall, Extended Memory has shown itself to be an effective Bayesian statistical technique to support information processing in a multi-target environment.

CHAPTER II
OPTIMAL SURVEILLANCE AGAINST A
MOVING TARGET IN A FALSE CONTACT ENVIRONMENT

A common and important ASW problem is the allocation of surveillance assets so as to obtain, and maintain over time, a specified degree of target localization. The purpose of this chapter is to describe an idealized surveillance situation in which effort is applied in a false contact environment with the objective of obtaining localization information on a moving target. The problem is to allocate our surveillance assets so as to obtain by some specified future time the best possible estimate for the target's location.

The results contained in this chapter can be viewed as an extension of the approach taken in Chapter III of reference [a]. In reference [a], a number of single-stage look-ahead surveillance policies were formulated and then compared using Monte Carlo simulation techniques. In this current work, we develop a dynamic programming solution for the optimal multi-stage surveillance policy and compare it with a number of other surveillance policies using analytic techniques. As a consequence of this we have been able to confirm a number of conjectures which were made in reference [a] on the basis of Monte Carlo studies. Additionally, our analytic techniques have yielded new insight into the moving target problem and have resulted in some new conjectures.

In the first three sections of this chapter we describe the general nature of the surveillance problem which we consider. The components of the surveillance problem are defined in terms of the sensor response capability and certain assumptions concerning target motion. The Bayesian updating of target location probability distributions to process contact data according to our sensor and target motion assumptions is discussed in the fourth section.

The fifth section is directed toward evaluating the effectiveness of a given surveillance plan against a specified target. This section indicates a number of different surveillance measures of effectiveness, and develops the fundamental recursion relationships which will be heavily exploited in subsequent sections.

The evaluation of optimal surveillance plans is described in the sixth and seventh sections. Such plans are shown to be the solution of a certain dynamic

programming problem. Unfortunately, however, the computational effort required to determine the solution of the dynamic programming problem is immense. For this reason we consider, in the next section, a class of suboptimal surveillance plans, called stationary plans, which are in many cases computationally easier to determine than the optimal plan.

In the ninth section we compare, by way of numerical examples, the optimal surveillance plan to a number of stationary surveillance plans. One of these stationary plans, called the 1-stage look-ahead maximum-information-gain plan, is shown to have a number of very desirable properties. In particular it appears to provide near optimal target localization over a variety of measures of surveillance effectiveness.

One striking feature of the numerical examples considered is the rapid convergence of surveillance effectiveness as the length of time of the surveillance operation increases. Accordingly, the asymptotic behavior of optimal surveillance plans is the subject of the tenth section. The existence of a limiting surveillance effectiveness is established under a number of different hypotheses. Unfortunately we have as yet been unable to determine explicitly the value of this limit. Although the asymptotic results contained in section ten are of interest in their own right, the methods developed here are a particularly important step in placing a theory of optimal surveillance on a sound theoretical basis.

The next three sections are concerned with a special type of surveillance sensor called a homogeneous sensor. The importance of such sensors is based on the fact that it is reasonable to model many operational surveillance sensors as homogeneous sensors. Moreover, our numerical examples indicate that the optimal surveillance plan for such a sensor, when employed against a target which satisfies certain motion assumptions, is in fact a stationary plan of a particularly simple type.

This chapter concludes finally with a section where our various results are summarized.

The Tactical Situation

Suppose we are interested in performing a surveillance operation against a stochastically moving target which is located in one of N cells, C_1, \dots, C_N . The precise cell containing the target is unknown, but at the beginning of the operation a probability distribution for the target's location has been established. We desire to allocate our surveillance assets so as to obtain, at the end of the surveillance operation, a posterior probability distribution for the target's position which localizes the target as much as possible.

The surveillance operation itself is performed in a sequence of discrete time stages, each of duration Δ , by a single movable surveillance sensor.

We are permitted only K stages, and we must use the sensor so as to obtain the best possible estimate for the target's position after these K stages. At the start of the surveillance operation, the fixed time at which we must obtain the best estimate of the target's position, the horizon, is $K\Delta$ units in the future.

At the beginning of each stage, one or more cells are chosen and then investigated for a total amount of time at most Δ . The choice of cell or cells, in general, depends upon the number of stages remaining in the operation, (i. e., the amount of time remaining until the horizon), the surveillance capability of our sensor, our target motion assumptions, and the current target location probability distribution. The sensor response results are then used to update the target location probability distribution. Target motion is assumed to take place at the end of each stage and results in a new target location distribution. The next stage begins with this new target location probability distribution and is the beginning of a surveillance operation involving one fewer stages. The phasing of these various activities is indicated in Figure II-1.

Observe that the surveillance operation described above is dynamic in that the allocation of effort at each stage depends upon the results of the previous stage. This is comparable to, say, a VP operation where flights are flown daily. The results of each day's search together with the target motion assumptions are then used to decide the allocation of effort for the next day's flights.

Sensor Response Assumptions


Because of the possibility of false contacts, a sensor response in a given cell does not necessarily imply that the target is in that cell. Similarly the lack of a sensor response in a given cell does not necessarily mean that the target is not in the cell being investigated. With this in mind we define r_{ij} , $i, j = 1, 2, \dots, N$ to be the conditional probability, given that the target is in cell C_i , that investigating cell C_j for the length of time Δ will result in a sensor response. Moreover we will assume that if the target is in cell C_i , then investigating cell C_j for the length of time t , $0 \leq t \leq \Delta$, will result in a sensor response with probability $r_{ij}t/\Delta$.

The $N \times N$ matrix

$$R = \begin{bmatrix} r_{11} & r_{12} & \dots & r_{1N} \\ r_{21} & r_{22} & \dots & r_{2N} \\ \vdots & \vdots & \dots & \vdots \\ r_{N1} & r_{N2} & \dots & r_{NN} \end{bmatrix}$$

FIGURE II-1

PHASING OF VARIOUS ACTIVITIES IN A SURVEILLANCE OPERATION



is called the sensor response matrix; in our model it completely characterizes the surveillance capability of the sensor. Note that, if there is no possibility of false contacts, then R is a diagonal matrix.

In order to simplify our calculations we will suppose that there can be at most one sensor response in any single surveillance stage. Additionally we will suppose that sensor responses from different time periods are statistically independent.

Target Motion Assumptions

Target motion is assumed to be a Markov process which takes place at the end of each surveillance stage. Let d_{ij} , $i, j = 1, 2, \dots, N$, be the conditional probability that the target will move to cell i given that it originated in cell j . The following $N \times N$ matrix, called the transition matrix,

$$M = \begin{bmatrix} d_{11} & d_{12} & \dots & d_{1N} \\ d_{21} & d_{22} & \dots & d_{2N} \\ \vdots & \vdots & \vdots & \vdots \\ d_{N1} & d_{N2} & \dots & d_{NN} \end{bmatrix},$$

completely describes the target motion assumptions. Observe that $\sum_i d_{ij} = 1$ and that in the case of a stationary target, the transition matrix is simply the identity matrix, $M = I$.

If* $X = (x_1, x_2, \dots, x_N)^T$ is the target location distribution before target motion, then $Y = MX$ is the target location distribution after a single stage of target motion. Moreover, assuming that no surveillance is performed between motion steps, the target location distribution after k -motion steps is $Y = M^k X$.

Mathematical Structure of K-Stage Surveillance Operations

In this section we formulate the mathematical structure of a K -stage surveillance operation performed on N cells, C_1, C_2, \dots, C_N . We will assume that the stochastic structure of the surveillance operation is completely

* We use the notation X' to denote the transpose of a vector X .

known, that is, that the target motion matrix M and the sensor response matrix R are known.

It is useful to introduce the positive orthant in N -space

$$V_N = \{(x_1, x_2, \dots, x_N)' : x_j \geq 0, j = 1, 2, \dots, N\}.$$

For each $X = (x_1, \dots, x_N)' \in V_N$, we define $\|X\|_1 = x_1 + \dots + x_N$. Observe that if x_k is the probability that the target is in cell C_k , $k = 1, 2, \dots, N$, then the vector $X = (x_1, \dots, x_N)'$ is an element of the $N-1$ dimensional simplex

$$\mathcal{S}_{N-1} = \{X \in V_N : \|X\|_1 = 1\}.$$

We will refer to \mathcal{S}_{N-1} as the probability state space. Every target location probability distribution on N cells C_1, C_2, \dots, C_N can be represented as a point of the probability state space.

We now define a K -stage surveillance plan on the cells C_1, C_2, \dots, C_N to be a function $\varphi : \mathcal{S}_{N-1} \times \{1, 2, \dots, K\} \rightarrow V_N$ such that $\|\varphi(X, k)\|_1 \leq 1$ for $(X, k) \in \mathcal{S}_{N-1} \times \{1, 2, \dots, K\}$. In order to explain this definition suppose that at stage k of the surveillance operation, $1 \leq k \leq K$, the current target location probability distribution is $X \in \mathcal{S}_{N-1}$. Let $\varphi(X, k) = (\varphi_1(X, k), \dots, \varphi_N(X, k))$. The surveillance plan φ then requires that we allocate to cell C_l the amount of surveillance effort $\varphi_l \Delta$, $l = 1, 2, \dots, N$. The condition $\|\varphi(X, k)\|_1 \leq 1$ is the constraint that the total amount of surveillance effort available at state k is Δ .

We suppose that the response capability of our surveillance sensor is defined by the $N \times N$ response matrix $R = [r_{ij}]$. Target motion is assumed to be Markovian and is defined by the transition matrix $M = [d_{ij}]$. The diagonal matrices $T_j = \text{diag}(r_{1j}, r_{2j}, \dots, r_{Nj})$, $j = 1, 2, \dots, N$, will prove particularly useful in the following development. Observe that T_j completely characterizes the sensor capability of investigating cell C_j .

At the beginning of stage k of our surveillance operation let the (prior) target location probability distribution be $X \in \mathcal{S}_{N-1}$. We will now determine the various (posterior) target location probability distributions possible at the end of stage k , and indicate the probabilities with which they will occur. Suppose that $\varphi(X, k) = (\varphi_1(X, k), \dots, \varphi_N(X, k))$ so that we are to allocate to cell C_l the amount of surveillance effort $\varphi_l(X, k)\Delta$. It follows then that the probability of obtaining a sensor response from cell l is $\varphi_l(X, k) \|T_l X\|_1$. Moreover it follows from Bayes' theorem that the posterior target location probability distribution given a sensor response in cell C_l is $T_l X / \|T_l X\|_1$.

Since we are assuming at most one sensor response in each surveillance stage, it follows that the probability of obtaining no response is $1 - \sum_{l=1}^N \varphi_l(X, k) \|T_l X\|_1$. Additionally the posterior target location probability distribution given no sensor response is

$$\left(1 - \sum_{l=1}^N \varphi_l(X, k) T_l X\right) / \left(1 - \sum_{l=1}^N \| \varphi_l(X, k) T_l X \|_1\right).$$

In order to obtain the (posterior) target location probability distributions at the end of stage k it is now only necessary to apply the target motion matrix M to each of the above distributions. We thus define

$$U_l(X, k) = \begin{cases} M \frac{\left(1 - \sum_{j=1}^N \varphi_j(X, k) T_j X\right)}{1 - \sum_{j=1}^N \| \varphi_j(X, k) T_j X \|_1} & l = 0. \\ M \frac{T_l X}{\| T_l X \|_1}, & l = 1, 2, \dots, N. \end{cases} \quad (\text{II-1})$$

Note that $U_0(X, k)$ is the posterior target location distribution at the end of stage k given that the prior distribution at the beginning of stage k is X and that there were no sensor responses during stage k. Similarly $U_l(X, k)$, $l = 1, 2, \dots, N$, is the posterior target location probability distribution at the end of stage k given that the prior distribution at the beginning of stage k is X and that there was a sensor response from cell C_l . It is interesting to observe for $l = 1, 2, \dots, N$ that $U_l(X, k)$ does not depend on the amount of surveillance effort applied to cell C_l and so in particular is independent of the surveillance stage k. This results from the linearity of our sensor response assumptions.

So as to simplify our notation further, define

$$\theta_l(X, k) = \begin{cases} 1 - \sum_{j=1}^N \|\varphi_j(X, k) T_j X\|, & l = 0 \\ \|\varphi_l(X, k) T_l X\|, & l = 1, 2, \dots, N. \end{cases} \quad (\text{II-2})$$

The quantity $\theta_0(X, k)$ is the probability of obtaining no sensor response at stage k given that the prior target location distribution at the beginning of stage k is X . Similarly, for $l = 1, 2, \dots, N$, $\theta_l(X, k)$ is the probability of obtaining a sensor response from cell C_l at stage k given that the prior target location distribution at the beginning of stage k is X .

Observe next that if X is the prior target location probability distribution at the beginning of stage k , then MX is the expected posterior target location probability distribution at the end of stage k . Indeed we have

$$MX = \sum_{l=0}^N \theta_l(X, k) U_l(X, k).$$

Consider now a surveillance operation lasting K stages. Let X_0 be the initial target location probability distribution and let X_l , $l = 1, 2, \dots, K$, be the target location probability distribution at the end of stage l . Each X_l is a random vector and the mapping $X_{l-1} \rightarrow X_l$ defines a discrete time Markov process on the probability state space \mathcal{S}_{N-1} . Indeed if for each $X \in \mathcal{S}_{N-1}$ and each Borel measurable set $A \subset \mathcal{S}_{N-1}$ we define for $k = 1, 2, \dots, K$,

$$p_k(X, A) = \sum_{l=0}^N \theta_l(X, k) \mu_{U_l(X, k)}(A), \quad (\text{II-3})$$

where $\mu_Y(\cdot)$ is the unit point mass concentrated at Y , then $p_k(\cdot, \cdot)$ is the stochastic transition function for stage k of the Markov process.

A major question in the theory of surveillance is the distribution of the random variable X_l , $l = 1, 2, \dots, K$. Indeed, the distribution of the X_l indicates the relative likelihood of the various target location probability distributions that can arise from a given surveillance plan. Define inductively

$$p^{(1)}(X_0, A) = p_1(X_0, A)$$

$$p^{(l)}(X_0, A) = \int_{\mathcal{S}_{N-1}} p^{(l-1)}(Y, A) p_l(X_0, dY), \quad l = 2, 3, \dots, K.$$

It follows that $p^{(l)}(X_0, \cdot)$ is the probability distribution for the target location probability distribution at the end of stage l , given that the initial target location probability distribution was X_0 . It is an immediate consequence of equation (II-3) that

$$M^l X_0 = \int_{\mathcal{S}_{N-1}} Y p^{(l)}(X_0, dY),$$

that is the mean target location probability distribution at stage l , given that the initial distribution was X_0 , is $M^l X_0$.

Evaluating Surveillance Plans

In order to determine the effectiveness of a surveillance operation, it is first necessary to establish, for each target location probability distribution, a measure of the localization information implied by that distribution. There are a tremendous number of possibilities for such a measure of effectiveness. Suppose for example that $X = (x_1, \dots, x_N)' \in \mathcal{S}_{N-1}$ is the current target's location probability distribution. Our best single-cell estimate for the target location is that cell which contains the target with highest probability. Indeed if $x_{i_1} \geq x_{i_2} \geq \dots \geq x_{i_N}$, then the target is most likely to be located in cell C_{i_1} , and the probability that this estimate is correct is x_{i_1} . If this is to be our measure of effectiveness, it is clear then that our surveillance effort should be applied so as to maximize, at the end of the surveillance operation, the probability that the target is in the highest probability cell. Other closely related measures of effectiveness would be, for $n = 1, 2, \dots, N$, the probability that the target is located in the first n highest probability cells $x_{i_1} + \dots + x_{i_n}$. We are thus led to define the functions f_1, \dots, f_N on the probability state space \mathcal{S}_{N-1} by $f_n(X) = x_{i_1} + \dots + x_{i_n}$.

As another example of a possible measure of surveillance effectiveness, suppose that, at the end of a surveillance operation, we will search for the target with a specified sensor. The goal of this search is to obtain a target detection, and we can measure the amount of localization information implied by a target location probability distribution in terms of parameters associated with this search. For example we can use the expected time to detect the target with the search sensor as one measure of the localization provided by a probability distribution. If the search has a limit on the total available effort, then we can use as our measure of effectiveness the probability of obtaining a detection within the specified time limit. Observe that, in the case of a perfect sensor* which is only permitted to search one cell, the

* A perfect sensor detects the target if and only if it searches in the cell which contains the target.

maximum probability of obtaining a detection is the probability that the target is located in the highest probability cell.

As a final example, we can suppose that the goal of our surveillance operation is to obtain a posterior target location probability distribution which contains the maximum information possible. Recall that if x_l is the probability that the target is located in cell C_l , $l = 1, 2, \dots, N$, the information content of the probability distribution $X = (x_1, \dots, x_N)'$ is

$$\mathcal{I}(X) = \ln N + \sum_{k=1}^N x_k \ln x_k.$$

Although we will be chiefly concerned with the measures of effectiveness given by the functions f_1, f_2, \dots, f_N defined above, it is generally useful to have a larger class of effectiveness measures. For this reason we define the space of objective functions C to be the space of all convex functions f defined on \mathcal{I}_{N-1} which satisfy the uniform Lipschitz condition: there is a constant $L > 0$ (which depends on f) such that if $X_1, X_2 \in \mathcal{I}_{N-1}$ then

$$|f(X_1) - f(X_2)| \leq L \|X_1 - X_2\|.$$

It is interesting to note that the convexity of $f \in C$ implies that f satisfies the above uniform Lipschitz condition on each compact subset of \mathcal{I}_{N-1} . See for example reference [e]. Thus our Lipschitz condition is only a restriction near the boundary of \mathcal{I}_{N-1} .

Because of the convexity assumption observe that for $X_1, X_2, \dots, X_m \in \mathcal{I}_{N-1}$ each objective function $f \in C$ satisfies

$$f\left(\sum_{k=1}^m \theta_k X_k\right) \leq \sum_{k=1}^m \theta_k f(X_k) \text{ whenever } \theta_k \geq 0, \theta_1 + \dots + \theta_m = 1.$$

In particular f attains its maximum at a vertex of \mathcal{I}_{N-1} , so that the greatest return provided by the objective function is when the target is localized to a single cell.

Suppose now that we wish to evaluate the K -stage surveillance plan ϕ using the measure of effectiveness given by the objective function $f \in C$. For each $X \in \mathcal{I}_{N-1}$ our concern is with the function

$$EV(\varphi, f)(X) = \begin{cases} \text{expected value of } f \text{ at the end of the} \\ \text{surveillance operation } \varphi \text{ given that} \\ \text{the initial target location distribution} \\ \text{is } X \in \mathcal{S}_{N-1}. \end{cases} \quad (\text{II-4})$$

The function EV indicates the amount of target localization, in terms of the objective f, that we can expect at the end of the surveillance operation φ .

The most convenient method of computing the function EV is to introduce, for each $k = 0, 1, 2, \dots, K$ and each $X \in \mathcal{S}_{N-1}$, the functions

$$\Lambda^k(\varphi, f)(X) = \begin{cases} \text{expected value of } f \text{ resulting from following} \\ \text{the last } k \text{ stages of plan } \varphi \text{ given that there} \\ \text{are } k \text{ stages remaining and that the target} \\ \text{location probability distribution at the} \\ \text{current stage (stage } K-k \text{) is } X. \end{cases}$$

The function $\Lambda^k(\varphi, f)$ indicates the expected target localization that can be achieved relative to the objective f given that there are k stages remaining in the surveillance operation. In particular $\Lambda^K(\varphi, f) = EV(\varphi, f)$.

We will compute the functions $\Lambda^k(\varphi, f)$, $k = 0, 1, 2, \dots, K$, inductively. It is evident that $\Lambda^0(\varphi, f) = f$. Suppose that $\Lambda^k(\varphi, f)$ is known and that we are to compute $\Lambda^{k+1}(\varphi, f)$. Let $\varphi = (\varphi_1(X, k), \dots, \varphi_N(X, k))$. If the prior at the beginning of stage $K-k$ is X , we are to allocate to cell C_l the amount of surveillance effort $\varphi_l(X, K-k)$. In the notation of equations (II-1) and (II-2), we will obtain a sensor response from investigating cell C_l with probability $\theta_l(X, K-k)$ and in this case the posterior probability distribution is $U_l(X, K-k)$. Similarly, the probability of obtaining no sensor response is $\theta_0(X, K-k)$ and the resulting posterior target location distribution is $U_0(X, K-k)$. It follows then from the definition of $\Lambda^{k+1}(\varphi, f)$ that

$$\Lambda^{k+1}(\varphi, f)(X) = \sum_{l=0}^N \theta_l(X, K-k) \Lambda^k(\varphi, f)(U_l(X, K-k)). \quad (\text{II-6})$$

It is possible to simplify equation (II-6) somewhat by extending $\Lambda^k(\varphi, f)(\cdot)$ to all of V_N by defining

$$\Lambda^k(\varphi, f)(Y) = \begin{cases} \left(\|Y\|_1 \right)^{-1} \Lambda^k(\varphi, f) \left(\frac{Y}{\|Y\|_1} \right) & \text{if } Y \in V_N, Y \neq 0 \\ 0 & \text{if } Y = 0. \end{cases}$$

The function $\Lambda^k(\varphi, f)(\cdot)$ defined in this manner is homogeneous in the sense that

$$\Lambda^k(\varphi, f)(\lambda Y) = \lambda \Lambda^k(\varphi, f)(Y), \quad \text{all } \lambda \geq 0, Y \in \mathcal{S}_{N-1}. \quad (\text{II-7})$$

Using now the relation (II-7) and the definitions (II-1) and (II-2), we can write equation (II-6) in the form

$$\begin{aligned} \Lambda^{k+1}(\varphi, f)(X) = & \Lambda^k(\varphi, f) \left(\sum_{l=1}^N \varphi_l(X, K-k) M(1-T_l) X \right) \\ & + \sum_{l=1}^N \varphi_l(X, K-k) \Lambda^k(\varphi, f)(MT_l X). \end{aligned} \quad (\text{II-8})$$

The recursion relationships (II-6) and (II-8) are fundamental in that they enable us to compute analytically the effectiveness of any given surveillance system. We will heavily exploit these relationships in the subsequent sections.

Optimal Surveillance Plans

The purpose of this section is to provide a method for determining the optimal K-stage surveillance plan for any objective function $f \in C$. A K-stage optimal surveillance plan φ^K is characterized by the condition

$$EV(\varphi^K, f) \geq EV(\varphi, f)$$

for any surveillance plan φ . This is to say that at the end of the surveillance operation (i. e., at the end of stage K) the optimal surveillance plan yields the greatest possible expected target localization relative to the measure of effectiveness f .

In order to determine the K-stage-optimal surveillance plan φ^K consider the functions Λ^l , $l = 0, 1, 2, \dots, K$ defined by equation (II-5). The iterative scheme (II-8) and the fact that $\Lambda^K(\varphi^K, f) = EV(\varphi^K, f)$ force immediately the

inequalities

$$\Lambda^l(\varphi^K, f) \geq \Lambda^l(\varphi, f), \quad l = 0, 1, 2, \dots, K, \quad (\text{II-9})$$

for all surveillance plans φ . With this we can now inductively compute the functions $\Lambda^l(\varphi^K, f)$, $l = 0, 1, 2, \dots, K$. Recall first that $\Lambda^0(\varphi^K, f) = f$. Now suppose that $\Lambda^k(\varphi^K, f)$ has been determined. It follows immediately from equations (II-8) and (II-9) that

$$\Lambda^{k+1}(\varphi^K, f)(X) = \max \{F^k(\alpha_1, \alpha_2, \dots, \alpha_N, X) : \alpha_1, \dots, \alpha_N \geq 0, \alpha_1 + \dots + \alpha_N \leq 1\}, \quad (\text{II-10})$$

where

$$F^k(\alpha_1, \dots, \alpha_N, X) = \Lambda^k(\varphi^K, f)\left(\sum_{l=1}^N M(I - \alpha_l T_l)X\right) + \sum_{l=1}^N \alpha_l \Lambda^k(\varphi^K, f)(MT_l X). \quad (\text{II-11})$$

It is important to observe that the functions $\Lambda^l(\varphi^K, f)(X)$, $l = 0, 1, 2, \dots, N$, are convex functions of $X \in \mathcal{S}_{N-1}$. To see this note first that $\Lambda^0(\varphi^K, f)(X) = f(X)$ is by assumption a convex function of X . We now argue by induction: suppose that $\Lambda^k(\varphi^K, f)$ is a convex function on \mathcal{S}_{N-1} . Then $\Lambda^k(\varphi^K, f)$ is a convex function on V_N and so, by equation (II-11), $F^k(\alpha_1, \dots, \alpha_N, X)$ is a convex function of $X \in \mathcal{S}_{N-1}$. It follows then from (II-10) that $\Lambda^{k+1}(\varphi^K, f)(X)$ is a convex function of $X \in \mathcal{S}_{N-1}$, the desired result.

Suppose now that, for a given $X \in \mathcal{S}_{N-1}$, the numbers $\alpha_1^*, \dots, \alpha_N^*$ satisfy

$$\Lambda^{k+1}(\varphi^K, f) = F^k(\alpha_1^*, \dots, \alpha_N^*, X), \quad \alpha_1^* + \dots + \alpha_N^* = \beta \leq 1.$$

It is evident then that a K -stage optimal surveillance plan φ^K can be achieved by allocating, at stage $K-k$, the fraction α_l^* of the available effort to cell C_l , i.e., $\varphi^K(X, K-k) = (\alpha_1^*, \dots, \alpha_N^*)$. We claim that the optimal allocation of surveillance effort can always be achieved by allocating the total available effort to a single cell. As we shall see, this is a consequence of the linearity of the sensor detection process.

Writing now

$$M(1 - \sum_{l=1}^N \alpha_l^* T_l) X = (1-\beta)MX + \beta (MX - \sum_{l=1}^N \frac{\alpha_l^*}{\beta} T_l X),$$

it follows from the convexity of $\Lambda^k(\varphi^K, f)(\cdot)$ that

$$\begin{aligned} F^k(\alpha_1^*, \dots, \alpha_N^*, X) &\leq (1-\beta) \Lambda^k(\varphi^K, f)(MX) \\ &\quad + \beta F^k\left(\frac{\alpha_1^*}{\beta}, \frac{\alpha_2^*}{\beta}, \dots, \frac{\alpha_N^*}{\beta}, X\right). \end{aligned}$$

But

$$\begin{aligned} \Lambda^k(\varphi^K, f)(MX) &\leq \Lambda^k(\varphi^K, f)\left(M\left(1 - \sum_{l=1}^N \frac{\alpha_l^*}{\beta} T_l\right) X\right) \\ &\quad + \sum_{l=1}^N \frac{\alpha_l^*}{\beta} \Lambda^k(\varphi^K, f)(T_l X) \\ &= F^k\left(\frac{\alpha_1^*}{\beta}, \dots, \frac{\alpha_N^*}{\beta}, X\right), \end{aligned}$$

and so we conclude that

$$F^k(\alpha_1^*, \dots, \alpha_N^*, X) \leq F^k\left(\frac{\alpha_1^*}{\beta}, \dots, \frac{\alpha_N^*}{\beta}, X\right),$$

with equality if and only if

$$\beta = 1 \text{ or } \Lambda^{k+1}(\varphi^K, f)(X) = \Lambda^k(\varphi^K, f)(MX). \quad (\text{II-12})$$

In order to interpret the second condition in (II-12), observe that if $X \in \mathcal{S}_{N-1}$ is the target location probability distribution with $k+1$ stages remaining in a surveillance operation, then MX is the probability distribution with k stages remaining assuming that no surveillance was performed. The second condition in (II-12) says then that the increment in the expected target localization produced by the optimal surveillance plan φ^K with $k+1$ stages remaining is the same as if no surveillance is performed at that stage. We conclude then that the optimal surveillance plan requires that we use the total amount of

surveillance effort available at a given stage ($\beta=1$), or else the available surveillance effort at the stage in question is insufficient to affect the result.

Without loss of generality we can now assume that in equation (II-10) $\alpha_1 + \dots + \alpha_N = 1$. Subject to this condition we then have

$$F^k(\alpha_1, \dots, \alpha_N, X) = \Lambda^k(\varphi^K, f) \left(\sum_{l=1}^N \alpha_l M(I-T_l)X \right) + \sum_{l=1}^N \alpha_l \Lambda^k(\varphi^K, f) (MT_l X).$$

But since $\Lambda^k(\varphi^K, f)(X)$ is a convex function of X , it follows that $F^k(\alpha_1, \dots, \alpha_N, X)$ is a convex function in the variables $\alpha_1, \dots, \alpha_N$ and so it achieves its maximum at one of the vertices of the convex simplex $\{(\alpha_1, \dots, \alpha_N) : \alpha_j \geq 0, \alpha_1 + \dots + \alpha_N = 1\}$. Thus as claimed the optimal surveillance plan can always be achieved by allocating the entire available effort to a single cell.

