

IDA

INSTITUTE FOR DEFENSE ANALYSES

**Bayesian Sensor Fusion for
Minimum-Cost I.D. Declaration**

James M. Ralston

June 1999

Approved for public release;
distribution unlimited.

IDA Paper P-3441

Log: H 98-003030

This work was conducted under contract DASW01 97 C 0056, Task CA-2-1617, for the Joint Theater Air and Missile Defense Organization (JTAMDO), Crystal Mall 3, Suite 511, 1931 Jefferson Davis Highway, Arlington, VA 22203. The publication of this IDA document does not indicate endorsement by the Department of Defense, nor should the contents be construed as reflecting the official position of that Agency.

© 1998, 1999 Institute for Defense Analyses, 1801 N. Beauregard Street, Alexandria, Virginia 22311-1772 • (703) 845-2000.

This material may be reproduced by or for the U.S. Government pursuant to the copyright license under the clause at DFARS 252.227-7013 (10/88).

PREFACE

This work was performed for the Joint Theater Air and Missile Defense Organization (JTAMDO) in partial fulfillment of the task, "Analyses of CID Procedures and Data in Support of Joint Theater Air and Missile Defense." It was presented at the 1998 Combat Identification Systems Conference (CISC) and is published in the conference proceedings.

ACKNOWLEDGMENTS

The author wishes to thank Kevin Hopkins of SAIC for sharing his work in developing CID systems for AEGIS, and also to acknowledge with gratitude the many valuable comments made by his colleagues at IDA: Kent Haspert, Roger Sullivan, and Michael Tuley.

CONTENTS

I. INTRODUCTION.....	1
II. ANALYSIS	3
A. The Sensor Performance Matrix.....	3
B. Combat Identification System States	4
C. I.D. Fusion Rules.....	6
D. Cost of Identification Errors.....	10
E. Least-Cost I.D. Designations	13
F. Declaration Probability	17
III. IMPLEMENTATION	19
IV. CONCLUSIONS.....	21
References	R-1
Glossary.....	GL-1

FIGURES

1.	The Combat I.D. Process.....	2
2.	Graphing CIS Performance	7
3.	Ranked Likelihood Ratios of Example CIS States.....	9
4.	Probabilities of Ranked CIS States	9
5.	Identification System Operating Characteristic	10
6.	Example of I.D. Error Cost Estimate	11
7.	Setting Declaration Thresholds	12
8.	Least-Cost I.D. Declarations [$p(F) = 0.9$]	14
9.	Least-Cost I.D. Declarations [$p(F) = 0.99$]	15
10.	$p(F) = p(H) = 0.5$. $c(\text{friend} \text{Hostile}) = 1.0$	16
11.	Comparative ISOCs of Original and Augmented CIS	17
12.	Declaration Probability with Constrained Errors	18

TABLES

1.	Performance Matrix of Sensor (1).....	3
2.	Performance Matrices of Hypothetical CID Sensors	4
3.	Computed Declaration Thresholds for Assumed Costs and Priors	13
4.	Declaration Probabilities for Hypothetical CIS and Costs	14
5.	Performance Matrices of Additional Hypothetical Sensors	16

I. INTRODUCTION

The process of assigning a "hostile" or "friend" identification (I.D.) to a specific target can be thought of as one of making an identification decision about an unknown target track based not only on all of the information our sensors furnish about the unknown but also on all the information we have about the functioning of the system of I.D. sensors itself. The optimum declaration will also be sensitive to the a priori probability that a given unknown track is friend or hostile as well as to the costs of making different kinds of errors in assigning identification. In fact, all currently observed sensor fusion rules or rules of engagement reflect these factors to some degree, in that they involve some subjective and often intuitive judgment or declared policy about the relative reliability or priority of different I.D. sensors and procedures, the consequences of making I.D. errors, and the nature of the threat environment. Subjective and intuitive judgments may be based on a great deal of direct experience and often lead to valid conclusions. Such judgments, however, and the processes by which they arise, are often difficult to study, quantitatively justify, or teach to others. In this paper we present an analytic approach to developing I.D. sensor fusion and declaration procedures that are sensitive to all key factors and which may be considered "optimum" in the sense of leading to the lowest total expected costs. This analysis is based on the application of a Bayes criterion in accord with standard techniques of statistical decision theory (Ref. 1). It represents a further development of the concept of *objective rules of engagement* introduced in an earlier paper (Ref. 2). In this paper we introduce the added factor of the cost of I.D. declaration error, using this concept to identify which of the objective rules leads to the lowest operational cost when considering errors of all kinds.

Achieving reliable identification of potential targets requires that identification data from many different sources be collected and effectively interpreted. Doing so requires first that a system be developed for collecting the outputs of several distinct sources and making them available at a single processing node. This is a formidable system-integration task, particularly when it must deal with merging sensors hosted on different platforms and possibly from different military Services. Although the challenges of implementing such an I.D. system integration should be understood, this paper will proceed from the assumption that the integration has been achieved. Figure 1 represents

the combat identification (CID) process: an unknown target is evaluated by a number of I.D. discriminants (here called *sensors*, regardless of the I.D. means employed). These individual evaluations are presented to a single fusion site where they are evaluated in accord with prescribed fusion rules leading to several possible I.D. declarations. For the purposes of this paper we have adopted the nomenclature employed in a developmental system for integrating I.D. sensors in AEGIS. The five designations are

- hostile
- unknown-assume enemy (uae)
- unknown evaluated (uev)
- unknown-assume friend (uaf)
- friend.

We assume that all unknowns can be divided into true-friend and true-hostile categories, and that the declarations represent shades of gray appropriate to the confidence of the I.D. process. In general, only the “hostile” declaration will lead to direct engagement of the target.

