

LEVEL:

12

DSA Report No. 314

AD A106391

**AN INFORMATION FUSION SYSTEM FOR WARGAMING
AND INFORMATION WARFARE APPLICATIONS**

Final Report

DTIC
OCT 26 1981
H

Prepared Under Contract No. N00014-80-C-0567 for:

Office of Naval Research
800 North Quincy Street
Arlington, Virginia 22217

DISTRIBUTION STATEMENT A
Approved for public release
Distribution Unlimited

28 August 1981



DECISION-SCIENCE APPLICATIONS, INC.

1500 WILSON BOULEVARD, SUITE 510, ARLINGTON, VIRGINIA 22200

81 0

DSA REPORT 88-314

(3)

AN INFORMATION FUSION SYSTEM FOR WARGAMING
AND INFORMATION WARFARE APPLICATIONS.

FINAL REPORT . . .

SECRET
H

George E. Pugh
David F. Noble

Prepared under Contract No. N00014-80-C-0567 for

Office of Naval Research
900 North Quincy Street
Arlington, Virginia 22217

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

28 August 1981



DECISION-SCIENCE APPLICATIONS, INC.

1800 WILSON BOULEVARD SUITE 810 ARLINGTON, VIRGINIA 22200

CONTENTS

<u>SECTION</u>		<u>PAGE</u>
1.0	INTRODUCTION	1
	1.1 Advantages of New Analytical Methodology	2
	1.2 Potential Applications	4
	1.3 The Design Concept	5
	1.4 Suggested Development Program	6
	1.5 Organization of this Report	7
2.0	BACKGROUND AND RESEARCH APPROACH	9
	2.1 The CSAP Project	9
	2.2 Initial Study Recommendations	10
	2.3 The Research Approach	11
3.0	OPERATING CONTEXT FOR INFORMATION FUSION SYSTEM	13
	3.1 Simulation of Intelligence Processes	15
	3.2 Information Processing Functions to be Simulated	19
	3.3 Other Complexities of Ocean Surveillance Intelligence	20
4.0	OVERVIEW OF DESIGN ARCHITECTURE	25
	4.1 The Essential Role of Historical Time	25
	4.2 Separation of Tracking and Identification Functions	30
	4.3 Impact of Tactics and Deception	31
	4.3.1 The Effect of Tactical Correlations	32
	4.3.2 Nonuniformity of Tactical Probabilities	33
	4.3.3 A Game Theory Analogy	34

Accession For
NTIS GRA&I
DTIC TAB
Unannounced
Justification For
By _____
Distribution
Availability Codes
Dist _____
A

CONTENTS (Cont.)

<u>SECTION</u>		<u>PAGE</u>
	4.3.4 Intelligence Methods as a Design Pattern	34
	4.4 The Detailed Design Concept	35
<u>ANNEX</u>		
A	CONCEPTUAL DESIGN FOR THE HISTORICAL CORRELATOR	37
B	CONCEPTUAL DESIGN FOR THE SHIP IDENTITY INFERENCE SYSTEM	65

FIGURES

<u>NO.</u>		<u>PAGE</u>
3-1a	Satellite Observation at 0400 (Received for Processing at 0600)	14
3-1b	Reports at 0800 Plotted on Same Arbitrary Coordinate Grid	14
3-2a	Satellite Observation at 0800 (Received for Processing at 1000)	16
3-2b	Bear D Observation at 1000 (Received at 1100)	16
3-3	Illustrating Relationship of Physical and Intelligence Events in an Intelligence Simulation	18
4-1	Illustrating the Essential Role of Historical Time	27
4-2	Comparison of Method of Historical Reinterpretation (Lower Diagram) with Direct Update of Situation Perception (Upper Diagram)	29
A-1	Illustrating Relationship Linkages in Knowledge Representation	48

1.0 INTRODUCTION

This report describes a new design concept for an ocean surveillance information fusion system that was developed under the Navy's Counter Surveillance Assessment Program (CSAP). The design concept utilizes a remarkably simple and computationally efficient new deductive algorithm that should permit real time tactical information fusion either in a wargame environment or in automated support of shipboard information fusion. The algorithm can also be applied in the context of an information warfare decision aid, to assist in evaluating what an opponent could learn from alternative courses of action that might be selected for cover and deception tactics, or emission control (EMCON) alternatives.

The new deductive algorithm achieves its computational efficiency by means of deductive procedures which closely parallel the methods used by human intelligence analysts. When these intelligence methods are formalized for application in a computer, the resulting system is, of course, far faster and more accurate than human analysis; and it is also free from the biases and personal prejudice that is unavoidable in human intelligence analysis. Consequently in a shipboard tactical environment the system should substantially speed up the intelligence cycle, and it should also contribute to greater accuracy in the assessment of the tactical situation by providing an unbiased second opinion against which the human analysts could check their own conclusions. In a wargaming environment, where the system is used to simulate human deductive processes, it may be necessary to deliberately degrade the system's performance to match the more limited computational capacity of human analysts.

In addition to its computational speed, the new deductive algorithm has another unexpected advantage over the traditional formal methods of mathematical inference. When the traditional formal methods are applied in an environment of deliberate

tactical deception, the calculated conclusions often appear to be more naive (i.e., more vulnerable to deceptive tactics) than an experienced intelligence analyst. Because the new algorithm is based on fundamental principles similar to those used in human intelligence analysis, it tends to duplicate the well justified conservatism of an experienced analyst.

These unique advantages of the present design concept can be viewed as an unexpected by-product of the design objectives that were required for the CSAP program. Because of the counter surveillance goals of the CSAP program, the effort was focused, not on the general problem of information fusion, but rather on the assessment of what Soviet ocean surveillance and intelligence systems might be able to learn from the electronic emissions of U.S. ships. The initial design effort, therefore, focused on the formulation of computer algorithms capable of simulating those human analytical processes that are required in an intelligence network, in order to integrate the available data into a coherent situation assessment.

Because the original design objective was simply to simulate (rather than to improve) the normal human intelligence processes, the analysis led quite naturally to a design concept that is quite different than would have resulted from a standard application of the methods of mathematical inference.

1.1 ADVANTAGES OF NEW ANALYTICAL METHODOLOGY

The main advantages of the new approach relative to the more conventional mathematical methods can be summarized as follows:

1. Computational efficiency. The traditional Bayesian methods of mathematical inference require very large combinatorial calculations that can make it difficult or infeasible to solve practical problems on even the largest computers. Because human intelligence analysts must, of necessity, work with limited

computational resources, the basic methods they use for information fusion tend to be far more efficient in the use of computational resources. When these common-sense methods are formalized in computer algorithms, the resulting system is capable of carrying out the basic information fusion operations not only very rapidly, but also far more accurately and comprehensively than is possible for human analysts. In addition, the compactness and efficiency of the new algorithm opens the possibility of automatic information fusion systems that could be used aboard ships to assist intelligence officers in providing a more up-to-date situation assessment and to help in evaluating information warfare alternatives.

2. Reduced vulnerability to deception. One of the common problems that is encountered in the application of the standard methods of mathematical inference is that they are not designed to deal with deliberate tactical deception. As a consequence, they can be unrealistically naive in their interpretation of the available evidence. Human intelligence analysts, on the other hand, are aware of the possibility of deception and tend to be more conservative in drawing firm conclusions from a limited amount of evidence. The use of mathematical algorithms based on human analytical methods makes it possible to produce conservative results that are less vulnerable to deception than the standard methods of mathematical inference.
3. Analytical simplicity. The resulting design concept has one other unexpected advantage. The mathematical methodology is analytically very simple and is thus easy to generalize to incorporate a wide variety of practical considerations that would be very difficult to incorporate in the traditional methods of

mathematical inference. This analytical simplicity is particularly unexpected because human methods of deduction seem subjectively to be very ad hoc and disorganized. Thus, the discovery of a simple mathematical generalization capable of simulating these mental processes came as a major surprise. The basic simplicity of the method reduces the complexity of the required computer code and should allow much simpler and more economical information fusion systems.

