

ESC-TR-2010-083

**Project Report
MD-26**

Bioinspired Resource Management for Multiple-Sensor Target Tracking Systems

H.C. Lambert
D. Sinno

20 June 2011

Lincoln Laboratory
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LEXINGTON, MASSACHUSETTS



Prepared for the Missile Defense Agency under Air Force Contract FA8721-05-C-0002.

Approved for public release; distribution is unlimited.

20110627117

MIT LINCOLN LABORATORY

0134



G18D279023N

This report is based on studies performed at Lincoln Laboratory, a center for research operated by Massachusetts Institute of Technology. This work was sponsored by the Missile Defense Agency/BC under Contract F48721-05-C-0002.

This report may be reproduced to satisfy needs of U.S. Government agencies.

The Missile Defense Agency Public Affairs Office has reviewed this report, and it is releasable to the National Technical Information Service, where it will be available to the general public, including foreign nationals.

This technical report has been reviewed and is approved for publication.

FOR THE COMMANDER



Gary Tutungian
Administrative Contracting Officer
Plans and Programs Directorate
Contracted Support Management

Non-Lincoln Recipients

PLEASE DO NOT RETURN

Permission has been given to destroy this document when it is no longer needed.

Massachusetts Institute of Technology
Lincoln Laboratory

**Bioinspired Resource Management for Multiple-Sensor
Target Tracking Systems**

*H.C. Lambert
D. Sinno
Group 36*

Project Report MD-26

20 June 2011

Approved for public release; distribution is unlimited.

Lexington

Massachusetts

This page intentionally left blank.

ABSTRACT

We present an algorithm, inspired by self-organization and stigmergy observed in biological swarms, for managing multiple sensors tracking large numbers of targets. We devise a decentralized architecture wherein autonomous sensors manage their own data collection resources and task themselves. Sensors cannot communicate with each other directly; however, a global track file, which is continuously broadcast, allows the sensors to infer their contributions to the global estimation of target states. Sensors can transmit their data (either as raw measurements or some compressed format) only to a central processor where their data are combined to update the global track file. We outline information-theoretic rules for the general multiple-sensor Bayesian target tracking problem. We provide specific formulas for problems dominated by additive white Gaussian noise. Using Cramér–Rao lower bounds as surrogates for error covariances, we illustrate, using numerical scenarios involving ballistic targets, that the bioinspired algorithm is highly scalable and performs very well for large numbers of targets.

This page intentionally left blank.

ACKNOWLEDGMENTS

The authors wish to thank Drs. K-P. Dunn and D. A. O'Connor of MIT Lincoln Laboratory for stimulating discussions and helpful suggestions.

This page intentionally left blank.

TABLE OF CONTENTS

	Page
Abstract	iii
Acknowledgments	v
List of Illustrations	ix
List of Tables	ix
1. INTRODUCTION	1
2. RELATIVE INFORMATION GAIN	5
2.1 Sequential Bayesian Estimation	5
2.2 Entropy and Divergence	6
3. DYNAMIC TARGET SELECTION	9
3.1 Self-Organization	9
3.2 Stigmergy	10
3.3 Data Flow	12
4. ADDITIVE WHITE GAUSSIAN NOISE	15
4.1 Sequential Bayesian Estimation	15
4.2 Entropy and Divergence	18
5. DETERMINISTIC MOTION	21
5.1 Cramér–Rao Lower Bounds	22
6. NUMERICAL RESULTS	25
7. SUMMARY AND DISCUSSION	29
REFERENCES	31

This page intentionally left blank.

LIST OF ILLUSTRATIONS

Figure No.		Page
1	"Forgetting" factor.	11
2	Data flow diagram: general noise models.	12
3	Data flow diagram: additive white Gaussian noise.	18
4	Average number of targets serviced.	27
5	Average number of targets serviced versus time.	28

LIST OF TABLES

Table No.		Page
1	Simulation parameters.	26

This page intentionally left blank.

1. INTRODUCTION

A major objective of multiple-sensor target tracking systems is to achieve certain levels of state estimation accuracy for as many targets as possible. In applications where a single sensor is tasked with collecting measurements on a single target, the management of the sensor's available resources does not pose any challenge. However, in situations where a small number of sensors are directed to collect measurements on a large number of targets, the decision process that determines which sensors are to collect measurements on which targets during the next observation time interval can be daunting. The sensor management problem becomes more difficult as the number of targets relative to the number of sensors increases. The problem is further exacerbated when targets and sensors are not arranged in a favorable geometric constellation. Ideally, one wishes to "service" as many targets as possible; that is, one would like to ensure that certain levels of accuracy, or "qualities of service," can be achieved in estimating the states of as many targets as possible. In this report, we present a target selection algorithm that is inspired by principles governing biological swarm intelligence.

In multiple-sensor target tracking systems, data collected by tasked sensors are often correlated and fused to form a global database of target state estimates and error covariances, referred to as the "global track file" (e.g., see [1] and [2]). Sensor data can be processed either at a central processing node or in a distributed fashion, giving rise to a plethora of data correlation and fusion architectures (e.g., see [3]). In this report, we pursue a mainly *decentralized* approach to sensor management that would be readily adaptable to a bioinspired implementation. However, our architecture is not devoid of central processing. In a way, it can be regarded as a hybrid. In our implementation, sensors are not allowed to directly communicate with each other, yet a global track file, maintained by a central processor, is continuously broadcast, and its contents are therefore assumed to be always available to all sensors. The sensors operate locally and autonomously by tasking their own resources based on the information available in the global track file. In our architecture, sensors make adjustments to their tasking decisions based on an assessment of their *anticipated* relative contributions to the global track file over a forthcoming data collection time interval; thus, the global track file provides an *indirect* means for sensors to communicate with each other. Sensors transmit their data only to the central processor, where they are correlated and fused to update the global track file.

A decentralized approach to sensor management has many benefits. Foremost, it is *robust* in that the overall system performance is not significantly degraded by the failure of a few member sensors. More importantly, a decentralized implementation is *scalable*; that is, the complexity of the implementation does not increase significantly as more sensors are added to the system. Since all tasking decisions are made locally by the sensors themselves, the computing and communication bandwidth requirements are less stringent than they would be in a purely centralized approach. Many examples of decentralized approaches to sensor management can be found in the literature. For example, in [4], the authors apply a novel game-theoretic approach to dynamic target selection. Of relevance to our work is [5], wherein the authors employ an information-theoretic approach to decentralized sensor management.

Bioinspired approaches to solving combinatorial optimization problems follow either or both of two main principles observed in nature: “self-organization” and “stigmergy.” For a full and lucid account on swarm intelligence and its applicability to solving combinatorial optimization problems, we refer the reader to [6]. There we learn that in biological swarms, such as schools of fish or flocks of birds, individual members, without any awareness of actions taken by other members, are capable of unconsciously arriving at a meaningful, and often very complex, global behavior by merely following a set of very simple rules. This phenomenon is referred to as “self-organization” or “emergent behavior.” Self-organization is not limited to biological systems and is also observed in other complex decentralized systems such as free-market economies. An application of self-organization to sensor management was developed in [7] (see also [8]). We extend and formalize the rules considered in [7] and [8] using the information-theoretic approaches employed in [5], [9], and [10].

In biological swarms, simple local rules alone may not be sufficient to provide the means for the system as a whole to reach a particular global goal. Although individual members in swarms are not consciously aware of actions taken by other members, they can nevertheless communicate with each other *indirectly* through the impact they lay on their environment. As individuals impact their environment through their actions, the environment, in turn, influences individual behavior, thereby fostering a learning mechanism. This phenomenon is referred to as “stigmergy” or “sematectonic communication.” A prime example in nature is that of ant trails. Frequently trodden paths to existing food sources are doused with pheromone sprayed by passing ants. Pheromone traces, in turn, encourage further traffic. Once a particular food source has been exhausted, ants, following simple rules, begin searching for new sources of nourishment. At this point, the trail to the empty food source becomes less frequented, and extant pheromone traces begin to evaporate, thereby discouraging further visitation. Ants who find a new food source by leaving behind new traces of pheromone then encourage other ants to visit the new path, and the process begins anew. The application of *both* self-organization and stigmergy to sensor management was first explored in [7]. By employing the information-theoretic formulation of [10], we extend and formalize ideas originally developed in [11], [12], and [7] to allow sensors to communicate with each other indirectly via the global track file that represents the “environment” in our application.

Throughout this report, we assume data association to be perfect. In many real-world problems, poor data association can be a major limiting factor to the reliable performance of target tracking systems, and its impact on sensor management must therefore be brought to bear (e.g., see [1] and [2]). Nevertheless, even in cases of perfect data association, sensor management can be a daunting task when the number of sensors is significantly smaller than the number of targets. Although *absolute* performance predictions would be optimistic, we believe that the formulation we present in this report can at least shed light on the *relative* performance of candidate target selection algorithms. We have also ignored the impact of sensor biases that can have an adverse effect on data association. We defer to a future investigation the impact of poor data association and sensor biases on the performance of the target selection algorithm presented in this report.

We begin in Section 2 by reviewing the sequential Bayesian formulation of the single-target tracking problem. In Section 2, we also present the Rényi α -entropy and α -divergence [13] that are extensively utilized in our information-theoretic approach (cf. [9] and [10]). The bioinspired target

selection algorithm for the general Bayesian formulation is presented in Section 3. In Section 4, we present the specific formulation of the multiple-sensor target tracking problem for systems dominated by additive white Gaussian noise. In Section 5, we focus on targets whose motion can be described by a deterministic system dynamic model. Deterministic dynamics allow a simple recursive formulation of the Cramér–Rao lower bound on the error covariance associated with the estimated target state [14]. In Section 6, we use the Cramér–Rao lower bounds derived in Section 5 as surrogates for target state estimation error covariances to evaluate the performance of the target selection algorithm outlined in Section 3 for an example scenario involving ballistic targets. Lastly, in Section 7, we summarize our results and discuss how the bioinspired target selection algorithm may be further improved for specific applications.

This page intentionally left blank.

2. RELATIVE INFORMATION GAIN

In this report, we adopt an information-theoretic approach to target selection wherein we favor targets whose measurements, once collected during an upcoming observation period, would result in the largest reduction in the uncertainties associated with their state estimates (cf. [5], [9], and [10]). We begin this section by reviewing the recursive Bayesian formulation of the single-target tracking problem. We then provide general expressions for the information-theoretic measures utilized by the bioinspired target selection logic outlined in Section 3.

2.1 SEQUENTIAL BAYESIAN ESTIMATION

Within the Bayesian framework, we seek to estimate the probability density function of a target's state vector, $\mathbf{x}_k \in \mathbb{R}^n$, valid at time t_k , conditioned on a sequence of noisy measurements [15]. Useful statistical measures, such as the mean and the covariance, are then obtained directly from the conditional density. We limit the discussion to the class of Markov processes. The general nonlinear system dynamic model prescribing the temporal evolution of the target state vector from time t_{k-1} to time t_k is assumed to be given by

$$\mathbf{x}_k = \varphi_{k-1}(\mathbf{x}_{k-1}, \mathbf{w}_{k-1}), \quad (1)$$

where \mathbf{w}_{k-1} is a random vector sampled from the known probability density function $p(\mathbf{w}_{k-1})$. The general nonlinear observation model mapping the target state space to the measurement space is, in turn, assumed to be given by

$$\mathbf{z}_k = \mathbf{h}_k(\mathbf{x}_k, \mathbf{v}_k), \quad (2)$$

where \mathbf{v}_k is a random vector sampled from the known probability density function $p(\mathbf{v}_k)$. The noise vectors \mathbf{w}_{k-1} and \mathbf{v}_k in (1) and (2), respectively, are assumed to be white and mutually independent.