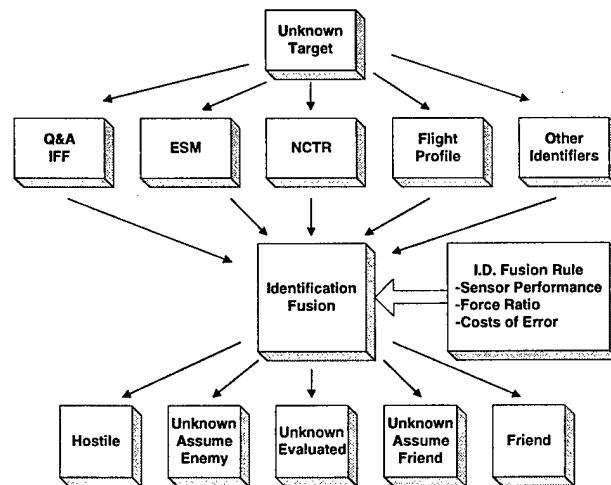


Figure 1. The Combat I.D. Process

II. ANALYSIS

A. THE SENSOR PERFORMANCE MATRIX

In the most general view, a combat identification system (CIS) comprises a set of independent means of making friend/foe identification. These means may include visual identification based on aircraft type and markings, judgments of whether an aircraft complies with current airspace control procedures (ACPs); third-party information from command and control (C²) nets; direct question-and-answer identification, friend or foe (IFF) sensor systems (Q&A); and possible noncooperative target recognition (NCTR) means. Whether or not these means involve specific hardware suites, we refer to them as sensors. The individual performance of each sensor in discriminating friend from foe can be expressed in a matrix of the form shown in Table 1. The two columns of this matrix, labeled "F" and "H," correspond to the true identity of the unknown target being evaluated. The n rows of the matrix, labeled "output state 1" through "output state n ," correspond to the n possible indications or outcomes of a test or observation employing that sensor. For example, an attempt at visual identification of an unknown aircraft might result in one of three indications: "friend," "hostile," or "cannot identify." In that case, the sensor performance matrix will have three rows, with each element of the sensor performance matrix representing the probability that a true-friend or true-hostile aircraft will be identified friend, identified hostile, or not be identified.

Table 1. Performance Matrix of Sensor (1)

Sensor (1)	True Friend (F)	True Hostile (H)
Output State (1)	$p(1 F)$	$p(1 H)$
•	•	•
•	•	•
•	•	•
Output State (n)	$p(n F)$	$p(n H)$

Other identification sensors may permit more than three distinct output indications. For example, advanced electronic IFF means may be able to categorize friendly indications into varying "degrees of friendship," depending on the level and

security of the reply received. As defined above, the sensor performance matrix elements are intended to embody objective data on the sensors in question. These probabilities could be obtained from exercises, tests, or analyses, and would be expected to vary with the nature and training of the forces engaged, as well as with combat theater and conditions. In general, we would expect the matrix elements to depend on the relative positions of the unknown and the identifying sensor. That dependence is not treated here, although it should be straightforward to include it if the I.D. system is integrated with a tracking system holding spatial information.

As a concrete example, consider a CIS comprising four distinct sources of I.D. information: Q&A IFF (e.g., Mark 10/12 IFF), a generic NCTR system, direct visual observation, and a track data network (e.g., Link 16). The sensor performance matrices given in Table 2 summarize the performance of this suite of sensors (values are arbitrarily chosen and do not represent the performance of any particular system). Note that for each sensor the sum of the probabilities in each column is unity. That is, each sensor gives some indication when applied to a target, even if the reply is “no reply” or “cannot identify.” In this analysis even null replies may contain useful information when fused with information from other sensors.

Table 2. Performance Matrices of Hypothetical CID Sensors

Q&A IFF			NCTR		
Sensor Output	True Identity		Sensor Output	True Identity	
	F	H		F	H
positive	0.6	0.1	f	0.4	0.1
no reply	0.4	0.9	h	0.05	0.3
			u	0.55	0.6

Visual Observation			C3		
Sensor Output	True Identity		Sensor Output	True Identity	
	F	H		F	H
f	0.5	0.1	f	0.7	0.05
h	0.1	0.5	h	0.05	0.2
u	0.4	0.4	u	0.25	0.75

B. COMBAT IDENTIFICATION SYSTEM STATES

Let N_s denote the total number of different sensors in a given combat identification system ($N_s = 4$ in the example above), and let i denote a particular sensor within that system. Then $1 \leq i \leq N_s$. Let n_i represent the number of allowed indicator states of the i th sensor (that is, the number of rows in its performance matrix) and let k_i denote a

specific output state of the i th sensor. Then $1 \leq k_i \leq n_i$. Unless some sensors are totally correlated, there will be a total of N distinct configurations of the overall CIS given by

$$N = \prod_{i=1}^{N_s} n_i \quad . \quad (1)$$

For the example above, $N = 54$. Each state of the CIS could be designated by listing the individual indications of each of the component sensor subsystems. That is, by a state vector $|k_1, \dots, k_{N_s}\rangle$. For convenience, however, we will define a single index, j , that runs over all N states of the system. Then $1 \leq j \leq N$. When we write $|j\rangle$ we refer to that configuration defined by the corresponding set of individual sensor states $k_1 \dots k_{N_s}$. Thus, the conditional probability that the overall CIS arrives at configuration $|j\rangle$ is denoted by $p(j|F)$ if the unknown is, in fact, a friend, and by $p(j|H)$ if the unknown is, in fact, hostile.

The computation of $p(j|F$ or $H)$ from the individual performance matrix elements $P(k_i|F$ or $H)$ requires consideration of the degree of correlation between sensors. For example, we would not reasonably expect that a ground-based visual observer's ability to correctly identify an aircraft would depend on the functioning of the aircraft's IFF transponder. Thus, there should be no significant correlation between visual identification and Q&A IFF. On the other hand, identification information on a C^2 net is derived from many sources, including possible visual sightings and possible previous attempts at electronic identification. In this case, some degree of correlation may exist between the indication of a local CIS sensor and the data held on a C^2 net, which may reflect prior observations using a similar sensor. Once this correlation is known, it may be included in straightforward fashion. Nevertheless, to simplify the following discussion and examples, we will assume that the degree of correlation between sensors is negligible.