Finally observe that the K -stage-optimal surveillance plan depends only on the number of stages remaining in the surveillance plan. Thus if $K_1 < K_2$ we have

$$\Lambda^k(\varphi^{K_1}, f) = \Lambda^k(\varphi^{K_2}, f), \quad k = 0, 1, 2, \dots, K_1.$$

Evaluating Optimal Surveillance Plans Numerically

In the previous section we described a method for determining a K -stage-optimal surveillance plan φ^K for any objective function $f \in C$ and any sensor response and target motion matrices R and M , respectively. Unfortunately, the method requires the solution of a dynamic programming problem, and like such solutions in general, it requires an enormous amount of computational effort to implement. In particular, as the number of cells increases, the dimension of the probability state space \mathcal{S}_{N-1} increases.

For example, consider a surveillance problem on N cells, C_1, C_2, \dots, C_N . Recall that, in order to determine $\Lambda^{k+1}(\varphi^K, f)$, we must know the values of $\Lambda^k(\varphi^K, f)$ on \mathcal{S}_{N-1} . In order to have the values of $\Lambda^k(\varphi^K, f)$ available on a computer, we must quantize the state space \mathcal{S}_{N-1} into a finite number of points at which the values of $\Lambda^k(\varphi^K, f)$ are stored. The values of $\Lambda^k(\varphi^K, f)$ at other points are then estimated by an interpolation scheme.

Suppose that we decide to partition each coordinate axis of \mathcal{S}_{N-1} into J equally spaced points. Our probability state space is then quantized into the points $(i_1, i_2, \dots, i_N)J^{-1}$ where i_1, \dots, i_N are nonnegative integers such that $i_1 + \dots + i_N = J$. Observe that $1/J$ is the coordinate resolution. It is easy to

see that the number of points in such a quantization is

$$\binom{J+N-1}{J}.$$

Ideally a coordinate resolution of .01 or better is highly desirable. Observe, however, from Table II-1 that such a resolution is practical for at most three cells. To solve a four-cell optimal surveillance problem clearly requires a resolution not much finer than .05. Unfortunately, such a resolution is barely adequate to guarantee a good approximation to the functions $\Lambda^k(\varphi^K, f)$, since it provides only 1.5 significant decimal places. Considerable programming care must thus be exercised to insure valid results, even for as few as four cells.

There are a number of special features of the functions $\Lambda^k(\varphi^K, f)$ which can be exploited so as to ease their evaluation. First note that these functions are convex, so that a piecewise linear interpolation with vertices at the points $(i_1, \dots, i_N)J^{-1}$ will provide an upper bound. We now would like to estimate the accuracy of this piecewise linear approximation.

Recall that the space of objective functions C is the space of all convex functions on \mathcal{S}_{N-1} which are uniformly Lipschitz continuous in the sense that if $f \in C$ there exists $L > 0$ such that

$$|f(X_1) - f(X_2)| \leq L \|X_1 - X_2\|_1, \quad \text{all } X_1, X_2 \in \mathcal{S}_{N-1}.$$

Each function $f \in C$ can be extended to a function defined in the positive orthant in N -space

$$V_N = \{Y = (y_1, \dots, y_N) : y_j \geq 0, j = 1, 2, \dots, N\}$$

by setting

$$f(Y) = \begin{cases} \|Y\|_1 f\left(\frac{Y}{\|Y\|_1}\right), & Y \in V_N, Y \neq 0 \\ 0 & Y = 0. \end{cases}$$

TABLE II-1

NUMBER OF POINTS IN A QUANTIZATION OF \mathcal{Y}_{N-1}

		Coordinate Resolution		
		.1	.05	.01
Number of Cells	2	11	21	101
	3	66	231	5,151
	4	286	1,771	176,851
	5	1,001	10,626	4,598,126
	6	3,003	53,130	96,560,646

Suppose now that

$$|f(X_1) - f(X_2)| \leq L \|X_1 - X_2\|_1, \quad \text{all } X_1, X_2 \in \mathcal{S}_{N-1}.$$

If $Y_1, Y_2 \in V_N$, $Y_1, Y_2 \neq 0$, it is easy to see that

$$|f(Y_1) - f(Y_2)| \leq L(2 + \max\{f(X) \mid X \in \mathcal{S}_{N-1}\}) \|Y_1 - Y_2\|_1$$

and so f is also equicontinuous on V_N . The smallest constant L such that

$$|f(Y_1) - f(Y_2)| \leq L \|Y_1 - Y_2\|, \quad \text{all } Y_1, Y_2 \in V_N$$

is called the modulus of continuity of f (on V_N). The modulus of continuity of f provides a usable estimate for the fineness of a quantization of \mathcal{S}_{N-1} required to achieve a piecewise linear approximation for f with a specified accuracy. The following theorem relates the modulus of continuity of the functions $\Lambda^k(\varphi^K, f)$ to that of f , and so provides estimates for the accuracy of approximations for $\Lambda^k(\varphi^K, f)$.

Theorem II-1. Let φ^K be the K -stage optimal surveillance plan for the objective function $f \in C$. If the modulus of continuity of f on V_N is L_0 , then $\Lambda^k(\varphi^K, f)$ is Lipschitz continuous with modulus of continuity L_k , $k = 1, 2, \dots, K$. Additionally

$$L_0 \geq L_1 \geq \dots \geq L_K.$$

Proof. Suppose that the sensor response matrix is R and the motion matrix is M . The proof proceeds by induction on k . First note that the result is trivially valid for $k = 0$. We assume now that the result is valid for $0, 1, 2, \dots, k$; we must verify it for $k+1$.

Now since the optimal surveillance plan can always be achieved by allocating the entire available effort to a single cell, it follows from equation (II-11) that

$$\Lambda^{k+1}(\varphi^K, f)(X) = \max_{l=1, 2, \dots, N} \Lambda^k(\varphi^K, f)(MT_l X) + \Lambda^k(\varphi^K, f)(M(I-T_l)X).$$

Now let $X_1, X_2 \in V_N$ and suppose that

$$\Lambda^{k+1}(\varphi^K, f)(X_1) \geq \Lambda^{k+1}(\varphi^K, f)(X_2).$$

If

$$\Lambda^{k+1}(\varphi^K, f)(X_1) = \Lambda^k(\varphi^K, f)(MT_l X_1) + \Lambda^k(\varphi^K, f)(M(I-T_l)X_1)$$

it follows that

$$\begin{aligned} & \left| \Lambda^{k+1}(\varphi^K, f)(X_1) - \Lambda^{k+1}(\varphi^K, f)(X_2) \right| \\ & \leq \left| \Lambda^{k+1}(\varphi^K, f)(X_1) - \Lambda^k(\varphi^K, f)(MT_l X_2) - \Lambda^k(\varphi^K, f)(M(I-T_l)X_2) \right| \\ & = \left| \Lambda^k(\varphi^K, f)(MT_l X_1) + \Lambda^k(\varphi^K, f)(M(I-T_l)X_1) \right. \\ & \quad \left. - \Lambda^k(\varphi^K, f)(MT_l X_2) - \Lambda^k(\varphi^K, f)(M(I-T_l)X_2) \right|. \end{aligned}$$

Using now our induction hypothesis, we obtain

$$\begin{aligned} \left| \Lambda^{k+1}(\varphi^K, f)(X_1) - \Lambda^{k+1}(\varphi^K, f)(X_2) \right| & \leq L \{ \| MT_l (X_1 - X_2) \|_1 \\ & \quad + \| M(I-T_l)(X_1 - X_2) \|_1 \}. \end{aligned}$$

But $\| MU \|_1 \leq \| U \|_1$ for any vector $U \in \mathbb{R}^N$, so that

$$\begin{aligned} \left| \Lambda^{K+1}(\varphi^{K+1}, f)(X_1) - \Lambda^{K+1}(\varphi^{K+1}, f)(X_2) \right| & \leq L \{ \| T_l (X_1 - X_2) \|_1 \\ & \quad + \| (I-T_l)(X_1 - X_2) \|_1 \} \\ & = L \| X_1 - X_2 \|_1 \end{aligned}$$

and the desired result is verified.

It is easy to construct examples in which $L_0 = L_1 = \dots = L_K$. Indeed, consider a two cell surveillance problem against a stationary target involving a uninformative sensor response matrix

$$R = \begin{bmatrix} c & c \\ c & c \end{bmatrix}, \quad 0 < c < 1.$$

Then for any K-stage surveillance plan φ (including the K-stage optimal plan) $\Lambda^k(\varphi, f) = f$, $k=0, 1, 2, \dots, K$.

Stationary Surveillance Plans

In most cases the K-stage optimal surveillance plan depends strongly on the horizon. Thus it may happen that the surveillance policy designed to obtain the best possible estimate at one specified future time is suboptimal for a different horizon.

Suppose however that we are interested in the maintenance of target localization over time relative to some objective function $f \in C$. In particular such a surveillance plan does not end at any specified horizon, but rather continues indefinitely over time.

Since the K-stage optimal surveillance plans φ^K depend on the number of stages until the horizon and not on the number of stages since the beginning of the operation, such plans are generally not extendable as optimal plans to horizons different than the initially specified horizon. It is thus desirable to find physically realizable surveillance plans which are nearly optimal for any horizon.

The simplest surveillance plans which are extendable to any possible horizon are the stationary surveillance plans. A K-stage surveillance plan φ defined on N cells C_1, C_2, \dots, C_N is stationary if it depends only on the current target location probability distribution $X \in \mathcal{P}_{N-1}$ and not on the current stage of the operation. Such a plan is a function $\varphi: \mathcal{P}_{N-1} \times \{1, 2, \dots, K\} \rightarrow V_N$ such that $\|\varphi\| \leq 1$ which is of the form $\varphi(X, 1) = \varphi(X, 2) = \dots = \varphi(X, K)$.

Such a plan is obviously extendable, as a stationary surveillance plan, to any number of stages K.

Suppose that φ is a stationary surveillance plan and that $f \in C$ is our localization

objective function. If X is the prior target location distribution, the expected degree of target localization obtained by φ after k stages is $\Lambda^k(\varphi, f)(X)$. Thus the expression

$$\frac{1}{K+1} \sum_{k=0}^K \Lambda^k(\varphi, f)(X)$$

is the average expected target localization during the first K -stages. If K is sufficiently large, this expression represents the long term expected target localization one can expect to maintain with φ .

For each k let φ^K be the K -stage optimal surveillance plan relative to $f \in C$. Observe that

$$\Lambda^k(\varphi, f) \leq \Lambda^k(\varphi^K, f).$$

Thus although the K -stage optimal surveillance plans φ^K may not be physically realizable for the long-term maintenance of target localization, they do provide an upper bound for the expected localization that stationary plans can achieve. This fact is useful in determining how close to optimal a particular stationary surveillance plan localizes a given target.

Examples Comparing Optimal and Stationary Surveillance Plans

In this section we compare numerically the localization achieved by the optimal surveillance plan to that which can be obtained by several stationary plans. The particular surveillance system considered is operating against a target located in one of four cells.

We assume that our surveillance system has the sensor response matrix

$$R = \begin{bmatrix} .15 & .01 & .01 & .09 \\ .01 & .15 & .01 & .09 \\ .01 & .01 & .15 & .09 \\ .01 & .01 & .01 & .1 \end{bmatrix}$$

← target in cell C_1

↑
investigate in cell C_2

Thus if we investigate cell C_2 and the target is located in cell C_1 , the probability of obtaining a sensor response is .01. Note that the sensor capability achieved by investigating any of the first three cells is the same. Additionally, if the prior target location distribution is uniform, $X = (.25, .25, .25, .25)$, a contact obtained from investigating any of the first three cells will localize the target to that specific cell with probability .83.

Investigating cell C_4 with this sensor is not as informative about the target's location as is investigation of any of the other cells. To see this, observe that if the prior target location distribution is uniform, a contact obtained from investigating cell C_4 will localize the target to that cell with probability only .27. This is in comparison with the corresponding value .87 obtained by investigating any of the first three cells.

We also assume that the target motion matrix is given by

$$M = \begin{bmatrix} .85 & .05 & .05 & .03 \\ .05 & .85 & .05 & .03 \\ .05 & .05 & .85 & .03 \\ .05 & .05 & .05 & .91 \end{bmatrix} .$$

Thus if the target is located in one of the first three cells it will remain there with probability .85 and move to any one of the other cells with probability .05. Similarly if the target is located in cell C_4 it will remain there with probability .91 and move to any one of the other cells with probability .03. It can be shown that, for each $X \in \mathcal{S}_3$

$$\lim_{k \rightarrow \infty} M^k X = (.214, .214, .214, .358).$$

Thus, if no surveillance effort is applied, the target location distribution will asymptotically approach (.214, .214, .214, .358). In particular, on the basis of our target motion assumptions, the target, after a sufficiently long period of time, is most likely to be located in cell C_4 , and this estimate for the target's position will be correct with probability .358.

The effect of surveillance is to increase our knowledge about the target's position beyond that which can be achieved from our target motion assumptions alone. The amount of increase achieved by a particular surveillance plan is a measure of the effectiveness of the surveillance system when deployed according

TABLE II-2

PRIOR PROBABILITY DISTRIBUTION AND NO SURVEILLANCE LOCATION DISTRIBUTIONS
FOR VARIOUS NUMBERS OF STAGES OF TARGET MOTION IN FOUR CASES OF INTEREST

Note: Highest probability cell is underlined

Prior	Case I	Case II	Case III	Case IV
0	(1.0, 0.0, 0.0, 0.0)	(0.0, 0.0, 0.0, 1.0)	(.7, .1, .1, .1)	(.2, .2, .2, .4)
1	(<u>1.0</u> , 0.0, 0.0, 0.0)	(0.0, 0.0, 0.0, <u>1.0</u>)	(<u>.7</u> , .1, .1, .1)	(.2, .2, .2, <u>.4</u>)
5	(.729, .089, .089, .093)	(.056, .056, .056, .833)	(.534, .150, .150, .167)	(.204, .204, .204, .380)
10	(.437, .175, .175, .213)	(.128, .128, .128, .617)	(.354, .197, .197, .253)	(.209, .209, .209, .374)
15	(.294, .208, .208, .289)	(.174, .174, .174, .479)	(.265, .213, .213, .308)	(.212, .212, .212, .364)
20	(.244, .216, .216, .325)	(.195, .195, .195, .415)	(.233, .216, .216, .334)	(.213, .213, .213, .361)
30	(.225, .216, .216, .342)	(.205, .205, .205, .384)	(.222, .216, .216, .346)	(.214, .214, .214, .358)
40	(.216, .215, .215, .354)	(.212, .212, .212, .363)	(.215, .215, .215, .355)	(.214, .214, .214, .358)
40	(.214, .214, .214, .358)	(.214, .214, .214, .358)	(.214, .214, .214, .358)	(.214, .214, .214, .358)

to the specified plan. To illustrate this point, we compare the effectiveness of a number of different surveillance plans against four hypothetical targets which are distinguished only by having different prior probability distributions. The four prior target location distributions, together with their corresponding no-surveillance posterior distributions, are indicated in Table II-2.

The surveillance plans which we consider here are a K-stage-optimal surveillance plan and the three stationary surveillance plans. The K-stage-optimal surveillance plan has as its objective function the probability that the target is located in the highest probability cell. The stationary plans are referred to as the 1-stage look-ahead maximum-information-gain policy, the 3-stage look-ahead maximum-information-gain policy, and the highest-probability-cell policy.

The 1-stage look-ahead maximum-information-gain policy, denoted φ_{I_1} , allocates its effort so as to maximize the expected information content of the posterior distribution after 1 stage of surveillance and target motion. For $X = (x_1, x_2, x_3, x_4)' \in \mathcal{S}_3$, let $\mathcal{J}(X)$ be the information content of X,

$$\mathcal{J}(X) = \sum_{j=1}^4 x_j \ln x_j + \ln 4,$$

and let φ^* be the one stage optimal surveillance plan for the objective function \mathcal{J} . Then φ^* satisfies

$$\Lambda^1(\varphi^*, \mathcal{J}) \geq \Lambda^1(\varphi, \mathcal{J})$$

for all 1-stage surveillance plans φ . We then define the stationary surveillance plan φ_{I_1} by

$$\varphi_{I_1}(X, k) = \varphi^*(X, 1) \quad \text{for all } k = 1, 2, 3, \dots$$

The 3-stage look-ahead maximum-information-gain plan, denoted φ_{I_3} is defined similarly. Let φ^* be the three-stage optimal surveillance plan for the measure of effectiveness \mathcal{J} . The plan φ^* then satisfies

$$\Lambda^3(\varphi^*, \mathcal{J}) \geq \Lambda^3(\varphi, \mathcal{J})$$

* Note that, although \mathcal{J} fails to satisfy a uniform Lipschitz condition and hence $\mathcal{J} \notin C$, we can still easily define $\Lambda^1(\cdot, \mathcal{J})$.

for every surveillance plan φ . We then set

$$\varphi_{13}(X, k) = \varphi^*(X, 1), \quad k = 1, 2, 3, \dots$$

Thus φ_{13} is the stationary surveillance plan which at each stage allocates effort according to the first stage of φ^* .

Finally, because of its great intuitive appeal, we consider the highest probability-cell policy. This surveillance plan allocates at each stage the entire available effort to one of the cells which currently contains the target with highest probability.

For each initial target location distribution and each surveillance plan under study, we have graphed, in Figures II-2, II-3, II-4, and II-5, the probability of correctly specifying the target's location at the end of the operation as a function of the number of stages in the surveillance operation. Our measure of effectiveness is the probability that the target is located in the highest probability cell. Observe that for each K , as expected, the K -stage optimal surveillance policy yields the greatest probability of correctly specifying the target's position at the end of the surveillance operation. Similarly, if no surveillance is performed, we achieve the least amount of target localization. The two maximum information gain policies perform almost as well as does the optimal policy, and both perform better than the highest probability cell policy.

In particular, in Case I (see Figure II-2), the initial target location probability distribution is $(1, 0, 0, 0)$ so that, with probability 1, the target is initially located in cell C_1 . Thus, as a triviality, any surveillance operation lasting zero stages can always correctly specify the target location at the end of the operation. For the surveillance operations considered here, however, the probability of correctly specifying the target's location decreases if the length of the operation is between 1 and 10 stages. The reason for this is the conflict between target motion and the applied surveillance effort. From Table II-2 we see that, in Case I, target motion during the first ten stages decreases rapidly our information about the target's position. During the first ten stages the available surveillance effort is insufficient to overcome the loss of target localization caused by target motion.

If no surveillance is applied, observe also, from Table II-2, that after about 13 stages of target motion, the most likely cell to contain the target changes from cell C_1 to cell C_4 . Additionally, after 15 stages of target motion the amount of target localization implied by the motion assumptions above actually increases. This is reflected in Figure II-2 by an increase in the localization


FIGURE II-2
COMPARISON OF SURVEILLANCE POLICIES

(Case I)

Notes: (1) Sensor Response Matrix (2) Target Motion Matrix

$$R = \begin{bmatrix} .15 & .01 & .01 & .09 \\ .01 & .15 & .01 & .09 \\ .01 & .01 & .15 & .09 \\ .01 & .01 & .01 & .10 \end{bmatrix} \quad M = \begin{bmatrix} .85 & .05 & .05 & .03 \\ .05 & .85 & .05 & .03 \\ .05 & .05 & .85 & .03 \\ .05 & .05 & .05 & .91 \end{bmatrix}$$

- (3) Surveillance objective function: probability target is in the highest probability cell.
- (4) Motion vector converges to (.214, .214, .214, .358)
- (5) Initial probability distribution (1.0, 0, 0, 0)



K Number of Stages in Surveillance Operation

FIGURE II-3

COMPARISON OF SURVEILLANCE POLICIES

(Case II)

Notes: (1) Sensor Response Matrix (2) Target Motion Matrix


$$R = \begin{bmatrix} .15 & .01 & .01 & .09 \\ .01 & .15 & .01 & .09 \\ .01 & .01 & .15 & .09 \\ .01 & .01 & .01 & .10 \end{bmatrix}$$

$$M = \begin{bmatrix} .85 & .05 & .05 & .03 \\ .05 & .85 & .05 & .03 \\ .05 & .05 & .85 & .03 \\ .05 & .05 & .05 & .91 \end{bmatrix}$$

(3) Surveillance objective function: probability target is in the highest probability cell.

(4) Motion vector converges to (.214, .214, .214, .358)

(5) Initial probability distribution (.0, .0, .0, 1.0)



K - Number of Stages in Surveillance Operation

FIGURE II-4

COMPARISON OF SURVEILLANCE POLICIES

(Case III)

Notes: (1) Sensor Response Matrix (2) Target Motion Matrix

$$R = \begin{bmatrix} .15 & .01 & .01 & .01 & .03 \\ .01 & .15 & .01 & .09 & .03 \\ .01 & .01 & .15 & .09 & .03 \\ .01 & .01 & .01 & .10 & .91 \end{bmatrix}$$

$$M = \begin{bmatrix} .85 & .05 & .05 & .03 \\ .05 & .85 & .05 & .03 \\ .05 & .05 & .85 & .03 \\ .05 & .05 & .05 & .91 \end{bmatrix}$$

(3) Surveillance objective function: probability target is in the highest probability cell.

(4) Motion vector converges to (.214, .214, .214, .358)

(5) Initial probability distribution (.7, .1, .1, .1)




FIGURE II-5

COMPARISON OF SURVEILLANCE POLICIES

(Case IV)


Notes: (1) Sensor Response Matrix (2) Target Motion Matrix

$$R = \begin{bmatrix} .15 & .01 & .01 & .09 \\ .01 & .15 & .01 & .09 \\ .01 & .01 & .15 & .09 \\ .09 & .09 & .09 & .10 \end{bmatrix} \quad M = \begin{bmatrix} .85 & .05 & .05 & .03 \\ .05 & .85 & .05 & .03 \\ .05 & .05 & .85 & .03 \\ .05 & .05 & .05 & .91 \end{bmatrix}$$

(3) Surveillance objective function: probability target is in the highest probability cell.

(4) Motion vector converges to (.214, .214, .214, .358)

(5) Initial probability distribution (.2, .2, .2, .4)



K - Number of Stages in Surveillance Operation

achieved by each of the surveillance plans after 15 stages.

One very remarkable aspect of Figure II-2 is that, after 13 stages of surveillance effort, the highest probability cell policy performs only slightly better in localizing the target than does no surveillance at all. The reason for this is that after 13 stages, the target is most likely to be located in cell C_4 , and so the highest probability cell policy will allocate most of its effort to cell C_4 . But we have already indicated that the sensor response matrix implies that investigating cell C_4 is comparatively unproductive. Indeed the effectiveness of such a search is so low that each increment of effort adds little to the success probability. The K-stage-optimal surveillance policy and the maximum-information-gain policies compensate for this by almost never investigating cell C_4 . Indeed the high probabilities of correctly specifying the target's location attained by these policies are due, in large measure, to exploiting the differences in response probability achieved by applying effort to the various cells. These policies investigate cells C_1 , C_2 , and C_3 , and use a lack of sensor responses to localize the target in cell C_4 . The highest probability cell policy, on the other hand, incorporated no information about the sensor response capability and is thus independent of the response matrix R . It performs correspondingly poorly.

It is also interesting to note in Figures II-2, II-3, II-4, and II-5 that both the 1-stage and 3-stage look-ahead maximum-information-gain policies perform much better than the highest probability cell policy, and nearly as well as the K-stage-optimal surveillance plan. These features are common to all of the cases studied to date. A somewhat surprising observation is that the 3-stage look-ahead maximum-information-gain policy does not perform significantly better than does the 1-stage look-ahead maximum-information-gain policy. Thus, in terms of stationary plans, a well-chosen 1-stage look-ahead policy appears to be almost as good as a policy which looks further into the future. We have yet to find an example which violates this conjecture.

Recall that the K-stage-optimal surveillance plan depends strongly on the horizon, and thus this plan may be quite different for different horizons. The 1- and 3-stage look-ahead maximum-information-gain policies are particularly attractive because they depend only on the present target location probability distribution and not on the number of stages remaining in the operation. Additionally, as illustrated by Figures II-2 through II-5, these policies incorporate sensor response characteristics sufficiently well to give probabilities of success reasonably close to the theoretical maximum determined by the K-stage optimal surveillance plan.

One particularly important feature of the maximum-information-gain policies is that they are not only independent of the number of stages in the surveillance operation, but since they maximize the information content of the posterior distribution, they also do not depend on the measure of effectiveness.

Moreover, they appear to give good target localization for a number of different measures of effectiveness.

To illustrate this point, we have graphed in Figures II-6, II-7, II-8, and II-9 the degree target localization attained, in each of the previous four cases, when the measure of effectiveness is the probability that the target is located in the two highest probability cells. In these figures, the K-stage-optimal surveillance plan is chosen so as to maximize, at the end of K stages, the expected probability that the target is located in the two highest probability cells. Thus the upper curve in Figures II-6 through II-9 indicates a theoretical upper bound for the amount of target localization possible for this surveillance objective function. The 1- and 3-stage look-ahead maximum-information-gain policies and the highest-probability-cell policy are exactly the same allocations of effort used in Figures II-2 through II-5.

Note that, in each case presented, the maximum-information-gain policies yield near optimal two cell target localization. Also, after about 13 stages, the highest-probability-cell policy performs only slightly better than if no surveillance effort was expended at all. Again the reason for this is that after about 10 stages the highest probability cell policy is allocating most of its effort to the comparatively uninformative cell C_4 .

It is thus apparent, at least in the cases at hand, that the maximum-information-gain policies provide robust estimates for target location over a variety of measures of effectiveness. This conjecture has been verified in each case studied to date. This is an important result because such plans are computationally easy to determine, and since they require neither a prior specification of the number stages involved in the surveillance operation nor knowledge of the surveillance objective function.