Given a sequence of measurements

$$\mathbf{z}_{1:k} = \{\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_k\}, \quad (3)$$

we wish to estimate the posterior conditional density, $p(\mathbf{x}_k | \mathbf{z}_{1:k})$, of the target state vector, \mathbf{x}_k , valid at time t_k . Models (1) and (2), with their underlying assumptions of Markovianity and white and independent noise vectors, render a recursive formulation. Given the density, $p(\mathbf{x}_{k-1} | \mathbf{z}_{1:k-1})$, of the target state vector valid at time t_{k-1} conditioned on the sequence of measurements collected up to time t_{k-1} , the *prior* conditional density, $p(\mathbf{x}_k | \mathbf{z}_{1:k-1})$, of the target state vector valid at time t_k can be obtained from the Chapman–Kolmogorov equation:

$$p(\mathbf{x}_k | \mathbf{z}_{1:k-1}) = \int p(\mathbf{x}_k | \mathbf{x}_{k-1}) p(\mathbf{x}_{k-1} | \mathbf{z}_{1:k-1}) d\mathbf{x}_{k-1}, \quad (4)$$

where $p(\mathbf{x}_k | \mathbf{x}_{k-1})$ is the state transition density obtained from (1). The *posterior* density, $p(\mathbf{x}_k | \mathbf{z}_{1:k})$, conditioned on the sequence of measurements collected up to time t_k is, in turn, obtained from the Bayes rule:

$$p(\mathbf{x}_k | \mathbf{z}_{1:k}) = \frac{p(\mathbf{z}_k | \mathbf{x}_k) p(\mathbf{x}_k | \mathbf{z}_{1:k-1})}{p(\mathbf{z}_k | \mathbf{z}_{1:k-1})}, \quad (5)$$

where the likelihood function, $p(\mathbf{z}_k | \mathbf{x}_k)$, is obtained from (2), and

$$p(\mathbf{z}_k | \mathbf{z}_{1:k-1}) = \int p(\mathbf{z}_k | \mathbf{x}_k) p(\mathbf{x}_k | \mathbf{z}_{1:k-1}) d\mathbf{x}_k. \quad (6)$$

It follows that starting with an initial probability density function, $p(\mathbf{x}_0)$, of the target state vector, \mathbf{x}_0 , valid at a time t_0 prior to any observation, relations (4) and (5) can be applied recursively to update the conditional probability density function of the target state vector as new measurements become available.

2.2 ENTROPY AND DIVERGENCE

The contribution of each new measurement to the conditional probability density function of a target's state by means of the Bayes rule (5) amounts to a reduction in the uncertainty associated with the target's state estimate. A reduction in uncertainty can be equivalently regarded as a gain in *information*. The *Rényi α -entropy* provides a general scalar measure of uncertainty [10]:

$$\mathcal{H}_\alpha(\mathbf{z}_{1:k}) = \frac{1}{1-\alpha} \log \int p^\alpha(\mathbf{x}_k | \mathbf{z}_{1:k}) d\mathbf{x}_k, \quad (7)$$

where $0 < \alpha < 1$. In information theory, “information” is defined as the negative of entropy [16]; that is, $-\mathcal{H}_\alpha(\mathbf{z}_{1:k})$. Applying L'Hôpital's rule, it follows that as α approaches unity, the Rényi α -entropy (7) reduces to the *Shannon entropy*:

$$\mathcal{H}_1(\mathbf{z}_{1:k}) = \lim_{\alpha \rightarrow 1} \mathcal{H}_\alpha(\mathbf{z}_{1:k}) = - \int p(\mathbf{x}_k | \mathbf{z}_{1:k}) \log p(\mathbf{x}_k | \mathbf{z}_{1:k}) d\mathbf{x}_k. \quad (8)$$

A general scalar measure of the relative information between the prior and posterior probability density functions, $p(\mathbf{x}_k | \mathbf{z}_{1:k-1})$ and $p(\mathbf{x}_k | \mathbf{z}_{1:k})$, respectively, is, in turn, obtained from the Rényi α -divergence [10]:

$$\mathcal{D}_\alpha(\mathbf{z}_{1:k} \| \mathbf{z}_{1:k-1}) = \frac{1}{\alpha-1} \log \int \left[\frac{p(\mathbf{x}_k | \mathbf{z}_{1:k})}{p(\mathbf{x}_k | \mathbf{z}_{1:k-1})} \right]^\alpha p(\mathbf{x}_k | \mathbf{z}_{1:k-1}) d\mathbf{x}_k. \quad (9)$$

Again, applying L'Hôpital's rule, it follows that as α approaches unity, the Rényi α -divergence reduces to the *Kullback–Leibler divergence*:

$$\mathcal{D}_1(\mathbf{z}_{1:k} \| \mathbf{z}_{1:k-1}) = \lim_{\alpha \rightarrow 1} \mathcal{D}_\alpha(\mathbf{z}_{1:k} \| \mathbf{z}_{1:k-1}) = \int p(\mathbf{x}_k | \mathbf{z}_{1:k}) \log \frac{p(\mathbf{x}_k | \mathbf{z}_{1:k})}{p(\mathbf{x}_k | \mathbf{z}_{1:k-1})} d\mathbf{x}_k. \quad (10)$$

When $\alpha = 1/2$, we obtain the *Hellinger affinity*:

$$\mathcal{D}_{1/2}(\mathbf{z}_{1:k} \| \mathbf{z}_{1:k-1}) = -2 \log \int \sqrt{p(\mathbf{x}_k | \mathbf{z}_{1:k}) p(\mathbf{x}_k | \mathbf{z}_{1:k-1})} d\mathbf{x}_k, \quad (11)$$

where the integral on the right-hand side of (11) is recognized as the *Bhattacharyya coefficient*.

The Rényi α -divergence is a measure of dissimilarity between two distributions. Although it violates the symmetry and triangle inequality conditions, it is nevertheless helpful to think of

the Rényi α -divergence as a distance function. For certain special cases, the Rényi α -divergence can in fact be appropriately transformed to yield a proper distance function. For example, in the case of $\alpha = 1/2$, the *Hellinger distance*, which satisfies all the conditions required to be a proper mathematical metric, is given by

$$\mathfrak{H}(\mathbf{z}_{1:k} \parallel \mathbf{z}_{1:k-1}) = \sqrt{1 - \exp\left[-\frac{\mathcal{D}_{1/2}(\mathbf{z}_{1:k} \parallel \mathbf{z}_{1:k-1})}{2}\right]}, \quad (12)$$

where the exponential expression under the square root sign on the right-hand side of (12) is recognized as the *Bhattacharyya distance*.

To determine on which target a sensor should collect measurements during the next observation period, we adopt the general strategy pursued by [5] and [10] by favoring the target whose expected measurements would result in the largest gain in relative information; that is, we select the target whose Rényi α -divergence would be the largest. The choice of α can have an impact on performance. Values of $\alpha = 1$ and $\alpha = 1/2$, corresponding to the Kullback–Leibler divergence and Hellinger affinity, respectively, have been reported to provide favorable results [10]. The value of $\alpha = 1/2$ appears to serve as a better discriminant between two similar densities—such as the prior and posterior conditional probability density functions encountered in target tracking problems. The choice of α need not remain constant and can be allowed to vary adaptively. The prior and posterior conditional probability density functions of newly formed tracks are arguably less similar in shape than the prior and posterior conditional densities of mature tracks, thereby suggesting values of α closer to unity for newly formed tracks and values of α closer to $1/2$ for mature tracks. In the numerical example presented in Section 6, we show results for $\alpha = 1/2$. The optimal choice of α for the sensor management algorithm outlined in this report will be pursued as part of a future study.

This page intentionally left blank.

3. DYNAMIC TARGET SELECTION

The decentralized dynamic target selection logic we present is predicated upon favoring the target whose measurements, once collected during an upcoming observation period, would provide the largest *weighted* information gain. We choose the Rényi α -divergence as a measure of relative information gain. The weighting, in turn, enables the local sensor to learn from the environment (here, the global track file) the contribution of the other sensors with which it has no direct contact. Furthermore, prior to maximization, we apply a set of preselection rules that reduce the processing load. Since no measurements on any targets are available prior to any data collection, the *potential* information gain is obtained by computing the *expected value* of the Rényi α -divergence with respect to the probability density function of the measurements to be collected during the forthcoming observation period.

In order to expedite the formulation of the decentralized dynamic target selection logic, we introduce some new notation first. We denote the set of measurements originating from the ν th target (where $\nu = 1, \dots, N_t$) collected by the μ th sensor (where $\mu = 1, \dots, N_s$) within the time interval $[t_{k-1}, t_k)$ with

$$\mathbf{z}_k^{\mu\nu} = \left\{ \mathbf{z}_{\ell_\mu}^{\mu\nu} : t_{k-1} \leq t_{\ell_\mu} < t_k \right\}; \quad (13)$$

we denote the set of *all* measurements originating from the ν th target collected by the μ th sensor up to (but not including) time t_k with

$$\mathbf{z}_{1:k}^{\mu\nu} = \left\{ \mathbf{z}_1^{\mu\nu}, \mathbf{z}_2^{\mu\nu}, \dots, \mathbf{z}_k^{\mu\nu} \right\}; \quad (14)$$

and we denote the set of all measurements originating from the ν th target collected by *all* sensors up to (but not including) time t_k with

$$\mathbf{z}_{1:k}^\nu = \bigcup_{\mu} \mathbf{z}_{1:k}^{\mu\nu}. \quad (15)$$

Our bioinspired target selection logic applies the principles of self-organization and stigmergy that are often encountered in nature. In the sequel, we give detailed accounts of how we implement each principle within the context of sensor management.