With this assumption, the probability that an attempt to identify an unknown, but in fact friendly, aircraft will put the CIS into state, j , is

$$p(j|F) = \prod_{i=1}^{N_s} p(k_i|F) \quad . \quad (2)$$

The individual probabilities, $p(k_i|F)$, are obtained from the "Friend" columns of the appropriate sensor performance matrices. Similarly, the probability that a hostile aircraft will put the system in state, j , is:

$$p(j|H) = \prod_{i=1}^{N_s} p(k_i|H) \quad , \quad (3)$$

where the $p(k_i|H)$ are given by the “Hostile” columns of the sensor performance matrices. Because any unknown aircraft, whether friend or hostile, will put the CIS into some state (even if it is the “undetermined” state), we have

$$\sum_{j=1}^N p(j|F) = \sum_{j=1}^N p(j|H) = 1 \quad . \quad (4)$$

C. I.D. FUSION RULES

There is every reason to expect that some states of the CIS will involve conflicting indications from two or more individual sensors. The identification fusion rule (sometimes referred to as a “rule of engagement”) must resolve all such possible conflicts, specifying whether or not to declare a target “hostile” and hence engageable for each of the N states of the CIS. Historically, these fusion rules have tended to be expressed verbally. (For example: “Engage only if either visual I.D. or NCTR indicates hostile, and neither IFF nor C^2 indicates friend.”) An earlier paper (Ref. 2) explored some of the limitations of verbal rules in contrast with the much greater flexibility of the objective rules that can be defined mathematically. In this paper we address the question of how to choose from among the many possible objective rules those that minimize the total cost of misidentification.

Consider a case in which only two I.D. designations are used—friend and hostile. In that case, a complete I.D. fusion rule can be expressed as a vector, R , of length N . Each element of the vector is 1 if the “hostile” I.D. is assigned for that state and 0 if “friend” is assigned. For example, $R(j)$ might look like (1,0,1,1,0,...). Under this fusion rule, the probability of correctly assigning a hostile I.D. to a true-hostile is then

$$p(h|H) = \sum_{j=1}^N p(j|H) \cdot R(j) \quad , \quad (5)$$

while the probability of mistakenly assigning the “hostile” designation to a true-friend is

$$p(h|F) = \sum_{j=1}^N p(j|F) \cdot R(j) \quad . \quad (6)$$

The problem is to choose the fusion rule $R(j)$ to maximize the former while minimizing the latter. The total number of distinct possible fusion rules is 2^N , or $\sim 10^{16}$ for the example case. Clearly, it is impossible to enumerate and test all such possible rules. We will show that a small subset of all possible fusion rules, representing the best performance that can be obtained with a given sensor suite, can be defined and selected.

The problem of fusing the indication of identification sensors is expressed graphically in Figure 2. Ideally, we want $P(h|H) = 1$ with $P(h|F) = 0$, which implies

perfect identification of all targets. Although that performance is not attainable, it is possible to determine how close we may approach it with a given set of sensors, as defined by their sensor performance matrices. To begin, two fusion rules immediately suggest themselves: “never declare hostile—regardless of the identification system configuration,” and “always declare hostile.” Although trivial, these are rational rules in the following sense: “never declare hostile” guarantees zero fratricide, and no alternative rule can provide better defense effectiveness without increasing the risk of fratricide. Similarly, “always declare hostile” guarantees that all hostiles will be engaged, and no alternative rule can provide lower fratricide without also reducing defense effectiveness.

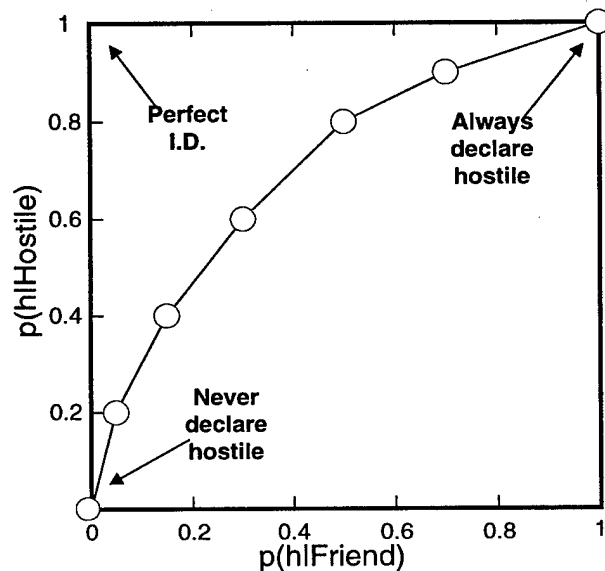


Figure 2. Graphing CIS Performance

Although neither of these rules may be appropriate to apply in most realistic scenarios, they provide two clear points of reference. Consider the “never declare hostile” rule, for which $R(j) = 0$ for all j . This is clearly the most conservative rule; the next most conservative rule is to engage in the single state, j , for which the likelihood ratio $P(j|H)/P(j|F)$ is largest. This gets us the maximum distance along the ordinate of Figure 2 (effectiveness), for the minimum distance along the abscissa (fratricide). This is also a rational rule of engagement, in that no other rule provides better effectiveness at the same fratricide, or vice versa. The next most conservative rule allows engagement on both this state and also on the system state with the next highest likelihood ratio $P(j|H)/P(j|F)$. This process is repeated, creating successively less conservative rules of engagements until the least conservative “always-engage” rule is reached [$R(j) = 1$ for all j]. The key point about this procedure is that the N distinct configurations of the CIS are mathematically tested using objective sensor performance data and then “turned on” in

decreasing order of their likelihood ratio. (In a looser sense, the likelihood ratio of each state is proportional to the “benefit/cost” ratio of declaring “hostile” on that state.) For a CIS with N states, there will be $N + 1$ points plotted. These points define a trajectory connecting the endpoints (0,0) and (1,1). Any point in this set represents a rational and objective rule of engagement in the sense that no alternative rule, applied to the same CIS, can provide higher effectiveness at the same or lower fratricide, or provide lower fratricide at the same or higher defense effectiveness. Within the set, each point represents an alternative trade-off between effectiveness and fratricide. Although, as pointed out above, the best trade-off among these alternatives will depend on combat requirements, the collective set of objective fusion rules completely and objectively characterizes the performance of the specific suite of identification sensors being analyzed. This trajectory of objective fusion rules summarizes the performance of an identification system in the same way that the “receiver operating curve” summarizes the detection/false-alarm performance of a detection system. Because of this similarity, the set of points defined above is denoted the *identification system operating characteristic* or ISOC.

Applying this process to the 4-sensor/54-state CIS we are using as an example, we first rank the 54 states in order of their hostile/friend likelihood ratio (see Figure 3). Large values of this ratio are more likely to result from a true-hostile target and small values are more likely to be caused by a true-friend.