1.2 POTENTIAL APPLICATIONS

The basic design concept developed here has a wide variety of potential applications, both in wargaming and in the development of tactical decision support systems. Because the basic concepts are new, we believe that the system should be developed initially for use in a wargaming environment. This should provide a relatively low cost environment in which the system could be tested, evaluated, and modified before attempting to develop any operational systems.

1. Wargaming applications. In the context of man/machine tactical wargames, the system could be used to provide real time automation of the basic information fusion processes. In this role, it should greatly reduce the burden of routine analysis required of human players, and it should provide a realistic context for assessing the potential tactical value of such a system in an operational environment. Such an automated information fusion capability could also provide the foundation for the development of "automated players" that could substitute for the blue and/or orange team in certain scenarios.
2. Tactical information fusion. Probably the most obvious and potentially important application of the design concept would be to provide computerized

assistance in the tactical analysis of real time sensor data. In this role the system should greatly speed up the intelligence cycle and could provide a "second opinion" against which human situation assessments could be checked.

3. Decision support for information warfare. An automatic information fusion system could be used as an aid in estimating what the opponent might already know from previous observations, and what he might be able to learn from new evidence. Thus, such a system could be used to support a wide variety of information warfare decisions dealing with issues such as: cover and deception; emissions control and ECM; surveillance and reconnaissance; and intelligence information. Because of the very critical military importance of such activities, and the large amount of information required to make good decisions, computerized assistance in these areas could be of great importance to the military commander.

1.3 THE DESIGN CONCEPT

The present design concept for a tactical information fusion system involves two basic processing systems:

1. An historical correlator
2. A ship identity inference system

The historical correlator is concerned with the time and space correlations between ship tracks and observed signal sources. This correlation of apparent signal sources with historical ship tracks is necessary so that all historical clues can be collected and organized in such a way that they can be fully exploited in the evaluation of ship identity. The ship identity inference system then uses the accumulated clues to

provide an estimate of the identity of the ship or ships associated with each track.

The new ship identity algorithm involves a logical extension of the principle of minimum cross-entropy. The algorithm provides a computationally feasible method for dealing with very large and complex problems in mathematical inference. The development and demonstration of this algorithm was one of the major achievements of this project, and the practical feasibility of the present design concept is a direct consequence of this new analytical method.

During the present contract, a fairly detailed basic design concept was developed for both components. In addition, a laboratory prototype of the ship identity inference system was developed to test and evaluate the performance of the ship identity algorithm.

Obviously, a considerable amount of additional work will be required to convert the present design concept into a practical operating system that can be tested and evaluated. However, we believe the most difficult theoretical problems have already been solved, so the remainder of the development should involve relatively little technological risk.

1.4 SUGGESTED DEVELOPMENT PROGRAM

Because a practical system that could provide efficient tactical fusion of sensor data with intelligence information would be of great importance to the Navy, we believe that development work on the concept should proceed on a priority basis. On the other hand, because many of the concepts are new, it would be a mistake to try to move too rapidly toward an operational system. The required development time cannot be accelerated beyond certain reasonable limits by increasing the

level of effort. What is required is an orderly development effort by a few well qualified analysts to provide an initial well structured operating system.

If a decision is made to proceed with the development of such a tactical information fusion system, we believe the next step should be to implement the overall concept within the context of a tactical naval warfare simulation where the usefulness of the concept can be demonstrated, and where inevitable weaknesses in the original system design can be identified and corrected before attempting an operational system.

1.5 ORGANIZATION OF THIS REPORT

This report includes a main paper and two annexes. The design concept for the historical correlator is developed in Annex A, and the design concept for the ship identity inference system is developed in Annex B. The main paper is limited to a discussion of the broad architectural concepts and the relationship between the two systems. The remainder of the main paper includes three sections: a brief review of the background and research plan; a general discussion of the ocean surveillance environment and associated analytical problems; and finally, a discussion of the broad architectural concept for the information fusion system.

2.0 BACKGROUND AND RESEARCH APPROACH

Funding for the present study was provided by NAVELEX (under the Navy's Counter Surveillance Analysis Program, CSAP) through the Office of Naval Research. DSA's basic responsibility under the contract was to assist NAVELEX in integrating advanced decision science methodologies into the Counter Surveillance Analysis Program. In particular, the objective was to determine how specific decision science developments which had been pioneered by DSA under ONR sponsorship, such as:

1. The Theory of Measures of Effectiveness;
2. Electronic Warfare Decision Aiding; and
3. Computer Simulation of Deductive Logic and Decision Processes

could be most effectively applied within the context of the Counter Surveillance Analysis Program (CSAP).

2.1 THE CSAP PROJECT

Broadly speaking, the Navy's CSAP project was designed to achieve three major objectives:

- Identify U.S. Navy vulnerabilities associated with signal emissions.
- Evaluate promising U.S. cover and deception options in terms of both tactics and equipment.
- Identify procedures to effectively exploit weaknesses and inefficiencies in Soviet C3 systems.

To provide a systematic way of addressing these issues the CSAP project envisioned the development of a large tactical simulation in which detailed models of U.S. emitters and Soviet ocean surveillance systems would be used to identify U.S. vulnerabilities, and to evaluate promising cover and deception options.

In pursuing this program, however, CSAP project officers anticipated some difficult practical problems in relating simulation results to broader CSAP objectives. Although the simulation could provide information concerning signals emitted and detected, there appeared to be no practical way to relate these results to the higher level objectives of mission outcome or combat effectiveness. Therefore, to assess the practical significance of the signal detections, it appeared to be necessary to expand the simulation to include human decision processes, so that the effects of combat outcomes could be observed. But such an introduction of human players into the simulation would inevitably make the simulation slow, costly, and nonreproducible. As a consequence, consideration was being given to the possibility of automating human decision processes within the simulation.

2.1 INITIAL STUDY RECOMMENDATIONS

The DSA study effort began with an initial exploratory phase which was followed by the design and development effort. Following the exploratory study DSA made the following general observations and recommendations:

1. That the gap between the specific simulation results and the higher level objectives (mission outcome) could be most effectively bridged by introducing an intermediate measure of performance--which could be related downward to the simulation results, and upward to the higher level objectives.
2. That the "situation perception" available to the opposing military commanders could provide such an intermediate performance measure. (Obviously, information warfare options that are successful in providing the commander an accurate situation assessment, and which can simultaneously deny an accurate assessment to the opponent, are generally more likely to lead to favorable mission outcomes).

3. That to efficiently use such a criterion it would be necessary to simulate the way multiple signal clues and intelligence sources can be combined by intelligence analysts to provide an overall situation assessment.
4. That such a capability to simulate human deductive processes would also be required to provide a computerized foundation for modeling human decision processes (if it were later decided to proceed with a full automation of the player decision processes).
5. That the remainder of the DSA study, therefore, should focus on the development of a conceptual design for a computerized system capable of simulating the kinds of deductive logic normally used by intelligence officers to convert an ensemble of clues into a "situation perception."

These general recommendations (which were presented at the Phase I project review in August 1980) were accepted by the NAVELEX project director, and DSA's remaining effort on the project was directed as recommended above.

2.3 THE RESEARCH APPROACH

DSA's research plan for the design effort included three major elements:

1. Development of a benchmark ocean surveillance scenario.
2. Theoretical analysis of key problems in integrating sensor evidence into a situation assessment.
3. Development of a conceptual design for the information fusion process.