3.1 SELF-ORGANIZATION

In our application, sensors following simple local target selection rules induce the multiple-sensor target tracking system to arrive at a global database of estimated target state vectors that meet their desired qualities of service. Here, we consider three simple rules. We begin by selecting targets that have not yet met their qualities of service (QoS) at time t_{k-1} :

$$\mathbb{S}_1^\mu = \left\{ \nu \mid \mathcal{H}_\alpha(\mathbf{z}_{1:k-1}^\nu) > \text{QoS}_k^\nu \right\}. \quad (16)$$

If \mathbb{S}_1^μ is empty, then no further action is required as all targets have already met their desired qualities of service. From the list of remaining targets (if any), we then select those that *would*

meet their qualities of service if new measurements were collected on them and were subsequently incorporated by the μ th sensor to update their local state estimates:

$$\mathbb{S}_2^\mu = \left\{ \nu \mid \nu \in \mathbb{S}_1^\mu, \hat{\mathcal{H}}_\alpha(\mathbf{z}_k^{\mu\nu}, \mathbf{z}_{1:k-1}^\nu) \leq \text{QoS}_k^\nu \right\}, \quad (17)$$

where

$$\hat{\mathcal{H}}_\alpha(\mathbf{z}_k^{\mu\nu}, \mathbf{z}_{1:k-1}^\nu) = \int p(\mathbf{z}_k^{\mu\nu} \mid \mathbf{z}_{1:k-1}^\nu) \mathcal{H}_\alpha(\mathbf{z}_k^{\mu\nu}, \mathbf{z}_{1:k-1}^\nu) d\mathbf{z}_k^{\mu\nu} \quad (18)$$

is the *expected* entropy of the posterior probability density function, and $p(\mathbf{z}_k^{\mu\nu} \mid \mathbf{z}_{1:k-1}^\nu)$ denotes the probability density function of the measurements originating from the ν th target expected to be collected by the μ th sensor. We revert to \mathbb{S}_1^μ if \mathbb{S}_2^μ is empty. Lastly, from the list of desired targets (if any), we select the target whose measurements would yield the largest gain in relative information:

$$\mu_\nu = \arg \max_\nu \left(\alpha \hat{\Delta}_k^{\mu\nu} \right), \quad (19)$$

where

$$\alpha \hat{\Delta}_k^{\mu\nu} = \int p(\mathbf{z}_k^{\mu\nu} \mid \mathbf{z}_{1:k-1}^\nu) \mathcal{D}_\alpha(\mathbf{z}_k^{\mu\nu}, \mathbf{z}_{1:k-1}^\nu \parallel \mathbf{z}_{1:k-1}^\nu) d\mathbf{z}_k^{\mu\nu} \quad (20)$$

is the *expected* local information gain.

3.2 STIGMERGY

The simple local rules considered thus far do not employ any learning. Although sensors cannot communicate with each other directly, they can nevertheless infer the impact of their measurements on the global track file relative to those of other sensors by comparing their local information gain achieved for a particular target during the *last* observation period with the global information gain that was achieved for the same target by the multiple-sensor system as a whole. If the local information gain achieved by the μ th sensor for the ν th target is the same or nearly the same as the global information gain on the ν th target, then the μ th sensor has been the major contributor to updating the global state estimate of the ν th target. Under this circumstance, the μ th sensor should continue collecting data on the ν th target as it is very likely that the other sensors are unable to either observe the ν th target or have been obliged by their local rules to attend to other targets. On the other hand, if the local information gain achieved by the μ th sensor for the ν th target is significantly smaller than the global information gain on the ν th target, then the μ th sensor could not have been a major contributor and should therefore be induced to collect measurements on some other target in urgent need of attention.

Inspired by biology, the global track file serves as the “environment” in our application and provides an indirect method of communication between the sensors. To include the environment’s impact on the target selection logic, we modify the local rule (19) to read as

$$\mu_\nu = \arg \max_\nu \left(w_{k-1}^{\mu\nu} \alpha \hat{\Delta}_k^{\mu\nu} \right), \quad (21)$$

where the weight, $w_{k-1}^{\mu\nu}$, with $0 < w_{k-1}^{\mu\nu} \leq 1$, serves as a *learning index*. Assuming that all targets yield the same expected information gain, $\alpha \hat{\Delta}_k^{\mu\nu}$, relation (21) implies that targets with weights

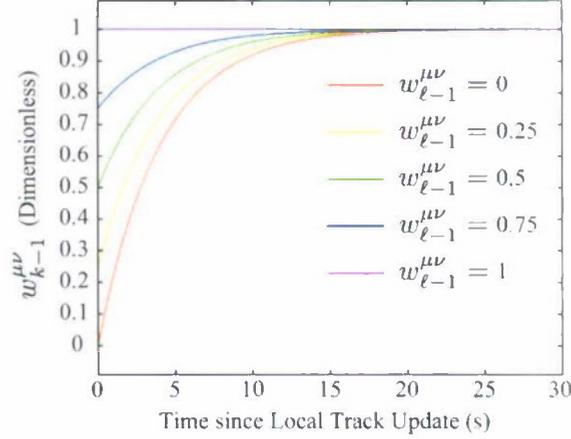


Figure 1. Heuristic function modeling the “forgetting” effect in the learning index, $w_{k-1}^{\mu\nu}$, when the last local track update occurs at a time $t_{\ell-1}$ earlier than the last global track update time, t_{k-1} .

closer to unity will be favored over targets with weights closer to zero. When there exists a local update at time t_{k-1} , we choose the learning index to be given by

$$w_{k-1}^{\mu\nu} = \frac{\alpha \Delta_{k-1}^{\mu\nu}}{\alpha \Delta_{k-1}^{\nu}}, \quad (22)$$

where

$$\alpha \Delta_{k-1}^{\mu\nu} = \mathcal{D}_{\alpha}(\mathbf{z}_{1:k-1}^{\mu\nu}, \mathbf{z}_{1:k-2}^{\nu} \parallel \mathbf{z}_{1:k-2}^{\nu}) \quad (23)$$

is the actual gain in local information obtained by the μ th sensor on the ν th target at time t_{k-1} , and

$$\alpha \Delta_{k-1}^{\nu} = \mathcal{D}_{\alpha}(\mathbf{z}_{1:k-1}^{\nu} \parallel \mathbf{z}_{1:k-2}^{\nu}) \quad (24)$$

is the global information gain for the ν th target at time t_{k-1} , which includes measurements collected by all sensors (including the μ th sensor). If $\alpha \Delta_{k-1}^{\mu\nu}$ is close to $\alpha \Delta_{k-1}^{\nu}$, then $w_{k-1}^{\mu\nu}$ is close to unity. Under this circumstance, the μ th sensor is perceived as a major contributor to updating the global state estimate of the ν th target and thus favors to continue collecting data on the ν th target. On the other hand, if $\alpha \Delta_{k-1}^{\mu\nu}$ is significantly smaller than $\alpha \Delta_{k-1}^{\nu}$, then the μ th sensor could not have been a major contributor to updating the global state estimate of the ν th target and thus should attend to other targets.

When the last local update occurred at some time $t_{\ell-1} < t_{k-1}$, we choose the following heuristic relation:

$$w_{k-1}^{\mu\nu} = w_{\ell-1}^{\mu\nu} + (1 - w_{\ell-1}^{\mu\nu}) \left[1 - \exp\left(-\frac{t_{k-1} - t_{\ell-1}}{\tau}\right) \right], \quad (25)$$

where τ is a globally specified time constant determining the rate at which the learning index approaches unity—thereby inducing a “forgetting” effect—when there are no local measurement updates (see Figure 1).

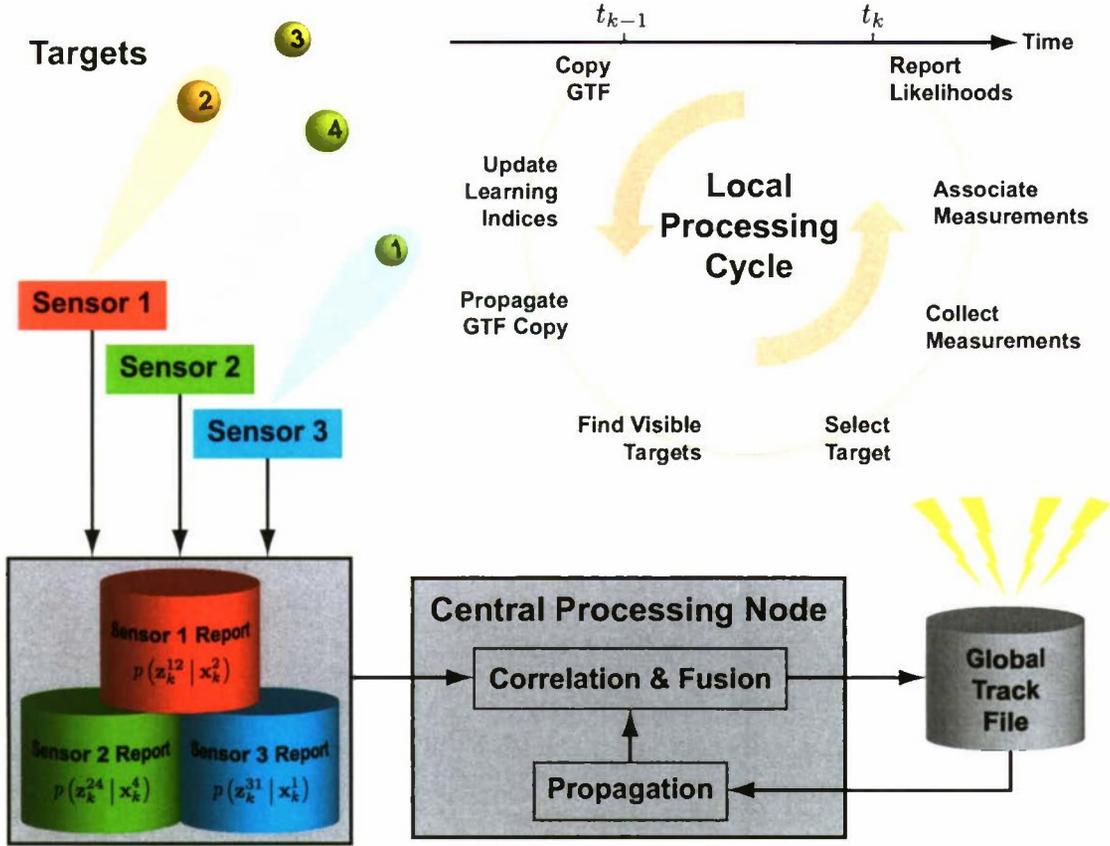


Figure 2. Data flow diagram for the distributed sensor management architecture for problems with general system dynamic and observation models.

3.3 DATA FLOW

The data flow diagram for the distributed sensor management architecture presented in this report is shown in Figure 2. In this architecture, autonomous sensors are responsible for allocating their own data collection resources based on (1) the information available in the global track file (GTF) and (2) the expected relative information to be gained during the forthcoming data collection time interval, $t_{k-1} \leq t < t_k$. The local processing cycle is detailed in the box in the upper right corner of Figure 2. While sensor management occurs strictly locally at the sensor level, track maintenance occurs globally at the central processing node. In other words, our architecture does not require maintenance of local track files. The global track file is populated and maintained only with reported likelihoods valid during the data collection time interval, $t_{k-1} \leq t < t_k$.

As discussed earlier, in this report we assume perfect data association; that is, we assume there is no ambiguity as to which measurement is originated from which target. In dense target environments, data association imposes a serious challenge and cannot be ignored. Although not pursued any further in this report, in our sensor management architecture, data association must

be performed both at the local and global levels. At the global level, the measurement likelihoods reported from multiple sensors must be correlated with propagated global tracks prior to fusion. At the sensor level (see top-right box in Figure 2), once measurements have been collected, they must be associated with the propagated copies of the global tracks. For the scenario shown on the left-hand side of Figure 2, each sensor tracks only one target at a time. For example, sensor 1 has allocated its resources to collect measurements solely on target 2. In this case, local data association would not be necessary. However, if target 2 in Figure 2 represented a dense target *complex* instead (that is, a *family* of closely-spaced objects as opposed to a single object), then the multiple measurements originating from the complex must be associated with the propagated copies of the global tracks that form the cluster.

This page intentionally left blank.