The individual probabilities, $p(j|H)$ and $p(j|F)$, are plotted in Figure 4 in the same order as in Figure 3. Although the likelihood ratios are monotonic, the individual state probabilities fluctuate widely with the ranked state index. There is a general trend, however, for state probability to increase with ranked index for friends and to decrease for hostiles.

Once the states are ordered in accord with their likelihood ratios, we now know in which order to take the partial sums used to yield the ISOC. The points of the ISOC $[x(m),y(m)]$ are parameterized by the variable limit of summation, m , and we have

$$\begin{aligned} x(m) &= \sum_{j=1}^m p(j|F) \\ y(m) &= \sum_{j=1}^m p(j|H) \end{aligned} \tag{7}$$

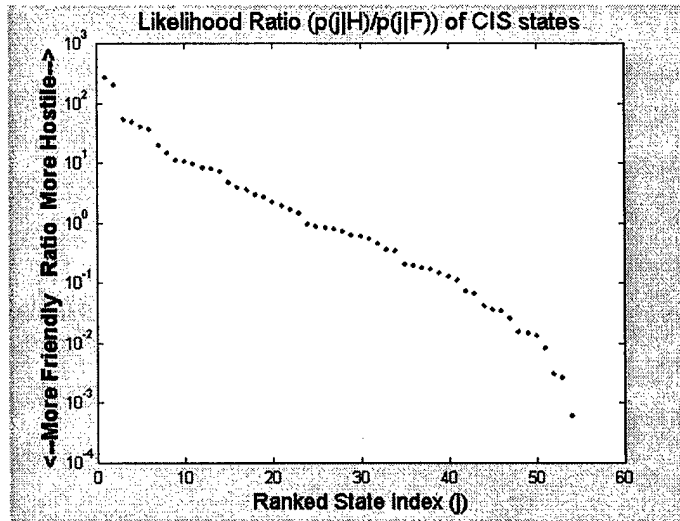


Figure 3. Ranked Likelihood Ratios of Example CIS States

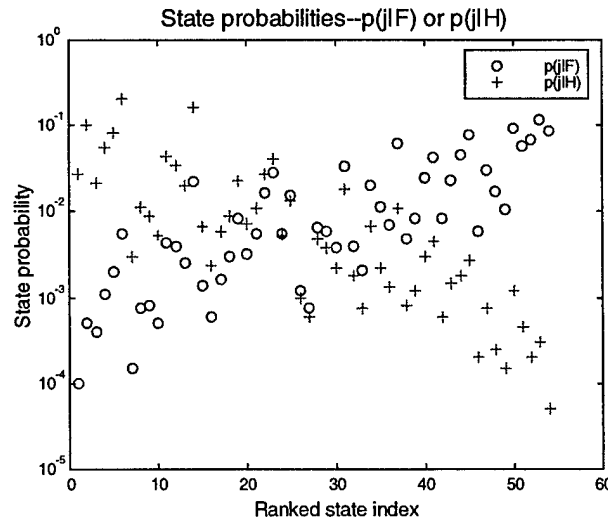


Figure 4. Probabilities of Ranked CIS States

Figure 5 plots the ISOC for the example suite of four sensors. Each point of the ISOC is a complete I.D. fusion rule. Together, the $N + 1$ ISOC points comprise the complete set of objective I.D. fusion rules. They are defined solely by an objective characterization of the performance of the sensor suite. They reflect no verbal, policy or traditional inputs or constraints, and represent the best discrimination performance that a given CIS can yield in distinguishing friend from hostile. At this point of the analysis all of these alternative rules may be considered equally valid. Although they differ in degree of “conservativeness,” none is clearly better than any other, and no means has been provided to determine which of the rules represents the best operating point for the CIS.

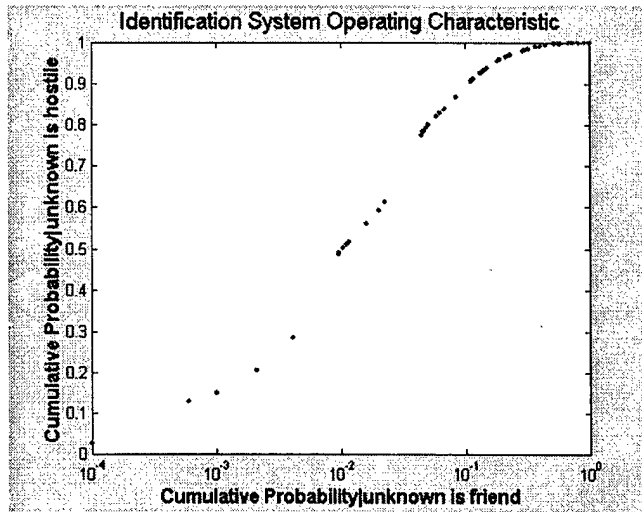


Figure 5. Identification System Operating Characteristic

To determine that, we require additional information about the anticipated ratio of encountering true-friends and true-hostiles in the theater of operation and about the costs of making identification errors of different kinds.

D. COST OF IDENTIFICATION ERRORS

Consider a situation in which only two true identifications are possible—true-friend and true-hostile. In this case we assume that neutrals and allies would be regarded as friends and all adversary forces would be lumped together as hostile. Because a particular unknown may be suspected of hostile intent without accruing enough evidence to warrant being declared hostile, and hence engageable, we consider the range of I.D. declarations indicated in Figure 1. Ideally, all true-hostiles would be assigned the “hostile” declaration and all true-friends assigned “friend” declaration. Any other declaration would incur a cost of error which may be large or small. The values assigned to these costs represent subjective judgments on the part of command authorities about the relative consequences of different I.D. failures based on both quantitative and non-quantitative factors appropriate to a specific time and place. Figure 6 displays an example of such costs; the values chosen are arbitrary and for illustration only. For this example we assume that the worst error is to declare a true-friend “hostile.” Although assigning a true-friend the “unknown-assume enemy” label is also undesirable, it would not necessarily lead to immediate engagement, and so the cost might be much less. The costs of the other, less hostile, designations are similarly less and the cost of the correct “friend” I.D. is zero. For true hostiles we assume that the most costly designation is