The purposes of the benchmark scenario were:

1. To make it easier to relate theoretical issues in mathematical inference to the practical problems of naval tactics and intelligence; and
2. To ensure that the resulting design concept would be compatible both with realistic cover and deception tactics, and with the practical problems of integrating multiple information sources within a realistic ocean surveillance scenario.

The scenario was structured to include most of the major types of information sources and deductive logic that would be encountered in an ocean surveillance system. Because the scenario was not designed as a final product for the study, it is not included in this report. However, the scenario proved to be extremely important in focusing theoretical attention on the problems of making valid deductive inferences in the context of deceptive tactics.

Whereas our initial theoretical analysis had focused primarily on the problems of computational feasibility, the emphasis of the scenario on tactics and deception raised very serious issues concerning the practical validity and usefulness of Bayesian methods in such a context. This in turn led to the consideration of more conservative methods of inference that promised to be less vulnerable to deception.

3.0 OPERATING CONTEXT FOR INFORMATION FUSION SYSTEM

Ocean surveillance information fusion is concerned primarily with the integration of multisensor information into an overall situation perception. The integration of multisensor information, as it is accomplished by human intelligence analysts, involves a deductive process in which a series of clues are combined to provide an overall interpretation of the situation.

Before discussing the deductive methods themselves it may be helpful to review some of the types of clues that typically will need to be processed within an ocean surveillance intelligence network. To illustrate the problem, we will consider a hypothetical series of observations such as might need to be processed by a Soviet ocean surveillance system in a limited war scenario.

- Following notification by intelligence sources that a U.S. Task Force is underway, satellite reconnaissance observations are received that show a group of ships in a configuration as illustrated in Fig. 3-1a. This satellite observation (which was made at 0400) arrives at the processing center at 0600. From a theoretical perspective, the significant clues in this observation include not only the data on ship location, but also clues concerning the probable size of the ship (or ships) at each location as suggested by the size of blobs.
- At 0800 a Bear D reports a detection of an SPS-52 but (to minimize the risk of being attacked) it avoids moving close enough to sight the ships. Also at 0800, a picket submarine reports sighting a specific U.S. destroyer. Fig. 3-1b illustrates how this 0800 information might look when plotted on the same coordinates as were used in Fig. 3-1a . Notice that

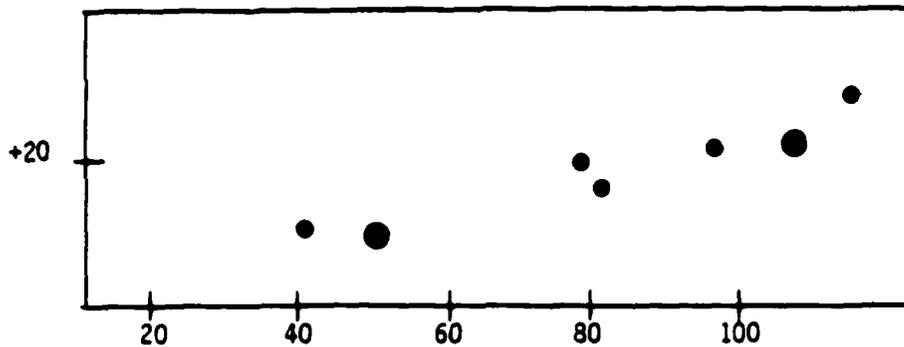


Figure 3-1a. Satellite Observation at 0400 (Received For Processing at 0600). Scale in nautical miles on an arbitrary grid.

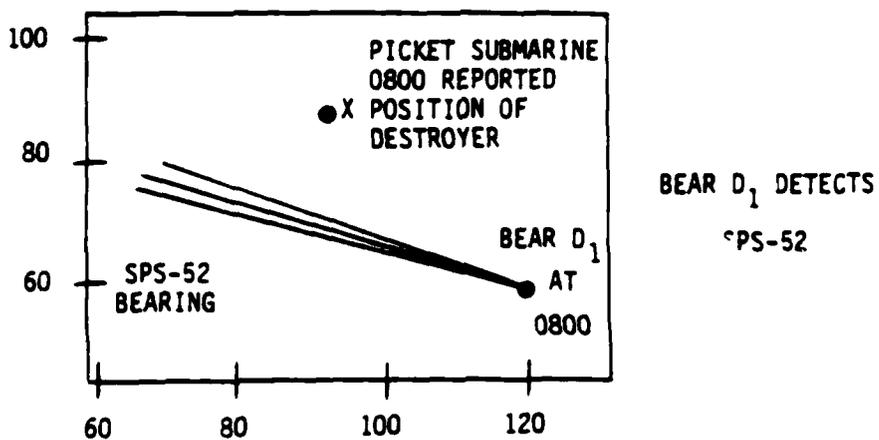


Figure 3-1b. Reports at 0800 Plotted on Same Arbitrary Coordinate Grid

the new information needs to be correlated with the previous satellite observation to provide much useful information.

- Subsequently, a report from an Elint satellite is received indicating six SPS-10s radiating near the location of the satellite observation. This confirms the existence of a U.S. Task Force in the area.
- At 1000 a second satellite report is received based on observations taken at 0800. Fig. 3-2a illustrates the type of information that might be received. Notice the ships seem to be moving to the northwest, but the configuration is not exactly the same, raising some ambiguity concerning the actual correspondence of these ships to those observed previously. By 1100 an Elint satellite report is received which indicates that six SPS-10s and five SPS-40s were radiating in the area at 0900. At 1100 another Bear D report is received indicating an SPS-48 radiating at 1000, at the bearing shown in Fig. 3-2b.

As this type of information accumulates, intelligence analysts try to integrate the various clues with intelligence estimates concerning the ships in the area and the analysts' a priori knowledge of the electronic equipment on various U.S. ships. Their goal is to develop an assessment of the situation in terms of the actual ships (and ship types) involved and the tactical deployment of the ships. Our design objective therefore is to develop computer algorithms that can be used to simulate the mental processes involved in developing such a situation assessment.

3.1 SIMULATION OF INTELLIGENCE PROCESSES

Our objective in the present project was to develop a conceptual design for computer simulation of these mental

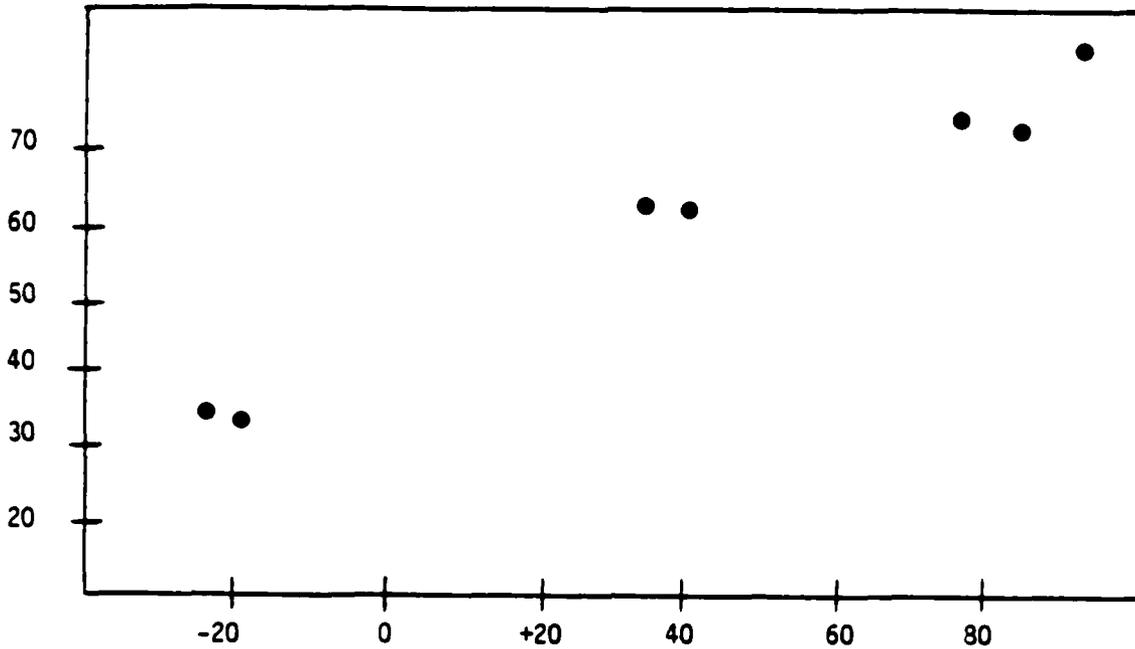


Figure 3-2a. Satellite Observation at 0800 (Received For Processing at 1000). Displayed on same coordinate grid as Fig. 1.