4. ADDITIVE WHITE GAUSSIAN NOISE

Although somewhat specialized, the sources of system dynamic and measurement uncertainties in many target tracking problems are often modeled to be additive, white, and Gaussian. Such modeling has the advantage of allowing the computation of the prior and posterior conditional probability density functions needed for the sequential Bayesian estimation of the target state to be reduced to the computation of their first two statistical moments—mean and covariance.

4.1 SEQUENTIAL BAYESIAN ESTIMATION

Under the additive white Gaussian noise assumption, the nonlinear system dynamic model (1) for the ν th target reduces to

$$\mathbf{x}_k^\nu = \boldsymbol{\varphi}_{k-1}^\nu(\mathbf{x}_{k-1}^\nu) + \mathbf{w}_{k-1}^\nu, \quad (26)$$

where

$$\mathbf{w}_{k-1}^\nu \sim \mathcal{N}(\mathbf{w}_{k-1}; \mathbf{0}, \mathbf{Q}_{k-1}^\nu), \quad (27)$$

is a white zero-mean Gaussian random vector with covariance \mathbf{Q}_{k-1}^ν . The function $\mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}) = \frac{1}{\sqrt{|2\pi\boldsymbol{\Sigma}|}} \exp\left[-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1}(\mathbf{x} - \boldsymbol{\mu})\right]$ denotes the probability density function of a Gaussian random variable, \mathbf{x} , with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$. Similarly, under the additive white Gaussian noise assumption, the nonlinear observation model (2) for the μ th sensor and the ν th target reduces to

$$\mathbf{z}_k^{\mu\nu} = \mathbf{h}_k^\mu(\mathbf{x}_k^\nu) + \mathbf{v}_k^{\mu\nu}, \quad (28)$$

where

$$\mathbf{v}_k^{\mu\nu} \sim \mathcal{N}(\mathbf{v}_k; \mathbf{0}, \mathbf{R}_k^{\mu\nu}) \quad (29)$$

is a white zero-mean Gaussian random vector with covariance $\mathbf{R}_k^{\mu\nu}$.

Since the process noise vector, \mathbf{w}_{k-1}^ν , is additive, white, and Gaussian, it follows from (26) that the state transition probability density function for the state vector of the ν th target from time t_{k-1} to time t_k is given by

$$p(\mathbf{x}_k^\nu | \mathbf{x}_{k-1}^\nu) = \mathcal{N}(\mathbf{x}_k^\nu; \boldsymbol{\varphi}_{k-1}^\nu(\mathbf{x}_{k-1}^\nu), \mathbf{Q}_{k-1}^\nu). \quad (30)$$

Starting with a Gaussian probability density function for the ν th target's state vector, \mathbf{x}_{k-1}^ν , valid at time t_{k-1} , conditioned on the sequence of measurements, $\mathbf{z}_{1:k-1}$:

$$p(\mathbf{x}_{k-1}^\nu | \mathbf{z}_{1:k-1}^\nu) = \mathcal{N}(\mathbf{x}_{k-1}^\nu; \hat{\mathbf{x}}_{k-1|k-1}^\nu, \mathbf{P}_{k-1|k-1}^\nu), \quad (31)$$

multiplying (30) and (31), employing a first-order Taylor expansion of the system dynamic function, $\boldsymbol{\varphi}_{k-1}^\nu(\cdot)$, about the state estimate, $\hat{\mathbf{x}}_{k-1|k-1}^\nu$, and integrating the result with respect to \mathbf{x}_{k-1}^ν over \mathbb{R}^n , it follows that the *prior* conditional probability density function of the state vector of the ν th target valid at time t_k can be approximated with a Gaussian function:

$$p(\mathbf{x}_k^\nu | \mathbf{z}_{1:k-1}^\nu) \simeq \mathcal{N}(\mathbf{x}_k^\nu; \hat{\mathbf{x}}_{k|k-1}^\nu, \mathbf{P}_{k|k-1}^\nu), \quad (32)$$

where $\hat{\mathbf{x}}_{k|k-1}^\nu$ and $\mathbf{P}_{k|k-1}^\nu$ denote the predicted state estimate and error covariance of the ν th target, respectively. It can be shown that they are given by [17]

$$\hat{\mathbf{x}}_{k|k-1}^\nu = \varphi_{k-1}^\nu \left(\hat{\mathbf{x}}_{k-1|k-1}^\nu \right) \quad (33)$$

and

$$\mathbf{P}_{k|k-1}^\nu = \hat{\Phi}_{k-1}^\nu \mathbf{P}_{k-1|k-1}^\nu \left(\hat{\Phi}_{k-1}^\nu \right)^\top + \mathbf{Q}_{k-1}^\nu, \quad (34)$$

where

$$\hat{\Phi}_{k-1}^\nu = \frac{\partial \varphi_{k-1}^\nu \left(\hat{\mathbf{x}}_{k-1|k-1}^\nu \right)}{\partial \mathbf{x}_{k-1}} \quad (35)$$

is the state transition Jacobian matrix.

Similarly, since the measurement noise vector, $\mathbf{v}_k^{\mu\nu}$, is additive, white, and Gaussian, it follows from (28) that the likelihood function associated with the measurement vector, $\mathbf{z}_k^{\mu\nu}$, originating from the ν th target and collected by the μ th sensor is given by

$$p \left(\mathbf{z}_k^{\mu\nu} | \mathbf{x}_k^\nu \right) = \mathcal{N} \left(\mathbf{z}_k^{\mu\nu}; \hat{\mathbf{z}}_k^{\mu\nu}, \mathbf{R}_k^{\mu\nu} \right), \quad (36)$$

where

$$\hat{\mathbf{z}}_k^{\mu\nu} = \mathbf{h}_k^\mu \left(\hat{\mathbf{x}}_{k|k-1}^\nu \right) \quad (37)$$

denotes the predicted measurement vector. Multiplying (32) and (36), employing a first order Taylor expansion of the measurement function, $\mathbf{h}_k^\mu(\cdot)$, about the predicted state estimate, $\hat{\mathbf{x}}_{k|k-1}^\nu$, and integrating the result with respect to \mathbf{x}_k^ν over \mathbb{R}^n , it follows that

$$p \left(\mathbf{z}_k^{\mu\nu} | \mathbf{z}_{1:k}^\nu \right) \simeq \mathcal{N} \left(\mathbf{z}_k^{\mu\nu}; \hat{\mathbf{z}}_k^{\mu\nu}, \hat{\mathbf{S}}_k^{\mu\nu} \right), \quad (38)$$

where

$$\hat{\mathbf{S}}_k^{\mu\nu} = \hat{\mathbf{H}}_k^\mu \mathbf{P}_{k|k-1}^\nu \left(\hat{\mathbf{H}}_k^\mu \right)^\top + \hat{\mathbf{R}}_k^{\mu\nu} \quad (39)$$

is the innovation covariance matrix, and

$$\hat{\mathbf{H}}_k^\mu = \frac{\partial \mathbf{h}_k^\mu \left(\hat{\mathbf{x}}_{k|k-1}^\nu \right)}{\partial \mathbf{x}_k} \quad (40)$$

is the measurement Jacobian matrix evaluated at $\hat{\mathbf{x}}_{k|k-1}^\nu$. In (39), $\hat{\mathbf{R}}_k^{\mu\nu}$ denotes the *anticipated* error covariance matrix associated with the measurement originating from the ν th target to be collected by the μ th sensor at time t_k . We assume that it is possible to compute the measurement error covariance matrix associated with a given target prior to its observation.

Substituting expressions (32), (36), and (39) into (5), it can be shown that the *posterior* conditional probability density function can also be approximated with a Gaussian function:

$$p \left(\mathbf{x}_k^\nu | \mathbf{z}_{1:k}^{\mu\nu} \right) \simeq \mathcal{N} \left(\mathbf{x}_k^\nu; \hat{\mathbf{x}}_{k|k}^{\mu\nu}, \mathbf{P}_{k|k}^{\mu\nu} \right), \quad (41)$$

where $\hat{\mathbf{x}}_{k|k}^{\mu\nu}$ and $\mathbf{P}_{k|k}^{\mu\nu}$ are the updated state estimate and error covariance, respectively, of the ν th target obtained using the measurement collected by the μ th sensor at time t_k . It can be shown that they are given by [17]

$$\hat{\mathbf{x}}_{k|k}^{\mu\nu} = \left(\mathbf{P}_{k|k}^{\mu\nu}\right)^{-1} \left[\left(\mathbf{P}_{k|k-1}^{\nu}\right)^{-1} \hat{\mathbf{x}}_{k|k-1}^{\nu} + \left(\hat{\mathbf{H}}_k^{\mu}\right)^{\text{T}} \left(\hat{\mathbf{R}}_k^{\mu}\right)^{-1} \mathbf{z}_k^{\mu\nu} \right] \quad (42)$$

and

$$\left(\mathbf{P}_{k|k}^{\mu\nu}\right)^{-1} = \left(\mathbf{P}_{k|k-1}^{\nu}\right)^{-1} + \left(\hat{\mathbf{H}}_k^{\mu}\right)^{\text{T}} \left(\hat{\mathbf{R}}_k^{\mu}\right)^{-1} \hat{\mathbf{H}}_k^{\mu}. \quad (43)$$

Equations (42) and (43) can be rewritten as

$$\hat{\mathbf{x}}_{k|k}^{\mu\nu} = \hat{\mathbf{x}}_{k|k-1}^{\nu} + \mathbf{K}_k^{\mu\nu} \left(\mathbf{z}_k^{\mu\nu} - \hat{\mathbf{H}}_k^{\mu} \hat{\mathbf{x}}_{k|k-1}^{\nu} \right) \quad (44)$$

and

$$\mathbf{P}_{k|k}^{\mu\nu} = \mathbf{P}_{k|k-1}^{\nu} - \mathbf{K}_k^{\mu\nu} \hat{\mathbf{S}}_k^{\mu\nu} \left(\mathbf{K}_k^{\mu\nu}\right)^{\text{T}}, \quad (45)$$

respectively, where

$$\mathbf{K}_k^{\mu\nu} = \mathbf{P}_{k|k-1}^{\nu} \left(\hat{\mathbf{H}}_k^{\mu}\right)^{\text{T}} \left(\hat{\mathbf{S}}_k^{\mu\nu}\right)^{-1} \quad (46)$$

is the Kalman gain matrix [17]. Equations (44) and (45) are recognized, respectively, as the state estimate and error covariance update equations of the extended Kalman filter. The state estimate and error covariance prediction equations of the extended Kalman filter, in turn, are given by equations (33) and (34), respectively.