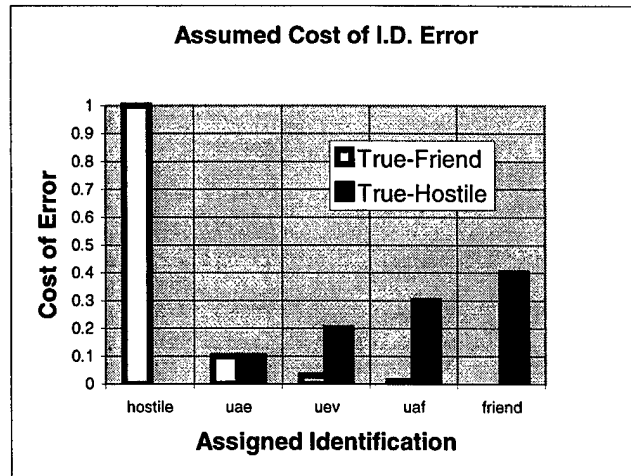


Figure 6. Example of I.D. Error Cost Estimate

“friend” with lesser costs assigned to the more ambiguous declarations and zero cost to the correct “hostile” declaration. For this particular example we assume that the consequences of possibly engaging a true-friend exceed those of allowing a true-hostile to pass unengaged. This reflects various assumptions about the lethality of our own defense systems, the effectiveness of hostile weapons, and the political and military consequences of engagement errors, all of which may vary widely in different theaters and circumstances.

The expected cost of evaluating an unknown target can be expressed by a *loss function*, L . This loss reflects the prior probabilities, $p(F)$ or $p(H)$, that a given unknown target is either true-friend or true-hostile, before any I.D. sensors are consulted. These prior probabilities may be estimated using various intelligence resources. The expected cost of evaluating a single unknown then depends on the probability that the unknown is friend or hostile, the probabilities that a true-friend or true-hostile will receive the various I.D. designations shown in Figures 1 and 6, and the cost of the consequent designations. The loss function expressing these costs is

$$L = p(F)[p(h|F)c(h|F) + p(uae|F)c(uae|F) + p(uev|F)c(uev|F) + p(uaf|F)c(uaf|F) + p(f|F)c(f|F)] + p(H)[etc \dots] \quad (8)$$

Here, $p(F)$ and $p(H)$ are the prior encounter probabilities, and the factors of the form $c(\text{“I.D. designation”} | F \text{ or } H)$ are the postulated costs of each accurate or mistaken I.D. declaration given the unknown’s true I.D. (cf. Fig. 6). To minimize L , we must vary the only parameters at our disposal, the probabilities of assigning various I.D. designations to friends and hostiles. This process is indicated schematically in Figure 7.

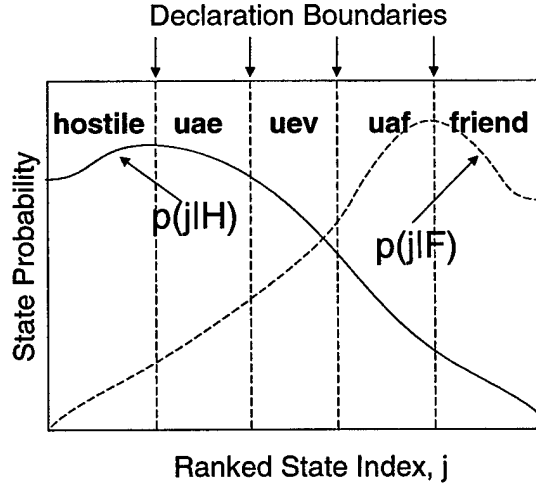


Figure 7. Setting Declaration Thresholds

This figure is a smoothed idealization of the state probability distributions shown in Figure 4, with the addition of the thresholds between I.D. declarations. The probability, for example, of declaring a true-hostile as “hostile” is the area of the $p(j|H)$ distribution within the “hostile” threshold area. We adjust the probabilities of various declarations by moving the threshold boundaries. Consider just the boundary between “hostile” and “uae.” At the minimum of the loss function the differential change of the loss with respect to changes in declaration probabilities is zero, so we have

$$dL = 0 = p(F)[c(h|F)dp(h|F) + c(uae|F)dp(uae|F)] + p(H)[c(h|H)dp(h|H) + c(uae|H)dp(uae|H)] \quad (9)$$

When we move this threshold, conservation of probability requires that

$$\begin{aligned} dp(h|F) &= -dp(uae|F) \\ dp(h|H) &= -dp(uae|H) \end{aligned} \quad (10)$$

Substitution into Eq. 9 yields, for the loss function minimum,

$$\frac{dp(h|H)}{dp(h|F)} = -\frac{p(F)}{p(H)} \cdot \frac{[c(h|F) - c(uae|F)]}{[c(h|H) - c(uae|H)]} \quad (11)$$

Because this is a discrete problem, the probability differentials, $dp(h|F)$ or H , are quantized, and hence so are their ratios. For CIS state j , therefore,

$$\frac{dp(h|H)}{dp(h|F)} = \frac{p(j|H)}{p(j|F)} \quad (12)$$

The first threshold criterion for CIS state likelihood ratios between the “hostile” and “uae” declarations is therefore

$$\frac{p(j|H)}{p(j|F)} = T_1 = -\frac{p(F)}{p(H)} \cdot \frac{[c(h|F) - c(uae|F)]}{[c(h|H) - c(uae|H)]} \quad (13)$$

If the likelihood ratio of the observed CIS state exceeds T_1 , declare the unknown "hostile," otherwise declare "uae" (or other, depending on the next threshold, etc.). This process is repeated to develop similar expressions for the other three least-cost thresholds demarcating the "uae," "uev," "uaf," and "friend" declarations. The resulting expressions for these declaration thresholds are

$$\begin{aligned}
 T_2 &= -\frac{p(F)}{p(H)} \cdot \frac{[c(uae|F) - c(uev|F)]}{[c(uae|H) - c(uev|H)]} \\
 T_3 &= -\frac{p(F)}{p(H)} \cdot \frac{[c(uev|F) - c(uaf|F)]}{[c(uev|H) - c(uaf|H)]} \\
 T_4 &= -\frac{p(F)}{p(H)} \cdot \frac{[c(uaf|F) - c(f|F)]}{[c(uaf|H) - c(f|H)]}
 \end{aligned} \tag{14}$$

Note that the least-cost declaration thresholds depend only on the assumed costs and the prior probabilities of encountering true-hostile and true-friend. They do not depend on sensor performance. The rank ordering of states, on the other hand, depends solely on sensor performance. Thus the problem has separated into independent parts.