Finally, it is of interest to observe the asymptotic behavior of the various surveillance plans as the number of stages becomes large. In Figures II-10 through II-13, we have graphed the functions

$$\max_{X \in S_3} \Lambda^k(\varphi, f_1)(X), \text{ and } \min_{X \in S_3} \Lambda^k(\varphi, f_1)(X)$$

for $k = 0, 1, 2, \dots, 20$, where $f_1(X)$ = probability that the target is located in the highest probability cell, and φ is respectively the K-stage-optimal surveillance plan, the 1- and 3-stage look-ahead maximum-information-gain policies, and the highest-probability-cell policy. By definition, the range of possible payoffs resulting from the surveillance plans φ lies between the upper and lower bounds given by the above functions. The striking feature here

FIGURE II-6

COMPARISON OF SURVEILLANCE POLICIES

(Case **a**)

Notes: (1) Sensor Response Matrix

$$R = \begin{bmatrix} .15 & .01 & .01 & .09 \\ .01 & .15 & .01 & .09 \\ .01 & .01 & .15 & .09 \\ .01 & .01 & .01 & .10 \end{bmatrix}$$


(2) Target Motion Matrix

$$M = \begin{bmatrix} .85 & .05 & .05 & .03 \\ .05 & .85 & .05 & .03 \\ .05 & .05 & .85 & .03 \\ .05 & .05 & .05 & .91 \end{bmatrix}$$

(3) Surveillance objective function: probability target is in the highest two probability cells

(4) Motion vector converges to (.214, .214, .214, .358)

(5) Initial probability distribution (1.0, 0, 0, 0)



K = Number of Stages in Surveillance Operation

FIGURE H-7

COMPARISON OF SURVEILLANCE POLICIES

(Case IIa)

Notes: (1) Sensor Response Matrix

$$R = \begin{bmatrix} .15 & .01 & .01 & .09 \\ .01 & .15 & .01 & .09 \\ .01 & .01 & .15 & .09 \\ .01 & .01 & .01 & .10 \end{bmatrix}$$


(2) Target Motion Matrix

$$M = \begin{bmatrix} .85 & .05 & .05 & .03 \\ .05 & .85 & .05 & .03 \\ .05 & .05 & .85 & .03 \\ .05 & .05 & .05 & .91 \end{bmatrix}$$

(3) Surveillance objective function: probability target is in the highest two probability cells

(4) Motion vector converges to (.214, .214, .214, .358)

(5) Initial probability distribution (.0, .0, .0, 1.0)



K = Number of Stages in Surveillance Operation

FIGURE II-8

COMPARISON OF SURVEILLANCE POLICIES

(Case IIIa)




FIGURE II-9

COMPARISON OF SURVEILLANCE POLICIES

(Case IVa)

Notes: (1) Sensor Response Matrix

$$R = \begin{bmatrix} .15 & .01 & .01 & .09 \\ .01 & .15 & .01 & .09 \\ .01 & .01 & .15 & .09 \\ .09 & .09 & .09 & .10 \end{bmatrix}$$

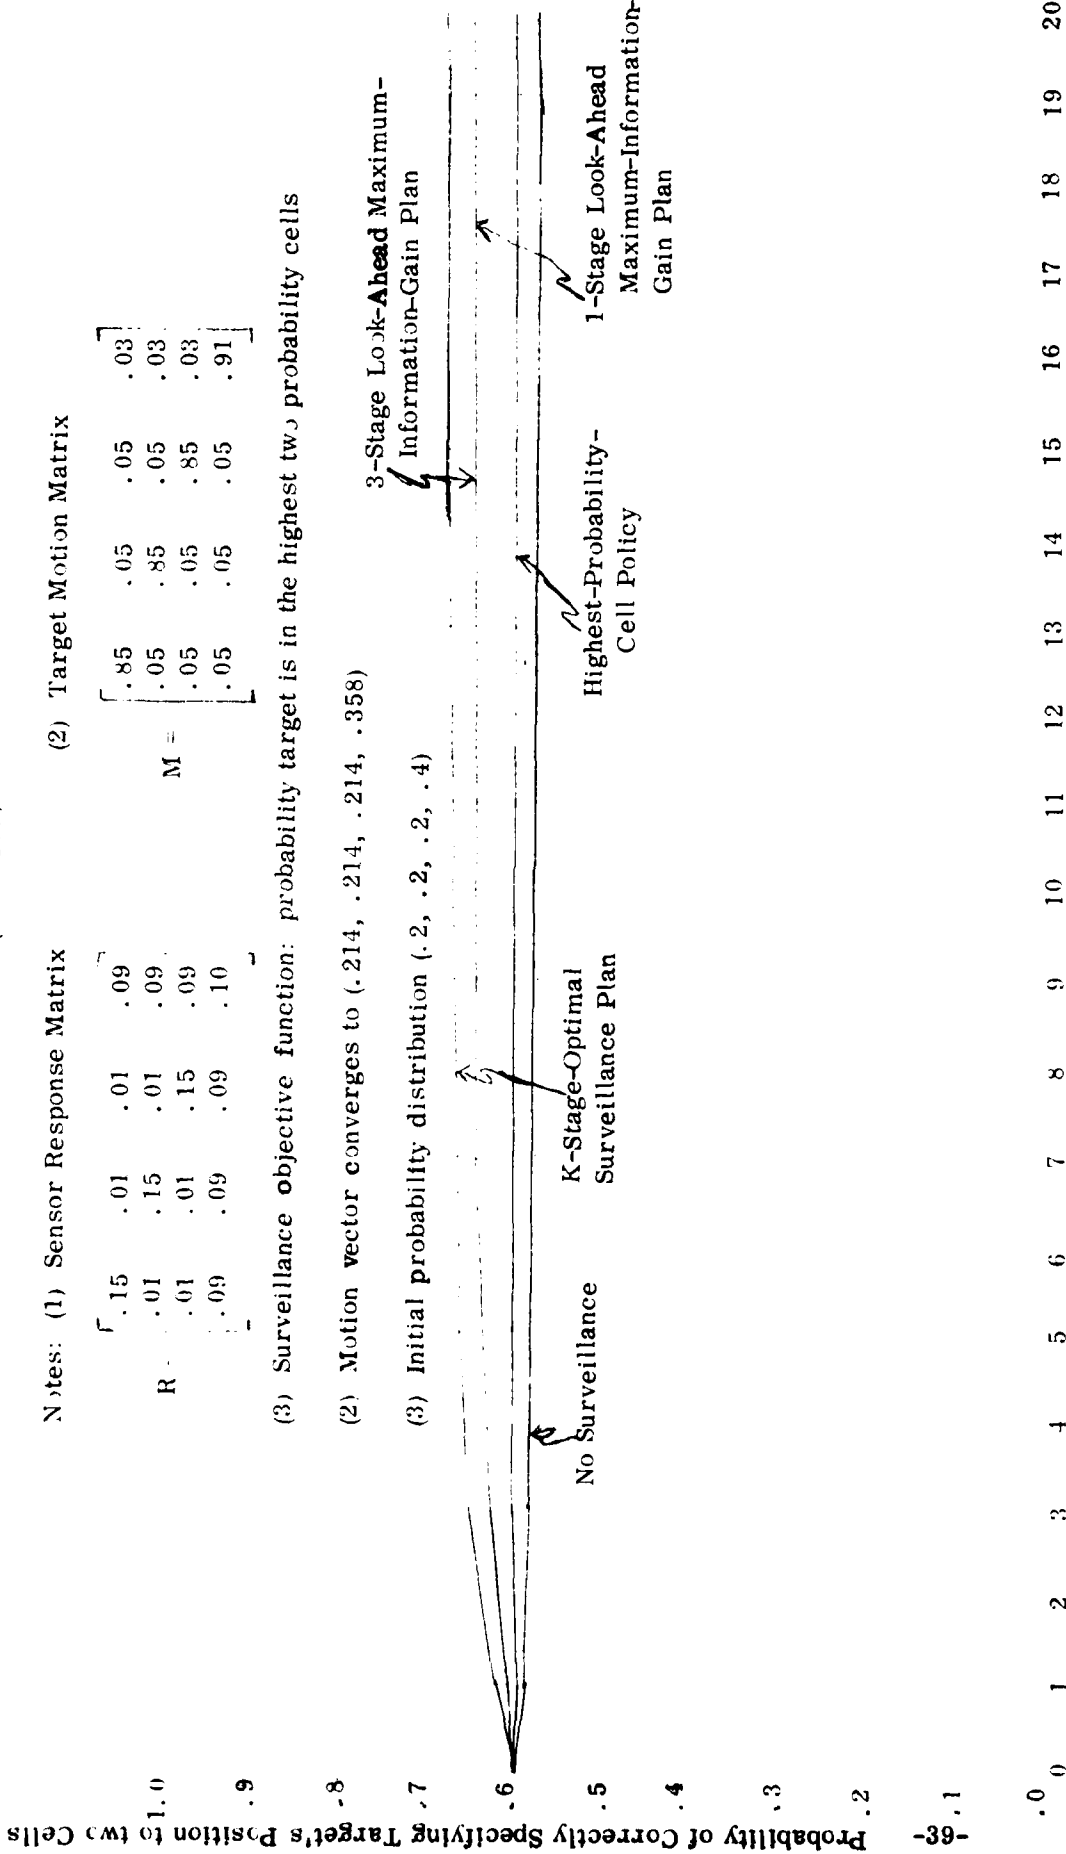
(2) Target Motion Matrix

$$M = \begin{bmatrix} .85 & .05 & .05 & .03 \\ .05 & .85 & .05 & .03 \\ .05 & .05 & .85 & .03 \\ .05 & .05 & .05 & .91 \end{bmatrix}$$

(3) Surveillance objective function: probability target is in the highest two probability cells

(2) Motion vector converges to (.214, .214, .214, .358)

(3) Initial probability distribution (.2, .2, .2, .4)



K - Number of Stages in Surveillance Operation

FIGURE II-10
ENVELOPE OF SURVEILLANCE EFFECTIVENESS FOR THE


K-STAGE OPTIMAL SURVEILLANCE POLICY

Notes: (1) Sensor Response Matrix (2) Target Motion Matrix

$$R = \begin{bmatrix} .15 & .01 & .01 & .09 \\ .01 & .15 & .01 & .09 \\ .01 & .01 & .15 & .09 \\ .01 & .01 & .01 & .10 \end{bmatrix}$$

$$M = \begin{bmatrix} .85 & .05 & .05 & .03 \\ .05 & .85 & .05 & .03 \\ .05 & .05 & .85 & .03 \\ .05 & .05 & .05 & .91 \end{bmatrix}$$

- (3) Surveillance objective function: $f = \text{probability target is in the highest probability cell.}$
- (4) $\varphi^K = K\text{-stage-optimal surveillance policy}$
- (5) Motion vector converges to (.214, .214, .214, .358)



K Number of Stages in Surveillance Operation

FIGURE II-11

ENVELOPE OF SURVEILLANCE EFFECTIVENESS FOR THE 1-STAGE

LOOK-AHEAD MAXIMUM-INFORMATION-GAIN POLICY

Notes: (1) Sensor Response Matrix (2) Target Motion Matrix


$$R = \begin{bmatrix} .15 & .01 & .01 & .01 & .09 \\ .01 & .15 & .01 & .01 & .09 \\ .01 & .01 & .15 & .01 & .09 \\ .01 & .01 & .01 & .15 & .10 \end{bmatrix}$$

$$M = \begin{bmatrix} .85 & .05 & .05 & .03 \\ .05 & .85 & .05 & .03 \\ .05 & .05 & .85 & .03 \\ .05 & .05 & .05 & .91 \end{bmatrix}$$

(3) Surveillance objective function: f = probability target is in the highest probability cell.

(4) φ = 1-stage look-ahead maximum-information-gain policy

(5) Motion vector converges to (.214, .214, .214, .358)



K : Number of Stages in Surveillance Operation

FIGURE II-12

ENVELOPE OF SURVEILLANCE EFFECTIVENESS FOR THE

3-STAGE LOOK-AHEAD MAXIMUM-INFORMATION-GAIN POLICY

Notes: (1) Sensor Response Matrix (2) Target Motion Matrix

$$R = \begin{bmatrix} .15 & .01 & .01 & .09 \\ .01 & .15 & .01 & .09 \\ .01 & .01 & .15 & .09 \\ .01 & .01 & .01 & .10 \end{bmatrix}$$

$$M = \begin{bmatrix} .85 & .05 & .05 & .03 \\ .05 & .85 & .05 & .03 \\ .05 & .05 & .85 & .03 \\ .05 & .05 & .05 & .91 \end{bmatrix}$$

(3) Surveillance objective function: $f =$ probability target is in the highest probability cell.

(4) $\varphi =$ 3-stage look-ahead maximum-information-gain policy.

(5) Motion vector converges to (.214, .214, .214, .358)





FIGURE II-13
 ENVELOPE OF SURVEILLANCE EFFECTIVENESS FOR THE
 HIGHEST-PROBABILITY-CELL POLICY

Notes: (1) Sensor Response Matrix (2) Target Motion Matrix

$$R = \begin{bmatrix} .15 & .01 & .01 & .09 \\ .01 & .15 & .01 & .09 \\ .01 & .01 & .15 & .05 \\ .01 & .01 & .01 & .10 \end{bmatrix}$$

$$M = \begin{bmatrix} .85 & .05 & .05 & .03 \\ .05 & .85 & .05 & .03 \\ .05 & .05 & .85 & .03 \\ .05 & .05 & .05 & .91 \end{bmatrix}$$

- (3) Surveillance objective function: f = probability target is in the highest probability cell.
- (4) φ = highest-probability-cell policy
- (5) Motion vector converges to (.214, .214, .214, .358).



is the apparent rapid convergence in each case to constant values which depend on the surveillance plan but not on the prior target location probability distribution. This indicates that precise knowledge of the initial target location distribution is unimportant to the long term ability of a surveillance system to localize a target. In view of this apparent rapid convergence of the expected localization to the limit values, it is valuable to be able to characterize the limit values in terms of the surveillance plan, the sensor response matrix, and the target motion matrix. Unfortunately we have not yet been able to push matters that far. In the next section, however, we establish the existence of the specified limit for the K-stage-optimal surveillance plans under a variety of different hypotheses.

Asymptotic Behavior of the K-Stage Optimal Surveillance Plan

One of the most striking features common to the examples discussed in the previous section is the rapid convergence, as $K \rightarrow \infty$, of the functions $\Lambda^K(\varphi, f)(\cdot)$. Additionally, it is of interest to note that, in all of the previous examples, the limiting values

$$\lim_{K \rightarrow \infty} \Lambda^K(\varphi, f)(X) \quad (\text{II-13})$$

depend only on the surveillance plan φ and the measure of effectiveness f and not on the initial target location probability distribution X . The purpose of this section is to discuss various issues related to the asymptotic behavior of the functions $\Lambda^K(\varphi^K, f)(\cdot)$ for K-stage optimal surveillance plans φ^K .

First note that the limit (II-13) need not always exist for the optimal surveillance plan, and indeed it is quite easy to construct counterexamples. To see this, consider a two-cell surveillance problem involving a completely uninformative response matrix R : for some c , $0 < c < 1$,

$$R = \begin{bmatrix} c & c \\ c & c \end{bmatrix}.$$

Also suppose that the Markovian target motion matrix M is given by

$$M = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}.$$

Such a motion matrix interchanges the position of the target between the two cells in a single stage.

Suppose now that, for $X = (x_1, x_2)' \in \mathcal{S}_1$, our objective function f is given by

$$f(X) = ax_1 + bx_2$$

where a and b are positive real numbers, $a \neq b$. Since our response matrix is completely uninformative, any surveillance plan is optimal, and we will thus assume that the surveillance plan φ^K is given by $\varphi^K(X, k) = (\frac{1}{2}, \frac{1}{2})$ for all $X \in \mathcal{S}_1$ and all $K = 1, 2, 3, \dots, k = 1, 2, \dots, K$.

Using now the definitions for the posterior distributions U_l given in (II-1) we see, for $l = 0, 1$, and 2 , that

$$U_l(X, K) = MX, \quad K = 1, 2, 3, \dots$$

Similarly, if $\theta_l(X, K)$, $l = 0, 1, 2$, are the event probabilities defined in equation (II-2), we have, for all $X \in \mathcal{S}_1$, and $K = 1, 2, 3, \dots$ that

$$\theta_l(X, K) = \begin{cases} 1-c & \text{if } l = 0 \\ c/2 & \text{if } l = 1, 2. \end{cases}$$

It follows then from equation (II-8) that

$$\Lambda^k(\varphi^k, f) = f(M^k X).$$

But note that $M^{2K+1} = M$ and that $M^{2K} = I$, $K = 1, 2, 3, \dots$. Thus if $X = (x_1, x_2)' \in \mathcal{S}_1$,

$$\Lambda^K(\varphi^K, f)(X) = \begin{cases} ax_1 + bx_2 & \text{if } K \text{ is even} \\ ax_2 + bx_1 & \text{if } K \text{ is odd,} \end{cases}$$

so that the limit (III-13) will exist if and only if $X = (\frac{1}{2}, \frac{1}{2})$.

In spite of this counterexample, observe that, in the case at hand, every subsequence of $\{\Lambda^K(\varphi^K, f)(\cdot)\}_{K=1}^{\infty}$ contains a uniformly convergent subsequence. Similarly the Cesaro limit of the functions $\Lambda^k(f, \varphi)(\cdot)$ exists for $X \in \mathcal{S}_1$.
Indeed

$$\lim_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K \Lambda^k(f, \varphi^k)(X) = a + b.$$

Recall that the Cesaro limit here indicates the long-term average target localization obtainable using the objective function f and the surveillance parameters R , M , and φ^k . In Theorems II-2 and II-3 below we establish these two facts for an arbitrary surveillance objective function.

In the remainder of this section we will consider a surveillance operation on N cells C_1, C_2, \dots, C_N . Recall that our space of objective functions C is the space of all convex functions on \mathcal{S}_{N-1} which are uniformly Lipschitz continuous on \mathcal{S}_{N-1} . We assume that each function $f \in C$ has been extended to a function defined on the positive orthant V_N of \mathbb{R}^N , by setting

$$f(Y) = \begin{cases} \|Y\|_1 f\left(\frac{Y}{\|Y\|_1}\right) & \text{if } Y \in V_N, Y \neq 0 \\ 0 & \text{if } Y = 0. \end{cases}$$

Recall that if $f \in C$ then there exists $L > 0$ such that, for all $Y_1, Y_2 \in V_N$,

$$|f(Y_1) - f(Y_2)| \leq L \|Y_1 - Y_2\|_1.$$

Theorem II-2. Let φ^K , $K = 0, 1, 2, 3, \dots$ be the K -stage optimal surveillance plans for the objective function $f \in C$. Then the family $\{\Lambda^K(\varphi^K, f)(\cdot)\}_{K=0}^{\infty}$ is an equicontinuous family of functions and so every subsequence from this family contains a uniformly convergent subsequence.

Proof. Suppose our sensor response matrix is R and that our target motion matrix is M . Choose a constant $L > 0$ such that if $X_1, X_2 \in V_N$ then

$$|f(X_1) - f(X_2)| \leq L \|X_1 - X_2\|_1.$$

It then follows from Theorem II-1 that if $X_1, X_2 \in \mathcal{S}_N$, then

$$|\Lambda^k(\varphi^k, f)(X_1) - \Lambda^k(\varphi^k, f)(X_2)| \leq L \|X_1 - X_2\|_1$$

and so the family $\{\Lambda^k(\varphi^k, f)\}$ is clearly equicontinuous.

As an immediate consequence of this theorem observe that every convergent subsequence of $\{\Lambda^k(\varphi^k, f)\}$ is in fact uniformly convergent.

The following theorem establishes the existence of the Cesaro limit in general. Note that this theorem is valid for every sensor response matrix R and every target motion matrix M . The theorem states that there is a long term average expected target localization for every prior target location distribution.

Theorem II-3. Let φ^k , $k = 0, 1, 2, \dots$ be the k -stage optimal surveillance policy for the objective function $f \in C$. Then the sequence of functions

$$\left\{ \frac{1}{K} \sum_{k=1}^K \Lambda^k(\varphi^k, f) \right\}_{K=1}^{\infty}$$

is uniformly convergent.

To establish this result we need first the following lemma.

Lemma. Let φ be the stationary surveillance on N cells given by

$$\varphi(X, k) = (1/N, 1/N, \dots, 1/N) \quad X \in S_{N-1}, \quad k = 0, 1, 2, \dots$$

If $f \in C$ then the sequence of functions

$$\left\{ \frac{1}{K} \sum_{k=1}^K \Lambda^k(\varphi, f) \right\}_{K=1}^{\infty}$$

is uniformly convergent.

Proof of Lemma. We first show that the family $\{\Lambda^k(\varphi, f)\}$ is an equicontinuous family. Choose a constant L such that

$$|f(X_1) - f(X_2)| \leq L \|X_1 - X_2\|_1, \quad X_1, X_2 \in V_N.$$

We will show, by induction on k , that

$$|\Lambda^k(\varphi, f)(X_1) - \Lambda^k(\varphi, f)(X_2)| \leq L \|X_1 - X_2\|_1.$$

The result is clear for $k = 0$. Using equation (II-8), we obtain

$$\begin{aligned} & |\Lambda^{k+1}(\varphi, f)(X_1) - \Lambda^{k+1}(\varphi, f)(X_2)| \\ & \leq \frac{1}{N} \sum_{l=1}^N |\Lambda^k(\varphi, f)(MT_l X_1) - \Lambda^k(\varphi, f)(MT_l X_2)| \\ & \quad + |\Lambda^k(\varphi, f)[M(I - \frac{1}{N} \sum_{l=1}^N T_l)X_1] - \Lambda^k(\varphi, f)[M(I - \frac{1}{N} \sum_{l=1}^N T_l)X_2]|. \end{aligned}$$

It follows now from our induction hypothesis that

$$\begin{aligned} |\Lambda^{k+1}(\varphi, f)(X_1) - \Lambda^{k+1}(\varphi, f)(X_2)| & \leq L \left\{ \frac{1}{N} \sum_{l=1}^N \|MT_l(X_1 - X_2)\|_1 \right. \\ & \quad \left. + \left\| M(I - \frac{1}{N} \sum_{l=1}^N T_l)(X_1 - X_2) \right\|_1 \right\}. \end{aligned}$$

It is easy to see that, for any vector $Y \in \mathbb{R}^N$, $\|M_Y\|_1 \leq \|Y\|_1$, whence we have

$$\begin{aligned} & |\Lambda^{k+1}(\varphi, f)(X_1) - \Lambda^{k+1}(\varphi, f)(X_2)| \\ & \leq L \left\{ \frac{1}{N} \sum_{l=1}^N \|T_l(X_1 - X_2)\|_1 + \left\| (I - T_l)(X_1 - X_2) \right\|_1 \right\}. \end{aligned}$$

Now recall that $T_l = \text{diag}(r_{1l}, r_{2l}, \dots, r_{Nl})$ and let $(y_1, y_2, \dots, y_N) = X_1 - X_2$. Since $0 \leq r_{jl} \leq 1$, it follows that

$$\begin{aligned} & \|T_l(X_1 - X_2)\|_1 + \|(I - T_l)(X_1 - X_2)\|_1 \\ &= \sum_{j=1}^N |r_{jl} y_j| + |(1 - r_{jl}) y_j| = \|X_1 - X_2\|_1. \end{aligned}$$

Thus we conclude that

$$|\Lambda^{k+1}(\varphi, f)(X_1) - \Lambda^{k+1}(\varphi, f)(X_2)| \leq l \cdot \|X_1 - X_2\|_1,$$

and so $\{\Lambda^k(\varphi, f)\}$ is an equicontinuous family.

Now set

$$G_K(f)(X) = \frac{1}{K} \sum_{k=1}^K \Lambda^k(\varphi, f)(X).$$

The family of functions $\{G_K(f)\}$ is clearly also an equicontinuous family, and so every subsequence contains a uniformly convergent subsequence. We thus need only show that every convergent subsequence from the family $\{G_K(f)\}$ converges to the same limit.

Suppose now that the sequences $\{G_{K_{j_1}}(f)\}_{j_1=1}^{\infty}$ and $\{G_{K_{j_2}}(f)\}_{j_2=1}^{\infty}$ are uniformly convergent. Set

$$\mathcal{A} = \{g \in C : \{G_{K_{j_1}}(g)\}_{j_1=1}^{\infty} \text{ and } \{G_{K_{j_2}}(g)\}_{j_2=1}^{\infty} \text{ are uniformly convergent sequences}\}.$$

Clearly $f \in \mathcal{A}$, and so \mathcal{A} is nonvoid.

For each $g \in \mathcal{A}$ define functionals H_1 and H_2 on \mathcal{A} by

$$H_1(g) = \lim_{j_1 \rightarrow \infty} G_{K_{j_1}}(g), \text{ and } H_2(g) = \lim_{j_2 \rightarrow \infty} G_{K_{j_2}}(g).$$

For $g \in \mathcal{A}$ the functions $H_1(g)$ and $H_2(g)$ are convex uniformly Lipschitz continuous functions on S_{N-1} and hence in C . Additionally, for each $X \in S_{N-1}$, $H_1(\cdot)(X)$ and $H_2(\cdot)(X)$ are linear functionals on \mathcal{A} .

Define now an operator on the space \mathcal{A} as follows: for each $g \in \mathcal{A}$ set

$$\Phi(g)(X) = \frac{1}{N} \sum_{l=1}^N g(MT_l X) + g(M(I - \frac{1}{N} \sum_{l=1}^N T_l)X).$$

Note that Φ is a linear operator on \mathcal{A} and that

$$\Lambda^k(\varphi, g) = \Phi^k(g) = \Phi \cdots \Phi(g), \quad g \in \mathcal{A}.$$

We now claim that $\Phi(g) \in \mathcal{A}$ whenever $g \in \mathcal{A}$. Clearly, $\Phi(g) \in C$, and we need only show that the sequences $\{G_{K_{j_i}}(\Phi g)\}_{j_i=1}^\infty$ are uniformly convergent. But, for $i = 1, 2$,

$$G_{K_{j_i}}(\Phi g) = G_{K_{j_i}}(g) + \frac{1}{K_{j_i}} [\Lambda^{K_{j_i}+1}(\varphi, g) - \Lambda^1(\varphi, g)].$$

Thus $\Phi(g) \in \mathcal{A}$, and indeed $H_i(g) = H_i(\Phi g)$, $i = 1, 2$. It follows then that $G_{K_{j_i}}(g) \in \mathcal{A}$, for $i = 1, 2, 3, \dots$.

We also claim that $H_i(g) \in \mathcal{A}$ and indeed that H_i is a continuous operator of \mathcal{A} into itself. To see that $H_i(g) \in \mathcal{A}$ whenever $g \in \mathcal{A}$, note that since $G_{K_{j_i}}(g)$ converges uniformly to $H_i(g)$ it follows from the linearity of Φ that $\Phi G_{K_{j_i}}(g)$ converges uniformly to $\Phi H_i(g)$. But $\Phi G_{K_{j_i}}(g) = G_{K_{j_i}}(\Phi g)$, and so $\Phi H_i(g) = H_i(\Phi g) = H_i(g)$; whence, $G_K(H_i(g)) = H_i(g)$, $i = 1, 2, K = 1, 2, 3, \dots$. Thus clearly $H_i(g) \in \mathcal{A}$.

Let \mathcal{A} have the topology induced by the supremum norm. To see that H_i is a continuous mapping of \mathcal{A} into itself, let $g_1, g_2 \in \mathcal{A}$ satisfy

$$\|g_1 - g_2\|_\infty = \sup_{X \in \mathcal{S}_{N-1}} |g_1(X) - g_2(X)| \leq \frac{\epsilon}{3}.$$

Note that if $X \in V_N$ then $|g_1(X) - g_2(X)| \leq (\epsilon/3) \|X\|_1$, so that

$$|\Phi(g_1)(X) - \Phi(g_2)(X)| \leq \frac{\epsilon}{3}, \quad X \in \mathcal{S}_{N-1}.$$

It follows then that

$$\|G_K(g_1) - G_K(g_2)\|_\infty \leq \frac{\epsilon}{3}, \quad K = 1, 2, 3, \dots$$

Now choose K_0 so large that if $K_{j_i} > K_0$, then

$$\|H_1(g_l) - G_{K_{j_i}}(g_l)\|_\infty \leq \frac{\epsilon}{3}, \quad i = 1, 2.$$

Using the triangle inequality it follows immediately that

$$\|H_1(g_1) - H_1(g_2)\|_\infty \leq \epsilon,$$

and so as claimed H_1 is a continuous map of \mathcal{A} into itself.

Now observe that

$$H_1(f) = H_1(G_{K_{j_2}}(f)) \quad \lim_{j_2 \rightarrow \infty} H_1(G_{K_{j_2}}(f)) = H_1(H_2(f))$$

since H_1 is continuous on \mathcal{R} . Similarly

$$H_1(f) = G_{K_{j_2}}(H_1(f)) = H_2(H_1(f)).$$

Interchanging the role of H_1 and H_2 above we obtain

$$H_1(f) = H_1 H_2(f) = H_2 H_1(f) = H_2(f)$$

and so every convergent subsequence from the family $\{G_{K_i}\}$ has the same limit, as claimed. This completes the proof of the Lemma.

Proof of Theorem II-3. For $k = 1, 2, 3, \dots$ let ϕ^k be the k -stage optimal surveillance plan for the measure of effectiveness function f . Recall that each ϕ^k is determined entirely by the number of stages remaining in the surveillance operation, so that, for all $X \in \mathcal{S}'_{N-1}$,

$$\phi^{k_1}(X, k_1 - l) = \phi^{k_2}(X, k_2 - l) \quad l = 0, 1, 2, \dots, \min\{k_1, k_2\}.$$

For each $j = 1, 2, 3, \dots$ let ψ_{jk} be the k -stage surveillance plan defined by

$$\psi_{jk}(X, l) = \begin{cases} \phi^k(X, l) & \text{if } k \leq j \\ \left(\frac{1}{N}, \frac{1}{N}, \dots, \frac{1}{N}\right) & l = 1, 2, \dots, k-j \\ \phi^j(X, l-k+j), l = k-j+1, \dots, k & \text{if } k > j. \end{cases}$$

Thus if $k > j$, ψ_{jk} allocates, during the first $k-j$ stages of the surveillance plan, the same amount of effort to each cell. During the remaining j stages of the plan, ψ_{jk} allocates effort according to the optimal surveillance plan ϕ^j .

Now define an operator $\Phi : C \rightarrow C$ by

$$\Phi(f)(X) = \frac{1}{N} \sum_{l=1}^N f(MT_l X) + f\left(M\left(I - \frac{1}{N} \sum_{l=1}^N T^l\right)X\right).$$

It is easy to see that

$$\Lambda^k(\psi_{jk}, f) = \begin{cases} \Lambda^k(\phi^k, f), & k = 1, 2, \dots, j \\ \Phi^{k-j}(\Lambda^j(\phi^j, f)), & k = j+1, j+2, \dots \end{cases}$$

Whence if

$$G_{K_j} = \frac{1}{K} \sum_{k=1}^K \Lambda^k(\psi_{jk}, f), \quad k = 1, 2, 3, \dots,$$

it follows that the sequence $\{G_{Kj}\}_{K=1}^{\infty}$ converges uniformly to a function $H_j \in C$, $j = 1, 2, 3, \dots$.

Suppose now that $j_1 < j_2$. It follows then that

$$\begin{aligned} \Lambda^{j_2}(\psi_{j_1 j_2}, f) &= \Phi^{j_2 - j_1}(\Lambda^{j_1}(\psi_{j_1}, f)) \\ &\leq \Lambda^{j_2}(\varphi^{j_2}, f) = \Lambda^{j_2}(\psi_{j_2 j_2}, f). \end{aligned}$$

In particular, for $j_1 < j_2 < k$, we have, by induction on k , that

$$\Lambda^k(\psi_{j_1 k}, f) \leq \Lambda^k(\psi_{j_2 k}, f)$$

whence $H_{j_1} \leq H_{j_2}$. Since $H_{j_2} \leq \|f\|_{\infty} \cdot \max_{x \in \mathcal{S}_{N-1}} |f(x)|$ it follows that the sequence H_j converges uniformly to a function $H \in C$.