The data flow diagram for the distributed sensor management architecture with system dynamic and observation models governed by additive white Gaussian noise is shown in Figure 3. The data flow is similar to that shown in Figure 2 for the general system dynamic and observation models with the exception that now only measurement vectors and their associated measurement error covariances need be reported. Compared to likelihoods, this is a significant reduction in the amount of information that the sensors need to report. As shown in Figure 3, in cases where there may be limited communication bandwidth, optionally, data can be further compressed. For example, given a sequence of measurements collected over the time interval $t_{k-1} \leq t < t_k$, for certain scenarios, it may be sufficient to report the parameter vector (along with its associated estimation error covariance) representing the fit to a polynomial of an appropriately chosen degree. The measurement vectors and their associated error covariances (or their compressed versions) are conflated and reported as a data bundle, denoted with $\mathbf{s}_k^{\mu\nu}$ for the μ th sensor and the ν th target in Figure 3:

$$\mathbf{s}_k^{\mu\nu} = \left\{ \mathbf{z}_k^{\mu\nu}, \mathbf{R}_k^{\mu\nu} \right\}, \quad (47)$$

where

$$\mathbf{z}_k^{\mu\nu} = \left\{ \mathbf{z}_{\ell_\mu}^{\mu\nu} : t_{k-1} \leq t_{\ell_\mu} < t_k \right\} \quad (48)$$

and

$$\mathbf{R}_k^{\mu\nu} = \left\{ \mathbf{R}_{\ell_\mu}^{\mu\nu} : t_{k-1} \leq t_{\ell_\mu} < t_k \right\}. \quad (49)$$

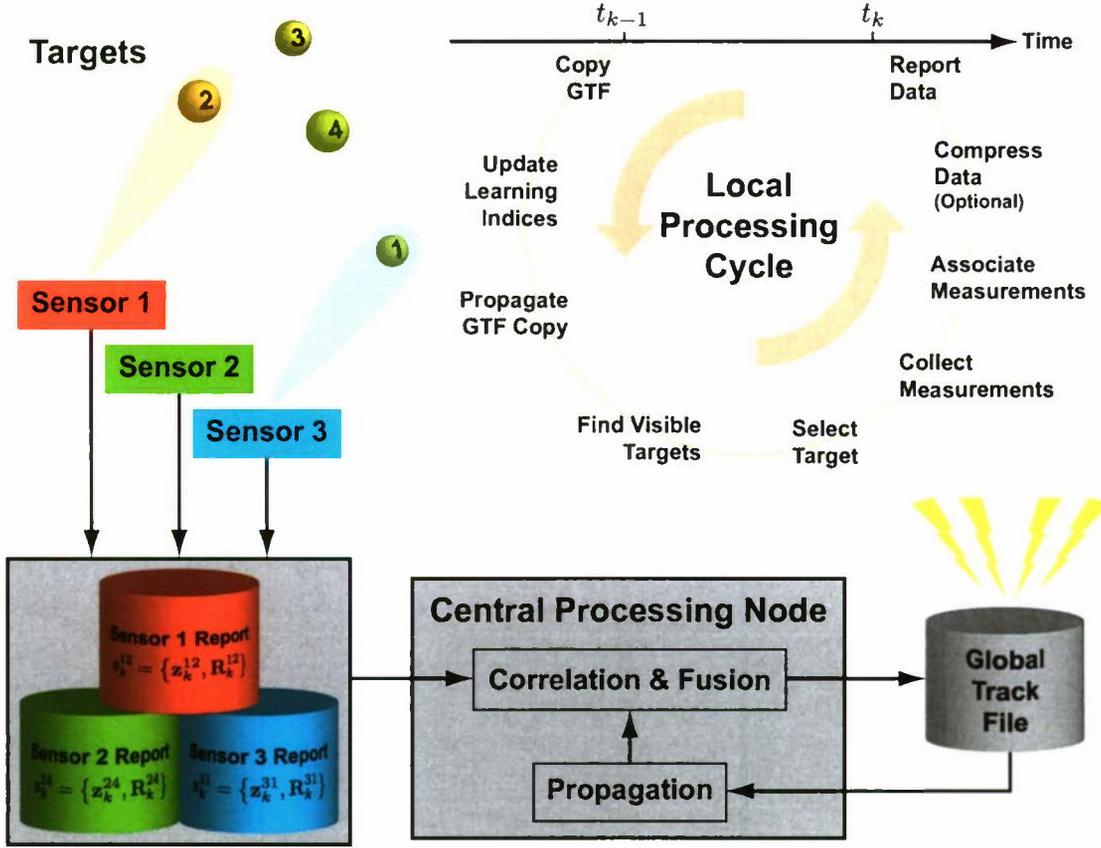


Figure 3. Data flow diagram for the distributed sensor management architecture for problems with system dynamic and observation models governed by additive white Gaussian noise.

4.2 ENTROPY AND DIVERGENCE

We next present expressions for the entropy and divergence. Substituting (41) into (7) and evaluating the integral analytically, it can be shown that the posterior α -entropy is given by [5]

$$\mathcal{H}_\alpha(\mathbf{z}_k^{\mu\nu}, \mathbf{z}_{1:k-1}^\nu) = \frac{1}{2} \log \left(\left| 2\pi\alpha^{\frac{1}{1-\alpha}} \mathbf{P}_{k|k}^{\mu\nu} \right| \right). \quad (50)$$

The *expected* posterior entropy needed for the dynamic target selection algorithm outlined in Section 3 is, in turn, obtained from substituting (50) and (38) into (18) and evaluating the integral analytically. It can be readily shown that it is given by

$$\hat{\mathcal{H}}_\alpha(\mathbf{z}_k^{\mu\nu}, \mathbf{z}_{1:k-1}^\nu) = \frac{1}{2} \log \left(\left| 2\pi\alpha^{\frac{1}{1-\alpha}} \mathbf{P}_{k|k}^{\mu\nu} \right| \right), \quad (51)$$

which is the same as (50). Similarly, substituting (41) into (9) and evaluating the integral analytically, it can be shown that the local information gain is given by [10]

$$\begin{aligned} \alpha \Delta_k^{\mu\nu} = & -\frac{1}{2(1-\alpha)} \log \left[\frac{|\mathbf{P}_{k|k-1}^\nu|^\alpha |\mathbf{P}_{k|k}^{\mu\nu}|^{1-\alpha}}{|\alpha \mathbf{P}_{k|k-1}^\nu + (1-\alpha) \mathbf{P}_{k|k}^{\mu\nu}|} \right] + \\ & \frac{\alpha}{2} \left\{ \left(\hat{\mathbf{x}}_{k|k}^{\mu\nu} - \hat{\mathbf{x}}_{k|k-1}^\nu \right)^\top \left[\alpha \mathbf{P}_{k|k-1}^\nu + (1-\alpha) \mathbf{P}_{k|k}^{\mu\nu} \right]^{-1} \left(\hat{\mathbf{x}}_{k|k}^{\mu\nu} - \hat{\mathbf{x}}_{k|k-1}^\nu \right) \right\}. \end{aligned} \quad (52)$$

The *expected* local information gain needed for the target selection algorithm outlined in Section 3 is, in turn, obtained from substituting (52) and (38) into (20) and evaluating the integral analytically. It can be shown that it is given by

$$\begin{aligned} \alpha \hat{\Delta}_k^{\mu\nu} = & -\frac{1}{2(1-\alpha)} \log \left[\frac{|\mathbf{P}_{k|k-1}^\nu|^\alpha |\mathbf{P}_{k|k}^{\mu\nu}|^{1-\alpha}}{|\alpha \mathbf{P}_{k|k-1}^\nu + (1-\alpha) \mathbf{P}_{k|k}^{\mu\nu}|} \right] + \\ & \frac{\alpha}{2} \text{tr} \left\{ \left(\mathbf{K}_k^{\mu\nu} \right)^\top \left[\alpha \mathbf{P}_{k|k-1}^\nu + (1-\alpha) \mathbf{P}_{k|k}^{\mu\nu} \right]^{-1} \mathbf{K}_k^{\mu\nu} \hat{\mathbf{S}}_k^{\mu\nu} \right\}. \end{aligned} \quad (53)$$

Using L'Hôpital's rule, it can be shown that as α approaches unity, equations (50) and (51) reduce to

$$\mathcal{H}_1(\mathbf{z}_k^{\mu\nu}, \mathbf{z}_{1:k-1}^\nu) = \frac{1}{2} \log \left(\left| 2\pi e \mathbf{P}_{k|k}^{\mu\nu} \right| \right) \quad (54)$$

and

$$\hat{\mathcal{H}}_1(\mathbf{z}_k^{\mu\nu}, \mathbf{z}_{1:k-1}^\nu) = \frac{1}{2} \log \left(\left| 2\pi e \mathbf{P}_{k|k}^{\mu\nu} \right| \right), \quad (55)$$

respectively; equations (52) and (53), in turn, reduce to

$$\begin{aligned} {}^1\Delta_k^{\mu\nu} = & -\frac{1}{2} \log \left(\left| \mathbf{P}_{k|k}^{\mu\nu} \right| \right) + \frac{1}{2} \log \left(\left| \mathbf{P}_{k|k-1}^\nu \right| \right) + \frac{1}{2} \text{tr} \left[\left(\mathbf{P}_{k|k-1}^\nu \right)^{-1} \mathbf{P}_{k|k}^{\mu\nu} - \mathbf{I}_n \right] + \\ & \frac{1}{2} \left[\left(\hat{\mathbf{x}}_{k|k}^{\mu\nu} - \hat{\mathbf{x}}_{k|k-1}^\nu \right)^\top \left(\mathbf{P}_{k|k-1}^\nu \right)^{-1} \left(\hat{\mathbf{x}}_{k|k}^{\mu\nu} - \hat{\mathbf{x}}_{k|k-1}^\nu \right) \right] \end{aligned} \quad (56)$$

and

$$\begin{aligned} {}^1\hat{\Delta}_k^{\mu\nu} = & -\frac{1}{2} \log \left(\left| \mathbf{P}_{k|k}^{\mu\nu} \right| \right) + \frac{1}{2} \log \left(\left| \mathbf{P}_{k|k-1}^\nu \right| \right) + \frac{1}{2} \text{tr} \left[\left(\mathbf{P}_{k|k-1}^\nu \right)^{-1} \mathbf{P}_{k|k}^{\mu\nu} - \mathbf{I}_n \right] + \\ & \frac{1}{2} \text{tr} \left[\left(\hat{\mathbf{S}}_k^{\mu\nu} \right)^{-1} \hat{\mathbf{H}}_k^\mu \mathbf{P}_{k|k-1}^\nu \left(\hat{\mathbf{H}}_k^\mu \right)^\top \right], \end{aligned} \quad (57)$$

respectively. In (56) and (57), \mathbf{I}_n denotes the $n \times n$ identity matrix where $n = \dim(\mathbf{x})$ is the dimension of the target state vector.