E. LEAST-COST I.D. DESIGNATIONS

To illustrate, we apply the threshold analysis above to the hypothetical four-sensor CID suite under the cost assumptions of Figure 6. For this example we assume the prior probability of encountering true-friends is 0.9 and 0.1 for hostiles. The application of Eqs. 13 and 14 for each threshold yields the likelihood ratio thresholds given in Table 3.

Table 3. Computed Declaration Thresholds for Assumed Costs and Priors

Declaration	hostile		uae		uev		uaf		friend
Threshold		81		6.3		1.8		0.9	

With these thresholds we can revisit the ranked states of Figure 3, but this time associate a specific I.D. declaration with each state of the system in accord with its likelihood ratio. In Figure 8 we show the same CIS states encoded with the I.D. declarations that minimize the total expected cost of identifying both friends and hostiles, under the specific assumptions made here about performance, force ratio, and cost.

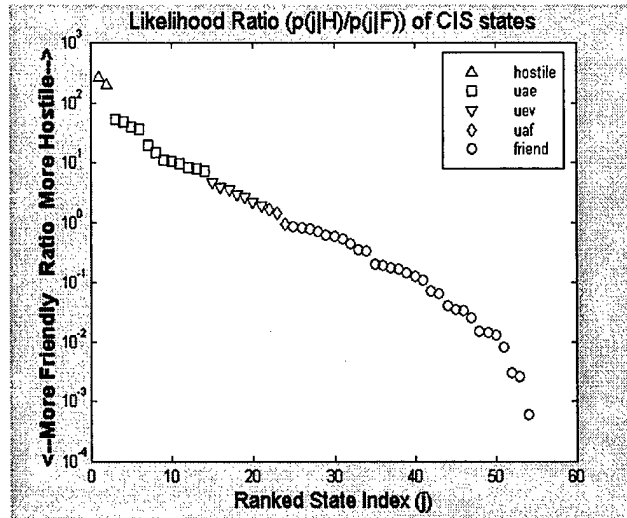


Figure 8. Least-Cost I.D. Declarations [$p(F) = 0.9$]

As we would expect from the modest performance assumed for the sensors in Table 2, the I.D. performance of the postulated CIS is far from perfect, even when the best sensor fusion, in the sense of lowest cost, is applied. The example CIS will classify true-friends and true-hostiles with the probabilities shown in Table 4, obtained by summing the state probabilities within the computed declaration boundaries. Notice that although 88 percent of true-friends are so declared, only ~13 percent of true-hostiles are so declared. Together, the fact that 90 percent of the potential targets are known to be true-friends, the high cost of mistakenly assigning friends the “hostile” designation, and the modest level of performance of our I.D. sensors mean that only the states with the highest likelihood ratios lead to “hostile” declarations. In effect, a high prior probability that an unknown is a friend can be overcome only by the strongest possible evidence of the contrary (see Figure 8.)

Table 4. Declaration Probabilities for Hypothetical CIS and Costs

I.D. Declaration	True Identification	
	Friend	Hostile
hostile	0.0006	0.1283
uae	0.0441	0.6470
uev	0.0236	0.0647
uaf	0.0491	0.0729
friend	0.8826	0.0873

An even clearer example of this results if the prior probability that unknowns are true-friends is 0.99, a situation which could arise after several days of a successful counter-air campaign. Changing prior probability will not change the likelihood ratios of any of the CIS states or their ranking, but it does change the declaration thresholds. Figure 9 shows the I.D. declarations resulting from this case. Note that the least-cost declaration is never “hostile” for any of the CIS states. This figure illustrates the case when the cost of I.D. error, the high frequency of friendly encounters, and the known limitations of I.D. sensors combine to make “never engage” an appropriate rule of engagement.

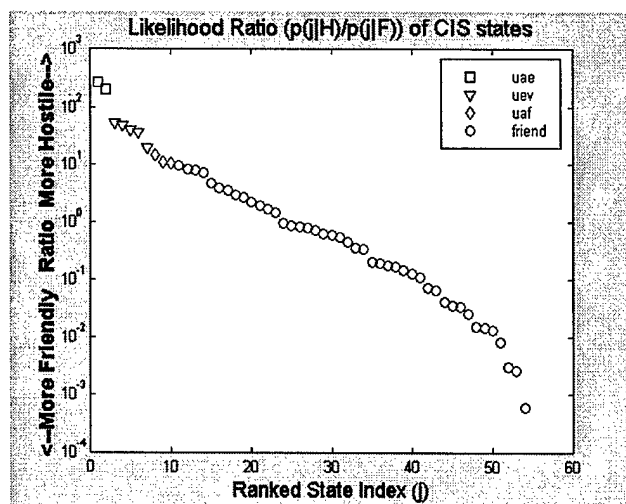


Figure 9. Least-Cost I.D. Declarations [$p(F) = 0.99$]

The contrary case may also be considered, in which we assume a more severe threat. The prior probability of encountering friends is dropped to 0.5, and, more significantly, the cost of assigning the “friend” declaration to a true-hostile is increased to 1.0—the same as the cost for misidentifying a true-friend as “hostile.” Sensor performance is assumed unchanged from the base level of Table 2. Not unexpectedly, in this case, shown in Figure 10, the number of CIS states leading to “hostile” declarations is greatly increased and those leading to “friend” classifications is correspondingly reduced.

Yet another variation arises from adding additional independent sensors to the four subsystems included in the baseline suite of Table 2. We postulate that I.D. information may be obtained by observing an unknown’s compliance with established airspace control procedures (ACPs) that constrain flight corridors, speeds, and altitudes of friendly aircraft. In addition, it may be possible to continuously track some fraction of