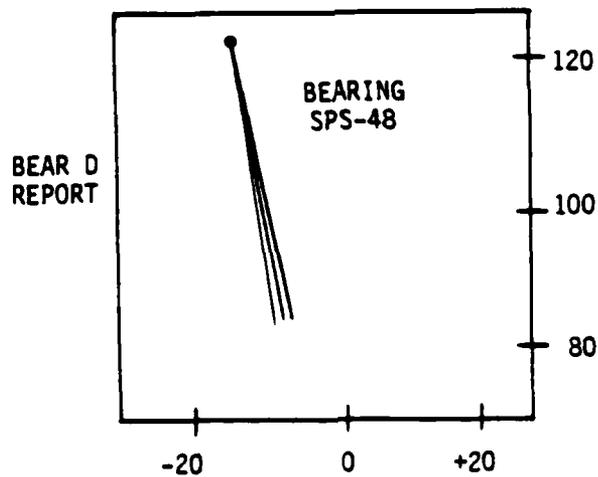


Figure 3-2b. Bear D Observation at 1000 (Received at 1100)

processes which could be incorporated as a part of an existing simulation of the physical processes.

Figure 3-3 illustrates the basic design approach for adding such intelligence processes to an existing ocean surveillance simulation. The lower line in the figure illustrates a standard simulation of the basic physical processes. In such a simulation, the physical status information is updated by a sequence of "events" as the simulation proceeds. In effect, the events operate on an "initial situation" to generate a "new situation." In order to add mental processes to such a simulation, we need to introduce some new types of events (which we will call "information processing" events) to simulate the mental processes. These "information processing" events are generated whenever the physical simulation provides new sensor information or communication messages which can affect the situation perception. The upper part of the figure shows that these mental events operate on the mental "situation perception" in much the same way that the physical events operate on the actual physical situation.

The information represented in the "situation perception" is typically very similar to the information represented in the physical status arrays, except that it tends to be incomplete and inaccurate. The actual quality of information in the situation perception will, of course, depend not only on the adequacy of the sensor data and intelligence information that is provided, but also on the information processing capabilities that are provided to integrate new information into an existing situation perception.

Such a simulation of intelligence processes, in principle, can operate with any desired number of independent "situation perceptions." For example, there could be a separate "situation perception" for each ship or for each command center--or in a

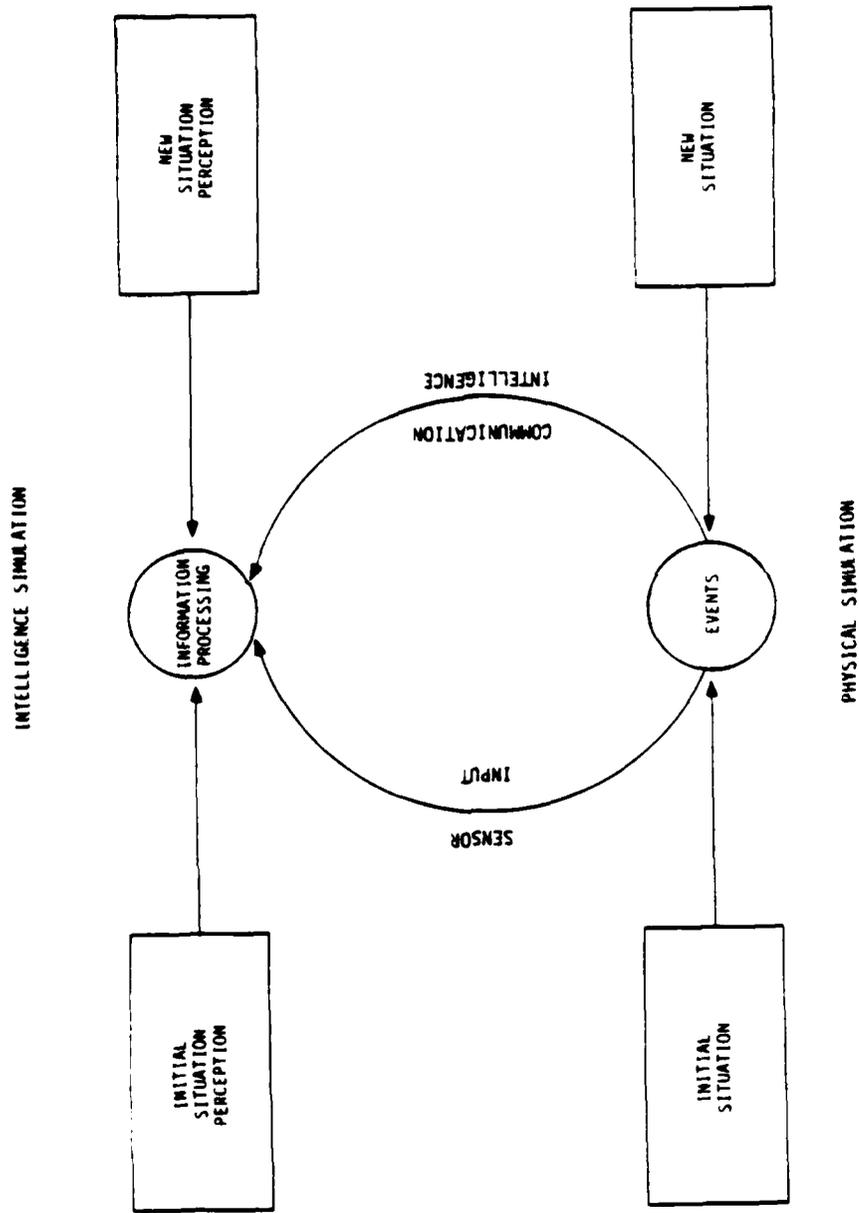


Figure 3-3. Illustrating Relationship of Physical and Intelligence Events in an Intelligence Simulation

simpler simulation there might be only two "situation perceptions," one for the blue forces and one for the red forces.

Although the simulation of mental processes can be treated, in theory, much like the simulation of physical processes, in practice it is usually much more difficult, because the mental processes are not as well understood and they do not obey simple physical laws. Because the actual mental processes are not well understood, it is necessary in practice to replace the actual mental processes by a computerized procedure that is capable of performing the same basic functions, but is sufficiently well structured that it can be defined within a computer program. The development of such a simplified representation of the mental processes is the most critical step in the development of a satisfactory mental simulation.

3.2 INFORMATION PROCESSING FUNCTIONS TO BE SIMULATED

The first step in developing such a representation is to understand in some depth the actual functions that are to be simulated. The real purpose of an ocean surveillance system is to provide an up-to-date understanding of the tactical situation, including the location and identity of all relevant ships in the area. The end product of the information processing procedure should, therefore, provide information analogous to what is contained in a war room situation map. Operationally, this is accomplished by a variety of high level deductive processes through which the totality of clues are combined to provide an overall situation assessment that is as accurate and up-to-date as possible.