Consider now the diagonal sequence of functions

$$G_{KK} = \frac{1}{K} \sum_{k=1}^K \Lambda^k(\psi_{K, k}, f) = \frac{1}{K} \sum_{k=1}^K \Lambda^k(\varphi^k, f).$$

It is easy to see now that this sequence converges uniformly to H and we have thus obtained the desired result.

The following two theorems establish limit results for the sequence of functions $\{\Lambda^k(\varphi^k, f)\}_{k=1}^{\infty}$ under various hypotheses concerning the target motion matrix. The first theorem is most useful in the case of a stationary target, for then the target motion transition matrix is the identity matrix. The second theorem is applicable to many moving target situations, in particular to those presented in the previous section.

Theorem II-1. Let φ^K , $K = 0, 1, 2, \dots$ be the K -stage-optimal surveillance plan for an objective function $f \in C$. Let M be the target motion transition matrix and suppose that $U = \{X \in \mathcal{S}_{N-1}; MX = X\}$. Then the sequence $\{\Lambda^K(\varphi^K, f)\}_{K=1}^{\infty}$ converges uniformly on U .

Proof. Let $X \in U_1$; then

$$\begin{aligned} \Lambda^{K+1}(\varphi^{K+1}, f)(X) &\geq \Lambda^K(\varphi^K, f)(MX) \\ &\geq \Lambda^K(\varphi^K, f)(X), \end{aligned}$$

and so the sequence $\{\Lambda^K(\varphi^K, f)\}_{K=1}^\infty$ is an increasing sequence and so uniformly convergent on U .

Theorem II-5. Let φ^K , $K = 0, 1, 2, \dots$ be the K -stage optimal surveillance plan for the objective function $f \in C$. If the target motion matrix M has a row with no zero entries and if every entry of the sensor response matrix is less than 1, then the sequence $\{\Lambda^K(\varphi^K, f)\}_{K=1}^\infty$ converges uniformly to a constant.

Proof. Let $M = [d_{ij}]$ and suppose that $d_{i0j} \geq \alpha > 0$, $j = 1, 2, \dots, N$. If $X = (x_1, \dots, x_N)$ and $Y = (y_1, \dots, y_N) \in \mathcal{P}_{N-1}$, consider

$$\|MX - MY\|_1 = \sum_i \left| \sum_j d_{ij} (x_j - y_j) \right|.$$

Observe that

$$\begin{aligned} \left| \sum_j d_{i0j} (x_j - y_j) \right| &= \left| \sum_j (d_{i0j} - \alpha) (x_j - y_j) \right| \\ &\leq \sum_j (d_{i0j} - \alpha) |x_j - y_j| \\ &\leq \theta \sum_j d_{i0j} |x_j - y_j| \end{aligned}$$

for some θ , $0 < \theta < 1$, whence

$$\begin{aligned} \|MX - MY\|_1 &\leq \sum_j (1 - d_{i0j} + \theta d_{i0j}) |x_j - y_j| \\ &\leq (1 - (1 - \theta)\alpha) \|X - Y\|_1. \end{aligned}$$

But $1-(1-\theta)\alpha < 1$ so that M is a contraction mapping of \mathcal{S}_{N-1} . Thus there is a unique $X_0 \in \mathcal{S}_{N-1}$ such that

$$\lim_{K \rightarrow \infty} M^K X = X_0, \quad \text{all } X \in \mathcal{S}_{N-1}.$$

Now observe that $MX_0 = X_0$, whence from equation (II-12) we obtain

$$\Lambda^{K+1}(\varphi^{K+1}, f)(X_0) \geq \Lambda^K(\varphi^K, f)(MX_0) = \Lambda^K(\varphi^K, f)(X_0).$$

Thus the sequence of numbers $\{\Lambda^K(\varphi^K, f)(X_0)\}$ converges.

We will now show that

$$\lim_{K \rightarrow \infty} \sup\{|\Lambda^K(\varphi^K, f)(X) - \Lambda^K(\varphi^K, f)(Y)| : X, Y \in \mathcal{S}_{N-1}\} = 0.$$

To see this, define for $K = 1, 2, 3, \dots$ numbers L_K such that

$$|\Lambda^K(\varphi^K, f)(X) - \Lambda^K(\varphi^K, f)(Y)| \leq L_K \|X - Y\|_1, \quad X, Y \in V_N.$$

The existence of the numbers L_K was established in Theorem II-1. Suppose now that

$$\begin{aligned} \Lambda^{K+1}(\varphi^{K+1}, f)(X) &= \Lambda^K(\varphi^K, f)(MT_l X) + \Lambda^K(\varphi^K, f)(M(I-T_l)X) \\ &\geq \Lambda^{K+1}(\varphi^{K+1}, f)(Y). \end{aligned}$$

We then have

$$\begin{aligned}
& \Lambda^{K+1}(\varphi^{K+1}, f)(X) - \Lambda^{K+1}(\varphi^{K+1}, f)(Y) \\
& \leq \left| \Lambda^K(\varphi^K, f)(MT_l X) - \Lambda^K(\varphi^K, f)(M(T_l Y)) \right| \\
& \quad + \left| \Lambda^K(\varphi^K, f)(M(I-T_l)X) - \Lambda^K(\varphi^K, f)(M(I-T_l)Y) \right| \\
& \leq L_K \left\{ \|MT_l(X-Y)\|_1 + \|M(I-T_l)(X-Y)\|_1 \right\}.
\end{aligned}$$

But if $R = [r_{ij}]$, then by our assumption on R , there is a $\beta > 0$ such that

$$d_{lj}^{(1-r_{lj})} \geq \beta > 0, \quad j = 1, 2, 3, \dots, N.$$

We then obtain, as above, for some appropriately chosen μ , $0 < \mu < 1$,

$$\begin{aligned}
\|M(I-T_l)(X-Y)\|_1 & \leq \sum_j | \sum_{ij} d_{ij}^{(1-r_{lj})} (x_j - y_j) | \\
& \leq \mu \sum_j (1-r_{lj}) (x_j - y_j).
\end{aligned}$$

Again using our assumption that $r_{lj} < 1$, we have

$$\begin{aligned}
|\Lambda^{k+1}(\varphi^{k+1}, f)(X) - \Lambda^{k+1}(\varphi^{k+1}, f)(Y)| & = \sum_j (\mu + r_{lj}^{(1-\mu)}) |x_j - y_j| \\
& \leq \theta_0 L_K \|X-Y\|_1
\end{aligned}$$

for some appropriately chosen constant θ_0 , $0 < \theta_0 < 1$. Observe that θ_0 is independent of K , whence, by induction

$$L_K \leq \theta_0^K L_1 \rightarrow 0 \quad (K \rightarrow \infty).$$

Now since the sequence of numbers $\{\Lambda^K(\phi^K, f)(X_{\phi})\}_{K=1}^{\infty}$ is convergent and since $L_K \rightarrow 0, (K \rightarrow \infty)$, it follows as claimed that $\{\Lambda^K(\phi^K, f)(\cdot)\}$ is uniformly convergent to a constant.

An Important Special Case: The Homogeneous Sensor

In this section we describe an important special surveillance sensor which is similar to that encountered in a number of operational situations.

Consider a sensor which is only able to detect a target if it investigates the cell which contains it. Occasionally, however, the sensor will produce a non-target-generated contact. These contacts may occur, for example, because of random acoustic fluctuations, electronic instabilities, or sensor operator error.

Suppose that, if the sensor investigates the cell which contains the target, the probability that the sensor will obtain a target-generated contact is β . Additionally, suppose that the probability of obtaining a non-target-generated contact, μ , is independent of the cell investigated.

We seek the sensor response matrix for the case at hand. Since, if we search in a cell which does not contain the target, the only possible sensor response must be non-target generated, it follows from the definition of r_{ij} that $r_{ij} = \mu$ when $i \neq j$. On the other hand, if we investigate the cell which does contain the target, then we assume a sensor response may be obtained in two mutually exclusive ways: either a sensor response may be target generated or it may be non-target generated. Setting $\lambda = \beta + \mu$, it follows that $r_{ij} = \lambda, i = 1, 2, \dots, N$. The sensor response matrix is thus the $N \times N$ matrix

$$R = \begin{bmatrix} \lambda & \mu & \mu & \dots & \mu \\ \mu & \lambda & \mu & \dots & \mu \\ \mu & \mu & \lambda & \dots & \mu \\ \cdot & \cdot & \cdot & \dots & \cdot \\ \mu & \mu & \mu & \dots & \lambda \end{bmatrix} .$$

Such a sensor is called a homogeneous surveillance sensor with parameters (λ, μ) , $\lambda \geq \mu$. The quantity $\lambda - \mu$ is the target-generated contact rate and μ is the non-target-generated contact rate. In terms of valid contacts (i.e., contacts coming from the cell containing the target regardless of whether

they are target generated), λ is the valid contact rate and μ is the false contact rate.

For $l = 1, 2, \dots, N$ we now compute $U_l(X, k)$, the posterior target location probability distribution at the end of stage k given that the location distribution at the beginning of stage k is $X = (x_1, \dots, x_N)$ and that there was a sensor response from cell C_l . Let $e_l = (0, 0, \dots, 0, 1, 0, \dots, 0)$, where the 1 occurs in the l^{th} coordinate. It follows then from equation (II-1) that

$$U_l(X, k) = M \left[\frac{(\lambda - \mu)x_l}{(\lambda - \mu)x_l + \mu} e_l + \frac{\mu}{(\lambda - \mu)x_l + \mu} X \right],$$

and it is instructive to interpret the terms on the right-hand side of this identity. Observe that

$$\frac{(\lambda - \mu)}{(\lambda - \mu)x_l + \mu}$$

is the conditional probability, given that a contact occurred from investigating cell C_l , that the contact was target generated; also e_l is the posterior target location probability distribution given that the contact was target generated (for then the target must be in cell C_l). Also,

$$\frac{\mu}{(\lambda - \mu)x_l + \mu}$$

is the conditional probability, given that a contact occurred from investigating cell C_l , that the contact was non-target generated. In this case, the posterior target location probability distribution is X , the same as the prior, since no new information has been gained.

Similarly, we can compute $U_0(X, k)$, the posterior target location at the end of stage k given that the prior distribution at the beginning of stage k is X and that there were no sensor responses during stage k . For $j = 1, 2, \dots, N$, let $\alpha_j(X, k)$ be the amount of surveillance effort allocated in cell C_j during stage k , and set $\kappa = (\lambda - \mu) \sum_{j=1}^N \alpha_j$. We then obtain, from equation (II-1), that

$$U_0(\mathbf{X}, k) = M \frac{\sum_{j=1}^N \varphi_j(\mathbf{X}, k) [X - \kappa x_j e_j]}{\sum_{j=1}^N \varphi_j(\mathbf{X}, k) [1 - \kappa x_j]},$$

and the probability of obtaining no sensor response is

$$(1 - \mu) \left[\sum_{j=1}^N \varphi_j(\mathbf{X}, k) [1 - \kappa x_j] \right].$$

Consider now a surveillance sensor with the response matrix

$$\begin{bmatrix} \epsilon & 0 & 0 & \dots & 0 \\ 0 & \epsilon & 0 & \dots & 0 \\ 0 & 0 & \epsilon & \dots & 0 \\ & & & \dots & \\ 0 & 0 & 0 & \dots & \epsilon \end{bmatrix}.$$

Such a sensor is a homogeneous surveillance sensor with parameters $(\epsilon, 0)$. Every contact for such a sensor is target generated, and the occurrence of such a contact completely localizes the target to a single cell. If for $j = 1, 2, \dots, N$, $\varphi_j(\mathbf{X}, k)\Delta$ is the amount of surveillance effort allocated to cell C_j during stage k by such a sensor, the posterior distribution at the end of stage k given no contact is

$$\frac{\sum_{j=1}^N \varphi_j(\mathbf{X}, k) [X - \epsilon x_j e_j]}{\sum_{j=1}^N \varphi_j(\mathbf{X}, k) [1 - \epsilon x_j]},$$

and the probability of obtaining no sensor response is $\sum_{j=1}^N \varphi_j(\mathbf{X}, k) [1 - \epsilon x_j]$.

It follows now that a homogeneous surveillance sensor with parameters (λ, μ) achieves the same posterior distribution given no contact as does a homogeneous surveillance sensor with parameters $((\lambda - \mu)/(1 - \mu), 0)$ when both are used in the same surveillance plan. Moreover the ratio of their probabilities of obtaining no contact is $1 - \mu$.

The K-Stage Optimal Surveillance Plan for a Homogeneous Sensor

Consider a homogeneous surveillance sensor with parameters (λ, μ) , $\lambda \geq \mu$, which is attempting to localize a target located in one of N cells C_1, C_2, \dots, C_N . Suppose also that the target motion matrix M is given by the circulant matrix of the form, with $0 \leq \delta \leq 1$,

$$M = \begin{bmatrix} 1 - \frac{N-1}{N}\delta & \frac{\delta}{N} & \frac{\delta}{N} & & \frac{\delta}{N} \\ \frac{\delta}{N} & 1 - \frac{N-1}{N}\delta & \frac{\delta}{N} & & \frac{\delta}{N} \\ \frac{\delta}{N} & \frac{\delta}{N} & 1 - \frac{N-1}{N}\delta & & \frac{\delta}{N} \\ & & & \dots & \\ \frac{\delta}{N} & \frac{\delta}{N} & \frac{\delta}{N} & & 1 - \frac{N-1}{N}\delta \end{bmatrix}$$

Such matrices were originally considered as target transition matrices for surveillance problems by Richardson, reference [a].

The parameter δ is called the dispersion constant of the target motion matrix. Observe that if $\delta = 0$, then the target is stationary, whereas if $\delta = 1$ the target distribution disperses to the uniform distribution in one step. Additionally if $X = (x_1, \dots, x_N)'$ is the initial target location probability distribution, assuming that no surveillance effort is expended, the posterior target location distribution after k stages is

$$M^k X = (1 - \delta)^k \left(X - \left(\frac{1}{N}, \dots, \frac{1}{N} \right) \right) + \left(\frac{1}{N}, \dots, \frac{1}{N} \right).$$

In particular, if $\delta > 0$, the target location probability distribution converges uniformly to the uniform distribution.

The optimal whereabouts search introduced by Kadane in reference [b] is the special case where $\lambda > 0$, $\mu = 0$, $\delta = 0$, and the surveillance objective function is the probability that the target is located in the highest probability cell. Thus, in the situation considered by Kadane, the target is stationary and a sensor response can occur only if the target is located in the cell being investigated (i. e., there are no false contacts). Kadane was able to show that the optimal allocation of surveillance effort in an optimal whereabouts search is always to investigate the second highest probability cell. See, for example, section 4.4 of reference [c]. Note in particular that this optimal allocation depends only on the current target location probability distribution and not on the horizon, and is thus a stationary surveillance plan. Because of this, the resulting surveillance plan yields uniformly optimal probabilities of localizing the target for any possible horizon.

We now view the K-stage-optimal surveillance plan for the homogeneous sensor and circulant target motion matrix as a generalization of the optimal whereabouts search. Remarkably enough, our numerical results appear to indicate that the optimal allocation of surveillance effort when $\mu \geq 0$, $\delta \geq 0$, and the objective function is the probability that the target is located in the highest probability cell is the same as for Kadane's optimal whereabouts search. Thus for any homogeneous surveillance sensor, when used against a target whose motion matrix is circulant, we conjecture that the K-stage-optimal surveillance plan is the stationary plan which allocates the available effort at each stage to the second highest probability cell.

Another possible generalization of Kadane's optimal whereabouts search is to consider a different surveillance objective function. Suppose for example that we have a homogeneous surveillance sensor with parameters (λ, μ) , $\lambda \geq \mu \geq 0$, which is attempting to localize a target whose stochastic motion matrix is circulant with dispersion constant $\delta \geq 0$. If the surveillance problem involves N cells, C_1, C_2, \dots, C_N , we can consider, for $1 \leq n \leq N-1$, the surveillance objective function given by the probability that the target is located in the n highest probability cells. In this case, our numerical results indicate that the K-stage-optimal surveillance plan is the stationary plan which allocates at each stage the entire effort to the (n+1)th highest probability cell.

On the basis of these numerical results, we conjecture that in the general case of a homogeneous surveillance sensor and a target whose stochastic motion matrix is circulant, the K-stage-optimal surveillance plan, relative to the objective function which gives the probability that the target is located in the n highest probability cells, is the stationary plan which allocates at each stage the available effort to the (n+1)th highest probability cell.¹ If this conjecture is established theoretically, it will have a number of important consequences. First note that such a surveillance plan depends only on the current target location probability distribution and not on the number of stages

¹ Since this has been written, a counterexample to this conjecture has been found for $n = 2$ by J. R. Weisinger.

in the operation. This plan would then result in uniformly optimal probabilities of localizing the target for each possible choice of horizon. Moreover, we feel that it is reasonable to model many operational situations with a homogeneous sensor. Since in this case the K-stage-optimal surveillance plan would give optimal results for every horizon, it has potential for widespread applications.

In order to understand this conjecture, recall that the goal of surveillance is target localization rather than target detection. The surveillance objective function defines what is meant by target localization. Thus for example if we need only localize the target to n out of N cells, then the appropriate objective function is the probability that the target is located in the n highest probability cells, for if we use our surveillance sensor so as to maximize this quantity, we will maximize the chance of localizing the target to the specified number of cells. If we now allocate our surveillance effort at each stage to the $(n+1)^{\text{th}}$ highest probability cell, we are effectively using our effort to minimize the chance that our best estimate of n cells which contain the target is wrong.

Examples Involving a Homogeneous Surveillance Sensor

Suppose that a homogeneous sensor with parameters (λ, μ) is used against a stochastically moving target located in one of four cells C_1, C_2, C_3, C_4 . We assume that the target motion matrix is a circulant matrix with dispersion constant δ .

In Figures II-14 and II-15 we compare the effect of the dispersion constant δ on the K-stage-optimal surveillance plan using a homogeneous sensor with parameters $(.1, .01)$. In Figure II-14 the initial target location probability distribution is $(.7, .1, .1, .1)$. Observe that, in the case of a stationary target ($\delta = 0$), the probability of specifying the target's location to a single cell increases, as surveillance effort is expended, monotonically to 1.0. On the other hand, if $\delta > 0$ then


$$\lim_{k \rightarrow \infty} M^k X = (.25, .25, .25, .25).$$

so that if no surveillance is performed the target distribution asymptotically becomes uniform. Again, the parameter δ determines how rapidly this convergence occurs and if $\delta = 1$, then the target location distribution disperses to a uniform distribution after a single stage. This is illustrated in Figure II-14 by the fact that, if $\delta = 1$, the probability of correctly specifying the target's location after one or more stages of surveillance is .25. The other two cases considered correspond to dispersion constants of .2 and .4.

FIGURE II-11

THE INFLUENCE OF TARGET MOTION ON THE K-STAGE-OPTIMAL SURVEILLANCE PLAN

- Notes: 1) Homogeneous surveillance sensor with parameters (1, .01)
 2) Circulant target motion matrix with dispersion constant δ .
 3) Surveillance objective function: probability target is in the highest probability cell.
 4) Prior probability distribution (.7, .1, .1, .1)




K - Number of stages in surveillance operation

FIGURE II-15

THE INFLUENCE OF TARGET MOTION OF THE K-STAGE-OPTIMAL SURVEILLANCE PLAN

- Notes:
- 1) Homogeneous surveillance sensor with parameters (-1, .01)
 - 2) Circulant target motion matrix with dispersion constant δ .
 - 3) Surveillance objective function: probability target is in highest probability cell.
 - 4) Prior probability distribution (.3, .3, .2, .2)



K - Number of stages in surveillance operation

In Figure II-15 we have graphed the influence of target motion on the K-stage-optimal surveillance plan using the same sensor parameters as in Figure II-14 but assuming that the prior probability distribution is (.3, .3, .2, .2). The cases considered correspond to dispersion constants of $\delta = 0$ (stationary target), $\delta = .2$, $\delta = .4$, and $\delta = 1$ (complete dispersion in a single stage).

Finally, in Figures II-16 and II-17, we indicate the influence of the false contact rate on the K-stage-optimal surveillance plan. We consider three cases involving homogeneous surveillance sensors with parameters (λ, μ) where $\lambda = .1$ and μ takes the values .001, .01, and .05, respectively. We also suppose that the target motion matrix is a circulant matrix with dispersion constant $\delta = .2$. In Figure II-16 we assume that the prior target location probability distribution is (.7, .1, .1, .1) whereas in Figure II-17 it is (.3, .3, .2, .2).

Conclusion

Our investigations have established the K-stage-optimal surveillance plan as the solution of a particular dynamic programming problem. Unfortunately solving this dynamic programming problem generally requires a tremendous amount of computational effort. Moreover, except perhaps in the case of a homogeneous surveillance sensor, the K-stage-optimal surveillance plan depends strongly on the time horizon. Thus in general there exists no uniformly optimal surveillance plan.

Our numerical examples have shown that the 1-stage look-ahead maximum-information-gain plan, in addition to being easily determined, is a good suboptimal surveillance plan over a wide variety of surveillance objective functions. Moreover, since this plan is a stationary plan, it can be trivially extended to a surveillance plan of arbitrary length.


Other surveillance policies, such as the highest-probability-cell policy, can provide substantially less target localization than either the K-stage-optimal surveillance policy or the 1-stage look-ahead maximum-information-gain policy. The reason for this is that the highest probability cell policy ignores all knowledge concerning the response characteristics of the surveillance sensor as well as knowledge concerning the target's motion. Thus the highest probability cell policy may allocate its effort to an unproductive cell, with a corresponding penalty in target localization.

Our asymptotic results have established the existence of the limiting average payoff for the K-stage-optimal surveillance plan. Additionally, under certain assumptions concerning the target motion matrix, the expected target localization achieved by the K-stage-optimal surveillance plan converges to a limit which is independent of the initial target location distribution. Thus, in this case, the K-stage-optimal surveillance plan asymptotically produces estimates for the target's location which are robust against errors in the prior distribution.

FIGURE II-16

THE INFLUENCE OF FALSE CONTACT RATE ON THE K-STAGE-OPTIMAL SURVEILLANCE PLAN

- Notes:
- 1) Homogeneous surveillance sensor with parameter (λ, μ) , $\lambda = 1$.
 - 2) Circulant target motion matrix with dispersion constant $\delta = .2$
 - 3) Surveillance objective function: probability target is located in the highest probability cell.
 - 4) Prior probability distribution (.3, .3, .2, .2)




K - Number of stages in surveillance operation

FIGURE II-17

THE INFLUENCE OF FALSE CONTACT RATE ON THE K-STAGE-OPTIMAL SURVEILLANCE PLAN

- Notes:
- 1) Homogeneous surveillance sensor with parameter (λ, μ) , $\lambda = .1$
 - 2) Circulant target motion matrix with dispersion constant $\delta = .2$
 - 3) Surveillance objective function: probability target is located in the highest probability cell.
 - 4) Prior probability distribution $(.7, .1, .1, .1)$



K = number of stages in surveillance operation

Similar results are apparently valid, from our numerical examples, for stationary surveillance plans.

Another consequence of our numerical examples has been to show that if a homogeneous surveillance sensor is deployed against a target whose motion matrix is of a special type, then the resulting K-stage-optimal surveillance plan is stationary and of a particularly simple type. For example, if the surveillance objective is to localize the target to k cells, the optimal plan appears to be to investigate that $(k+1)^{\text{th}}$ highest probability cell.

CHAPTER III

A STATISTICAL MODEL FOR PROCESSING ASW CONTACT INFORMATION TO ESTIMATE TARGET PATTERNS OF OPERATION

In this chapter we consider two major problems which an ASW planner must frequently face in the presence of sparse contacts of various types and quality. The first is to obtain an estimate for the track of a specified target, and the second is to make inferences about overall target behavior patterns on the basis of contact data from a number of different targets. The purpose here is to describe a Bayesian method for obtaining both of these types of estimates.

Unfortunately, the methods in this chapter have not yet been developed to the point where numerical examples have been computed, and so none are included.

Our approach is based on a parametric model for target motion. The object is to use the available contact data to obtain posterior estimates for the target's track as well as the parameters which describe target motion. A major consideration in what follows is the development of a parametric model for target motion which is rich enough to model real world situations but which is also computationally tractable.

The approach outlined below is most applicable to the case of a target in transit. Additionally, since we are obtaining Bayesian estimates we are required to assume the existence of patterns of motion for which we have a reasonable prior knowledge. This prior knowledge may take the form of past operational experience, or may be a consequence of certain operational or geographical constraints.

In the first section of this chapter, we describe our parametric model target motion. The model is based on the notion of a number of different target scenarios or basic target tracks. It is assumed in this section that the parameters for the motion model are completely known.

In the second section we suppose that the parameters for the target motion model are unknown but that prior information about their possible values has been quantified in the form of a probability distribution. This probability distribution is the key item required to perform the Bayesian updating for the parameters of the target motion model.

We describe in the third and fourth sections the calculations necessary to perform the Bayesian updating on the probability distributions for the target motion parameters so as to account for the information provided by a completely known target track. Although it is unlikely that such high quality information would ever be available, this is the simplest case to consider, and the calculations are suggestive of the direction taken in subsequent sections.

The fifth section describes our model for target contacts. We assume that each contact provides us with an estimate for the target's actual position together with a covariance matrix which represents the uncertainty in the contact data. The next section describes the calculations necessary to marry the parametric target motion model with contact data so as to obtain a Bayesian estimate for the track of the specified target. This estimate consists of posterior scenario weights, a mean target track for each scenario, and associated covariance matrices.

The calculations necessary to perform the Bayesian updating for the distribution on parameters of the target motion model so as to account for contact information on a given target are described in the last two sections of this chapter. The calculations here are strongly motivated by those given in the third and fourth sections.

A Parametric Model for Target Motion

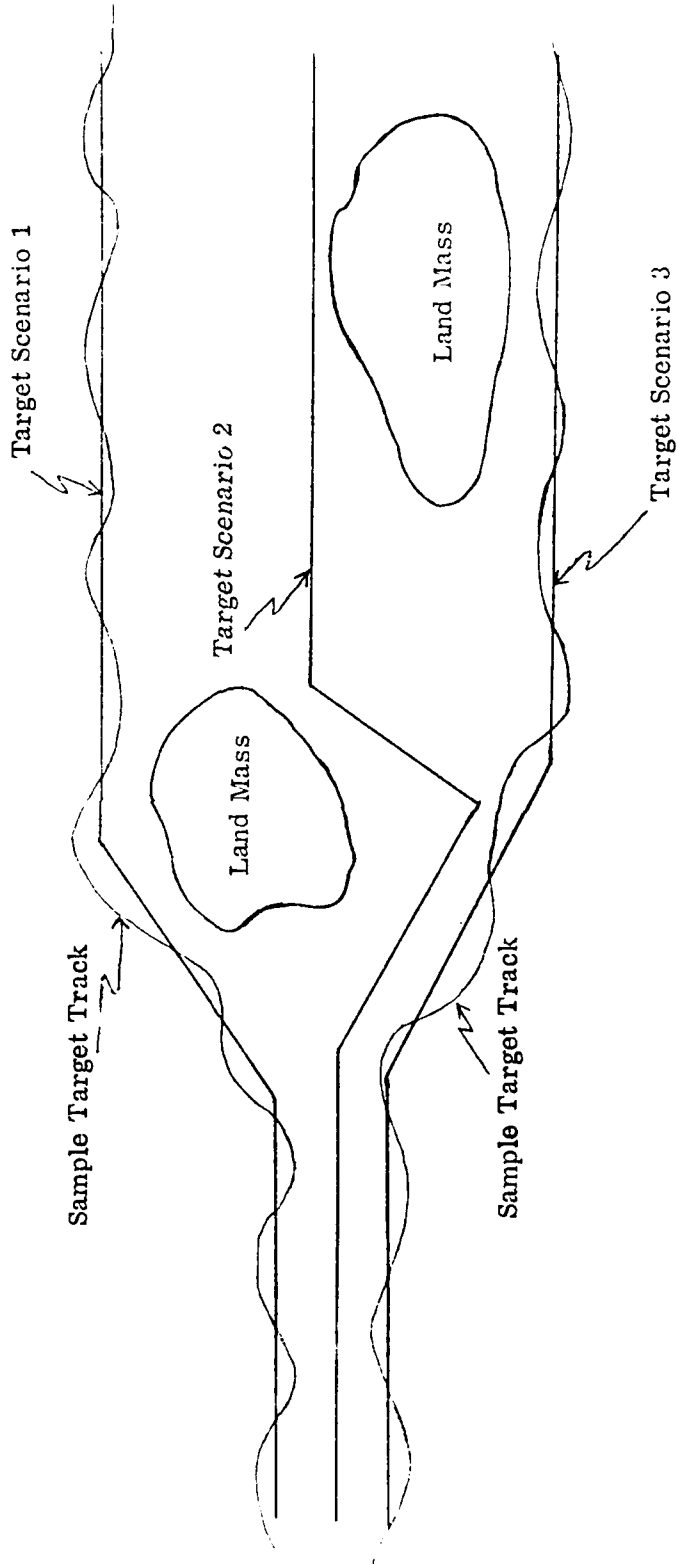
In this section we describe a parametric model for target motion. This model is based on the notion of target scenarios. We postulate the existence of a finite number of such scenarios, or operating plans, which a target might follow during a specified phase of its mission. Each scenario may be thought of as corresponding to a basic geometric pattern of target motion. We assume that each target chooses one of these basic patterns and follows it throughout the specified mission phase.