When $\alpha = 1/2$, equations (50) and (51) reduce to

$$\mathcal{H}_{1/2}(\mathbf{z}_k^{\mu\nu}, \mathbf{z}_{1:k-1}^\nu) = \frac{1}{2} \log \left(\left| \frac{\pi}{2} \mathbf{P}_{k|k}^{\mu\nu} \right| \right) \quad (58)$$

and

$$\hat{\mathcal{H}}_{1/2}(\mathbf{z}_k^{\mu\nu}, \mathbf{z}_{1:k-1}^\nu) = \frac{1}{2} \log \left(\left| \frac{\pi}{2} \mathbf{P}_{k|k}^{\mu\nu} \right| \right), \quad (59)$$

respectively; equations (52) and (53), in turn, reduce to

$$\begin{aligned} {}^{1/2}\Delta_k^{\mu\nu} = & -\log \left[\frac{\sqrt{|\mathbf{P}_{k|k-1}^\nu| |\mathbf{P}_{k|k}^{\mu\nu}|}}{\left| \frac{1}{2} (\mathbf{P}_{k|k-1}^\nu + \mathbf{P}_{k|k}^{\mu\nu}) \right|} \right] + \\ & \frac{1}{4} \left[(\hat{\mathbf{x}}_{k|k}^{\mu\nu} - \hat{\mathbf{x}}_{k|k-1}^\nu)^\top \left(\frac{\mathbf{P}_{k|k-1}^\nu + \mathbf{P}_{k|k}^{\mu\nu}}{2} \right)^{-1} (\hat{\mathbf{x}}_{k|k}^{\mu\nu} - \hat{\mathbf{x}}_{k|k-1}^\nu) \right] \end{aligned} \quad (60)$$

and

$${}^{1/2}\hat{\Delta}_k^{\mu\nu} = -\log \left[\frac{\sqrt{|\mathbf{P}_{k|k-1}^\nu| |\mathbf{P}_{k|k}^{\mu\nu}|}}{\left| \frac{1}{2} (\mathbf{P}_{k|k-1}^\nu + \mathbf{P}_{k|k}^{\mu\nu}) \right|} \right] + \frac{1}{4} \text{tr} \left[(\mathbf{K}_k^{\mu\nu})^\top \left(\frac{\mathbf{P}_{k|k-1}^\nu + \mathbf{P}_{k|k}^{\mu\nu}}{2} \right)^{-1} \mathbf{K}_k^{\mu\nu} \hat{\mathbf{S}}_k^{\mu\nu} \right], \quad (61)$$

respectively.

5. DETERMINISTIC MOTION

In the absence of process noise—that is, when the process noise vector, \mathbf{w}_{k-1} , in the system dynamic model (1) can be set equal to zero—target motion becomes deterministic. For example, in the special case of tracking ballistic targets, target motion is prescribed by well-understood forces that can be modeled to a high degree of accuracy (e.g., see [18, ch. 3]). Thus, for ballistic targets, motion can be safely assumed to be deterministic. When tracking N_t targets following deterministic trajectories, the nonlinear system dynamic model prescribing the temporal evolution of the ν th target from time t_{k-1} to time t_k becomes

$$\mathbf{x}_k^\nu = \varphi_{k-1}^\nu(\mathbf{x}_{k-1}^\nu), \quad (62)$$

where

$$\varphi_{k-1}^\nu(\mathbf{x}_{k-1}^\nu) = \mathbf{x}_{k-1}^\nu + \int_{t_{k-1}}^{t_k} \mathbf{f}_t^\nu(\mathbf{x}_t^\nu) dt \quad (63)$$

is a deterministic function prescribing the noise-free transition of the state vector of the ν th target from time t_{k-1} to time t_k . The function $\mathbf{f}_t^\nu(\cdot)$ in (63) models the underlying continuous-time forces acting on the ν th target; that is,

$$\dot{\mathbf{x}}_t^\nu = \mathbf{f}_t^\nu(\mathbf{x}_t^\nu), \quad (64)$$

where $\dot{\mathbf{x}}_t^\nu$ denotes the first derivative of \mathbf{x}_t^ν with respect to time¹. Here, we use the subscript t to distinguish continuous-time variables from discrete-time variables, which in turn we denote with the subscript k for time *index*; thus, using this notation, $\mathbf{x}_k = \mathbf{x}_{t_k}$. Applying (62) and (63) recursively, it follows that the state vector, \mathbf{x}_k^ν , of the ν th target, valid at time t_k , can be obtained directly from the target state vector, \mathbf{x}_0^ν , valid at a time t_0 prior to any data collection:

$$\mathbf{x}_k^\nu = \mathbf{x}_0^\nu + \int_{t_0}^{t_k} \mathbf{f}_t^\nu(\mathbf{x}_t^\nu) dt. \quad (65)$$

Thus, in case of deterministic motion, target dynamics are completely characterized by the initial state vector, \mathbf{x}_0^ν , valid at time t_0 .

In the absence of process noise, the covariance \mathbf{Q}_{k-1}^ν in equation (34), which we use to propagate the error covariance of the ν th target's error covariance between measurements, vanishes. Hence, starting with the updated *global* state estimate, $\hat{\mathbf{x}}_{k-1|k-1}^\nu$, and error covariance, $\mathbf{P}_{k-1|k-1}^\nu$, it follows that for a set of M_k^μ measurements expected to be collected by the μ th sensor on the ν th target within the forthcoming observation period, $t_{k-1} < t \leq t_k$, the covariance propagation and update equations—equations (34) and (43), respectively—can be combined and applied recursively to yield

$$\begin{aligned} \left(\mathbf{P}_{k|k}^{\mu\nu}\right)^{-1} &= \left(\mathbf{P}_{k-1|k-1}^\nu\right)^{-1} + \\ &\sum_{\ell=1}^{M_k^\mu} \left[\mathbf{H}_{t_{\ell\mu}}^\mu \left(\hat{\mathbf{x}}_{k-1|k-1}^\nu \right) \hat{\Phi}^\nu(t_{\ell\mu}, t_{k-1}) \right]^T \left(\hat{\mathbf{R}}_{t_{\ell\mu}}^{\mu\nu} \right)^{-1} \mathbf{H}_{t_{\ell\mu}}^\mu \left(\hat{\mathbf{x}}_{k-1|k-1}^\nu \right) \hat{\Phi}^\nu(t_{\ell\mu}, t_{k-1}), \end{aligned} \quad (66)$$

¹ Assuming that the function \mathbf{f} has continuous first partial derivatives, it can be shown that the transition matrix (35) can be obtained from solving $\dot{\Phi} = \mathbf{F}\Phi$, where $\mathbf{F} = \partial\mathbf{f}/\partial\mathbf{x}$ (e.g., see [19]).

where t_{ℓ_μ} , with $t_{k-1} < t_{\ell_\mu} \leq t_k$, denotes the time of a particular measurement expected to be collected by the μ th sensor within the upcoming observation period. Thus, the anticipated M_k^μ measurements to be collected by the μ th sensor on the ν th target during the forthcoming observation period, $t_{k-1} < t \leq t_k$, can potentially produce the local gain:

$$\Delta \hat{\mathbf{J}}_k^{\mu\nu} = \sum_{\ell=1}^{M_k^\mu} \left[\mathbf{H}_{t_{\ell_\mu}}^\mu \left(\hat{\mathbf{x}}_{k-1|k-1}^\nu \right) \hat{\Phi}^\nu(t_{\ell_\mu}, t_{k-1}) \right]^\top \left(\hat{\mathbf{R}}_{t_{\ell_\mu}}^{\mu\nu} \right)^{-1} \mathbf{H}_{t_{\ell_\mu}}^\mu \left(\hat{\mathbf{x}}_{k-1|k-1}^\nu \right) \hat{\Phi}^\nu(t_{\ell_\mu}, t_{k-1}) \quad (67)$$

in information relative to the *inverse* of the global error covariance, $\mathbf{P}_{k-1|k-1}^\nu$, valid at time t_{k-1} . In other words, equation (66) can be rewritten as

$$\hat{\mathbf{J}}_{k|k}^{\mu\nu} = \mathbf{J}_{k-1|k-1}^\nu + \Delta \hat{\mathbf{J}}_k^{\mu\nu}, \quad (68)$$

where $\hat{\mathbf{J}}_{k|k}^{\mu\nu} = \left(\hat{\mathbf{P}}_{k|k}^{\mu\nu} \right)^{-1}$ and $\mathbf{J}_{k-1|k-1}^\nu = \left(\mathbf{P}_{k-1|k-1}^\nu \right)^{-1}$ are the Fisher information matrices associated with the local and global updated error covariances valid at times t_k and t_{k-1} , respectively. Thus, expression (67) provides the *expected* gain in the Fisher information matrix at time t_k . Once actual measurements have been collected, the predicted measurement error covariances, $\hat{\mathbf{R}}_{t_{\ell_\mu}}^{\mu\nu}$, in (67) are replaced with actual measurement error covariances, $\mathbf{R}_{t_{\ell_\mu}}^{\mu\nu}$, in order to compute the actual Fisher information matrix; that is, at that point, $\Delta \hat{\mathbf{J}}_k^{\mu\nu}$ and $\hat{\mathbf{J}}_{k|k}^{\mu\nu}$ in (68) are replaced with their actual values, denoted by $\Delta \mathbf{J}_k^{\mu\nu}$ and $\mathbf{J}_{k|k}^{\mu\nu}$, respectively.

5.1 CRAMÉR–RAO LOWER BOUNDS

Since target dynamics in the absence of process noise are completely characterized by the state, \mathbf{x}_0 , valid at a particular time, t_0 , we only need to focus on estimating \mathbf{x}_0 . The state estimate valid at an any other time can subsequently be obtained from (62). The uncertainty in the state estimate valid at time t_0 reduces as we incorporate new measurements. We can express the recursion relation (68) in terms of $\hat{\mathbf{x}}_0^\nu$:

$$\hat{\mathbf{J}}_k^{\mu\nu}(\hat{\mathbf{x}}_0^\nu) = \mathbf{J}_{k-1}^\nu(\hat{\mathbf{x}}_0^\nu) + \Delta \hat{\mathbf{J}}_k^{\mu\nu}(\hat{\mathbf{x}}_0^\nu), \quad (69)$$

where $\hat{\mathbf{J}}_k^{\mu\nu}$ and \mathbf{J}_{k-1}^ν now denote the predicted local and actual global Fisher information matrices associated with the estimates of \mathbf{x}_0^ν at times t_k and t_{k-1} , respectively. Specifically, the local gain in the expected Fisher information matrix, $\Delta \hat{\mathbf{J}}_k^{\mu\nu}$, is given by

$$\Delta \hat{\mathbf{J}}_k^{\mu\nu}(\hat{\mathbf{x}}_0^\nu) = \sum_{\ell=1}^{M_k^\mu} \left[\mathbf{H}_{t_{\ell_\mu}}^\mu \left(\hat{\mathbf{x}}_0^\nu \right) \hat{\Phi}^\nu(t_{\ell_\mu}, t_0) \right]^\top \left(\hat{\mathbf{R}}_{t_{\ell_\mu}}^{\mu\nu} \right)^{-1} \mathbf{H}_{t_{\ell_\mu}}^\mu \left(\hat{\mathbf{x}}_0^\nu \right) \hat{\Phi}^\nu(t_{\ell_\mu}, t_0). \quad (70)$$

Again, once actual measurements have been collected and incorporated, $\hat{\mathbf{J}}_k^{\mu\nu}$, $\Delta \hat{\mathbf{J}}_k^{\mu\nu}$, and $\hat{\mathbf{R}}_{t_{\ell_\mu}}^{\mu\nu}$ in (69) and (70) are replaced with $\mathbf{J}_k^{\mu\nu}$, $\Delta \mathbf{J}_k^{\mu\nu}$, and $\mathbf{R}_{t_{\ell_\mu}}^{\mu\nu}$, respectively.