To accomplish this objective within a computerized system we will need to identify and simulate these higher level information processing steps that are used to integrate the lower level data into a comprehensive situation perception. Whereas the traditional signal processing activities tend to be concerned with the analysis of individual signals, the higher level

processing operations are concerned primarily with association processes involving multiple signals.

Table 3-1 illustrates the relationship between the traditional signal processing functions and the higher level information processing functions. The upper part of the table lists some of the standard signal processing functions. Such basic signal processing functions are usually included as a part of the physical simulation in a computer wargame, so for our present analysis we will assume that these functions exist, and that they can be used to provide required input data for the higher level information processing functions.

The lower portion of the table lists some of the major information association processes that are required to convert the lower level signals into higher level situation perception. These are the basic deductive processes that are addressed in the present design concept.

3.3 OTHER COMPLEXITIES OF OCEAN SURVEILLANCE INTELLIGENCE

The actual intelligence processes involve some additional complications which are not reflected in Table 3-1.

First, the information processing is accomplished in a distributed command and control system. Much of the input data available at each location consists of information that has been processed (at least in part) at some other location, and this opens the possibility of receiving the same basic information in different but redundant forms from different locations. The distributed processing environment also introduces communication delays within the processing operation.

Although this distributed network has very important practical effects on system performance, it was decided to postpone this complication until after a basic system design had been developed. The present design concept, therefore, has been

TABLE 3-1

RELATIONSHIP BETWEEN BASIC SIGNAL PROCESSING ACTIVITIES
AND HIGHER LEVEL ASSOCIATION PROCESSES

A. BASIC SIGNAL PROCESSING ACTIVITIES FOR INDIVIDUAL SIGNALS

- Signal search (ESM and radar)
- Signal detection (including detection of radar return)
- Signal analysis (develops signal attributes
--frequency, repetition rate, pulse width, modulation
form, azimuth, etc.)
- Signal classification (emitter type, familiar or new,
identification if possible)
- Signal tracking (updates azimuth, range, etc.)

B. HIGHER LEVEL SIGNAL ASSOCIATION PROCESSES INVOLVING
MULTIPLE SIGNALS

- Signal correlation (correlates multiple intercepts of
same emitter)
- Localization correlation (e.g., direction finding,
triangulation, or signal time of arrival differences)
- Recognition, localization, and tracking of platform sets
- Association of new emitters with known platform or
platform sets
- Decomposition of general platform sets into subsets or
individual ships
- Accumulation of historical signal clues for each
platform set
- Assessment of type, number, and identification of ships
that can explain emissions for each platform set
- Cross-correlation of evidence from all platform sets
to provide an integrated situation assessment

developed as if all of the processing were to take place within a single information processing center. Nevertheless, the system includes a capability to deal both with delayed information and with partially processed data. The system design, therefore, can be interpreted either as one processing center within an overall network, or as a single unified processing center. This approach should make it possible to extend the basic design concept to represent an overall intelligence network simply by adding the communication links, and by defining specific processing responsibilities for each specialized processing center.

A second complication involved in a real intelligence system arises because the intelligence functions go well beyond the specific deductive functions shown in Table 3-1, to provide a tactical interpretation of the situation. For example, in a situation where the concrete evidence indicates only the presence of a missile cruiser, an intelligence analyst is likely to conjecture the presence of destroyers and perhaps other typical elements of a naval task force. The intelligence analyst, however, will make a sharp distinction between what he "knows" (for example, that an emitter unique to a particular cruiser has been observed) and what he "thinks" (for example, that cruisers almost always are accompanied by destroyers, and are often deployed with aircraft carriers). This distinction between what he "knows" and what he "thinks" is useful both in communicating his knowledge and in interpreting new evidence.

The present system design effort has focused primarily on the development of information corresponding to what the intelligence analyst would say he "knows." The extension of the deductive process to draw conclusions concerning higher level tactical interpretations of the situation is a logical next step, which would require a new conceptual design effort. The information processing functions shown in Table 3-1 are focused

primarily on relatively concrete interpretations of the data,
they are not concerned with the higher level interpretations of
the tactical situation.

4.0 OVERVIEW OF DESIGN ARCHITECTURE

The first step in the development of algorithms that can simulate human analytical processes is to develop an accurate functional understanding of the mental processes that are to be simulated. Our analysis of the analytical processes required to maintain an accurate situation assessment showed that process is quite different than is commonly believed. The following sections discuss some of the major findings that emerged as a result of our critical review of these analytical processes.

4.1 THE ESSENTIAL ROLE OF HISTORICAL TIME

It is widely believed that people can update their understanding of the current situation directly, simply by correcting it to take into account new information. However, when we tried to develop system designs based on this simple intuitive concept we found it was impossible to develop a logically consistent system. In order to provide satisfactory system performance we found it was essential to do much of the analysis within the context of "historical time." A little reflection on the functional role of a "situation perception" in a C^2 system shows why this is necessary.

The "situation perception" required by any C^2 system is action oriented--that is, its function is to support practical decisions. Since decisions can only influence the future, the situation perception is pertinent only insofar as it provides a basis for projecting future outcomes. On the other hand, because of inevitable time delays in the transmission and analysis of sensor data, the available situation perception is necessarily based on information that was collected in the past. The moving line that divides the past from the future, therefore, plays a very important role in the deductive processes of any C^2 system.

Although the goal of the system is to provide the best possible understanding of the present situation, this

understanding is inevitably based on an analysis or interpretation of past events. The C2 system, therefore, is continually involved in the refinement of its interpretation of events in the recent past. As new evidence concerning past events becomes available, it can sometimes require major reinterpretations of the earlier events.

Figure 4-1 illustrates this general concept with a single simplified example. The figure plots historical time from left to right and shows a boundary "now" which separates the past from the future. All available information is derived from events that happened in the past, whereas decisions can only influence the future. In order to develop an accurate assessment of the present situation, and to project possible future outcomes, it is necessary to develop a satisfactory interpretation of events that happened in the past.

Because of the unavoidable limitations of a two-dimensional figure, the plot in Fig. 4-1 shows ship positions plotted one-dimensionally, on the vertical axis. At the time illustrated, three sets of ship tracks (A, D, and E) appear to be of current interest. However, the identity of the ships associated with these tracks depends on historical information. The upper track which was most recently radiating an SPS-10 was previously observed radiating an SPS-43. Thus it seems reasonable to conclude that the ship or ships associated with this track are equipped with both an SPS-43 and an SPS-10.

Tracks D and E, on the other hand, pose a more difficult problem. Depending on how we interpret the ambiguous intersection of the tracks we can reach quite different conclusions. If we were to assume a simple crossing of the tracks then we might conclude: that track E includes an SPS-43 and an SPS-52; and that track D includes an SPS-48. On the other hand, if we consider the possibility of a deceptive maneuver it is possible that track E and track C are associated with the same