Each scenario is determined by a mean target track and corresponding covariance matrix. Once a scenario has been chosen for a given target, the target must move roughly according to the mean track of the scenario. The target, however, is permitted to operate with some deviation from the mean track. For example, the target may move faster or more slowly than the specified mean track, or it may vary its course along the mean track. The extent of these perturbations in target motion is determined by the covariance matrix associated with the scenario.

An example of some basic tracks which might be used to define a scenario is given in Figure III-1. Additionally, a number of target tracks drawn from the target scenarios are indicated in Figure III-1.

FIGURE III-1

EXAMPLE OF TARGET SCENARIOS AND SAMPLE TARGET TRACKS



In order to formulate precisely our parametric model for target motion, suppose that there are J possible target scenarios: S_1, S_2, \dots, S_J . Let \mathcal{J} be the random variable with range in the set $\{1, 2, \dots, J\}$ which specifies the scenario a target will follow. Thus the event $\{\mathcal{J} = j\}$ represents all target tracks which follow scenario S_j . We let p_j be the probability that a target will follow scenario S_j :

$$\Pr \{ \mathcal{J} = j \} = p_j.$$

We now suppose that a target track is specified by the position of the target at times $t_1 < t_2 < \dots < t_\tau$. Thus each target track is a sequence $Z = (z_1, \dots, z_\tau)$ of target positions $z \in \mathbb{R}^2$. Let Z be the 2τ dimensional random vector which specifies a target's track. The conditional distribution function of Z , given that it represents the track of a target operating according to scenario S_j , is assumed to be a multivariate normal distribution with mean μ_j and covariance matrix Σ_j^{-1} .

In order to simplify the following discussion, we will use the notation $f(\cdot | \mu, \Sigma)$ to indicate a normal density function with mean μ and covariance matrix Σ^{-1} . The matrix Σ is called the precision matrix of the distribution. In this notation we have, for a measurable subset U of $\mathbb{R}^{2\tau}$,

$$\Pr \{ Z \in U | \mathcal{J} = j \} = \int_U f(z | \mu_j, \Sigma_j) dz.$$

The Parametric Target Motion Model in the Face of Uncertainty

We now assume that the parameters in the target motion model described above are uncertain, but that we have prior knowledge about their possible values which we will quantify as probability distributions. In particular, we assume imprecise knowledge about:

- i) the vector of probabilities $p = (p_1, p_2, \dots, p_J)$ which describes the distribution of the random variable \mathcal{J} ,
- ii) the mean target paths $\mu_1, \mu_2, \dots, \mu_J$, and
- iii) the covariance matrices $\Sigma_1^{-1}, \Sigma_2^{-1}, \dots, \Sigma_J^{-1}$.

We do assume, however, that the number of possible target scenarios, J , has already been established. Our goal now is to use contact data on a number of different targets to improve our knowledge of these parameters.

We now replace each of the above parameters with a random variable. Let Q be a random vector with values in the J -fold Cartesian product $[0, 1] \times \dots \times [0, 1]$, such that the coordinates of Q sum to 1. Let $\Gamma_1, \dots, \Gamma_J$ be random vectors with values in $\mathbb{R}^{2\tau}$, and let A_1, \dots, A_J be random matrices with values in the cone of positive definite, symmetric $2\tau \times 2\tau$ dimensional real-valued matrices. We will use Q to represent our state of knowledge concerning the vector of probabilities p ; $\Gamma_1, \dots, \Gamma_J$ to represent our state of knowledge concerning the mean target tracks μ_1, \dots, μ_J ; and A_1, \dots, A_J to represent our state of knowledge concerning the precision matrices $\Sigma_1, \dots, \Sigma_J$. Note that, although our primary interest is in obtaining estimates for the covariance matrices $\Sigma_1^{-1}, \dots, \Sigma_J^{-1}$, for technical reasons we use the random variables A_1, \dots, A_J to represent our state of knowledge about the precision matrices $\Sigma_1, \dots, \Sigma_J$.

Now let pr be the prior joint probability density function for the random quantities $Q, \Gamma_1, \dots, \Gamma_J$, and A_1, \dots, A_J . The prior density function pr must be chosen with considerable care. Indeed the ease with which we will ultimately obtain estimates for p, μ_1, \dots, μ_J , and $\Sigma_1, \dots, \Sigma_J$ depends strongly on the form of the density function pr . In particular, it is important that the posterior distribution for $Q, \Gamma_1, \dots, \Gamma_J$, and A_1, \dots, A_J , given a certain collection of contact data, be of the same form as the prior density pr . Finally, to facilitate computations, it is necessary that the density function pr have a simple form which can be completely described with as few parameters as possible.

We have yet to find a density function pr which completely satisfies all of the above described constraints. It is possible, however, to formulate a prior density function pr which fulfills our demands reasonably well.

We will assume that the parameters which describe a specified scenario are independent from one another as well as independent from the scenario weights. If we then let pr_0 be the marginal density function for Q , and pr_j , $j = 1, 2, \dots, J$, be the joint marginal density function for Γ_j and A_j , by our independence assumption we can write pr as the product

$$pr(q, \gamma_1, \dots, \gamma_J, a_1, \dots, a_J) = pr_0(q) \prod_{j=1}^J pr_j(\gamma_j, a_j). \quad (\text{III-1})$$

We now assume that pr_0 , the density function for Q , is a Dirchlet distribution with parametric vector $\alpha = (\alpha_1, \dots, \alpha_J)$, $\alpha_j > 0$, $j = 1, 2, \dots, J$. Using the notation $g(\cdot | \alpha)$ to indicate such a density function, $g(q | \alpha)$ then has the form

$$g(q | \alpha) = \frac{\Gamma(\alpha_1 + \dots + \alpha_J)}{\Gamma(\alpha_1) \dots \Gamma(\alpha_J)} q_1^{\alpha_1 - 1} q_2^{\alpha_2 - 1} \dots q_J^{\alpha_J - 1}, \text{ with } 0 \leq q_i \leq 1, \\ 1 \leq i \leq J.$$

We now describe the density functions pr_j , $j = 1, 2, \dots, J$. We assume that the conditional distribution of Γ_j , given that $A_j = a$, is a multivariate normal

distribution with mean vector $\hat{\gamma}_j \in \mathbb{R}^{2\tau}$ and covariance matrix $(v_j a)^{-1}$ where $v_j > 0$. Additionally we assume that the marginal distribution of A_j is a Wishart distribution with m_j degrees of freedom, $m_j > 2\tau - 1$, and precision matrix V_j , where V_j is a $2\tau \times 2\tau$ positive definite symmetric matrix. The Wishart distribution is extensively discussed in references [f] and [g].

We will use the notation $h(\cdot | m, V)$ to represent the density function of a Wishart distribution with m degrees of freedom and precision matrix V . In general then, for some appropriately chosen constant c , we have, for all $2\tau \times 2\tau$ positive definite symmetric matrices a ,

$$h(a | m, V) = c |V|^{m/2} |a|^{-(m-2\tau-1)/2} \exp\{-\frac{1}{2}\text{tr}(Va)\},$$

where $|V|$ = determinant of V . Thus, the joint density function for Γ_j and A_j has the form

$$\text{pr}_j(\gamma, a) = f(\gamma | \hat{\gamma}_j, v_j a) h(a | m_j, V_j), \quad j = 1, 2, \dots, J.$$

We now indicate the significance of the parameters which define the prior joint density pr . Note that the prior choice for the values of these parameters must be based on past experience or operational and geographic constraints.

The choice of the prior parametric vector α is governed by our prior estimates on the relative likelihoods of the scenarios S_1, \dots, S_J . Indeed if $\beta = \alpha_1 + \dots + \alpha_J$ then

$$E(Q) = \int q \text{pr}_0(q) dq = \alpha \beta^{-1},$$

and so the components of α are precisely our prior estimates for the relative likelihoods of the various scenarios.

In order to obtain the proper prior values for the parameters $\hat{\gamma}_j$ and V_j , we consider the prior conditional density, pr , on target tracks Z given that they are samples from scenario S_j . Observe that

$$\begin{aligned} \text{pr}(z | \mathcal{J} = j) &= \int \int \text{pr}(z | \mathcal{J} = j, \Gamma_j = \gamma, A_j = a) \text{pr}_j(\gamma, a) d\gamma da \\ &= \int f(z | \gamma, a) f(\hat{\gamma}_j, (v_j a)) h(a | m_j, V_j) d\gamma da \\ &= \int f(z | \hat{\gamma}_j, \frac{v_j}{1+v_j} a) h(a | m_j, V_j) da. \end{aligned} \tag{III-2}$$

To evaluate this last integral we need to define the t distribution with $m > 0$ degrees of freedom, location vector $\sigma \in \mathbb{R}^{2\tau}$, and positive definite symmetric $2\tau \times 2\tau$ precision matrix T , $t(x | m, \sigma, T)$. This density is defined for all $x \in \mathbb{R}^{2\tau}$ by the formula

$$t(x | m, \sigma, T) = c |T|^{\frac{1}{2}} \left[1 + \frac{1}{m} (x-\sigma)' T (x-\sigma)\right]^{-(m+2\tau)/2},$$

where c is a constant chosen so that the density integrates to 1. Observe that, for all sufficiently large m , the t distribution is approximately a normal distribution:

$$t(x | m, \sigma, T) \approx f(x | \sigma, T).$$

It is now an easy matter to evaluate the integral on the right hand side of equation (III-2) to obtain

$$\begin{aligned} \text{pr}(z | \mathcal{J} = j) &= t(z | m_j - 2\tau + 1, \hat{\gamma}_j, \frac{v_j}{1+v_j} (m_j - 2\tau + 1) V_j^{-1}) \\ &\approx f(z | \hat{\gamma}_j, \frac{v_j}{1+v_j} (m_j - 2\tau + 1) V_j^{-1}), \end{aligned} \tag{III-3}$$

where the last approximation is valid provided m_k is sufficiently large. Thus, our prior estimate for the probability distribution on target tracks which are associated with scenario S_j is approximately a multivariate normal distribution with mean track $\hat{\gamma}_j$ and covariance matrix

$$\hat{V}_j = \frac{v_j}{v_j + 1} (m_j - 2\tau + 1) V_j^{-1}.$$

In the theory of Bayesian statistical analysis the Bayes estimate for a parameter is the expectation of the random variable representing that parameter. Thus, the prior Bayes estimate for the vector of the probabilities p is

$$E(Q) = \alpha \beta^{-1}, \quad \beta = \alpha_1, \dots, \alpha_J,$$

and the prior Bayes estimate for the mean target track μ_j , of scenario S_j is

$$\hat{\gamma}_j = E(\Gamma_j), \quad j = 1, 2, \dots, J.$$

It is also possible to show that the prior Bayes estimates for the covariance matrix Σ_j^{-1} and precision matrix Σ_j are respectively

$$E(A_j^{-1}) = \frac{1}{m_j - 2\tau + 1} V_j,$$

$$= \frac{v_j}{v_j + 1} \hat{V}_j^{-1}, \text{ and}$$

$$E(A_j) = m_j V_j^{-1}.$$

See for example reference [g].

Estimating Target Motion Parameters When the Track and Motion Scenario Are Known

In this section we indicate how to perform a Bayesian update on the probability distribution for the parameters of our target motion model so as to reflect the information provided by a completely known target track. We thus assume that we are considering a target

- 1) which is following scenario S_l (so that $J = l$), and
- 2) whose track is $z = (z_1, z_2, \dots, z_\tau)$.

In this case we will be able to compute precisely the posterior distribution for the parameters of our target motion model. In the other cases which we will consider, it will be necessary to make a number of approximations to perform the updating.

Let

$$\text{pr}(q, \gamma_1, \dots, \gamma_J, a_1, \dots, a_J) = \text{pr}_0(q) \prod_{j=1}^J \text{pr}_j(\gamma_j, a_j)$$

be the joint prior probability distribution for the random quantities $q, \gamma_1, \dots, \gamma_J$ and A_1, \dots, A_J respectively. We are assuming, as in the case of

$$\text{pr}_0(q) = g(q | \alpha), \text{ and}$$

$$\text{pr}_j(\gamma, a) = f(\gamma | \hat{\gamma}_j, v_j a) h(a | m_j, V_j).$$

We seek the joint posterior density

$$\begin{aligned} & \text{pr}(q, \gamma_1, \dots, \gamma_J, a_1, \dots, a_J | z, \mathcal{J} = l) \\ &= \frac{\text{pr}(z | q, \gamma_1, \dots, \gamma_J, a_1, \dots, a_J, \mathcal{J} = l) \text{Pr}\{\mathcal{J} = l | q, \gamma_1, \dots, \gamma_J, a_1, \dots, a_J\}}{\text{pr}(z | \mathcal{J} = l) \text{Pr}\{\mathcal{J} = l\}} \\ & \quad \times \text{pr}_0(q) \prod_{j=1}^J \text{pr}_j(\gamma_j, a_j). \end{aligned} \quad (\text{III-4})$$

In order to evaluate this quantity, recall from the previous section that

$$\begin{aligned} & \text{pr}(z | q, \gamma_1, \dots, \gamma_J, a_1, \dots, a_J, \mathcal{J} = l) = f(z | \gamma_l, a_l), \\ & \text{Pr}\{\mathcal{J} = l | q, \gamma_1, \dots, \gamma_J, a_1, \dots, a_J\} = q_l, \\ & \text{pr}(z | \mathcal{J} = l) \approx f(z | \hat{\gamma}_l, \hat{V}_l), \text{ and} \\ & \text{Pr}\{\mathcal{J} = l\} = \alpha_l / \beta, \quad \beta = \alpha_1 + \dots + \alpha_J. \end{aligned} \quad (\text{III-5})$$

Substituting these quantities into equation (III-4) we obtain

$$\text{pr}(q, \gamma_1, \dots, \gamma_J, a_1, \dots, a_J | z, \mathcal{J} = l) \propto \beta \alpha_l^{-1} q_l f(z | \gamma_l, a_l) \text{pr}_0(q) \prod_{k=1}^K \text{pr}_k(\gamma_k, a_k).$$

Observe now that

$$\beta \alpha_l^{-1} q_l \text{pr}_0(q) = \beta \alpha_l^{-1} q_l g(q | \alpha) = g(q | \alpha^*)$$

where $\alpha^* = (\alpha_1, \dots, \alpha_l + 1, \dots, \alpha_J)$. Thus, the posterior distribution for the random variable Q is a Dirichlet distribution which depends only on the scenario followed by the target in question.

We now consider the product

$$\begin{aligned} f(z | \gamma, a) \text{pr}_l(\gamma, a) &= f(z | \gamma, a) f(\gamma | \hat{\gamma}_l, v_l, a) h(a | m_l, V_l) \\ &= f(\gamma | \hat{\gamma}_l^*, (v_l + 1)a) f(z | \hat{\gamma}_l, \frac{v_l}{1+v_l} a) h(a | m_l, V_l), \end{aligned} \quad (\text{III-6})$$

where

$$\hat{\gamma}_l^* = (v_l \hat{\gamma}_l + z) (v_l + 1)^{-1}.$$

Using the identity

$$\frac{v_l}{1+v_l} (z - \hat{\gamma}_l)' a (z - \hat{\gamma}_l) = \frac{v_l}{1+v_l} \text{tr} (z - \hat{\gamma}_l) (z - \hat{\gamma}_l)' a,$$

we can write the second normal density which occurs on the right hand side of equation (III-6) as

$$(2\pi)^{-2\tau} \frac{v_l}{1+v_l} a_l e^{-\frac{1}{2} \frac{v_l}{v_l+1} \text{tr} \{ (z - \hat{\gamma}_l) (z - \hat{\gamma}_l)' a \}}.$$

Combining this with the expression for the Wishart density function which appears in (III-6), we thus obtain

$$f(z | \gamma, a) \text{pr}_l(\gamma, a) \propto f(\gamma | \hat{\gamma}_l^*, (v_l + 1)a) h(a | m_l + 1, \hat{V}_l^*),$$

where

$$\hat{V}_l^* = V_l + \frac{v_l}{1+v_l} (z - \hat{\gamma}_l) (z - \hat{\gamma}_l)'.$$

We conclude finally that the posterior joint distribution for the random variables $Q, \Gamma_1, \dots, \Gamma_K, A_1, \dots, A_K$, given that the target's track is $z = (z_1, \dots, z_T)$ and that it is following scenario S_l , is

$$\begin{aligned} \text{pr}(q, \gamma_1, \dots, \gamma_J, a_1, \dots, a_J \mid z, \mathcal{J} = l) \\ = g(q \mid \alpha^*) \prod_{j=1}^J f(\gamma_j \mid \hat{\gamma}_j^*, v_j^* a_j) h(a_j \mid m_j^*, V_j^*), \end{aligned} \quad (\text{III-7})$$

where

$$\alpha^* = \alpha + (0, 0, \dots, 1, \dots, 0), \quad (1 \text{ appears in the } l^{\text{th}} \text{ component})$$

$$\hat{\gamma}_j^* = \begin{cases} \hat{\gamma}_j & \text{if } j \neq l \\ \frac{1}{v_l + 1} (v_l \hat{\gamma}_l + z) & \text{if } j = l, \end{cases}$$

$$v_j^* = \begin{cases} v_j & \text{if } j \neq l \\ v_l + 1 & \text{if } j = l, \end{cases}$$

$$m_j^* = \begin{cases} m_j & \text{if } j \neq l \\ m_l + 1 & \text{if } j = l, \end{cases}$$

$$V_j^* = \begin{cases} V_j & \text{if } j \neq l \\ V_l + \frac{v_l}{v_l + 1} (z - \hat{\gamma}_l) (z - \hat{\gamma}_l)' & \text{if } j = l. \end{cases}$$

In particular observe, in this case, that the posterior distribution for the target motion parameters has precisely the same form as does the prior.

Estimating Target Motion Parameters When the Track Is Known but the Motion Scenario Is Unknown

Suppose that we consider a target whose track $z = (z_1, z_2, \dots, z_T)$ is completely known, but we have yet to determine which of the motion scenarios S_1, S_2, \dots, S_J the target is following. In this case we would like to estimate the joint posterior distribution for the target motion parameters given the target track z . Observe that we can write

$$\begin{aligned} & \text{pr}(q, \gamma_1, \dots, \gamma_J, a_1, \dots, a_J \mid z) \\ &= \sum_{l=1}^J \text{pr}(q, \gamma_1, \dots, \gamma_J, a_1, \dots, a_J \mid z, \mathcal{J}=l) \text{Pr}\{\mathcal{J}=l \mid z\}. \end{aligned} \quad (\text{III-8})$$

The first factor of each term appearing on the right hand side of equation (III-8) has already been determined in equation (III-7). We thus need only evaluate the conditional probability that the target is following scenario S_l given that the target's track is z . To do this we employ Bayes theorem to write

$$\text{Pr}\{\mathcal{J}=l \mid z\} = \frac{\text{pr}(z \mid \mathcal{J}=l) \text{Pr}\{\mathcal{J}=l\}}{\sum_{j=1}^J \text{pr}(z \mid \mathcal{J}=j) \text{Pr}\{\mathcal{J}=j\}}.$$

Using now our parametric model for target motion, we can evaluate the quantities which appear on the right hand side of this identity. We thus obtain

$$\text{Pr}\{\mathcal{J}=l \mid z\} = \alpha_l f(z \mid \hat{\gamma}_l, \hat{V}_l) \left[\sum_{k=1}^K \alpha_k f(z \mid \hat{\gamma}_k, \hat{V}_k) \right]^{-1}. \quad (\text{III-9})$$

Combining equations (III-7), (III-8), and (III-9), we are able to compute the joint posterior distribution for our target motion parameters given only the target track z . Observe, however, that this posterior has a different form than the prior. Additionally, this posterior involves many more parameters than does the prior, and this unfortunate state of affairs can rapidly make the necessary calculations impractical.

The cause of this problem, of course, is the impossibility of assigning, without additional data, a given track to a unique scenario. One possible solution would be to use equation (III-9) to determine the highest probability scenario, and then assign the target track z to that scenario. The calculation of the posterior joint distribution for the target motion parameters would then be carried out as in the previous section. Such a scheme is particularly meaningful when the

scenarios are widely separated in space and the target track z fits one scenario significantly better than any other scenario.

Although we are unable to compute the posterior joint distribution $\text{pr}(q, \gamma_1, \dots, \gamma_J, a_1, \dots, a_J | z)$ without either an increase in the number of parameters or some maximum likelihood estimating scheme, nevertheless we can give a useful approximation for the marginal posterior distributions $\text{pr}_0(q | z)$. Suppose that $\text{pr}_0(q) = g(q | \alpha)$. We then have

$$\begin{aligned} \text{pr}_0(q | z) &= \sum_{j=1}^J \text{pr}_0(q | z, \mathcal{J} = j) \text{Pr} \{ \mathcal{J} = j | z \} \\ &= \sum_{j=1}^J g(q | \alpha^j) \text{Pr} \{ \mathcal{J} = j | z \}, \end{aligned} \tag{III-10}$$

where

$$\alpha^j = \alpha + (0, 0, \dots, 1, \dots, 0) \quad (1 \text{ appears in the } j^{\text{th}} \text{ coordinate}).$$

Observe, however, that if we set

$$\alpha^* = \sum_{j=1}^J \alpha^k \text{Pr} \{ \mathcal{J} = j | z \},$$

then the Dirchlet distribution $g(\cdot | \alpha^*)$ has the same mean as does the sum of Dirchlet distributions (III-10). Thus, in many applications, it may be reasonable to approximate the posterior distribution $\text{pr}_0(q | z)$ with the Dirchlet distribution $g(q | \alpha^*)$. This approximation of course prevents the increase in the number of parameters that occur in equation (III-10).

Target Contacts

It seldom happens that the complete track of a target is available for analysis. Indeed it may not even be known which of the various scenarios S_1, \dots, S_J a given target is following. Rather, it is much more common to have available a collection of target contact data on which all conclusions about target motion must be based. The purpose of this section is to describe our model for target contacts.

Suppose that the track of a specified target is $z = (z_1, \dots, z_T) \in \mathbb{R}^{2T}$, and let D_{tk} be the k^{th} contact event on the target at time t . Each contact event consists of

an estimated position for the target at time t and a known covariance matrix Δ_{tk} . We assume that the estimated position d_{tk} for the target, given that the contact occurred, is a random vector which depends only on the actual position z_t of the target at time t . Moreover, the distribution of the random vector d_{tk} is a bivariate normal distribution with mean z_t and precision matrix Δ_{tk} . Additionally, we assume that the random vectors d_{tk} are all mutually independent. Thus the joint probability distribution of the random vectors d_{tk} given that $Z = (z_1, z_2, \dots, z_\tau)$ is

$$\prod_{t,k} f(d_{tk} | z_t, \Delta_{tk}).$$

Here, for $z \in \mathbb{R}^2$ and Δ a 2×2 positive definite matrix, we use $f(\cdot | z, \Delta)$ to represent the bivariate normal density function.

Estimating the Track of a Target from Contact Data

Suppose that we are interested in obtaining an estimate for the track of a specified target from contact data. Recall that the track of this target is a random vector Z with values in $\mathbb{R}^{2\tau}$. Our prior estimate for the distribution of Z depends on the parameters which describe the uncertainty in our model for target motion. Thus let $Q, \Gamma_1, \dots, \Gamma_J$ and A_1, \dots, A_J are the random quantities which represent our state of knowledge about the vector of scenario weights $p = (p_1, \dots, p_J)$, the mean target paths μ_1, \dots, μ_J , and the scenario precision matrices $\Sigma_1, \dots, \Sigma_J$ respectively. We assume that the joint distribution for these quantities has the density function pr given by equation (III-1) where $pr_j(q) = g(q | \alpha)$ and, for $j = 1, 2, \dots, J$,

$$pr_j(\gamma, a) = f(\gamma | \hat{\gamma}_j, (v_j, a)) h(a | m_j, V_j).$$

It follows from equations (III-2) and (III-3) that our prior estimate for the distribution for the track Z is given by the density function pr :

$$\begin{aligned} pr(z) &= \sum_{j=1}^J pr(z | \mathcal{J} = j) Pr\{\mathcal{J} = j\} \\ &\approx \sum_{j=1}^J f(z | \hat{\gamma}_j, \hat{V}_j) \alpha_j \beta^{-1}, \end{aligned}$$

where $\beta = \alpha_1 + \dots + \alpha_J$ and

$$\hat{V}_j = \frac{v_j}{1+v_j} (m_j - 2\tau + 1) V_j^{-1}.$$

Suppose now that we have obtained contact events D_{tk} on the target. We thus have the position estimates $d_{tk} = \xi_{tk}$, with corresponding precision matrices Δ_{tk} , at various times t , $1 \leq t \leq \tau$. In order to simplify our notation, let D be the joint contact event

$$D = \{d_{tk} = \xi_{tk}, \text{ all } t, k\}.$$

We seek then the posterior distribution

$$\text{pr}(z | D) = \text{pr}(D)^{-1} \sum_{j=1}^J \text{pr}(D | z, \mathcal{J}=j) \text{pr}(z | \mathcal{J}=j) \text{Pr}\{\mathcal{J}=j\}. \quad (\text{III-11})$$

Since the distribution of d_{tk} depends only on the actual position of the target at time t , and not on the target scenario, we have

$$\text{pr}(D | z, \mathcal{J}=j) = \text{pr}(D | z) = \prod_{t,k} f(\xi_{tk} | z_k, \Delta_{tk}).$$

For each time t at which we have at least one contact on the target, we can write, by the process of completing the square, described in Appendix A,

$$\sum_k (\xi_{tk} - z_t)' \Delta_{tk} (\xi_{tk} - z_t) = (z_t - \sigma_t)' \Lambda_t (z_t - \sigma_t) + L_t$$

where L_t is a constant independent of $z = (z_1, \dots, z_\tau)$, Λ_t is the 2×2 positive definite symmetric matrix

$$\Lambda_t = \sum_k \Delta_{tk},$$

and σ_t is the 2-vector, independent of z , given by

$$\sigma_t = \Lambda_t^{-1} \left(\sum_k \Delta_{tk} \xi_{tk} \right).$$

Note that these calculations require nothing more difficult than inverting a 2 x 2 matrix. Also if there are no contacts at time t, let Λ_t be the 2 x 2 zero matrix, $\sigma_t = (0, 0)$, and $L_t = 0$.

Now let Λ be the $2\tau \times 2\tau$ block diagonal matrix

$$\Lambda = \text{diag} (\Lambda_1, \Lambda_2, \dots, \Lambda_\tau),$$

let $\sigma \in \mathbb{R}^{2\tau}$ be the vector

$$\sigma = (\sigma_1', \sigma_2', \dots, \sigma_\tau')',$$

and let $L = L_1 + \dots + L_\tau$. We can then write

$$\text{pr}(D | z, \mathcal{J} = j) = C \sqrt{\prod_{tk} |\Delta_{tk}|} e^{-\frac{1}{2}(z-\sigma)' \Lambda (z-\sigma)}, \quad (\text{III-12})$$

where C is the constant $(2\pi)^{-\mu} e^{-L}$, and μ is the total number of contacts. Observe that C is independent of the target track z as well as the scenario followed by the target.