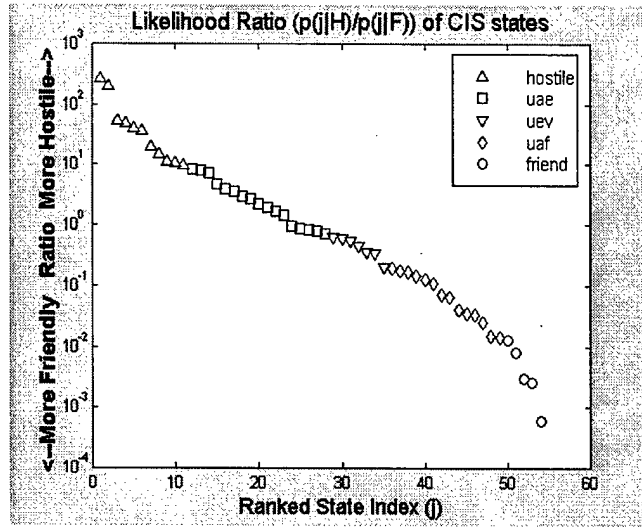


Figure 10. $p(F) = p(H) = 0.5$. $c(\text{friend}|\text{Hostile}) = 1.0$

both friend and hostile units from their points of origin. Finally, we assume that at least some of the true-hostile units can be observed undertaking hostile action. For these three additional sensors we arbitrarily postulate the performance shown in Table 5. Note that none of these additional sensors is assumed to have outstanding I.D. characteristics; the important assumption is that all of the sensors are independent of each other. The effect of adding these additional sensors to the CIS can be shown by comparing the I.D. System Operating Characteristics (ISOCs) of the original and augmented CISs, shown in Figure 11. At the operating point for which 50 percent of the true hostiles are declared hostile, the 4-sensor system has a probability of ~1 percent of declaring a true-friend hostile. The 7-sensor system, at the same level of hostile declaration, would declare ~0.2 percent of true-friends hostile. Thus, adding additional imperfect sensors can increase the probability of correct identification provided the new sensors are not correlated with the old.

Table 5. Performance Matrices of Additional Hypothetical Sensors

Airspace Control Procedures			Point of Origin			Hostile Action		
Sensor Output	True I.D.		Sensor Output	True I.D.		Sensor Output	True I.D.	
	F	H		F	H		F	H
<i>f</i>	0.6	0.2	<i>f</i>	0.4	0.05	<i>h</i>	0.05	0.4
<i>h</i>	0.2	0.5	<i>h</i>	0.3	0.4	<i>u</i>	0.95	0.6
<i>u</i>	0.2	0.3	<i>u</i>	0.3	0.55			

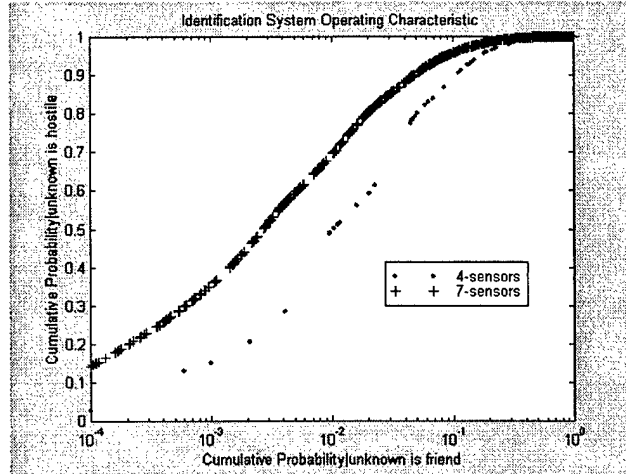


Figure 11. Comparative ISOCs of Original and Augmented CIS

F. DECLARATION PROBABILITY

In some cases, it may be required that any I.D. declaration have less than a stated probability of being in error. For example, it may be required that the probability of declaring a true-friend as “hostile” be less than 1 percent and that the probability of declaring a true-hostile as “friend” be similarly limited to less than 1 percent. Where these confidence criteria cannot be met then no I.D. declaration is made.

This constraint can be readily described with reference to Figures 4 and 7. Once we have ranked the states of the CIS in order of the hostile/friend likelihood ratio, we can declare unknowns to be hostile if they put the CIS into any state from the first or “most-hostile” state up to the state index, m_1 , for which

$$\sum_{j=1}^{m_1} p(j | F) = p(h | F) \leq T_1 \quad . \quad (15)$$

That is, we are willing to declare unknowns as hostile for all CIS states up to the point where the accumulated probability of mistakenly declaring a true-friend “hostile” reaches the threshold, T_1 . Similarly, we are willing to declare as friends unknowns that put the CIS into any state from the last or “most friendly” state down to the state index, m_2 , for which

$$\sum_{j=m_2}^N p(j | H) = p(f | H) \leq T_2 \quad , \quad (15)$$

which means we declare unknowns as friend for all CIS states down to the point where the accumulated probability of mistakenly declaring a true-hostile to be “friend” reaches the threshold, T_2 . In this way, we establish two threshold indices in Figure 4. The

probability that we will have insufficient confidence to declare I.D. is the probability weight between these threshold indices. Thus the probability of “non-declaration,” $p(nd)$, is given by

$$\begin{aligned}
 p(nd | F) &= \sum_{m_1 < j < m_2} p(j | F) \\
 p(nd | H) &= \sum_{m_1 < j < m_2} p(j | H) \\
 p(nd) &= p(nd | F)p(F) + p(nd | H)p(H) \\
 p(declare) &= 1 - p(nd)
 \end{aligned}
 \tag{17}$$

Note that the declaration probability depends on the prior force ratio. In Figure 12 we plot the declaration probability for both the original 4-sensor CIS and the augmented 7-sensor system over a range of values of the force ratio, $p(F)/p(H)$. This figure shows that although it is possible to constrain the probability of declaration errors to relatively low values, even with mediocre I.D. sensors, the price paid is a system that gives no I.D. declaration a large fraction of the time.