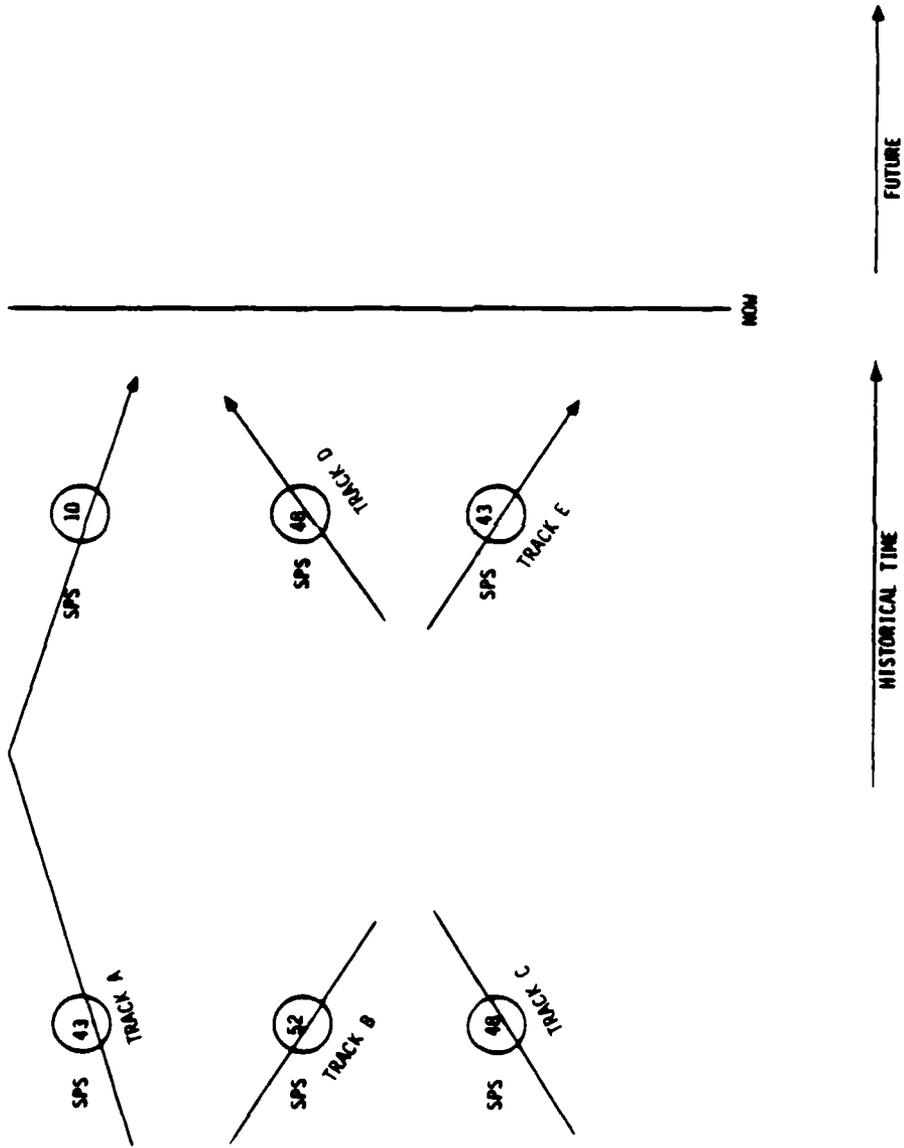


Figure 4-1. Illustrating the Essential Role of Historical Time

ships and that tracks D and B contain the same ships. In this case we could conclude: that track E includes both an SPS-43 and an SPS-48; and that track D includes both an SPS-48 and an SPS-52. Thus, the appropriate interpretation of the current situation depends critically on the interpretation of a past event. If new information were to become available which changed the interpretation of this historical event, it would necessarily change our interpretation of the present situation.

In terms of the design of a deductive system this observation is very important, because it means that the deductive system must maintain records of historical events which might need to be reinterpreted. If such records are not maintained there is no way that new information concerning those events can be utilized. From these observations it is apparent that a deductive system that directly updates the situation perception cannot operate satisfactorily. Instead the system must operate in terms of an interpretation of past events.

The basic operation of the system, therefore, must be concerned with the continuing refinement of an interpretation of recent past events, and with the updating of this historical interpretation to include more recent events as they occur. As events recede into the past, they become progressively less pertinent to current activities. They are less likely to be updated by new information, and they are less likely to be pertinent to the projection of future outcomes. Thus the interpretation of the older historical events can generally be relegated to a long-term memory archive.

Figure 4-2 illustrates the functional differences between the two methods of situation update. The upper diagram illustrates the approach which directly updates the situation perception. This approach precludes any reinterpretation of historical events and thus tends to "lock in" erroneous conclusions. The lower diagram illustrates the present design

concept in which historical events are available for reinterpretation as new information becomes available.

This analytical approach breaks the deductive process naturally into two components. An historical correlator which is responsible for maintaining and updating the historical interpretation, and a second component that can be utilized at any time to produce an up-to-date situation assessment. In the present system design this second component is identified as a Ship Identity Inference System because one of its most significant functions concerns the evaluation of ship identity.

4.2 SEPARATION OF TRACKING AND IDENTIFICATION FUNCTIONS

A second important finding of our critical review of the human analytical processes concerns the possibility of achieving a logical separation of the tracking and identification functions. As Fig. 4-1 illustrates, the clues that are pertinent to ship identification are typically accumulated over a period of time. In order to make the best possible evaluation of ship identity we need to take into account all of the available clues. Moreover, new information concerning old events can produce important changes in the way we associate identification "clues" with tracks.

This suggests that the analysis of ship identity can be logically separated from the processes involved in the correlation of clues with tracks. In the present system design, the historical correlator carries out all of the historical tracking and correlation analysis that is required to recognize ship tracks, to project ship tracks into the future, and to decide which identification "clues" are associated with which tracks. The ship identification inference system processes the accumulated identification clues for each track to develop a logically consistent estimate of ship identity.

The two systems are joined by a data interface which tabulates all of the "clues" associated with each track and provides, for each clue, an estimate of the probability that the clue is actually associated with the specific track.

The discovery that the deductive processes could be separated in this way into two relatively simple processors proved it to be a critical step in the development of a practical methodology. The basic division of responsibility between the two processors is summarized in the table below:

1. The Historical Correlator:

- Associates signals with emitters
- Recognizes, localizes, and tracks platform sets
- Correlates signals with platform sets
- Accumulates historical clues for each platform set

2. The Ship Identity Inference System:

- Determines combinations of hulls required to "explain" emissions from each platform set
- Combines clues from all platform sets to provide an integrated situation assessment

4.3 IMPACT OF TACTICS AND DECEPTION

A third important conclusion derived from our review of the human analytical processes concerns the fundamental importance of tactics and deception in the choice of an appropriate deductive logic. Most scientific applications of mathematical inference are concerned with situations in which the a priori probability of alternatives can be objectively evaluated. This is definitely not the case when one faces an active opponent who can select tactics from an astronomical set of possibilities. This observation, which is of fundamental importance in the design of the information fusion system, is worth considering in some

detail since the problem is rarely, if ever, addressed in text books dealing with the problems of mathematical inference.

4.3.1 The Effect of Tactical Correlations

Although any problem in mathematical inference can, in theory, be viewed as a problem in Bayesian inference, the tremendous variety of possible tactics (combined with the deliberate use of deceptive tactics) makes it totally infeasible to provide the a priori probabilities that are required for formal Bayesian inference.

To illustrate this point, consider the problem of ship identification in the context of deceptive maneuvers and deceptive electronic warfare tactics. Because ship maneuvers and electronic emissions are likely to be deliberately coordinated with deceptive intent, one cannot subdivide the problem into independent problems dealing with individual ships. Moreover, one cannot even define standardized a priori probabilities that any specific set of emitters on a ship will be turned on, since these probabilities will depend both on the tactic that is being employed and on the relationship to other ships.

Thus to properly apply Bayesian deductive logic one would have to supply a priori probabilities for every conceivable "operational plan"--where an "operational plan" implies a complete specification of the time-dependent position, and time-dependent emission patterns for all relevant ships! Even if we were to consider only ship positions and emission patterns at a given point in time, the number of alternatives would be astronomical. If there were 100 ships and one million ocean locations, the total number of position possibilities would be $(1,000,000)^{100}$ or 10^{600} . Moreover, for each position option (if we assume about 10 emitters per ship that can be either off or on) there are $(2^{10})^{100}$ or about 10^{300} emission alternatives. This yields a total of about 10^{900} possibilities. Obviously,

when we also consider the time dependence of position and emitter status, the number of conceivable "operational plans" is so astronomical that there is no possibility of considering them individually, even with the largest computer.