At this point it is convenient to compute the probability of obtaining the collection of contact data D given that the target is following scenario S_j . Indeed, from equation (III-12) we have

$$\begin{aligned} \text{pr}(D | \mathcal{J} = j) &= \int \text{pr}(D | z, \mathcal{J} = j) \text{pr}(z | \mathcal{J} = j) dz \\ &= C_j e^{-\frac{1}{2}(\sigma - \hat{\gamma}_j)' \hat{V}_j (\hat{V}_j + \Lambda)^{-1} \Lambda (\sigma - \hat{\gamma}_j)}, \end{aligned} \quad (\text{III-13})$$

where

$$C_j = \sqrt{|\hat{V}_j| |\hat{V}_j + \Lambda|^{-1} \prod_{tk} |\Delta_{tk}|} C.$$

Also, since the prior probability that the target will follow scenario S_j is $\text{Pr}\{\mathcal{J} = j\} = \alpha_j \beta^{-1}$, ($\beta = \alpha_1 + \dots + \alpha_J$), we have

$$\text{pr}(D) = \beta^{-1} \sum_{j=1}^J \text{pr}(D | \mathcal{J} = j) \alpha_j.$$

Consider now the product

$$\text{pr}(D | z, \mathcal{J} = j) \text{pr}(z | \mathcal{J} = j) = C \sqrt{\prod_{tk} |\Delta_{tk}|} e^{-\frac{1}{2}(z-\sigma)' \Lambda (z-\sigma)} f(z | \hat{\gamma}_j, \hat{V}_j).$$

In order to compute this density note that again by completing the square of a quadratic form, we can write

$$\begin{aligned} (z - V_j)' \hat{V}_j (z - V_j) + (z - \sigma)' \Lambda (z - \sigma) \\ = (z - \delta_j)' U_j (z - \delta_j) + (\sigma - \hat{\gamma}_j)' \hat{V}_j (\hat{V}_j + \Lambda)^{-1} \Lambda (\sigma - \hat{\gamma}_j), \end{aligned}$$

where

$$\delta_j = (\hat{V}_j + \Lambda)^{-1} (\hat{V}_j \hat{\gamma}_j + \Lambda \sigma).$$

then have

$$\begin{aligned} \text{pr}(D | z, \mathcal{J} = j) \text{pr}(z | \mathcal{J} = j) \\ = C_j f(z | \delta_j, U_j) e^{-\frac{1}{2}(\sigma - \hat{\gamma}_j)' \hat{V}_j (\hat{V}_j + \Lambda)^{-1} \Lambda (\sigma - \hat{\gamma}_j)}, \end{aligned} \quad \text{(III-14)}$$

where C_j is the constant defined above.

Finally we can compute the posterior distribution on target tracks given the collection of contact data D . Indeed combining equations (III-11) and (III-14) we obtain

$$\text{pr}(z | D) = \text{pr}(D)^{-1} \sum_{j=1}^J f(z | \delta_j, U_j) C_j e^{-\frac{1}{2}(\sigma - \hat{\gamma}_j)' \hat{V}_j (\hat{V}_j + \Lambda)^{-1} \Lambda (\sigma - \hat{\gamma}_j)}. \quad \text{(III-15)}$$

Estimating Target Motion Parameters from Contact Data When the Motion Scenario Is Known

Let the prior probability density function for the random quantities Q , T_1, \dots, T_J , A_1, \dots, A_J be given, as before, by

$$\text{pr}(q, \gamma_1, \dots, \gamma_J, a_1, \dots, a_J) = \text{pr}_0(q) \prod_{j=1}^J \text{pr}_j(\gamma_j, a_j),$$

where

$$\text{pr}_0(q) = g(q | \alpha), \text{ and}$$

$$\text{pr}_j(\gamma, a) = f(\gamma | \hat{\gamma}_j, v_j, a) h(a | m_j, V_j), \quad j = 1, 2, \dots, J.$$

Consider a target which is known to be following scenario S_l (so that $\mathcal{J} = l$). Suppose, however, that the target track $z = (z_1, \dots, z_\tau)$ is unknown but that we have obtained the contact events D_{tk} at various times t , $1 \leq t \leq \tau$. Each contact event consists of a mean target position estimate $d_{tk} = \xi_{tk}$ together with a corresponding covariance matrix Δ_{tk} . Let D be the joint detection event

$$D = \{d_{tk} = \xi_{tk}, \text{ all } t, j\}.$$

We seek the joint posterior distribution for the target motion parameters conditioned on the joint contact event D :

$$\text{pr}(q, \gamma_1, \dots, \gamma_J, a_1, \dots, a_J | D, \mathcal{J} = l) =$$

$$\frac{\text{pr}(D | q, \gamma_1, \dots, \gamma_J, a_1, \dots, a_J, \mathcal{J} = l) \text{Pr} \{ \mathcal{J} = l | q, \gamma_1, \dots, \gamma_J, a_1, \dots, a_J \}}{\text{pr}(D | \mathcal{J} = l) \text{pr} \{ \mathcal{J} = l \}} \quad (\text{III-16})$$

$$\cdot \text{pr}_0(q) \prod_{j=1}^J \text{pr}_j(\gamma_j, a_j).$$

We have already established the values of two of the factors which appear on the right side of (III-16). Indeed,

$$\text{Pr} \{ \mathcal{J} = l | q, \gamma_1, \dots, \gamma_K, a_1, \dots, a_K \} = q_l, \text{ and}$$

$$\text{Pr} \{ \mathcal{J} = l \} = \alpha_l \beta^{-1} \quad (\beta = \alpha_1 + \dots + \alpha_K).$$

Also, since each estimated position for the target d_{tk} is a random vector which depends only on the actual position z_t of the target at time t , the likelihood function $\text{pr}(D | q, \gamma_1, \dots, \gamma_J, a_1, \dots, a_J, \mathcal{J} = l)$ is independent of the parameters γ_j and a_j for $j \neq l$. We thus can write

$$\text{pr}(D | q, \gamma_1, \dots, \gamma_J, a_1, \dots, a_J, \mathcal{J} = l) = \text{pr}(D | \gamma_l, a_l, \mathcal{J} = l).$$

Substituting these quantities into equation (III-16) we obtain

$$\begin{aligned} \text{pr}(q, \gamma_1, \dots, \gamma_J, a_1, \dots, a_J | D, \mathcal{J} = l) \\ \propto \text{pr}(D | \gamma_l, a_l, \mathcal{J} = l) q_l \alpha_l^{-1} \beta \text{pr}_0(q) \prod_{j=1}^J \text{pr}_j(\gamma_j, a_j). \end{aligned} \quad (\text{III-17})$$

But observe that

$$\beta \alpha_l^{-1} q_l \text{pr}_0(q) = \beta \alpha_l^{-1} q_l g(q | \alpha) = \mathcal{J}(q | \alpha^*),$$

where $\alpha^* = (\alpha_1, \dots, \alpha_l + 1, \dots, \alpha_J)$. Thus the posterior distribution for the random variable Q is a Dirchlet distribution which depends only on the scenario followed by the target in question.

In order to complete the evaluation of (III-17) we need to compute the likelihood function $\text{pr}(D | \gamma_l, a_l, \mathcal{J} = l)$. Since the distribution of each d_{tk} depends only on the actual position of the target at time t , we have

$$\begin{aligned} \text{pr}(D | \gamma_l, a_l, \mathcal{J} = l) &= \int \text{pr}(D | z) \text{pr}(z | \gamma_l, a_l, \mathcal{J} = l) dz \\ &= \int \text{pr}(D | z) f(z | \gamma_l, a_l) dz. \end{aligned} \quad (\text{III-18})$$

Recall now from equation (III-12) that we can find a constant L , a vector $\sigma \in \mathbb{R}^{2\tau}$ and a $2\tau \times 2\tau$ block diagonal positive semidefnite matrix Λ such that

$$\text{pr}(D | z) = C \sqrt{\prod_{t,k} |\Delta_{tk}|} e^{-\frac{1}{2}(z-\sigma)' \Lambda (z-\sigma)},$$

where C is the constant $(2\pi)^{-\mu} e^{-L}$, and μ is total number of contacts. By the process of completing the square described in Appendix A, we can find a vector $\gamma \in \mathbb{R}^{2\tau}$ such that

$$\begin{aligned} (z-\sigma)' \Lambda (z-\sigma) + (z-\gamma_l)' a_l (z-\gamma_l) \\ = (z-\gamma)' (\Lambda + a_l) (z-\gamma) + (\sigma-\gamma_l)' \Lambda (a_l + \Lambda)^{-1} a_l (\sigma-\gamma_l). \end{aligned}$$

Thus, evaluating the integral in Equation (III-15) we obtain

$$\text{pr}(D | \gamma_l, a_l, \mathcal{J} = l) = C \sqrt{|a_l| |a_l + \Lambda|^{-1} \prod_{t,k} |\Delta_{tk}|} e^{-\frac{1}{2}(\sigma-\gamma_l)' a_l (a_l + \Lambda)^{-1} \Lambda (\sigma-\gamma_l)}.$$

In order to complete our evaluation of the posterior density given in (III-17), we must consider the product

$$\begin{aligned} \text{pr}(D | \gamma_l, a_l, \mathcal{J} = l) \text{pr}_l(\gamma_l, a_l) = \\ \text{pr}(D | \gamma_l, a_l, \mathcal{J} = l) f(\gamma_l | \hat{\gamma}_l, v_l, a_l) h(a_l | m_l, V_l). \end{aligned}$$

Note that, again by the process of completing the square of a quadratic form, we have

$$\begin{aligned} (\sigma-\gamma_l)' a_l (a_l + \Lambda)^{-1} \Lambda (\sigma-\gamma_l) + (\gamma_l - \hat{\gamma}_l)' v_l a_l (\gamma_l - \hat{\gamma}_l) \\ = (\gamma_l - \hat{\gamma}_l^*)' (v_l + 1) a_l T (\gamma_l - \hat{\gamma}_l^*) + (\sigma - \hat{\gamma}_l)' a_l U (\sigma - \hat{\gamma}_l) \quad \text{(III-19)} \\ = (\gamma_l - \hat{\gamma}_l^*)' (v_l + 1) a_l T (\gamma_l - \hat{\gamma}_l^*) + \text{tr} \{ U (\sigma - \hat{\gamma}_l) (\sigma - \hat{\gamma}_l)' a_l \}, \end{aligned}$$

where

$$T = (a_l + \Lambda)^{-1} \left(\frac{v_l}{v_l + 1} a_l + \Lambda \right),$$

$$U = \frac{v_l}{v_l+1} T^{-1} (a_l + \Lambda)^{-1} \Lambda, \text{ and}$$

$$\hat{\gamma}_l^* = \frac{1}{v_l+1} \left\{ v_l T^{-1} \hat{\gamma}_l + \left[\frac{v_l}{v_l+1} a_l + \Lambda \right]^{-1} \Lambda \sigma \right\}.$$

Thus for an appropriately chosen constant C_l , we can write

$$\begin{aligned} & \text{pr}(D | \gamma_l, a_l, \mathcal{J} = l) \text{pr}_l(\gamma_l, a_l) \\ &= C_l f(\gamma_l | \hat{\gamma}_l^*, (v_l+1) a_l, T) \sqrt{\frac{|a_l| \prod_{t,j} |\Delta_{tj}|}{|a_l + (1+1/v_l) \Lambda|}} h(a_l | m_l, V_l + U(\sigma - \hat{\gamma}_l)(\sigma - \hat{\gamma}_l)'). \end{aligned}$$

In order to continue our calculations, it is necessary at this point to make a number of important approximations. First, note that for each positive definite matrix a_l we can find a number $\theta(a_l)$ such that

$$\left[\theta(a_l) - \frac{v_l}{v_l+1} \right] \text{tr } a_l = [1 - \theta(a_l)] \text{tr } \Lambda.$$

Observe that if v_l is sufficiently large we then have

$$T = (a_l + \Lambda)^{-1} \left(\frac{v_l}{v_l+1} a_l + \Lambda \right) \approx \theta(a_l) I.$$

Here I is the $2\tau \times 2\tau$ identity matrix. It is easy to see that $v_l (v_l+1)^{-1} \leq \theta(a_l) \leq 1$. The quantity $\theta(a_l)$ is important in that it provides a measure for the amount of target localization information provided by the contact data D relative to that provided by the precision matrix a_l . Indeed if no contacts were obtained, so that $\Lambda = 0$, then $\theta(a_l) = v_l (v_l+1)^{-1}$. Similarly, if we obtained a contact at each time t , $t = 1, 2, \dots, \tau$, then $\theta(a_l) \rightarrow 1$ as the precision of the target contact data converges to the target's actual track.

Using now this approximation we can write

$$T \approx \theta(a_l) I,$$

$$U \approx \frac{v_l}{v_l + 1} \frac{1}{\theta(a_l)} (a_l + \Lambda)^{-1} \Lambda, \text{ and}$$

$$\hat{\gamma}^* = \frac{1}{v_l + 1} \left[\frac{v_l}{\theta(a_l)} \gamma_l + \left[\frac{v_l}{v_l + 1} a_l + \Lambda \right]^{-1} \Lambda \sigma \right].$$

Next note that, for each positive definite matrix a_l , there is a number $\varphi(a_l)$ such that

$$\frac{|a_l| \prod_{t,j} |\Delta_{tj}|}{|a_l + \Lambda(1 + \frac{1}{v_l})|} = |a_l| \varphi(a_l).$$

Since $|a_l + \Lambda(1 + \frac{1}{v_l})| \geq \prod_{t,j} |\Delta_{tj}|$, it follows that $0 \leq \varphi(a_l) \leq 1$. The quantity $\varphi(a_l)$ is a second measure for the amount of target localization information provided by the contact data D relative to that provided by the precision matrix a_l . Indeed, if there is no contact data then $\varphi(a_l) = 0$. Also if a contact was obtained at each time t , $t = 1, 2, \dots, \tau$, then $\varphi(a_l) \rightarrow 1$ as the precision of the target contact data converges to the target's actual track.

In order to obtain computationally useful results it is necessary to obtain estimates for the quantities T , U , $\hat{\gamma}^*$, θ , and φ which do not depend on the precision matrix a_l . To obtain such estimates, we will replace a_l with a certain expected value. Note that A_l^{-1} is a random matrix which represents the covariance matrix Σ_l^{-1} . Thus, the quantity

$$E(A_l^{-1}) = \frac{1}{m_l - 2\tau + 1} V_l$$

is the Bayesian estimate for the matrix Σ_l^{-1} . We will approximate the matrix a_l with the matrix

$$V_l^0 = E(A_l^{-1})^{-1} = (m_l - 2\tau + 1) V_l^{-1}.$$

Observe that $V_l = v_l (v_l + 1)^{-1} V_l^0$.

Now set $\theta_l = \theta(V_l^0)$ and $\varphi_l = \varphi(V_l^0)$. We then have the approximations

$$T \approx \theta_l I,$$

$$U \approx \frac{v_l}{v_l+1} \theta_l^{-1} (V_l + \Lambda)^{-1} \Lambda,$$

$$\hat{\gamma}_l^* \approx \frac{1}{v_l+1} [v_l \theta_l^{-1} \gamma_l + [\hat{V}_l + \Lambda]^{-1} \Lambda \sigma], \text{ and}$$

$$\varphi_l(a_l) \approx \varphi_l.$$

Substituting these approximations into equation (III-17) we conclude finally that the posterior distribution for the random quantities $Q, \Gamma_1, \dots, \Gamma_J$, and A_1, \dots, A_J , given the joint contact data D on a target and the fact that the target is following scenario S_l , satisfies

$$\text{pr}(q, \gamma_1, \dots, \gamma_J, a_1, \dots, a_J \mid D, \mathcal{J} = l) \approx$$

(III-20)

$$g(q \mid \alpha^*) \prod_{k=1}^J f(\gamma_k \mid \hat{\gamma}_k^*, v_k^* a_k) h(a_k \mid m_k^*, V_k^*),$$

where

$$\alpha^* = \alpha + (0, 0, \dots, 1, \dots, 0), \quad (1 \text{ appears in the } l^{\text{th}} \text{ component}),$$

$$\hat{\gamma}_j^* = \begin{cases} \gamma_j & \text{if } j \neq l \\ \frac{1}{v_l+1} \{v_l \theta_l^{-1} \gamma_l + [\hat{V}_l + \Lambda]^{-1} \Lambda \sigma\} & \text{if } j = l, \end{cases}$$

$$v_j^* = \begin{cases} v_j & \text{if } j \neq l \\ (v_l+1)\theta_l & \text{if } j = l, \end{cases}$$

$$m_j^* = \begin{cases} m_j & \text{if } j \neq l \\ m_l + \varphi_l & \text{if } j = l, \text{ and} \end{cases}$$

$$V_j^* = \begin{cases} V_j & \text{if } j \neq l \\ V_j + \frac{v_l}{v_l + 1} \theta_l^{-1} [V_l^0 + \Lambda]^{-1} \Lambda (\sigma - \hat{\gamma}_l) (\sigma - \hat{\gamma}_l)^{-1} & \end{cases}$$

Estimating Target Motion Parameters from Contact Data

Suppose now that we have obtained the joint contact event D on a given target, but we have yet to determine which of the motion scenarios S_1, S_2, \dots, S_J the target is following. This is the most likely form that the available data will take. We would like to estimate the joint posterior distribution for the target motion parameters from the contact data D . Observe that we can write

$$\begin{aligned} & \text{pr}(q, \gamma_1, \dots, \gamma_J, a_1, \dots, a_J | D) \\ &= \sum_{l=1}^J \text{pr}(q, \gamma_1, \dots, \gamma_J, a_1, \dots, a_J | D, J = l) \text{Pr} \{ \mathcal{J} = l | D \}. \end{aligned} \quad (\text{III-21})$$

The first factor of each term appearing in (III-21) has already been evaluated in equation (III-20).

To complete the evaluation of (III-21), we need only compute the conditional probability that the target is following scenario S_l given the contact data D . To do this we employ Bayes theorem to write

$$\text{Pr} \{ \mathcal{J} = l | D \} = \frac{\text{Pr}(D | \mathcal{J} = l) \text{Pr} \{ \mathcal{J} = l \}}{\sum_{j=1}^J \text{Pr}(D | \mathcal{J} = j) \text{Pr} \{ \mathcal{J} = j \}}. \quad (\text{III-22})$$

Recall now that $\text{Pr} \{ \mathcal{J} = l \} = \alpha_l \beta^{-1}$, $\beta = \alpha_1 + \dots + \alpha_J$. Also the probability of obtaining the contact data D given the target is following scenario S_l has been evaluated in equation (III-13). Thus combining equations (III-20), (III-21), (III-22), and (III-13) we are able to estimate the joint posterior distribution for our target motion parameters given only the target contact data D . Observe, however, as in the case where the target track is known but the motion scenario is not, this posterior has a different form than the prior and involves many more parameters than does the prior.

The cause of this problem is, much as before, the impossibility of determining solely from contact data which scenario a target is following. A possible solution to this problem would be to use equation (III-22) to determine the highest probability scenario, and then assign the target to that scenario. The calculation of the

posterior joint distribution for the target motion parameters would then be carried out as in the previous section. Again, such a scheme is particularly meaningful when the scenarios are widely separated in space and the contact data fit one scenario significantly better than any other scenario.

CHAPTER IV

ASW INFORMATION PROCESSING IN A MULTI-TARGET ENVIRONMENT

In this chapter we develop Bayesian statistical methods for processing ASW information in a multi-target environment. A potential application of this work is in ASW information processing systems designed to accept ASW information from diverse sources, including direct surveillance, in an effort to initiate and maintain simultaneous localizations on a number of submarine targets. Such systems can be used both to carry out ASW surveillance and to assess vulnerability to such surveillance by others.

The chapter is divided into a number of sections. The first section is an introduction and provides the necessary background to place the current work on information processing in perspective. The second section provides a brief description of ASWIPS (ASW Information Processing System), a small scale system developed as a testbed for information processing methodology. The third section involves a theoretical discussion of information processing algorithms. Finally, the fourth section presents some numerical results based on ASWIPS comparing a number of alternative information processing algorithms. Extended Memory processing is described in detail in Appendix B.

Background and Introduction

Suppose that a number of submarine targets are known to be operating in a more or less well defined ocean area and periodic estimates are to be generated for each of their positions. Potential sources of information to assist in target force localization include: water depth contours; submarine speed limitations; direct surveillance information such as observations of port arrivals and departures and ASW sensor detection data; geographic constraints imposed by submarine missions; and so on. We treat this localization problem in terms of ASW surveillance information processing systems, i. e., systems designed to accept as input information of the type described and produce as output estimates of the locations of the target submarines.

Our principal interest in this chapter is in Bayesian statistical methods for the systematic generation and updating of target location predictions in a multi-target

environment. By a multi-target environment we mean one in which, in general, there are situations in which it is not possible to determine the target that generated a given sensor contact. For example, it may be that the submarine targets are acoustically indistinguishable. In addition, it is assumed that individual targets are not confined to operate in disjoint subregions. Under these restrictions it is not possible to treat the multi-target surveillance problem as a composite of many independent single target problems even though the underlying target motions may be statistically independent. It is the information processing difficulties arising from this intrinsic multivariate nature of the problem that is the concern of this chapter.

A substantial portion of the discussion below will deal with specific multi-target Bayesian information processing algorithms and their comparative evaluation. A particular such algorithm called Extended Memory processing will emerge as the only one of the Bayesian approaches we consider that both makes accurate use of the observational data and is computationally practical.

In order to provide the Extended Memory methodology with a framework within which it can be implemented, tested, and compared with other processing approaches, we have put together a small scale developmental information processing system called ASWIPS (ASW Information Processing System). We next proceed to describe ASWIPS.

Structure of ASWIPS (ASW Information Processing System)

ASWIPS is a small scale information processing system of computer programs currently implemented on a PRIME 400 minicomputer designed to accept as input contact data on a small number of targets and using Bayesian statistical methods to generate target location estimates. The principal output of ASWIPS at the end of each processing update stage consists of probability distributions for the location of each active target.

Target locations are assumed to be discretized into $1^{\circ} \times 1^{\circ}$ cells within a $10^{\circ} \times 10^{\circ}$ overall planar grid. Target locations are thus specified in terms of probability distributions on this 100 cell grid. During the time between two successive updates a target can either remain in its initial cell or move to one of the eight adjacent cells (five cells, if the initial cell shares an edge in common with the region boundary; three cells, if it is a corner cell) in the two-dimensional grid. Each of these possible transitions is assumed to have equal probability. Thus the motion model has a very simple Markovian structure.

Observational information is input to ASWIPS in the form of probability distributions on (target-like) objects detected. Thus, the raw contact data are effectively assumed to have been preprocessed before being input to ASWIPS. Each object location distribution is assumed to correspond to an observation on one of the targets in the operating area.

The targets are assumed to be effectively indistinguishable in that sensor contacts cannot be uniquely associated with the target causing them. It should be noted, however, that in some cases such an identification may be possible by statistical inference. This would be the case when an object location distribution is out of sensor contact range of all but one of the prior target location distributions. In general, however, an object location distribution may overlap with more than one prior target location distribution in which case the identity of the target that caused the underlying responses is ambiguous. In such a case we will refer to the object location distribution (and the underlying response pattern) as unresolved.

The updating of target locations distributions is performed at the end of each update stage. First, the target location distributions as produced by ASWIPS at the end of the preceding processing stage are updated for target motion. This is done by applying the one-stage Markov transition operator for the motion model that we described earlier. The effect of this motion updating is to make the target location distributions more diffuse, reflecting the fact that target motion since the last update has increased the degree of our uncertainty about target locations.

The second phase of processing at the end of each update stage involves revising the target location distributions for the contact information obtained since the last update. We will reserve our discussion of the specifics of the updating algorithms used for the next section of this chapter.

It should be observed that just as target motion tends to diffuse the target location distributions, contact information tends to concentrate them. It is this continuing tug of war between the loss of information resulting from target motion and the gain of information resulting from contact information that controls the dynamics of the target localization capability of ASWIPS. Factors which control the direction that this information struggle will take include the number of targets, target speed, the size of the operating area, contact data rate, sensor detection ranges, etc. and of equal importance, the processing algorithms themselves. This brings us to the topic of the next section.

Discussion of Information Processing Algorithms

In this section we will discuss some of the mathematical aspects of information processing algorithms. This discussion will lay the groundwork for a comparison in the next section of a number of processing algorithms using ASWIPS.

Target location distributions. Suppose we use k to index processing update stages and assume that at the k^{th} update stage there are N submarines actively deployed in the operating area of interest. We then define

$X_n^{(k)}$ = grid cell location of the n^{th} target at the end of update stage k .

In the case of ASWIPS, target locations are discretized into cells so that $X_n^{(k)}$ is the index of the cell containing the n^{th} target at stage k . Under the presumption that these target locations are uncertain, the objective of information processing is to estimate these locations in terms of probability distributions. Thus we can think of

$$X^{(k)} = (X_1^{(k)}, X_2^{(k)}, X_3^{(k)}, \dots, X_N^{(k)})$$

as a random N -vector whose probability distribution we seek. Specifically, using all available information, we would like the information processor to estimate the joint frequency function

$$P^{(k)}(x_1, x_2, \dots, x_N) = \Pr \{ X^{(k)} = (x_1, x_2, \dots, x_N) \}, \quad (\text{IV-1})$$

where $x_n \in I$ for $n = 1, 2, \dots, N$ with I being the collection of all cell indices.

Before pursuing the question of how $P^{(k)}$ is to be estimated we consider the question of how many values it takes to specify this location frequency function. Consider a case in which the operating area grid contains approximately $1,000 \times 1,000$ cells and the number of targets, N equals 10. The number of possible states of the system, each of which corresponds to a joint location of the 10 targets, is $1,000^{10} = 10^{30}$ which is obviously prohibitively large from a computational standpoint.

Suppose now that it can be assumed that the N targets are operating independently. If we define the marginal target location distributions

$$p_n^{(k)}(x) = \Pr \{ X_n^{(k)} = x \},$$

then, assuming independence, we obtain the relationship,

$$P^{(k)}(x_1, x_2, \dots, x_N) = \prod_{n=1}^N p_n^{(k)}(x_n). \quad (\text{IV-2})$$

If the relationship in equation (IV-2) were valid and remained valid throughout information processing, the numerical size of the problem of estimating target location distributions would be sharply reduced. Going back to our example, we observe that each $p_n^{(k)}$, $i = 1, 2, \dots, 10$ involves the specification of up to 1,000 values. Thus the total number of values required to specify all target location distributions would be only 10,000, well within the capability of modern computers.

A problem unfortunately arises from the fact that even if the relationship in equation (IV-2) were valid at the outset of processing, i. e., before any sensor contacts were processed, and even if the individual target motions were statistically independent, equation (IV-2) in general would not continue to hold during the course of processing. A very simple example serves to illustrate this fact.

Suppose we consider a case in which the operating area consists of only two cells, each containing a single sensor which can detect a target only in its own cell. Suppose there are two targets operating independently each of which is a priori initially equally likely to be in either of the two grid cells. By assumption, then equation (IV-2) holds initially. Now suppose that during the first update period each of the two sensors reports a contact. Designate the two targets as T_1 and T_2 and the two grid cells as C_1 and C_2 . Then it is easily seen that there are only two possibilities: target T_1 is in cell C_1 and target T_2 is in cell C_2 , or vice versa. Based on our observational data alone, we cannot conclude with certainty which of these two equally likely possibilities is the correct one, although we can with certainty eliminate the possibility that the two targets are in the same cell. One now observes an interesting phenomenon. Whereas before the contacts were reported the target locations were statistically independent, after the contacts C_1 and C_2 occur, the target locations are completely correlated. Specifically, knowledge of the location of one target implies deterministic knowledge of the location of the other target. If target T_1 is in cell C_1 , then necessarily target T_2 is in cell C_2 -- thus the complete correlation between the two targets. We will refer to the problem of properly incorporating this contact-induced correlation between targets into information processing as the target coupling problem.

One might reasonably argue that the example we have given is extreme and that in more realistic cases the degree of correlation among target location distributions induced by the sensor data is likely to be small. Suppose for the time being we assume that the target location distributions can in fact be treated as approximately independent and we proceed naively on the basis of this assumption. This means that we will deal only with the individual marginal distributions $p_n^{(k)}$ for target location. The problem then becomes one of devising a procedure for updating these marginals to reflect sensor data. A standard and well-studied procedure for this kind of updating is based on Bayes' theorem on conditional probabilities.

In the next two subsections below we describe two possible target location distribution updating schemes based on the independence assumption and Bayes' theorem.

Sequential processing. We assume that the sensor contacts reported during an update stage have been clustered into sensor response patterns each generated by a single target. We associate with each response pattern the presence of a (target-like) object. Under sequential processing the target location distributions are collectively updated once for each such response pattern. Let $p_n^{(k,l)}(x)$ can be obtained recursively from $p_n^{(k,l-1)}(x)$.

Assume L response patterns were observed during the k^{th} stage and let

$$\lambda^{(k)}(x) = \Pr \{l^{\text{th}} \text{ observed target-like object is in cell } x \text{ at stage } k\},$$

$$l = 1, 2, \dots, L$$

be the associated observed object location distributions. We then let

$$V_n^{(k, l)} = \sum_{x \in I} p_n^{(k, l-1)}(x) \lambda_l^{(k)}(x), \quad n=1, 2, \dots, N \quad (\text{IV-3})$$

and

$$\alpha_n^{(k, l)} = \frac{V_n^{(k, l)}}{\sum_{j=1}^N V_j^{(k, l)}}. \quad (\text{IV-4})$$

The quantity $V_n^{(k, l)}$ is a measure of the fit between the l^{th} object location distribution in k^{th} stage processing and the target location distribution for the n^{th} target based on processing the first $l-1$ object location distributions. Probabilistically, $V_n^{(k, l)}$ is an estimate of the likelihood that the l^{th} object and the n^{th} target are in the same cell at the k^{th} stage. The quantity $\alpha_n^{(k, l)}$ is then an estimate of the probability that the l^{th} object and n^{th} target are one and the same.