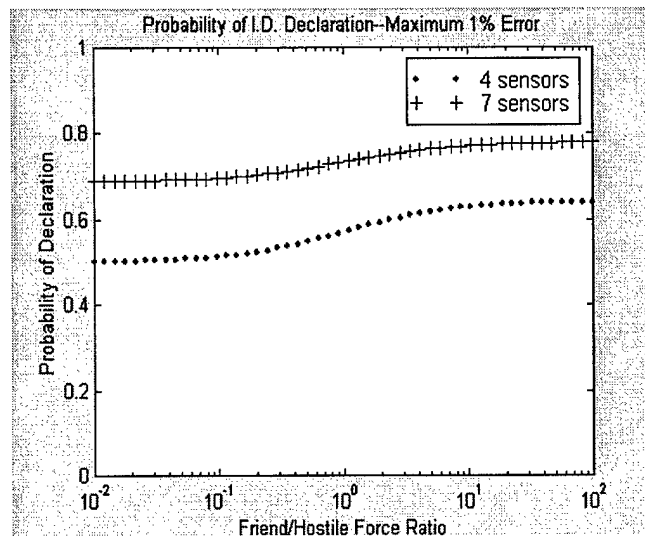


Figure 12. Declaration Probability with Constrained Errors

III. IMPLEMENTATION

Implementing this methodology as part of a real-time CIS requires that both objective and subjective data be supplied. Objective data include sensor performance and prior Hostile/Friend force ratio. The most critical, and probably most difficult, step in the implementation of this methodology in a CIS is obtaining accurate estimates of the performance of individual I.D. sensors. The most desirable way to determine sensor performance would be through detailed measurements, but this is unlikely to be practical given the wide range of circumstances that could conceivably be involved. The performance of all I.D. sensors is likely to depend on the range from the sensor to the unknown targets and possibly on other factors as well, such as target altitude and aspect angle. Although all of these variables can in principle be included in a measurement program, in practice it is unlikely that a complete set of measurements can be provided. Nevertheless, whether the sensor performance matrices are determined by measurement, analysis or some other method, such as the judgment of an expert panel, all subsequent analysis is conditioned on the validity of the values employed. Once an adequate set of performance matrices, possibly reflecting range and other dependency, is obtained, the second step is to estimate the prior force ratio. This would normally be accomplished as part of intelligence preparation of the battlefield (IPB).

Obtaining estimates of the cost of I.D. determination errors for both friends and hostiles is likely to rely heavily on subjective command judgments based on a variety of factors. These include estimates of the payloads carried by hostile penetrators, the vulnerability of their most likely targets, the effectiveness of hostile targeting and weapon delivery systems, the effectiveness of point-defense systems, the lethality of U.S. defense systems, and the military and political costs of fratricide. Although there is no reason in principle why some or all of these factors cannot be quantified and subjected to objective analysis, it is likely to be impossible to completely escape the need for subjective judgments.

The third step is to use the prior force ratio and misidentification costs to determine the least-cost likelihood ratio thresholds between I.D. declarations. This step need only be done once for each set of costs and force ratios. For each unknown encounter, the likelihood ratio of the CIS indication state is computed and compared to

the thresholds. No rank ordering is needed. The likelihood ratio of each state is computed as it arises. This process leads to the lowest cost I.D. declaration consistent with all information and assumptions.

IV. CONCLUSIONS

In this paper we have described a process of CID sensor fusion that is sensitive to known CID sensor performance, known ratio between friendly and hostile encounters, and the perceived costs of I.D. errors. This process can be used to study existing I.D. fusion rules, such as those embodied in the AEGIS and PATRIOT systems, and also to provide and evaluate alternative procedures. This method can be used to supplant or augment existing I.D. procedures for real-time I.D. sensor fusion. If implemented in this way, this method would facilitate the rapid adaptation of CIS rules to the specific circumstances of local theaters.

The most time-consuming part of this process is likely to be obtaining the objective sensor performance data. Although some data may be available from evaluation and readiness tests, the most desirable (and costly) procedure would be to conduct specifically instrumented tests. As an interim alternative, the consensus of an expert panel could provide the necessary data to allow further evaluations of these methods.

The following are among possible topics for additional research on this subject:

- The effect of partial correlation between I.D. sensors.
- Temporal and spatial variation of sensor performance.
- Ranges of uncertainty for prior force ratios and sensor performance.
- Sensitivity of optimum rules to assumed costs of error.
- Performance comparisons with other I.D. fusion rules.

REFERENCES

1. Harry L. Van Trees, "Detection, Estimation and Modulation Theory," John Wiley, 1968.
2. James M. Ralston, "Objective Rules of Engagement and Sensor Fusion in Combat Identification Systems," *Proc. of the 1989 Symposium on Command and Control Research*, pp. 153-8.

GLOSSARY

ACP	airspace control procedures
CID	combat identification
CIS	Combat Identification System
I.D.	identification
IFF	identification, friend or foe
IPB	intelligence preparation of the battlefield
ISOC	identification system operating characteristic
NCTR	noncooperative target recognition
Q&A	question and answer
uae	unknown-assume enemy
uaf	unknown-assume friend
uev	unknown evaluated

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 the 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 Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE June 1999	3. REPORT TYPE AND DATES COVERED Final — January 1998 – October 1998	
4. TITLE AND SUBTITLE Bayesian Sensor Fusion for Minimum-Cost I.D. Declaration			5. FUNDING NUMBERS DASW01-97-C-0056 CA-2-1617	
6. AUTHOR(S) James M. Ralston				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Institute for Defense Analyses 1801 N. Beauregard St. Alexandria, VA 22311-1772			8. PERFORMING ORGANIZATION REPORT NUMBER IDA Paper P-3441	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Joint Theater Air Missile Defense Organization Crystal Mall 3, Suite 511 1931 Jefferson Davis Highway Arlington, VA 22203			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 180 words) Current operational procedures for I.D. sensor fusion tend to be rigid, stereotypical, and based on unvalidated assumptions about the performance of specific I.D. discriminant sources, such as Q&A IFF, NCTR, etc. To be optimum, any I.D. sensor fusion procedure must be <i>explicitly</i> sensitive to several factors. These include (1) the objective performance of the specific sensors in use, (2) the costs of making different kinds of I.D. declaration error, and (3) the friend/hostile force ratio at the specific time and place of CID declaration. In this paper we describe an analytic methodology for developing and evaluating an I.D. sensor fusion procedure that is sensitive to these factors and which can be shown to yield the lowest risk I.D. declarations, in the sense of having the lowest expected cost due to errors of all kinds. This approach to fusion analysis allows the performance of individual sensors to be reflected directly in CID effectiveness and allows the connection between individual sensor performance and the overall system performance requirement to be traced. By permitting direct visibility of and access to the often "hidden" assumptions about force ratio and cost of error, it can promote the development of fusion rules suited to the joint operations and coalition forces environment. Operationally, it would permit I.D. declaration rules to be quickly adapted to the threat, forces, and sensor performance appropriate to a particular theater. It is possible that procedures based on this analysis could be applied within existing and developmental computer-based command and decision systems.				
14. SUBJECT TERMS Bayesian Fusion, combat identification, rules of engagement, sensor fusion			15. NUMBER OF PAGES 32	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT SAR	