4.3.2 Nonuniformity of Tactical Probabilities

Of course, the existence of an astronomical number of alternatives does not, in itself, make the Bayesian approach infeasible. In scientific problems (such as the derivation of the Boltzman distribution in thermodynamics) it has sometimes been possible to use Bayesian deductive methods (despite an astronomical number of possible states) by assuming equal a priori probabilities within well defined classes of system states. However, in the ocean surveillance problem this assumption of uniform probabilities is not tenable. Naval tactics do not involve random motions of ships; ship motion is always coordinated and deceptive tactics are frequently employed.

If we were to try to proceed with the Bayesian approach we would probably begin by classifying the totality of conceivable "operational plans" into a number of specific classes, and for each subclass we might assign a different a priori probability. But no matter how detailed we made our set of classifications, we would always find within each subclass a large number of unspecified tactical parameters; and we would have to find some way of assigning a priori probabilities for these unspecified tactical parameters.

The traditional Bayesian approach of using equal a priori probabilities is clearly inapplicable because only a tiny fraction of the "conceivable" operational plans make any tactical sense. Thus any simplification that treats all alternatives equally, without regard to their tactical relevance, is likely to lead us very far astray in our deductive inferences. For example, other things being equal, those alternatives that give away the least real information are likely to be the most probable. But

there is no obvious way to classify the "conceivable" plans either by their tactical relevance, or by the amount of information they give away. As a practical matter, therefore, the standard methods of Bayesian inference cannot be correctly applied, because it is not feasible to provide the required a priori information.

4.3.3 A Game Theory Analogy

The foregoing issues may be generally suggestive of a game theory problem in which each side seeks to maximize its own information while attempting to minimize the information obtainable by the opponent. Thus it seems reasonable to ask whether some form of game theory could be applied in the formulation of the deductive algorithms.

For example, we might assume on a priori grounds that the opponent's operational plans have been selected to minimize our knowledge. In a Bayesian formulation this might be equivalent to assuming a different a priori probability distribution over the possible "operational plans."

On the other hand real information warfare problems do not conform accurately to the basic premises of game theory. In particular, the opponent's real objective is usually unknown, and it changes predictably over time. Thus, as a practical matter, these tactical problems are not manageable using the basic premises of either game theory or Bayesian inference.

4.3.4 Intelligence Methods as a Design Pattern

Nevertheless, military commanders and intelligence analysts have developed a variety of heuristic and common-sense deductive methods, which allow them to address these problems in ways that seem reasonably satisfactory--except for the limitations of human intellectual speed, accuracy, processing capacity, and personal biases.

If computerized systems are to be effective in supporting information warfare decision processes it seems likely that they will need to incorporate some of the same basic simplifications that are traditionally used by military intelligence analysts. The present design concept for both systems follows this basic design philosophy.

4.4 THE DETAILED DESIGN CONCEPTS

Since the two system components, "The Historical Correlator" and "The Ship Identity Inference System" are entirely different in design and function, the designs for the two components are discussed separately in the following annexes. Annex A describes the design concept for the historical correlator, and Annex B describes the design concept for the ship identity inference system.

ANNEX A
CONCEPTUAL DESIGN FOR THE HISTORICAL CORRELATOR

A.1 DESIGN GOALS AND OBJECTIVES

The overall design concept for computerized fusion of ocean surveillance data, as outlined in the main paper, includes two major processing components: a historical correlator and a ship identity inference system. This annex develops the design concept of the historical correlator as we anticipate it would be implemented for the CSAP environment. Functionally, the purpose of the historical correlator is to generate and maintain ship tracks and to correlate received signals with the tracks. The logic required to assess the identity of specific ships is discussed separately in Annex B which describes the ship identity inference system.

The output of the historical correlator consists of the following major elements:

1. A list of current ship tracks with an assessment of the position, velocity, and the accuracy of these estimates.
2. For each ship track, a complete tally of all emitters and other identification clues believed to be associated with the track, together with an estimate of the probability of this association.

The tally of identity clues for each track provides the basic input data that is required by the ship identity inference system to provide an up-to-date assessment of ship identity.

Within existing intelligence systems most of the deductive processes outlined above are accomplished by human analysts working at different locations in a command and control network. Because of time delays in both analysis and communication, the

information that is available at any given time is different in the different processing centers; so to that extent, each individual center must operate with its own individual assessment of the situation. Obviously, a satisfactory model of such a network must take into account the time delays, the inaccuracies, and the loss of information that occurs within such a distributed system.

The design of a computerized simulation of the deductive processes in such a large distributed system is a substantial effort. Therefore, although the ultimate objective of the present design effort was to produce a system capable of simulating a distributed system such as the Soviet ocean surveillance system, our development strategy was to focus first on the design of a basic system that could model the generic deductive processes as they might occur within a single processing center.

Such a system could then be used either to model an individual processing center, or to model the operation of an idealized or unified surveillance system--as it might perform if it were not subject to the processing delays and communication limitations of the actual network. Once this basic deductive system is functioning, it should be a fairly straightforward step to link a number of such processing elements together via a communication net to provide a more realistic model of an actual surveillance network. This paper, therefore, is concerned with the design of a basic deductive system that could model a generic information processing center.

The design of this basic system can also be viewed as a two-step process. The first step is concerned with the development of a basic analytical system that is capable of duplicating the essential deductive processes without regard for the time delays and logical errors that are inevitable in human mental processes. After this basic system has been designed, the

second step is to assess the basic system's performance and to determine where its performance needs to be improved or degraded to provide a satisfactory simulation of actual intelligence processes. The specific design concept that is developed in this paper, therefore, is concerned with a generic representation of the analytical capabilities required for processing ocean surveillance data, without regard for the time delays and human inaccuracies that might be required to model the total network, or the quantitative performance of any specific processing center.

A.2 GENERAL OPERATING CONCEPT

The information or "situation perception" required by any C² system must be action oriented. That is, its purpose is to support practical decisions. Since decisions can only influence the future, the situation perception is pertinent only insofar as it provides a basis for projecting future outcomes. On the other hand, the situation perception is necessarily based on information that was collected in the past. The moving line that divides the past from the future, therefore, plays a very important role in the deductive processes of any C² system.

The C² system is continually involved in the refinement of its interpretation of events in the recent past. Moreover, because of time delays in communication, newly received information may provide new evidence that is pertinent to an earlier time--and this "new" evidence must be incorporated to update the interpretation of events which occurred at the earlier historical time.

The general operation of the deductive system, therefore, is concerned with the continuing refinement of an interpretation of recent past events; and with the updating of this historical understanding to include more recent events as they occur. As events recede into the past, they become progressively less pertinent to current activities. They are less likely to be

updated by new information, and they are less likely to be pertinent to the projection of future outcomes. Thus, the older historical events can generally be relegated to a long-term memory archive.

The present system design is concerned with the updating of an interpretation of recent historical events. Any time a decision is needed, this interpretation of recent history can be used either to project an assessment of the current situation (situation perception) or to project possible future outcomes for alternative courses of action.

At all times, the system maintains a store of "knowledge," which reflects its interpretation of recent history. As new information becomes available, it is incorporated to refine, revise, and update this store of historical knowledge. The design for the deductive system, therefore, provides a data structure for representing the pertinent store of knowledge; and it provides formal procedures for updating the knowledge as new information becomes available. Most of this annex is concerned with the design of an approach for representing historical knowledge, and for updating that knowledge as new information becomes available.