Now there are two possibilities: either the n^{th} target is the same as the l^{th} object or it is not. If it is not, then there is no reason to modify our current estimate of the location of the n^{th} target based on λ_l ; if the n^{th} target and the l^{th} object are the same, then a composite revised estimate of the location of the n^{th} target should be made combining the information contained in the prior distribution for the location of the n^{th} target with the information contained in the l^{th} object location distribution.

In mathematical terms, using Bayes' theorem we then obtain

$$p_n^{(k, l)}(x) = \{1 - \alpha_n^{(k, l)}\} p_n^{(k, l-1)}(x) + \alpha_n^{(k, l)} \frac{p_n^{(k, l-1)}(x) \lambda_l^{(k)}(x)}{V_n^{(k, l)}}. \quad (\text{IV-5})$$

The expressions in equations (IV-3) through (IV-5) provide the promised recursion. One proceeds in this fashion iteratively updating for each object location distribution associated with a given update. The entire series of updates for sensor information during a processing stage is carried out following an update for target motion at the beginning of each such processing stage.

Parallel processing. Under sequential processing the object location distributions during a processing stage lead to a series of updates to the target location distributions. An alternative mode of processing, which we call parallel processing, would attempt to perform a single update for all such object location distributions simultaneously.

To carry out the processing of target location distributions in parallel we introduce the notion of an assignment function. This is a map f such that

$$f: \{1, 2, \dots, L\} \rightarrow \{1, 2, \dots, N\},$$

where L is the number of object location distributions being processed in parallel. In general an assignment function need not be one-to-one since more than one object location might have been generated by the same target. In the current presentation, for simplicity, we will assume that assignment functions are one-to-one, so that in particular $L \leq N$. We then define

$\mathcal{F}^{(k)}$ = the class of all possible one-to-one assignments of k^{th} stage objects to targets.

Parallel processing is then carried out as follows. First set

$$\hat{w}_f^{(k)} = \prod_{l=1}^L \sum_{x \in S} p_{f(l)}^{(k-1)}(x) \lambda_l^{(k)}(x), \quad (\text{IV-6})$$

where $p_n^{(k-1)}$, as defined earlier, is the location distribution of the n^{th} target through $(k-1)^{\text{st}}$ stage processing and $\lambda_l^{(k)}$ is the l^{th} object location distribution observed during the k^{th} stage. Next normalize to obtain

$$w_f^{(k)} = \frac{\hat{w}_f^{(k)}}{\sum_{f \in \mathcal{F}^{(k)}} \hat{w}_f^{(k)}}. \quad (\text{IV-7})$$

The quantities $w_f^{(k)}$ given in equation (IV-7) are direct analogues of the quantities $\alpha_n^{(k, l)}$ defined by equation (IV-4) except that whereas in sequential processing object locations were matched with targets one at a time, under parallel processing a simultaneous assignment of all objects to targets is made.

The update of the target locations is then based on the following relationship:

$$p_n^{(k)}(x) = \sum_{f \in \mathcal{F}^{(k)}} w_f^{(k)} \frac{p_n^{(k-1)}(x) \lambda_{f^*(n)}^{(k)}(x)}{\sum_{y \in S} p_n^{(k-1)}(y) \lambda_{f^*(n)}^{(k)}(y)}, \quad (\text{IV-8})$$

where f^* is the inverse function of f , and

$$\lambda_{f^*(i)}^{(k)} = 1, \text{ if } f^*(i) \text{ is not defined.}$$

Considerable experimentation has been carried out using both the sequential and parallel methods of processing. Experience has shown that both approaches function reasonably well when the actual locations of the targets remain spatially separated so that there is little possibility for confusion of their identities. However, as the number of targets increases and correspondingly the likelihood that unresolved object location distributions will be generated also increases, the assumption that the target locations can be treated even approximately as independent seriously degrades target location prediction performance, the problem being more acute in the case of sequential processing than parallel processing.

The specific symptoms of the breakdown in processor performance include (1) possible permutation of target identities and (2) the doubling up of target location distributions. The permutation of target identities occurs when the estimated location distribution of target T_1 , for example, corresponds closely to the true position of target T_2 , and vice versa. The doubling up phenomenon occurs when estimated location distributions on two targets appear to be virtually identical when in fact there is only one actual target whose position is compatible with these distributions. In such a case, for example, the computed location distributions of both targets T_1 and T_2 may accurately fit the actual location of target T_1 while target T_2 is actually off in another part of the operating area with a location not covered by any of the computed target location distributions.

The two phenomena of target identity permutation and doubling up are related in that one may lead to the other; both are extremely damaging to proper prediction performance. Typical examples illustrating this kind of anomalous behavior using ASWIPS will be given in the next section.

In summary, our work has shown that it appears impossible to obtain a completely satisfactory information processing model based on the assumption of target location independence. As a result we have developed a third form of processing, called Extended Memory processing, which is based on the key observation that while the

target locations are, in general, not independent, they are conditionally independent given a specific assignment of all past sensor responses to the targets. This simply says that if one knew which object location distributions (or more fundamentally which raw sensor responses) were caused by which targets, i. e., if one knew the correct response pattern to target association, then the correlation in target locations induced by such observational data would disappear. In general, the correct such association is not known. However, the target locations can be effectively uncoupled by first conditioning on a possible response pattern to target association, then performing a conditional update of target locations based on such an association, and finally removing the conditioning by averaging over all possible such associations.

Observe that under Extended Memory processing, conditional target location distributions must be retained in processor memory for each association assigned a nonzero weight. Extended Memory processing thus requires a substantially expanded memory capability over the other processing algorithms and this, as we shall see in the next section, principally accounts for its improved prediction capability.

We remark that Extended Memory processing is a form of generalization of parallel processing. Under parallel processing the conditioning is based on an assignment to the targets of the response patterns observed during a single update period. Under Extended Memory processing the conditioning is based in effect on a simultaneous assignment to the targets of all response patterns observed to date.

The mathematical description of Extended Memory processing is somewhat more involved than that of sequential or parallel processing. Thus a detailed theoretical development of Extended Memory processing is relegated to Appendix B which also includes some comments on the computing and memory storage requirements imposed by this method of processing.

Comparison of Processing Algorithms Using ASWIPS

This section will be devoted to a series of numerical examples and comparisons of the various updating algorithms we described in the preceding section. All such numerical results have been obtained using the computer model ASWIPS.

Example 1. In our first example we assume that there are three targets operating in a 10 x 10 cellular grid according to the motion model described earlier under which during a given update stage a target is equally likely to move from its current cell to any one of the adjacent cells.

The assumed target tracks are shown in Figure IV-1 and the prior distributions for the three target locations are shown in Table IV-1 below. Circled target locations in Figure IV-1 indicate target detections at those locations. We assume that near

FIGURE IV-1

ACTUAL TARGET TRACKS

(Example 1)

Note: k indexes update stages with the corresponding target locations as indicated.
Circled target locations indicate detections




TABLE IV-1

INITIAL TARGET LOCATION DISTRIBUTIONS

(Example 1 Extended Memory Processing)

Target T₁

Note: The entry in each cell divided by 1000 gives the probability of finding the target in that cell.

	1	2	3	4	5	6	7	8	9	10
1	63.	47.	31.	47.	63.	0.	0.	0.	0.	0.
2	47.	35.	23.	35.	47.	0.	0.	0.	0.	0.
3	31.	23.	16.	23.	31.	0.	0.	0.	0.	0.
4	47.	35.	23.	35.	47.	0.	0.	0.	0.	0.
5	63.	47.	31.	47.	63.	0.	0.	0.	0.	0.
6	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
7	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
8	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
9	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
10	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.

TABLE IV-1 (continued)

Target T₂

	1	2	3	4	5	6	7	8	9	10
1	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
2	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
3	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
4	0.	0.	0.	20.	31.	41.	31.	20.	0.	0.
5	0.	0.	0.	31.	46.	61.	46.	31.	0.	0.
6	0.	0.	0.	41.	61.	82.	61.	41.	0.	0.
7	0.	0.	0.	31.	46.	61.	46.	31.	0.	0.
8	0.	0.	0.	20.	31.	41.	31.	20.	0.	0.
9	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
10	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.

TABLE IV-1 (continued)

Target T₃

	1	2	3	4	5	6	7	8	9	10
1	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
2	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
3	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
4	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
5	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
6	0.	0.	0.	0.	0.	22.	22.	22.	22.	22.
7	0.	0.	0.	0.	0.	44.	44.	44.	44.	44.
8	0.	0.	0.	0.	0.	67.	67.	67.	67.	67.
9	0.	0.	0.	0.	0.	44.	44.	44.	44.	44.
10	0.	0.	0.	0.	0.	22.	22.	22.	22.	22.

the end of the first update stage, two contacts are observed: one in the cell (5, 5) and a second in the cell (6, 6). The sensors involved are assumed to have very short range and, therefore, a contact guarantees the presence of a target in the cell containing the sensor. Thus two object location distributions are observed during stage 1: the first placing a target with unit mass in cell (5, 5), the second placing a target in the cell (6, 6).

After an update for motion during the first stage and processing for the first stage contacts, the posterior target locations under Extended Memory processing are as shown in Table IV-2.

One observes from Table IV-2 that both of the contacts are unresolved. The contacts in cells (5, 5) and (6, 6) could each have been caused by any of the three targets. Consequently, Table IV-2 shows a sharp concentration of target mass in cell (5, 5) in the location distribution of each target and a similar peak in cell (6, 6).

We now turn to the second processing stage. We assume that during this stage a single contact was obtained in cell (3, 8). Referring to Table IV-2 one observes that even allowing for target motion only T_2 could have caused this contact, and therefore it is known with certainty that T_2 is located in cell (3, 8) at the end of the second update. (Again we assume that the contact occurred at or near the end of the update period.) Now it becomes possible to fit some pieces of information together. T_2 is known to be in cell (3, 8). Since T_2 could have moved at most one cell since the last update, T_2 could not have caused either of the contacts in cells (5, 5) or (6, 6). Thus, virtually all of the ambiguity in the problem has retroactively been resolved by the latest contact. In particular, T_1 probably caused the contact in cell (5, 5), and T_3 , the contact in (6, 6), although there is a small chance of the alternative assignment. Allowing for one update period of additional target motion the locations of the three targets at the end of the second update period must be as shown in Table IV-3.

As noted earlier, this example is based on the use of Extended Memory processing in performing the updates for sensor contacts. Substantially different results are obtained using either sequential processing or parallel processing. These results are shown in Table IV-4 and IV-5. One observes that, while in all three types of processing, perfect localization is obtained on target T_2 , the degree of localization on the other two targets depends on the method of processing, and the correct degree of localization, as supported by the contact data, is given only by Extended Memory processing.

Example 2. Figure IV-2 below shows the actual track of two targets which form the basis of our second example. As in Example 1, the circled target locations correspond to sensor contacts on that target. Thus in the example there were five contacts on target T_1 , one each for five successive update periods. In contrast there is only one contact on target T_2 , that occurring during the first processing stage. Once again we assume that all sensor contacts occurred sufficiently close to the end of an update stage so that subsequent target motion during that stage is negligible.

TABLE IV-2

TARGET LOCATIONS DISTRIBUTIONS AFTER 1st UPDATE

(Example 1 Extended Memory Processing)

Target T₁

Note: The entry in each cell divided by 1000 gives the probability of finding the target in that cell.

	1	2	3	4	5	6	7	8	9	10
1	11.	13.	10.	11.	9.	5.	0.	0.	0.	0.
2	13.	16.	12.	14.	11.	6.	0.	0.	0.	0.
3	8.	10.	9.	10.	8.	5.	0.	0.	0.	0.
4	9.	11.	10.	11.	9.	5.	0.	0.	0.	0.
5	7.	9.	8.	9.	575.	4.	0.	0.	0.	0.
6	4.	5.	5.	5.	4.	128.	0.	0.	0.	0.
7	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
8	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
9	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
10	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.

TABLE IV-2 (continued)

Target I₂

	1	2	3	4	5	6	7	8	9	10
1	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
2	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
3	0.	0.	0.	1.	1.	1.	1.	1.	0.	0.
4	0.	0.	1.	2.	3.	4.	3.	2.	1.	0.
5	0.	0.	1.	3.	379.	6.	6.	3.	1.	0.
6	0.	0.	1.	4.	6.	510.	6.	4.	1.	0.
7	0.	0.	1.	3.	6.	6.	6.	3.	1.	0.
8	0.	0.	1.	2.	3.	4.	3.	2.	1.	0.
9	0.	0.	0.	1.	1.	1.	1.	1.	0.	0.
10	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.

TABLE IV-2 (continued)

TARGET T₃

	1	2	3	4	5	6	7	8	9	10
1	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
2	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
3	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
4	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
5	0.	0.	0.	0.	60.	3.	4.	4.	5.	4.
6	0.	0.	0.	0.	4.	380.	13.	13.	15.	11.
7	0.	0.	0.	0.	8.	17.	25.	25.	30.	21.
8	0.	0.	0.	0.	10.	20.	30.	30.	34.	25.
9	0.	0.	0.	0.	9.	18.	27.	27.	32.	23.
10	0.	0.	0.	0.	5.	10.	15.	15.	17.	12.

TABLE IV - 3
TARGET LOCATION DISTRIBUTIONS AFTER 2nd UPDATE

(Example 1 Extended Memory Processing)

Target T₁

Note: The entry in each cell divided by 1000 gives the probability of finding the target in that cell.

	1	2	3	4	5	6	7	8	9	10
1	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
2	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
3	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
4	0.	0.	0.	105.	105.	105.	0.	0.	0.	0.
5	0.	0.	0.	105.	111.	111.	6.	0.	0.	0.
6	0.	0.	0.	105.	111.	111.	6.	0.	0.	0.
7	0.	0.	0.	0.	6.	6.	6.	0.	0.	0.
8	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
9	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
10	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.

TABLE IV-3 (continued)

Target T₂

	1	2	3	4	5	6	7	8	9	10
1	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
2	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
3	0.	0.	0.	0.	0.	0.	0.	1000.	0.	0.
4	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
5	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
6	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
7	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
8	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
9	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
10	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.

TABLE IV-3 (continued)

Target T₃

	1	2	3	4	5	6	7	8	9	10
1	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
2	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
3	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
4	0.	0.	0.	6.	6.	6.	0.	0.	0.	0.
5	0.	0.	0.	6.	111.	111.	105.	0.	0.	0.
6	0.	0.	0.	6.	111.	111.	105.	0.	0.	0.
7	0.	0.	0.	0.	195.	105.	105.	0.	0.	0.
8	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
9	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
10	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.

TABLE IV-4

TARGET LOCATION DISTRIBUTIONS AFTER 2nd UPDATE

(Example 1 Sequential Processing)

Target T₁

	1	2	3	4	5	6	7	8	9	10
1	8	11	10	8	7	3	1	0	0	0
2	11	15	13	11	9	5	1	0	0	0
3	10	13	11	9	8	4	1	0	0	0
4	8	11	9	8	70	82	80	15	0	0
5	7	9	8	70	82	80	15	0	0	0
6	3	5	4	67	80	78	14	0	0	0
7	1	1	1	1	15	14	14	0	0	0
8	0	0	0	0	0	0	0	0	0	0
9	0	0	0	0	0	0	0	0	0	0
10	0	0	0	0	0	0	0	0	0	0

TABLE IV-4 (continued)

Target T₂

	1	2	3	4	5	6	7	8	9	10
1	0	0	0	0	0	0	0	0	0	0
2	0	0	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	1000	0	0
4	0	0	0	0	0	0	0	0	0	0
5	0	0	0	0	0	0	0	0	0	0
6	0	0	0	0	0	0	0	0	0	0
7	0	0	0	0	0	0	0	0	0	0
8	0	0	0	0	0	0	0	0	0	0
9	0	0	0	0	0	0	0	0	0	0
10	0	0	0	0	0	0	0	0	0	0

TABLE IV-4 (continued)

Target T₃

	1	2	3	4	5	6	7	8	9	10
1	0	0	0	0	0	0	0	0	0	0
2	0	0	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0	0	0
4	0	0	0	6	6	7	1	1	1	1
5	0	0	0	7	49	51	46	5	6	4
6	0	0	0	8	52	57	53	14	16	11
7	0	0	0	2	48	56	61	23	25	18
8	0	0	0	3	9	18	24	28	31	22
9	0	0	0	2	8	17	23	27	30	21
10	0	0	0	1	5	11	14	17	18	13

TABLE IV-5

TARGET LOCATION DISTRIBUTIONS AFTER 2 UPDATES

(Example 1 Parallel Processing)

Target T₁

	1	2	3	4	5	6	7	8	9	10
1	17	23	20	17	14	7	2	0	0	0
2	23	30	26	22	19	10	3	0	0	0
3	20	26	22	19	16	9	3	0	0	0
4	17	22	19	46	43	37	2	0	0	0
5	14	19	16	43	55	50	16	0	0	0
6	7	10	9	37	50	47	15	0	0	0
7	2	3	3	2	16	15	14	0	0	0
8	0	0	0	0	0	0	0	0	0	0
9	0	0	0	0	0	0	0	0	0	0
10	0	0	0	0	0	0	0	0	0	0

TABLE IV-5 (continued)

Target T₂

	1	2	3	4	5	6	7	8	9	10
1	0	0	0	0	0	0	0	0	0	0
2	0	0	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	1000	0	0
4	0	0	0	0	0	0	0	0	0	0
5	0	0	0	0	0	0	0	0	0	0
6	0	0	0	0	0	0	0	0	0	0
7	0	0	0	0	0	0	0	0	0	0
8	0	0	0	0	0	0	0	0	0	0
9	0	0	0	0	0	0	0	0	0	0
10	0	0	0	0	0	0	0	0	0	0

TABLE IV-5 (continued)

Target T₃

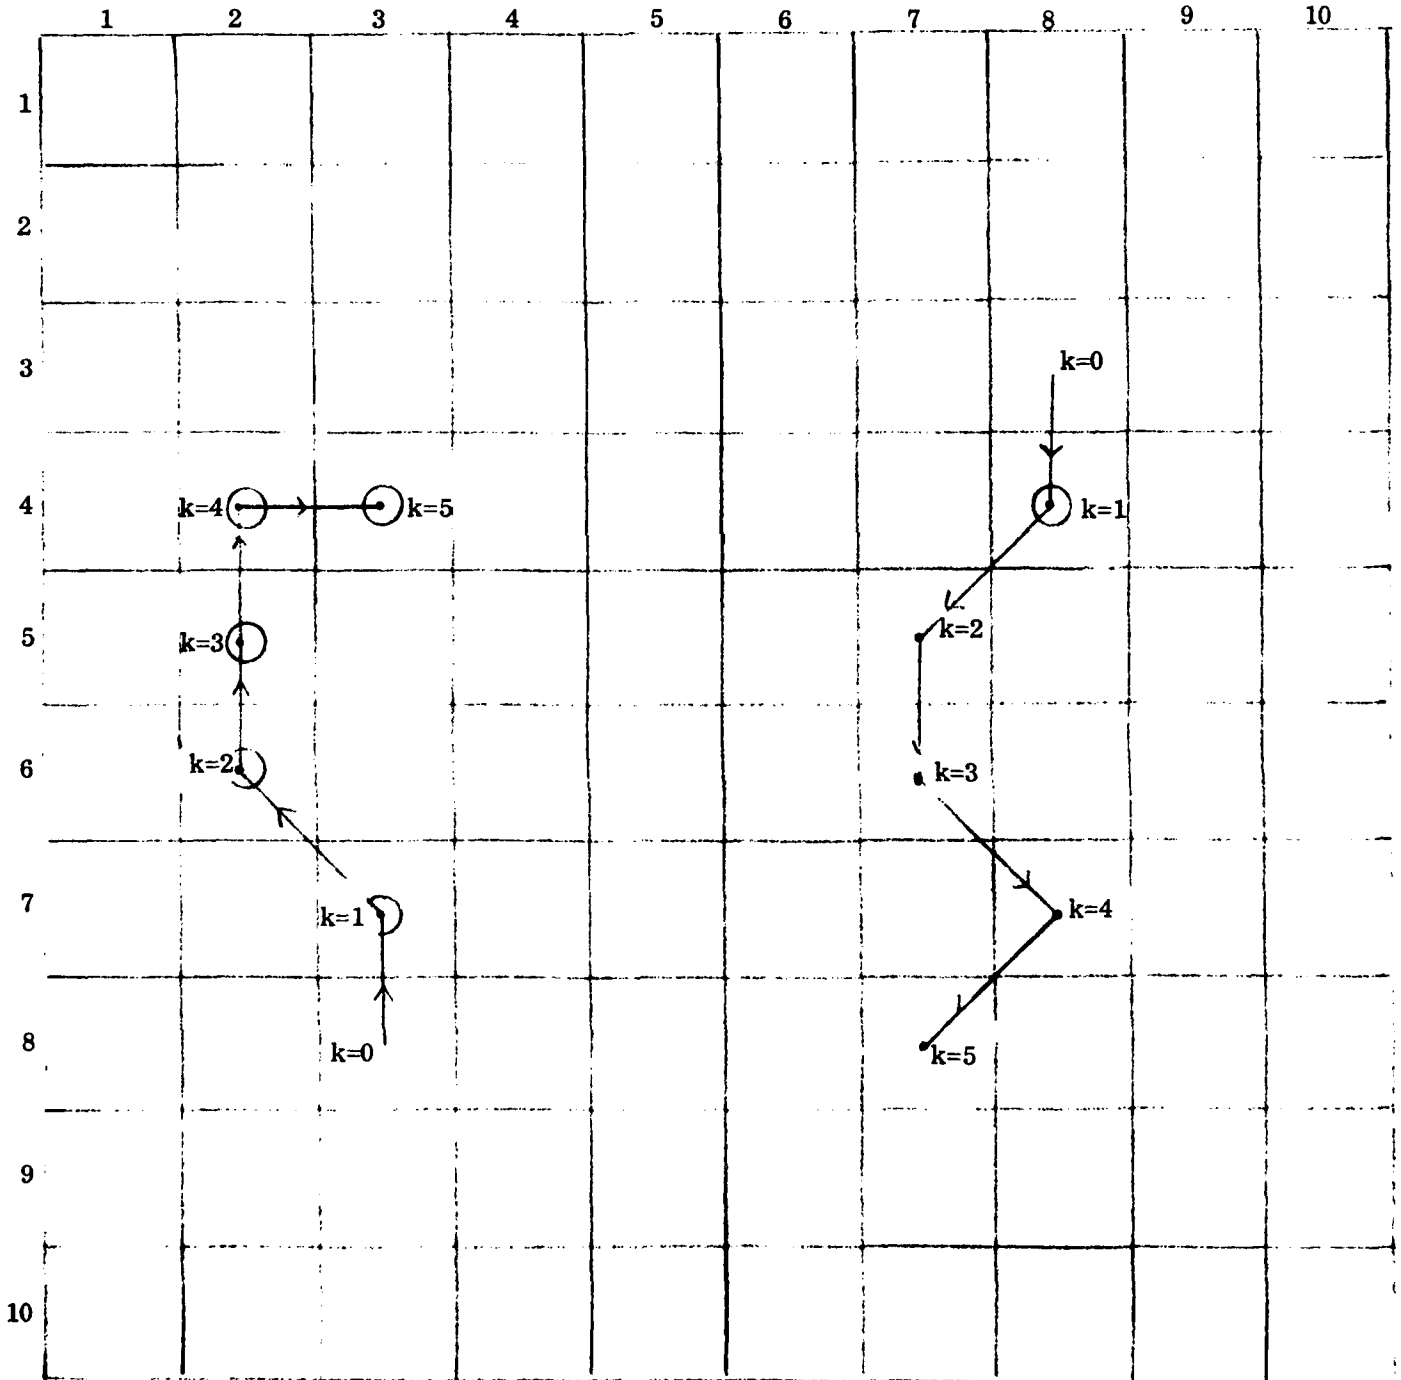
	1	2	3	4	5	6	7	8	9	10
1	0	0	0	0	0	0	0	0	0	0
2	0	0	0	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0	0	0	0
4	0	0	0	2	2	3	1	1	1	1
5	0	0	0	3	45	47	46	6	6	4
6	0	0	0	4	48	53	54	15	17	12
7	0	0	0	2	49	57	62	25	27	19
8	0	0	0	3	9	19	25	30	33	23
9	0	0	0	3	9	18	24	29	31	22
10	0	0	0	1	5	11	15	18	20	14

FIGURE IV-2

ACTUAL TARGET TRACKS

(Example 2)

Note: k indexes update stages with the corresponding target locations as indicated.
Circled target locations indicate detections.



We assume that at the initiation of processing essentially nothing is known about the prior locations of the two targets. Consequently, it is assumed that the prior location distribution of each target is uniform over the 10 x 10 cell grid with each target initially having probability .01 of being in any cell.

As shown in Figure IV-2, contacts are observed on each of the two targets during the first update stage: one placing a target in cell (7, 3) and the other placing a target in cell (4, 8). Because the prior target location distributions are uniform, there is no basis on which to show a preference in deciding which target caused which response. As a result, under either parallel processing or Extended Memory processing, .5 of the mass of each target would be assigned to each of the critical cells after processing these two contacts. Sequential processing, however, would introduce an artificial asymmetry into the problem and assign a total of somewhat more than a unit target mass to the cell containing the second contact processed and correspondingly somewhat less than a unit mass to the cell containing the sensor contact processed first. This behavior is simply another intrinsic flaw in sequential processing.

Table IV-6 shows the target location distributions after the fifth processing update using the Extended Memory processing technique. One first observes that the location distributions for targets T_1 and T_2 are identical, as in fact they should be in view of the complete symmetry in the nature of our information about their locations. A contact was observed in cell (4, 3) at the end of the fifth update period and Table IV-6 shows that this contact has been allocated equally to the two targets. No responses have been observed on target T_2 since the first processing stage. Thus each target map shows a second mode which is quite diffuse in contrast to the mode induced by the contact in cell (4, 3).

An overall assessment is that, at the end of the fifth update, a target (we do not know which one) can be placed with certainty in cell (4, 3). In addition, there appears to be a second target somewhere in the general northeastern portion of the operating area. One concludes that the results produced by Extended Memory processing are quite consistent with the underlying sensor contacts and appear to represent about the best localization that the available information supports.

Table IV-7 shows the analogous results for Example 2 based on parallel processing after the fifth update (the results under sequential processing are quite similar). Again the two target location distributions are identical because of the basic symmetry in the problem. In addition, each distribution shows a sharp mode in cell (4, 3) corresponding to the recent contact there, although parallel processing predicts an expected number of targets of approximately 1.07 in contrast to the 1.00 predicted by Extended Memory processing.

The most dramatic difference between Tables IV-6 and IV-7, however, lies in the distribution of mass outside cell (4, 3). Parallel processing predicts that a second target is likely to be found in the general vicinity of cell (4, 3) with extremely little likelihood given to the possibility that either of the two targets is to be found in the eastern half of the operating region.

TABLE IV-6
TARGET LOCATION DISTRIBUTIONS AFTER 5th UPDATE

(Example 2 Extended Memory Processing)

Target T
1

Note: The entry in each cell divided by 1000 gives the probability of finding the target in that cell.

	1	2	3	4	5	6	7	8	9	10
1	0.	0.	0.	0.	1.	3.	5.	7.	6.	4.
2	0.	0.	0.	1.	3.	8.	13.	16.	15.	9.
3	0.	0.	0.	1.	5.	12.	20.	24.	23.	14.
4	0.	0.	500.	1.	6.	14.	23.	28.	27.	17.
5	0.	0.	0.	1.	5.	12.	20.	24.	23.	14.
6	0.	0.	0.	1.	3.	8.	12.	15.	14.	9.
7	0.	0.	0.	0.	1.	3.	5.	6.	6.	4.
8	0.	0.	0.	0.	0.	1.	1.	1.	1.	1.
9	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
10	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.

TABLE IV-6 (continued)

Target T₂

	1	2	3	4	5	6	7	8	9	10
1	0.	0.	0.	0.	1.	3.	5.	7.	6.	4.
2	0.	0.	0.	1.	3.	8.	13.	16.	15.	9.
3	0.	0.	0.	1.	5.	12.	20.	24.	23.	14.
4	0.	0.	500.	1.	6.	14.	23.	28.	27.	17.
5	0.	0.	0.	1.	5.	12.	20.	24.	23.	14.
6	0.	0.	0.	1.	3.	8.	12.	15.	14.	9.
7	0.	0.	0.	0.	1.	3.	5.	6.	6.	4.
8	0.	0.	0.	0.	0.	1.	1.	1.	1.	1.
9	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
10	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.