When the estimated state, $\hat{\mathbf{x}}_0^\nu$, of the ν th target in (69) and (70) is replaced with the *true* state, \mathbf{x}_0^ν , the inverse of the updated Fisher information matrix provides the Cramér–Rao lower bound

on the error covariance [14]. When computing Cramér–Rao lower bounds, the terms involving $(\hat{\mathbf{x}}_{k|k}^{\mu\nu} - \hat{\mathbf{x}}_{k|k-1}^\nu)$ in the expressions for the α -divergence that were presented in Section 4 all vanish. Subsequently, the expected value of the α -divergence, ${}^\alpha\tilde{\Delta}_k^{\mu\nu}$, becomes the same as the actual α -divergence, ${}^\alpha\Delta_k^{\mu\nu}$, and therefore $\hat{\mathbf{J}}_k^{\mu\nu}$ and $\Delta\hat{\mathbf{J}}_k^{\mu\nu}$ become the same as $\mathbf{J}_k^{\mu\nu}$ and $\Delta\mathbf{J}_k^{\mu\nu}$, respectively. Furthermore, the error covariances, $\mathbf{P}_{k|k}^{\mu\nu}$ and $\mathbf{P}_{k|k-1}^\nu$ in the expressions for the α -entropy and α -divergence can be replaced with $(\mathbf{J}_k^{\mu\nu})^{-1}$ and $(\mathbf{J}_{k-1}^\nu)^{-1}$, respectively.

Without committing to a specific tracking filter, we can evaluate the overall performance of the target selection logic outlined in Section 3 using the Cramér–Rao lower bounds derived in this section² (cf. [21]). Once a particular sensor selects to collect measurements on a particular target, the local Fisher information matrix is computed according to

$$\mathbf{J}_k^{\mu\nu} = \mathbf{J}_{k-1}^\nu + \delta_{\nu\nu_\mu} \Delta\mathbf{J}_k^{\mu\nu_\mu}, \quad (71)$$

where ν_μ denotes the index of the target selected by the μ th sensor. The global Fisher information matrix is, in turn updated according to

$$\mathbf{J}_k^\nu = \mathbf{J}_{k-1}^\nu + \sum_{\mu} \delta_{\nu\nu_\mu} \Delta\mathbf{J}_k^{\mu\nu_\mu}. \quad (72)$$

For each local measurement cycle, the μ th sensor begins the target selection process by accessing the global track file that is broadcast to all sensors at regular time intervals. The sensor’s local processor then selects only those targets that are predicted to be within its field of regard during the forthcoming observation period, $t_{k-1} < t \leq t_k$. Next, the local processor updates the learning indices of all visible targets according to (22) or (25) depending on whether or not a particular target track was locally updated during the last observation period. Subsequently, the local processor selects a target according to the rules outlined in Section 3. Once a target has been selected, the local gain, $\Delta\mathbf{J}_k^{\mu\nu_\mu}$, for the selected target is reported to the central processor. The central processor, in turn, collects similar information from all sensors and updates the global track file according to (72).

²For ballistic targets, it has been shown that the batch least-squares filter provides efficient state estimates with error covariances that match the Cramér–Rao lower bound (e.g., see [20]).

This page intentionally left blank.

6. NUMERICAL RESULTS

We illustrate the performance of the bioinspired target selection algorithm outlined in Section 3 using simulation. Twenty-one different scenarios are examined, with target numbers ranging from 1 to 100. Our example scenarios utilize three independently operating narrowband IR cameras on board three separate stationary platforms, each of which can track up to 100 ballistic targets. Each camera is allowed to pan up to 360 degrees in the horizontal direction, whereas the panning movement in the vertical direction is restricted to lie within 0 and 40 degrees. The ballistic targets originate at the same time, $t_0 = 0$, from multiple sources randomly distributed over a fixed rectangular area and terminate at the same location. The targets' times of flight are chosen to be proportional to the ground range between their points of origin and the final destination; thus, different targets arrive at different times. The wide spatial separation of the targets together with the sensors' narrow fields of view guarantee that at most one target can be tracked by a sensor at any given time—thus, eliminating the need for any local data association algorithm. However, the problem of *global* data association remains. To compute the entropies and divergences needed by the target selection algorithm outlined in Section 3, we employ the Cramér–Rao lower bounds of Section 5 as surrogates for the error covariances associated with the state estimates of the ballistic targets valid at time $t_0 = 0$.

At regular time intervals, a central processor collects data reported by each IR camera—here, in the form of $\Delta \mathbf{J}_k^{\mu\nu}$ discussed in Section 5—and updates the global track file using equation (72). The global track file, the contents of which are broadcast to all sensors at regular time intervals, is assumed to have been already populated with the initial values of all targets' Fisher information matrices—for example, obtained by means of a separate set of sensors operating at a previous processing epoch—prior to the targets' entries into the fields of regard of any of the three IR cameras. We do not account for any search or acquisition time; thus, once a sensor has tasked itself to collect data on a target different than the one currently being tracked, local target-to-target switching is assumed to occur instantaneously³.

We vary the complexity of the resource allocation problem by starting with a single target and subsequently increasing the number of targets, in steps of 5, up to 100 (i.e., 1, 5, 10, 15, . . . , 95, 100 targets). Each numerical experiment is repeated 10 times, and results are averaged to evaluate performance. The simulation parameters are listed in Table 1. The global-track-file update occurs every 2 seconds; thus, each sensor, with a measurement sampling time of 1 second, can potentially produce up to two measurements per target (i.e., $M_k^\mu = 2$). Our goal is to ensure that for as many targets as possible the state estimates meet a particular level of accuracy as quantified by the α -entropy. Targets whose α -entropies fall below a certain threshold are referred to as “serviced”; the *threshold* itself is referred to as the “quality of service.” We have chosen $\alpha = 1/2$, in which case the α -divergence becomes the Hellinger affinity (see Section 2). We defer the examination of the impact of other values of α on system performance to a future study. Since no process noise is included in the problem, the value of the global entropy for a particular target does not increase once a sufficient number of measurements have been incorporated to lower it below a certain level.

³Although not pursued here, one can readily account for the time lost in acquiring a given target by excluding the measurements that would otherwise have been included in the sum (70).

TABLE 1. Simulation Parameters

PARAMETER	VALUE
Maximum Number of Targets	100
Number of Sensors	3
Local Sampling Time	1 s
Global Sampling Time	2 s
Measurement Noise Standard Deviation	500 μ rad
Quality of Service	$25 \log(\text{m}^6/\text{s}^3)$
Forgetting Time Constant, τ	10 s
Number of Monte Carlo Trials	10

The average number of targets serviced as a function of the total number of targets in a given scenario is shown in Figure 4. Averages are computed over 10 Monte Carlo trials. The launch locations of the ballistic targets are randomly varied over a rectangular region from one Monte Carlo trial to the next; however, in all trials, all targets terminate at the same location. The locations of the three sensor platforms also remain the same throughout the simulation. The dashed line in Figure 4 corresponds to the ideal case when all targets meet their qualities of service at all times. However, as the number of targets increases, degradation in performance would be expected. The green curve labeled as “bio inspired” in Figure 4 is obtained by employing the target selection algorithm outlined in Section 3. We see that the bioinspired algorithm performs very well. We only observe a relatively small degradation in performance when the number of targets nears 100.

For reference, we also compare the performance of the bioinspired algorithm with two other target selection algorithms. The red curve labeled as “uniformly random” in Figure 4 is similar to the bioinspired algorithm, a decentralized sensor-level target selection algorithm. This algorithm randomly selects a target from a list of targets predicted to be visible within the upcoming observation period. The targets are assigned equal weights; thus, the random selection is uniform. As is evident from Figure 4, for small number of targets (up to around 25 in our example), the performance of the uniformly random algorithm is nearly perfect—as is indeed the case for the other algorithms. This result suggests that at least for a small number of targets, an algorithm as simple as the uniformly random target selection algorithm can be employed reliably to meet all resource allocation requirements, and therefore, there may be no need to resort to a more expensive solution.

In addition to the decentralized, uniformly random target selection algorithm, we also compare the performance of the bioinspired algorithm to that of a greedy, centrally processed sensor-to-target assignment algorithm. The blue curve labeled as “greedy” in Figure 4 shows results obtained using

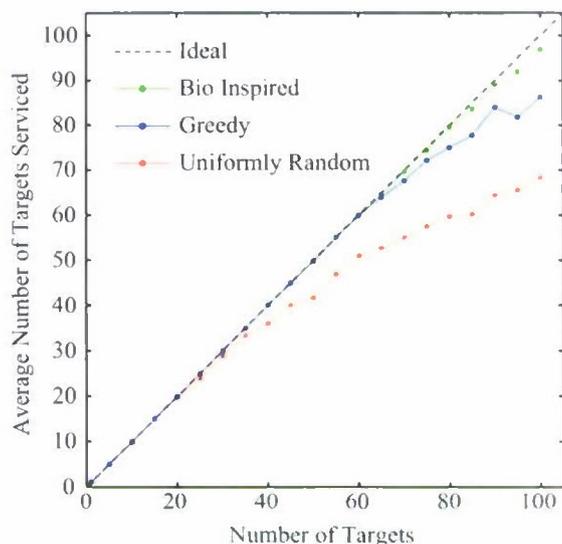


Figure 4. Average number of targets serviced versus actual number of targets. Averages are computed over 10 Monte Carlo trials. See text for an explanation of the different curves

this algorithm. Our centrally processed algorithm assigns sensors to targets by solving a linear assignment problem. The centrally processed algorithm is “greedy” in that at any time interval, (1) it assigns sequentially one sensor per target, and (2) it solves the assignment problem only for the forthcoming observation period without any regard to past assignments or prediction of future performance. The former restriction prevents the simultaneous coverage of targets by multiple sensors, whereas the target selection algorithm of Section 3, although not “consciously,” does allow for simultaneous multiple-sensor coverage. The lack of simultaneous multiple-sensor coverage over a given observation period in the greedy algorithm may partially, although not entirely, explain the poorer performance of the algorithm for a larger number of targets (greater than around 60 targets in our example). We would expect a more sophisticated centralized algorithm that fully takes into account the stochastic nature of the scenario to perform at least as good as or better than the bioinspired target selection algorithm. However, the main advantages of the target selection logic of Section 3 are: (1) it is scalable, and (2) it requires less processing and communication bandwidth. Thus, in many applications, a quick cost-benefit analysis may reveal its implementation to be more advantageous as compared to an otherwise optimal, centrally processed algorithm. A complete comparative analysis of various centralized and decentralized algorithms is beyond the scope of this report and is deferred to a future study.