TABLE IV - 7
TARGET LOCATION DISTRIBUTIONS AFTER 5th UPDATE

(Example 2 Parallel Processing)

Target T₁

	1	2	3	4	5	6	7	8	9	10
1	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
2	0.	0.	0.	0.	0.	1.	1.	1.	1.	1.
3	32.	34.	32.	2.	0.	1.	1.	1.	1.	1.
4	38.	41.	536.	4.	1.	1.	1.	2.	2.	1.
5	44.	50.	41.	7.	1.	1.	1.	1.	1.	1.
6	13.	18.	12.	6.	2.	1.	1.	1.	1.	1.
7	9.	12.	8.	5.	1.	1.	0.	0.	0.	0.
8	3.	4.	3.	2.	1.	0.	0.	0.	0.	0.
9	1.	2.	2.	1.	1.	0.	0.	0.	0.	0.
10	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.

TABLE IV-7 (continued)

Target T₂

	1	2	3	4	5	6	7	8	9	10
1	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.
2	0.	0.	0.	0.	0.	1.	1.	1.	1.	1.
3	32.	34.	32.	2.	0.	1.	1.	1.	1.	1.
4	38.	41.	536.	4.	1.	1.	1.	2.	2.	1.
5	44.	50.	41.	7.	1.	1.	1.	1.	1.	1.
6	13.	18.	12.	6.	2.	1.	1.	1.	1.	1.
7	9.	12.	8.	5.	1.	1.	0.	0.	0.	0.
8	3.	4.	3.	2.	1.	0.	0.	0.	0.	0.
9	1.	2.	2.	1.	1.	0.	0.	0.	0.	0.
10	0.	0.	0.	0.	0.	0.	0.	0.	0.	0.

Example 2 demonstrates what we have come to refer to as the quiet target problem. The quiet target problem arises from the propensity of an information processor employing either sequential or parallel processing to overreact to a string of responses on a noisy target spanning a number of update stages when, at the outset, the initial location distribution for that target had significant overlap with the location distributions of one or more other targets that did not generate sensor responses during the same period. The overreaction that takes place is the gradual buildup in the area of the contact string of the probability mass not only of the target that is actually causing the contacts, but also of all other "quiet" targets that initially had some chance of being the responsible target. The net result is that the processor attempts to explain the string of contacts by placing with virtual certainty the locations of all of the targets it can in the area, rather than by simply attributing the chain of contacts to a single target.

Exactly this sort of situation typically arises in applications when one target is intrinsically more noisy than another. In this case the noisier one tends to produce a higher frequency of contacts and eventually the doubling up phenomenon described earlier takes place. This means that little or no overall probability is assigned to the area near the quieter target. Eventually this quiet target, however, will cause a contact, in which case the decision of the processor about the identity of the quiet target may be driven by extremely small tail probabilities in the distributions of a number of targets whose actual locations may be far from the location of the contact. The result is a very dramatic shift in the predicted location of such targets. This then is a typical source of the permutation problem which we also described earlier.

In short, it should be clear from our two examples that Extended Memory processing shows substantially improved target location prediction performance over the other two processing algorithms. Future developmental work on Extended Memory processing will concentrate on generalizing its use to include negative information and false targets.

REFERENCES

- [a] ASW Information Processing and Optimal Surveillance in a False Target Environment, Daniel H. Wagner, Associates Report to Office of Naval Research, by H. R. Richardson, October 31, 1973.
- [b] "Optimal Whereabouts Search," J. B. Kadane, Operations Research, Vol. 19, 1975, pp. 894-904.
- [c] Theory of Optimal Search, L. D. Stone, Academic Press, 1975.
- [d] "Information Theory and the Optimal Detection Search," W. H. Barker, Operations Research, Vol. 25, No. 2, March-April 1977, pp. 304-314.
- [e] Convex Functions, A. W. Roberts and D. E. Varberg, Academic Press, 1973.
- [f] Optimal Statistical Decisions, M. H. DeGroot, McGraw-Hill, 1970.
- [g] Distributions in Statistics: Continuous Multivariate Distributions, N. L. Johnson and S. Kotz, John Wiley and Sons, Inc., 1972.
- [h] Multidimensional Gaussian Distributions, K. S. Miller, John Wiley and Sons, Inc., 1964.

APPENDIX A

AUXILIARY FORMULAS FOR EVALUATING GAUSSIAN INTEGRALS

This appendix gives the formulas of multivariate Gaussian analysis and matrix relations which are needed in the analysis. Reference [h] is a standard reference for the Gaussian analysis.

Expressions for Multivariate Gaussian Quantities

The multivariate Gaussian density function on N-dimensional space is

$$f(x) = k \exp -\frac{1}{2}\{x'P^{-1}x\} \quad (A-1)$$

where P is the covariance matrix of x, and k is a constant so chosen that the integral of f over \mathbf{R}^N equals 1. It is well known that

$$\frac{1}{k} = \int_{\mathbf{R}^N} \exp -\frac{1}{2}\{x'P^{-1}x\} dx = (2\pi)^{N/2} \sqrt{|P|} \quad (A-2)$$

where $|P|$ is the determinant of P. Clearly, it is true that for any N-vector d,

$$\int_{\mathbf{R}^N} \exp -\frac{1}{2}\{(x-d)'P^{-1}(x-d)\} dx = \int_{\mathbf{R}^N} \exp -\frac{1}{2}\{x'P^{-1}x\} dx, \quad (A-3)$$

as can be seen by a simple change of variables.

Matrix Identities

Some useful identities for symmetric square matrices will now be given.

Completing the square. The first one is known as completing the square. Let

$$F(z, \gamma, \sigma) = (z-\gamma)' A(z-\gamma) + (z-\sigma)' B(z-\sigma)$$

be a second order homogeneous relationship in z where A is a positive definite matrix and B is a positive semi-definite matrix.

We claim that

$$F(z, \gamma, \sigma) = (z-\delta)' C(z-\delta) + (\gamma-\sigma)' D(\gamma-\sigma)$$

where

$$C = A+B,$$

$$D = B - B(A+B)^{-1}B$$

$$= A - A(A+B)^{-1}A$$

$$= A(A+B)^{-1}B,$$

$$\delta = C^{-1}(A\gamma + B\sigma)$$

$$= \gamma + (A+B)^{-1}B(\sigma - \gamma).$$

In order to establish the relationship, note that

$$F(z, \gamma, \sigma) = (z-\delta)' C(z-\delta) + L$$

where L is independent of z . Setting $z = \delta$ we then have

$$L = (\delta - \gamma)' A(\delta - \gamma) + (\delta - \sigma)' B(\delta - \sigma)$$

$$= (\gamma - \sigma)' B(A+B)^{-1} A A^{-1} A(A+B)^{-1} B(\gamma - \sigma)$$

$$+ (\gamma - \sigma)' A(A+B)^{-1} B(A+B)^{-1} A(\gamma - \sigma).$$

But observe that

$$\begin{aligned} B(A+B)^{-1} A &= (I - A(A+B)^{-1}) A \\ &= A - A(A+B)^{-1} A \\ &= A(A+B)^{-1} B, \end{aligned}$$

whence we obtain, as claimed,

$$\begin{aligned} L &= (\gamma - \sigma)' A(A+B)^{-1} (B + B A^{-1} B) (A+B)^{-1} A (\gamma - \sigma) \\ &= (\gamma - \sigma)' A(A+B)^{-1} B (I + A^{-1} B) (I + A^{-1} B)^{-1} (\gamma - \sigma) \\ &= (\gamma - \sigma)' A(A+B)^{-1} B (\gamma - \sigma). \end{aligned}$$

In particular, if B is invertible, then

$$L = (\gamma - \sigma)' (A^{-1} + B^{-1})^{-1} (\gamma - \sigma).$$

Inverse of a partitioned matrix. Let a matrix M and its inverse M^{-1} be partitioned into submatrices

$$M = \left[\begin{array}{c|c} A & B \\ \hline B' & C \end{array} \right], \quad (A-9)$$

$$M^{-1} = \left[\begin{array}{c|c} D & E \\ \hline E' & F \end{array} \right], \quad (A-10)$$

where the dimensions of A and D are the same, but possibly different from the dimensions of C and F. Define auxiliary matrices S and T by

$$S = A^{-1} B, \quad (A-11)$$

$$T = C - B'S. \quad (A-12)$$

Then

$$D = A^{-1} + ST^{-1}S', \quad (\text{A-13})$$

$$E = -ST^{-1}, \quad (\text{A-14})$$

$$F = T^{-1}. \quad (\text{A-15})$$

Note also that

$$|M| = |A| / |F| = |A| \cdot |T|. \quad (\text{A-16})$$

Equations (A-13) to (A-16) can be used directly to find M^{-1} and $|M^{-1}|$ given A^{-1} , B, C, and $|A|$. Alternately, given M^{-1} (i. e., D, E, F) and $|M^{-1}|$ these same relations can be solved to yield $|A|$ and A^{-1} . For this latter purpose, observe from (A-14) and (A-15) that

$$T = F^{-1}$$

$$S = -EF^{-1}$$

so that

$$A^{-1} = D - E(EF^{-1})' = D - EF^{-1}E', \quad (\text{A-17})$$

$$|A^{-1}| = 1/|M| \cdot |F| = |M^{-1}| / |F|. \quad (\text{A-18})$$

Dimension reduction. Let

$$\mathcal{J}(y) = (y-a)' M^{-1} (y-a) \quad (\text{A-19})$$

where M and a are given. Using the partitioning of (A-9), let M, M^{-1} , a, and y be partitioned as:

$$M = \begin{bmatrix} D & | & E \\ \hline E' & | & F \end{bmatrix}, \quad M^{-1} = \begin{bmatrix} A & | & B \\ \hline B' & | & C \end{bmatrix}, \quad a = \begin{bmatrix} \alpha \\ \hline \beta \end{bmatrix}, \quad y = \begin{bmatrix} y_\alpha \\ \hline y_\beta \end{bmatrix},$$

where the dimensions of D, A, α , and y_α are compatible. Then (A-19) can be

re-expressed as:

$$\mathcal{L}(y) = (c_{\alpha} - \sigma)' \Sigma^{-1} (y_{\alpha} - \sigma) + \tau \quad (\text{A-20})$$

where

$$\Sigma = D - EF^{-1}E', \quad (\text{A-21})$$

$$\sigma = \alpha + EF^{-1}(y_{\beta} - \beta), \quad (\text{A-22})$$

$$\tau = R + (y_{\beta} - \beta)' F^{-1} (y_{\beta} - \beta), \quad (\text{A-23})$$

$$|\Sigma| = |M| / |F|. \quad (\text{A-24})$$

These are all in terms of entries in M and a.

The point of recasting (A-19) into the form (A-20) is to reduce the dimension of the variable y_{α} (prior to integrating over y_{α}). Observe that σ is linear and τ is quadratic in y_{β} .

Another matrix inverse result.* Let a matrix A and its inverse $B = A^{-1}$ be partitioned into submatrices as follows:

$$A = \begin{pmatrix} A_{11} & A_{12} & A_{13} \\ A_{21} & A_{22} & A_{23} \\ A_{31} & A_{32} & A_{33} \end{pmatrix},$$

$$B = A^{-1} = \begin{pmatrix} B_{11} & B_{12} & B_{13} \\ B_{21} & B_{22} & B_{23} \\ B_{31} & B_{32} & B_{33} \end{pmatrix}.$$

* This section is based on work done by R. V. Kohn and S. S. Brown.

Let

$$\Delta = \begin{pmatrix} 0 & 0 & 0 \\ 0 & \Delta_{22} & 0 \\ 0 & 0 & 0 \end{pmatrix} .$$

If B and $|B|$ are known, what is $(A + \Delta)^{-1}$ and its determinant? Now

$$(A + \Delta)^{-1} = (A(I + A^{-1}\Delta))^{-1} = (I + B\Delta)^{-1}A^{-1}, \quad (\text{A-25})$$

where I is the identity matrix. Furthermore,

$$(I + B\Delta) = \begin{pmatrix} I_{11} & B_{12}\Delta_{22} & 0 \\ 0 & I_{22} + B_{22}\Delta_{22} & 0 \\ 0 & B_{32}\Delta_{22} & I_{33} \end{pmatrix}, \quad (\text{A-26})$$

where I_{11} , I_{22} , and I_{33} are identity matrices.

Let

$$\beta = (I_2 + B_{22}\Delta_{22})^{-1}. \quad (\text{A-27})$$

By direct evaluation

$$|(I + B\Delta)^{-1}| = |\beta|,$$

so that

$$|A + \Delta|^{-1} = |\beta| \cdot |B|. \quad (\text{A-28})$$

The inverse of (β) is, by inspection

$$(I+A^{-1}\Delta)^{-1} = \begin{pmatrix} I_{11} & -B_{12}\Delta_{22}\beta & 0 \\ 0 & \beta & 0 \\ 0 & -B_{32}\Delta_{22}\beta & I_{33} \end{pmatrix} = I + \begin{pmatrix} 0 & -B_{12}\Delta_{22}\beta & 0 \\ 0 & \beta - I_{22} & 0 \\ 0 & -B_{32}\Delta_{22}\beta & 0 \end{pmatrix} .$$

Now

$$\beta - I_{22} = (I_{22} - \beta^{-1})\beta = -B_{22}\Delta_{22}\beta .$$

Hence,

$$(I+A^{-1}\Delta)^{-1} = I - \begin{pmatrix} 0 & B_{12}\Delta_{22}\beta & 0 \\ 0 & B_{22}\Delta_{22}\beta & 0 \\ 0 & B_{32}\Delta_{22}\beta & 0 \end{pmatrix} ,$$

and $\Delta_{22}\beta = (I_2 + B_{22} \Delta_{22})^{-1} = (\Delta_{22}^{-1} + B_{22})^{-1}$ so that (A-25) becomes

$$(A + \Delta)^{-1} = B - \begin{pmatrix} B_{12} \\ B_{22} \\ B_{32} \end{pmatrix} (\Delta_{22}^{-1} + B_{22})^{-1} (B_{12} \ B_{22} \ B_{32}) . \quad (A-29)$$

APPENDIX B

EXTENDED MEMORY PROCESSING METHODOLOGY

In this appendix we develop the mathematical theory underlying the information updating algorithm we have called Extended Memory processing.

As the examples in the fourth section of Chapter IV clearly demonstrated, a Bayesian updating scheme, assuming target location independence in a multi-target environment and operating solely on the marginal target location distributions, generally runs into serious trouble when confronted with unresolved sensor responses. A sensor response (or response pattern when preprocessing is employed) is said to be resolvable if one can, at least in a Bayesian statistical sense, infer uniquely the identity of the target that generated the response. In ASW applications the resolvable responses may be limited to port arrivals and departures and any sensor response such that the prior target location distributions at the time of the response place only one target within possible detection range of the associated sensor(s).

In the first section below we introduce the notion of a response pattern-to-target association distribution and describe the use of such distributions in updating target location distributions. The updating of the response pattern-to-target association distributions themselves is described in the second section.

Response Pattern-to-Target Association (Scenario) Distributions

We begin by indexing the response patterns C_1, C_2, \dots in order of processing (usually approximate chronological order). Included in this list are all target port departures and arrivals. Also index the targets T_1, T_2, \dots in order of port departure. We assume here that all contacts are on valid targets. The generalization to false targets is reasonably straightforward.

We next define the random variables

$$N_i = \text{index of the target that caused response pattern } C_i, i = 1, 2, \dots$$

and the random vector

$$S_i = (N_1, N_2, \dots, N_i).$$

Suppose for the moment that the N_i were not random variables but deterministically known at the time each response pattern was reported to the processor. Assume, for example, that the radiated noise characteristics of the various targets are sufficiently different to permit this degree of target identification. The target coupling problem would then completely disappear since the target location distributions updated for both motion and responses, but conditioned on a specific association of responses with targets, would be conditionally mutually independent. In making this statement, we of course assume that the targets may be treated as moving independently.

Unfortunately, there are many cases in practice in which the identity of the target that generated a specific response pattern cannot be unambiguously inferred. Nevertheless, our statement about the conditional independence of the target location distributions given an assignment of all past responses to targets points the way toward a possible solution to the target coupling problem. The fundamental idea in Extended Memory processing is simply to produce a separate estimate of each currently active target's location distribution for each possible scenario, i. e., for each specific assignment of all past response patterns to targets. Weighted averages of these conditional location distributions, the weights being the probabilities associated with each scenario, would then give current estimates of each target's (unconditioned) location distribution.

To make all of this more precise requires some work. The first step is to define the probability distribution on scenarios after processing the i^{th} joint response pattern:

$$F_i(n_1, n_2, \dots, n_i) = \Pr \{N_1 = n_1, N_2 = n_2, \dots, N_i = n_i\}. \quad (\text{B-1})$$

We suppose that by one means or another, estimates of the F_i are available for each index i . A Bayesian procedure for generating such estimates will, in fact, be described in the next section. We now describe a recursive procedure for developing posterior estimates of (i. e., updating) the target location distributions. This procedure makes use both of the current distribution on scenarios and the prior target location distributions.

Associated with each response pattern index i , there is the index k_i of the information processing update stage at which C_i is processed. In current ASWIPS processing, k_i is simply the index of the first update following the occurrence of C_i . Also define i_k to be the index of the last response pattern processed through the end of the k^{th} stage. Thus, $C_{i_{k+1}}, C_{i_{k+2}}, \dots, C_{i_{k+1}}$ is the sequence of response

patterns processed during the $(k+1)^{\text{st}}$ stage update. If $i_{k+1} = i_k$, then no response patterns were observed during the period between the k^{th} and $(k+1)^{\text{st}}$ updates.

We now define the conditional target location weights

$$p_n^{(k)}(x | S_{i_k} = (n_1, n_2, \dots, n_{i_k})) = \Pr \left\{ \begin{array}{l} \text{target } n \text{ is in cell } x \text{ at} \\ \text{time of the } k^{\text{th}} \text{ stage update} \end{array} \middle| N_1 = n_1, N_2 = n_2, \dots, N_{i_k} = n_{i_k} \right\}, \quad (\text{B-2})$$

where it is to be understood in this definition that these weights $p_n^{(k)}$ reflect target motion through the end of the k^{th} update stage as well as all response patterns through C_{i_k} .

The key to Extended Memory is the following identity expressing the conditional independence of the target locations given the scenario:

$$\Pr \left\{ X^{(k)} = (x_1, x_2, \dots, x_N) \middle| S_{i_k} \right\} = \prod_{n=1}^N \Pr \left\{ X_n^{(k)} = x_n \middle| S_{i_k} \right\}, \quad (\text{B-3})$$

where, as in Chapter IV

$$X_n^{(k)} = \text{grid cell location of } n^{\text{th}} \text{ target at end of update } k,$$

$$X^{(k)} = (X_1^{(k)}, X_2^{(k)}, \dots, X_N^{(k)}).$$

Suppose now that $i_{k+1} > i_k$ so that at least one contact was reported between the k^{th} and $(k+1)^{\text{st}}$ updates and we wish to update the weights $p_n^{(k)}$ to a revised set $p_n^{(k+1)}$ accordingly. The case $i_{k+1} = i_k$ requires only target motion and negative information (not considered here) updates.

Let $p_n^{(k+1)}(\cdot | S_{i_k})$ denote the conditional target location distributions obtained from the $p_n^{(k)}(\cdot | S_{i_k})$ updated for $(k+1)^{\text{st}}$ stage motion and negative information. The next step is to generate the $p_n^{(k+1)}(\cdot | S_{i_{k+1}})$ from the $p_n^{(k+1)}(\cdot | S_{i_k})$.

Fix a target index n and suppose

$$S_{i_k} = (n_1, n_2, \dots, n_{i_k})$$

and

$$S_{i_{k+1}} = (n_1, n_2, \dots, n_{i_k+1}, \dots, n_{i_{k+1}}).$$

If none of the indices $n_{i_{k+1}}, \dots, n_{i_{k+1}}$ is equal to n , then $p_n^{(k+1)}(\cdot | S_{i_k} = (n_1, n_2, \dots, n_{i_k})) = p_n^{(k+1)}(\cdot | S_{i_{k+1}} = (n_1, n_2, \dots, n_{i_k}, n_{i_{k+1}}, \dots, n_{i_{k+1}}))$. If one or more of these indices is equal to n , then denote the ordered list of such indices by $\nu_1^{(n)}, \nu_2^{(n)}, \dots, \nu_{m_n}^{(n)}$. Thus, the given association $S_{i_{k+1}}$ assigns to target n precisely the m response patterns

$$C_{\nu_1^{(n)}}, C_{\nu_2^{(n)}}, \dots, C_{\nu_{m_n}^{(n)}}$$

out of those observed between the k^{th} and $(k+1)^{\text{st}}$ updates.

Let

$$\lambda_{\nu_1^{(n)}}, \lambda_{\nu_2^{(n)}}, \dots, \lambda_{\nu_{m_n}^{(n)}}$$

be the object location distributions corresponding to the response patterns

$$C_{\nu_1^{(n)}}, C_{\nu_2^{(n)}}, \dots, C_{\nu_{m_n}^{(n)}}.$$

Then

$$p_n^{(k+1)}(x | S_{i_{k+1}} = (n_1, n_2, \dots, n_{i_k}, n_{i_{k+1}}, \dots, n_{i_{k+1}})) = \frac{\left(\prod_{l=1}^{m_n} \lambda_{\nu_l^{(n)}}(x) \right) \cdot p_n^{(k+1)}(x | S_{i_k} = (n_1, n_2, \dots, n_{i_k}))}{\sum_y \left(\prod_{l=1}^{m_n} \lambda_{\nu_l^{(n)}}(y) \right) \cdot p_n^{(k+1)}(y | S_{i_k} = (n_1, n_2, \dots, n_{i_k}))} \quad (\text{B-4})$$

Equation (B-4) gives the updating formula for the conditional target location distributions. It remains to show how the unconditioned target location distributions are obtained. For this we define

$$p_n^{(k)}(x) = \Pr \{ \text{target } n \text{ is in cell } x \text{ at time of } k^{\text{th}} \text{ stage update} \},$$

where, as in the conditional target location distributions, it is to be understood that the $p_n^{(k)}$ reflect target motion through the end of the k^{th} update as well as all response patterns through C_{i_k} . Then, letting \mathcal{N}_k denote the set of possible k -vector values

of S_{i_k} , we set

$$p_n^{(k+1)}(x) = \sum_{(n_1, n_2, \dots, n_{i_{k+1}}) \in \mathcal{N}_{k+1}} F_{i_{k+1}}(n_1, n_2, \dots, n_{i_{k+1}}) p_n^{(k+1)}(x | S_{i_{k+1}} = (n_1, n_2, \dots, n_{i_{k+1}})). \quad (\text{B-5})$$

Thus, the unconditioned target location distribution for target n is an overlay of the conditional target location distributions weighted by probabilities $F_{i_{k+1}}$ of the corresponding scenarios. Observe that equation (B-5) has the effect of relating the current assessment of the location of target n to the scenario weights not just of the latest round of response patterns, $C_{i_{k+1}}, \dots, C_{i_{k+1}}$, but to those of all past response patterns. Thus, changes at the $(k+1)$ st update in the scenario weights of response patterns first recorded at earlier updates can have significant impact on the $(k+1)$ st stage target location distributions. This memory feature is an essential advantage of Extended Memory processing over both sequential processing and parallel processing.

Before turning to the Bayesian updating of the scenario distribution, a few comments are in order on the computational feasibility of the target location distribution updating procedure we have just described.

It would appear superficially that the number of conditional target location distributions which must be stored in computer memory and updated at stage k is on the order of $N \cdot 2^{i_k}$, where N is the number of active targets at stage k and, as defined earlier, i_k is the number of response patterns processed through stage k . Since i_k grows more or less linearly with time (assuming a roughly constant data rate), 2^{i_k} can become quite large for a lengthy processing period. With $N = 7$ and an average of 2 response patterns per update, $N \cdot 2^{i_k} = 7 \cdot 2^{100}$ after 50 updates. This number is even larger than the $(1000)^7$ elements of the state space associated with the multivariate joint location distribution for 7 targets in a grid of 1000 cells, making the situation look grim, indeed.

Fortunately, the situation in ASW application is not nearly as bad as the extreme we have described. First of all, many response patterns are resolvable, i. e., uniquely assignable to an identifiable target. Those response patterns which are unresolved generally are associable with one of two targets or, perhaps on rare occasions, with one of three or four targets. Another factor is that it frequently happens that even when a response pattern C_j is unresolved, it is uniquely associable with an earlier response pattern C_i ($i < j$) in the sense that C_i (within the accuracies of Bayesian inference) must have been caused by the same target (whatever its true identity) that caused C_j . This situation arises when $F_j(\dots, n_i, \dots, n_j) = 0$ for $n_i \neq n_j$. In such a case C_i is more or less redundant and can be effectively dropped (by summing on n_i) from the list of response patterns in favor of C_j .

A final factor in reducing the dimensionality of the problem relates to target motion. Suppose response pattern C_i was observed and processed at time t_k . Then as time passes, the current location distribution of the target that generated C_i diffuses until a point in time is reached at which that distribution is spread over such a large part of the target operating region that its usefulness in target localization is virtually nil. At this point C_i can reasonably be discarded from the list of active response patterns (again by summing over n_i in the current scenario distribution).

To put all of this in perspective, consider a situation with 7 active targets in which only the last 8 response patterns are carried in the scenario distribution, say C_1, C_2, \dots, C_8 . Suppose C_4, C_7 , and C_8 are resolvable with $N_4 = 3$, $N_7 = 6$, and $N_8 = 1$. In addition, assume C_1, C_2 , and C_5 are unresolved relative to T_2 and T_5 . Finally suppose that C_3 and C_6 are unresolved relative to T_4 and T_7 but that C_3 and C_6 are necessarily due to different targets. The number of conditional target location distributions carried in computer memory would then be: $3 + 2 \cdot 2^3 + 2 \cdot 2 = 23$, which is a far cry from $7 \cdot 2^8 = 1,792$. While this is only a representative case, it does tend to support the conclusion that the updating scheme we have described is computationally practical.

We now finally turn to the problem of the Bayesian updating of the distribution on scenarios.

The Bayesian Updating of the Distribution on Scenarios

We assume that the first response pattern C_1 is a port departure by target T_1 and that C_1 is the only pattern processed at the first update. These restrictions can be modified to correspond to any initial conditions including the allowance for holdover targets from earlier (prior to the start of the processing period) port departures. However, the case $F_1(1) = 1$ corresponds to the current setup in ASWIPS, so we make the stated specialization.

Suppose, inductively, that, the $F_{i_k}(n_1, n_2, \dots, n_{i_k})$ have been computed for every $(n_1, n_2, \dots, n_{i_k}) \in \mathcal{A}_k$. We now wish to generate $F_{i_{k+1}}$. If $i_{k+1} = i_k$, so that no response patterns were observed between the k^{th} and $(k+1)^{\text{st}}$ updates, then $F_{i_{k+1}} = F_{i_k}$, their domains of definition, in particular, being identical.

Suppose then that $i_{k+1} > i_k$ with (following our earlier notation) the response patterns $C_{i_k+1}, C_{i_k+2}, \dots, C_{i_{k+1}}$ generated between the k^{th} and $(k+1)^{\text{st}}$ updates. Consider a particular assignment of these $i_{k+1} - i_k$ responses to the active targets, i. e., assume

$$N_i = n_i \quad \text{for } i_k < i \leq i_{k+1}.$$

Let λ be the object location distribution for C_i , $i_k < i \leq i_{k+1}$. Then

$$F_{i_{k+1}}(n_1, n_2, \dots, n_{i_{k+1}}) = \alpha \cdot \prod_{n=1}^N \sum_x \left(\prod_{\substack{i=i_{k+1} \\ \ni n_i=n}}^{i_{k+1}} \lambda_i(x) \right) \cdot p_n^{(k+1)}(x | S_{i_k} = (n_1, n_2, \dots, n_{i_k})) F_{i_k}(n_1, n_2, \dots, n_{i_k}) \quad (\text{B-6})$$

where, as earlier, the $p_n^{(k+1)}(x | S_{i_k})$ are the conditional target location weights prior to the processing of C_i , $i_k < i \leq i_{k+1}$ but current to the end of the $(k+1)^{\text{st}}$ update stage for target motion. In equation (B-6), α is a renormalization constant to insure that

$$\sum_{(n_1, n_2, \dots, n_{i_{k+1}})} F_{i_{k+1}}(n_1, n_2, \dots, n_{i_{k+1}}) = 1.$$

This completes our description of the basic Extended Memory Bayesian processing procedure.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Information Processing to Maintain Localization in ASW Surveillance		5. TYPE OF REPORT & PERIOD COVERED Final
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) William H. Barker and Barry Belkin		8. CONTRACT OR GRANT NUMBER(s) N00014-76-C-0676
9. PERFORMING ORGANIZATION NAME AND ADDRESS Daniel H. Wagner, Associates Station Square One Paoli, PA 19301		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS Office of Naval Research Department of the Navy Arlington, VA 22217		12. REPORT DATE July 29, 1977
		13. NUMBER OF PAGES 159
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) sequential search and surveillance multiple targets Bayesian information processing false contacts moving targets		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) See Abstract page.		