In Figure 4, the average number of targets serviced were computed when the three sensors’ opportunities to observe all targets had been exhausted. For a subset of scenarios considered, plots of the average number of targets serviced as a function of time are shown in Figure 5. In addition to the three curves that were considered in Figure 4, we have also included a curve labeled as “serviceable” in Figure 5. Since not all targets are within the fields of regard of all sensors at all times, even the existence of an optimal resource allocation algorithm could not guarantee that the quality of service can be met for all targets. This is a purely geometrical restriction. Nevertheless,

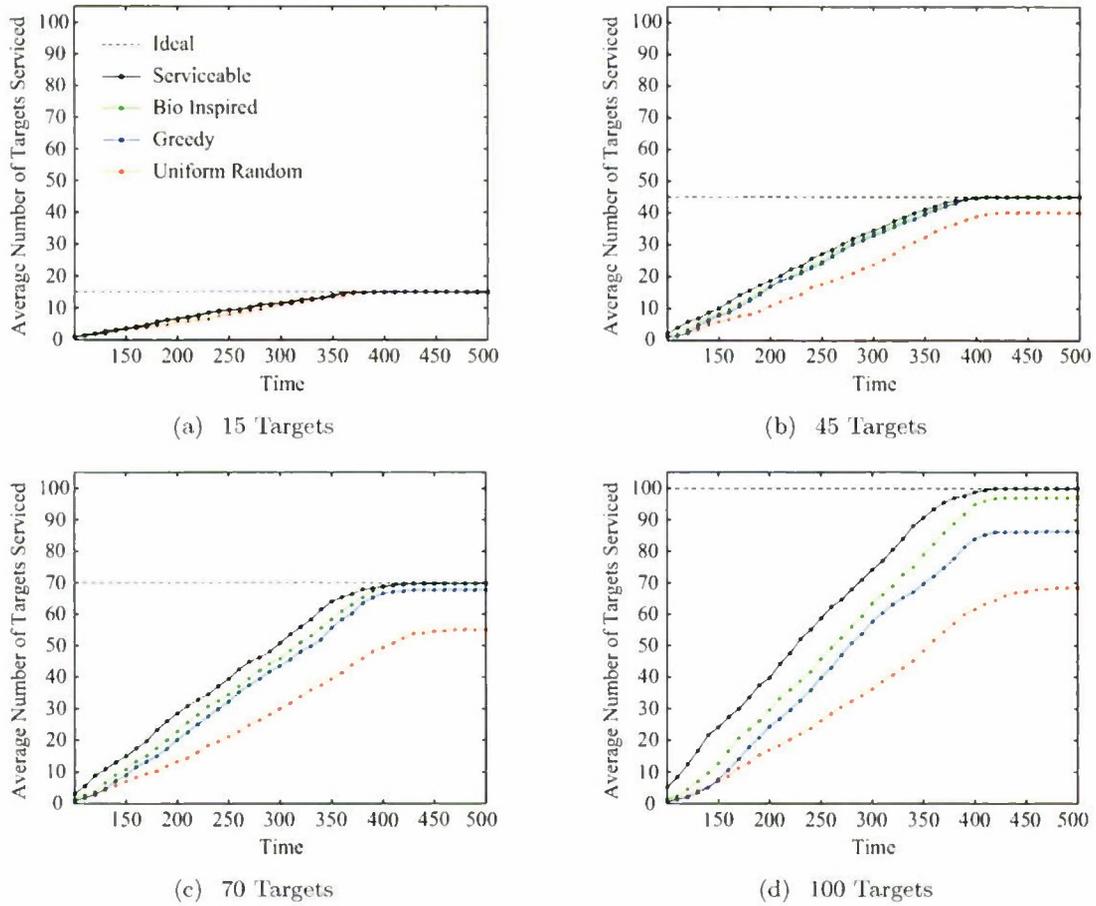


Figure 5. Average number of targets serviced versus time. See text for an explanation of the different curves.

the “serviceable” curves provide a loose upper bound on performance. As is evident in Figure 5, it is not after around 400 s that eventually all targets enter the fields of regard of all sensors in our example. From the temporal behavior shown in Figure 5, it is clear that the target selection algorithm outlined in Section 3 performs better than the other two algorithms in that larger number of targets meet their qualities of service earlier. The superior performance of the bioinspired target selection algorithm becomes even more evident as the number of targets increases.

7. SUMMARY AND DISCUSSION

In this report, we presented a bioinspired algorithm for managing the resources of multiple sensors tracking multiple targets within the framework of a decentralized sensor management architecture. Our target selection logic follows the principles of self-organization and stigmergy observed in biological swarms. In our implementation, sensors are self-tasking and are not required to communicate with each other directly. Nevertheless, sensors can infer from the global track file, which is continuously broadcast by a central processing node, their relative contributions to achieving certain levels of global state estimation accuracy for particular targets.

We devised local information-theoretic target selection rules for the general Bayesian formulation of the target tracking problem. We also provided specific formulas for problems dominated by additive white Gaussian noise. We rely on the Rényi α -entropy and α -divergence as measures of absolute and relative information, respectively. We have not examined the impact of the particular choice of α . In our examples, following [10], we have chosen $\alpha = 1/2$. We defer to a future study the examination of the impact of the particular choice of α on the performance of the bioinspired target selection algorithm.

For reference, we have compared the performance of the bioinspired target selection algorithm with those of other suboptimal algorithms using numerical examples involving ballistic targets (no process noise). To demonstrate performance, we have used Cramér–Rao lower bounds as surrogates for target state estimation error covariances. Targets following deterministic trajectories—such as ballistic targets—render a simple recursive formulation of their Cramér–Rao lower bounds. We have shown that the bioinspired algorithm is highly scalable and performs very well for large numbers of targets. Our numerical examples were devised such that data association would pose no challenge to tracking performance. Poor data association can in general impact sensor management. In a future study, we will examine the impact of finite process noise, poor data association, and sensor biases on the performance of the bioinspired target selection algorithm.

The target selection rules outlined in Section 3 are by no means the only rules one could devise. The rules could be extended, and the bioinspired target selection algorithm may be optimized for specific applications. For example, in the case of deterministic dynamics, such as the example scenarios considered in Section 6, the predictable nature of target motion may be exploited to improve the formulation of the learning index, w_{k-1} . Some knowledge of other sensors' observation capabilities, if available, can also be exploited to improve the formulation of the learning index. Although in a decentralized sensor management architecture sensors are not required to communicate with each other directly, certain information about them, such as their locations and fields of regard, can nevertheless be readily broadcast by the central processing node with little added cost. We defer the examination of such extensions of the algorithm to a future study.

This page intentionally left blank.

REFERENCES

- [1] Y. Bar-Shalom and X. R. Li. *Multi-Target Multi-Sensor Tracking: Principles and Techniques*. Storrs, Connecticut: YBS Publishing, 1995.
- [2] S. Blackman and R. Popoli. *Design and Analysis of Modern Tracking Systems*. Boston: Artech House, 1999.
- [3] A. B. Poore, S. M. Gadaleta, and B. J. Slocumb, "Multiple hypothesis correlation in track-to-track fusion management," in *Handbook on Modelling for Discrete Optimization*. New York: Springer, 2006, vol. 88, no. II, pp. 341–371.
- [4] X. Li, G. Chen, E. Blasch, J. Patrick, C. Yang, and I. Kadar, "A geometric feature-aided game-theoretic approach to sensor management," in *Proceedings of the International Conference on Information Fusion*, Seattle, Washington, July 2009, pp. 1155–1162.
- [5] J. Manyika and H. Durrant-Whyte, *Data Fusion and Sensor Management: A Decentralized Information-Theoretic Approach*. New York: Ellis Horwood, 1994.
- [6] E. Bonabeau, M. Dorigo, and G. Theraulaz, *Swarm Intelligence: From Natural to Artificial Systems*. New York: Oxford University Press, 1999.
- [7] D. Sinno, "Self-organizing networks (SONets) with application to target tracking," in *Adaptive Sensor Array Processing (ASAP) Workshop*, March 2004.
- [8] B. S. Weir and T. M. Sokol, "Radar coordination and resource management in a distributed sensor network using emergent control," in *Proceedings of SPIE*, R. Suresh, Ed., vol. 7350, 2009.
- [9] C. Kreucher, K. Kastella, and A. O. Hero, "Multi-target sensor management using alpha-divergence measures," in *Proceedings of the International Conference on Information Processing in Sensor Networks (ISPN) 2003*, F. Zhao and L. Guibas, Eds. Berlin: Springer-Verlag, April 2003, pp. 209–222.
- [10] A. O. Hero, C. M. Keucher, and D. Blatt, "Information theoretic approaches to sensor management," in *Foundations and Applications of Sensor Management*, A. O. Hero, D. Castanon, D. Cochran, and K. Kastella, Eds. New York: Springer, 2008, pp. 33–57.
- [11] D. Sinno, "Attentive management of configurable sensor systems," Ph.D. dissertation, Arizona State University, May 2000.
- [12] D. Sinno and D. Kreithen, "A constrained joint optimization approach to dynamic sensor configuration," in *Thirty-Sixth Asilomar Conference on Signals, Systems, and Computers*, November 2002, pp. 1179–1183.

- [13] A. Rényi, “On measure of entropy and information,” in *Berkeley Symposium on Mathematics, Statistics and Probability*, vol. 1, 1961, pp. 547–561.
- [14] J. H. Taylor, “The Cramér–Rao estimation error lower bound computation for deterministic nonlinear systems,” *IEEE Transactions on Automatic Control*, vol. AC-24, pp. 343–344, April 1979.
- [15] A. H. Jazwinski, *Stochastic Processes and Filter Theory*. New York: Academic Press, 1970.
- [16] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. New York: John Wiley & Sons, 1991.
- [17] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation*. New York: John Wiley & Sons, 2001.
- [18] O. Montenbruck and E. Gill, *Satellite Orbits*. Berlin: Springer, 2000.
- [19] L. Perko, *Differential Equations and Dynamical Systems*, 2nd ed. New York: Springer, 1996.
- [20] C. B. Chang, “Ballistic trajectory estimation with angle-only measurements,” *IEEE Transactions on Automatic Control*, vol. AC-25, no. 3, June 1980.
- [21] M. L. Hernandez, T. Kirubarajan, and Y. Bar-Shalom, “Multisensor resource deployment using posterior Cramér–Rao bounds,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 40, no. 2, pp. 399–416, April 2004.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE 20 June 2011			2. REPORT TYPE Project Report		3. DATES COVERED (From - To)	
4. TITLE AND SUBTITLE Bioinspired Resource Management for Multiple-Sensor Target Tracking Systems					5a. CONTRACT NUMBER FA8721-05-C-0002	
					5b. GRANT NUMBER	
					5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Hendrick C. Lambert and Dana Sinno, Group 36					5d. PROJECT NUMBER 1293	
					5e. TASK NUMBER 0	
					5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) AND ADDRESS(ES) MIT Lincoln Laboratory 244 Wood Street Lexington, MA 02420-9108					8. PERFORMING ORGANIZATION REPORT NUMBER MD-26	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) MDA/BC Von Braun II, Bldg 5222 Martin Rd Redstone Arsenal, AL 35898					10. SPONSOR/MONITOR'S ACRONYM(S)	
					11. SPONSOR/MONITOR'S REPORT NUMBER(S) ESC-TR-2010-083	
12. DISTRIBUTION / AVAILABILITY STATEMENT Distribution unlimited; approved for public release.						
13. SUPPLEMENTARY NOTES						
14. ABSTRACT We present an algorithm, inspired by self-organization and stigmergy observed in biological swarms, for managing multiple sensors tracking large numbers of targets. We devise a decentralized architecture wherein autonomous sensors manage their own data collection resources and task themselves. Sensors cannot communicate with each other directly; however, a global track file, which is continuously broadcast, allows the sensors to infer their contributions to the global estimation of target states. Sensors can transmit their data (either as raw measurements or some compressed format) only to a central processor where their data are combined to update the global track file. We outline information-theoretic rules for the general multiple-sensor Bayesian target tracking problem. We provide specific formulas for problems dominated by additive white Gaussian noise. Using Cramér-Rao lower bounds as surrogates for error covariances, we illustrate, using numerical scenarios involving ballistic targets, that the bioinspired algorithm is highly scalable and performs very well for large numbers of targets.						
15. SUBJECT TERMS						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON	
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified	Same as report	44	19b. TELEPHONE NUMBER (include area code)	

This page intentionally left blank.