A.3 GENERAL CONCEPTS FOR REPRESENTING AND UPDATING HISTORICAL KNOWLEDGE

Computer techniques for representing events are rather well developed as a result of experience with a wide variety of computer simulations. In an event-sequenced simulation, the state of the simulation at any time is stored in a "status information" file which contains the essential data concerning the location and current activity of all objects within the simulation. Any changes in this status information which take place during the simulation are characterized as "events." Within the simulation specific "event processors" are provided which calculate the outcome of interactions between objects in the simulation and accomplish the necessary changes in the status

information. Such a simulation can, therefore, be viewed as a time sequence of historical states for (each object in the simulation) in which the states are linked by the historical events that are associated with the change of state.

The present design uses the same basic concept to represent historical knowledge. The historical knowledge concerning each object (as it is interpreted within the deductive process) is represented by a time sequence of historical states that are linked by historical events. As the deductive process proceeds, the interpretation of any specific historical occurrence can change, and appropriate changes must then be made in the chain of "historical states" and "historical events" that represent that particular segment of history.

The deductive process itself proceeds much like an event-sequenced simulation. As new information arrives from sensors or communication links, the information triggers a sequence of deductive processing events which produce changes in the "knowledge" base. Specifically, these "deductive events" produce changes in the interpretation and representation of historical events. Thus, in discussing the simulation, it is necessary to distinguish between the deductive events which take place in real time in the simulation and the historical events as they are interpreted and updated within the simulation. For simplicity, we will usually use the word "event" to refer to real time events such as sensor events, communication events, and deductive events that occur in real time in the simulation. The interpretation of history (or historical knowledge) that is updated will be referred to either as an "historical event" or an "historical state."

A.4 ENTITIES REPRESENTED IN HISTORICAL KNOWLEDGE

Before considering a specific structure for representing historical knowledge, it is appropriate to ask what types of entities will need to be represented within such a deductive

interpretation of history. From the point of view of the command and control system, the primary concern is to have accurate information on the location and activities of potentially hostile naval forces. As a practical matter, however, this means that the system will need to consider in its deductive processes all ships that potentially could be confused with such naval forces.

Most, if not all, of such information that is obtained by a C² system is derived ultimately from sensor data. The sensor may be a radar, a passive electronic sensor, an electronic intelligence intercept system, or a human agent whose eyes or ears serve as the sensor. To identify the kinds of entities required within the simulation, it is useful to follow a typical chain of logic that leads from the sensor data to an estimate of ship activity, position, and identity.

A typical chain of logic involves the following sequence of objects:

- Sensor
- Signal
- Emitter
- Platform
- Ship

The actual chain of logic is roughly as follows: a "sensor" (which can be either active or passive) receives a "signal" that seems potentially significant. This leads to a conjecture that some kind of "emitter" is present. The emitter could be a radio transmitter, an active radar, the reflecting hull of a ship that returns a radar echo, or a warm hull that is detected by an infrared sensor. The nature of the received signal provides evidence concerning the type of emitter and the characteristics of the emitter. It also provides some evidence concerning the location of the emitter--but the evidence is likely to be incomplete and subject to some probable error. If multiple signals can be combined in the deductive process, the error and uncertainty in the localization of the emitter can be reduced.

For any given localization of observed "emitters," the next logical conjecture is that the emitters are associated with "platforms." Typically, the platforms will be ships, submarines, or aircraft. In some cases, it may not be clear whether two emitters are on the same platform or on different platforms. Where one or more emitters are observed, one can be reasonably sure that there are one or more platforms, but often one cannot be sure of the number of platforms involved. Thus, the observation of emitters will generally lead to a conjecture about the presence of a "platform set," in which the actual number of platforms may be uncertain.

Finally, as a result of a priori knowledge concerning the electronic order of battle for the ships of interest and prior knowledge about the probable location of such ships, it may then be possible to draw conclusions or inferences about the identity of specific ships in the platform set by exploiting the information about the emitters which have been correlated with the platform set.

Thus, a rather standardized logical chain of deductions leads from the sensor observations to the conclusions or inferences about the location of ships. In order to duplicate this kind of logic within an automated system, it seems clear that the automated system must be able to deal with the same basic classes of entities.

The design for the deductive system, therefore, provides for the following basic kinds of entities which can be included within any interpretation of history.

1. Sensor. Any sensor system whether active or passive; electronic, acoustic, or optical (including human eyes and ears); which intercepts signals emitted by opposing forces.

2. Signal. Any signal received by a sensor which provides information pertinent to the location, identity, or activity, of a platform. A signal is typically associated both with a specific emitter and a specific sensor.
3. Emitter. Any active or passive emitter that produces signals that can be detected by a sensor.
4. Platform set. A set of one or more platforms which are conjectured to be associated with a set of observed emitters.
5. Ship. A ship (military or commercial) of a specific type and class, identified by a specific name or hull number. Major properties of specific ships; such as size, speed, shape, electronic equipment, and acoustic signatures; can be assumed to be known in advance and will be utilized to assist in the deductive processes.

Within a single self-contained deductive system, the deductive process would probably always follow the same basic logical sequence from sensor to ship. However, in a distributed system a significant portion of the information available at any processing center is information that was collected and analyzed at another location. This information can enter such a system via communication links, and such communicated information, also, can deal with any of the entity types. It can introduce new information about signals, emitters, platform sets, ships, or the relationships among these entities.

The ability to send and receive such messages, of course, cannot be fully exercised until a full network of simulated processing centers is placed in operation. However, even in an initial version it is important to be able to respond to information received from an outside source that has already been

processed to some level. This is particularly important in order to demonstrate the ability of the system to deal appropriately with unreliable information, or with information which can be either redundant with its own deductions or in conflict with them. Although the repertoire of capabilities to deal with such messages necessarily will be limited in the initial system design, some illustrative capabilities should be provided that can be expanded as the system becomes more complete.

A.5 ASSOCIATION LINKAGES FOR KNOWLEDGE REPRESENTATION

The human ability to carry out rapid analysis and updating of knowledge depends on a complex set of memory associations which permit rapid recall of pertinent related information. The updating and analysis of the historical knowledge within a computerized deductive system presents a very similar problem, and requires a similar solution in the form of a network of linkages which will permit efficient retrieval of relevant stored information.

The basic organizational concept for the system is as follows. Each entity in the simulation has its own data file which contains the specific attributes, or descriptors, required to describe that specific entity. In addition, each entity also has a standardized skeletal record which contains all of the association linkages required to retrieve related information. Most of the required information retrieval processes can then be accomplished quite mechanically by processing the linkage records without regard to the actual data files--except where the information in the data files is actually needed for the analysis. This approach also makes it possible to postpone detailed decisions about the specific attributes to be included in each data file until the specific deductive algorithms are defined.

The knowledge representation provides for two basic types of linkage--relationship linkages and time-sequence linkages.

The relationship linkages reflect the interactions (and relationships) between historical entities that pertain at any specific point in historical time. The time-sequence linkages make it possible to trace the sequence of states (and events) for any historical entity either forward or backward in historical time.

The ensemble of these linkages provides a basic skeleton of associations which allows the system to efficiently locate, retrieve, and update any information pertinent to a specific deductive process.

A.5.1 TIME-SEQUENCE LINKAGES

During the period of time that an entity is under observation, one can expect more or less continuous changes in the estimated localization of the entity (i.e., in the estimated position and velocity and the estimated accuracy of this information). Moreover, during the period of observation the localization of the entity will typically be derived from a sequence of sensor observations. These two aspects of the history of an entity are provided for in two basic linked lists: "an observation list," which points to the basic sensor observations or signals; and a localization history, which contains the interpretation of these observations in terms of a time-dependent localization history for the object of interest.

An adequate representation of history, however, will require more than a simple history for each entity. The entities may occasionally undergo transformations in which their identity seems to change. For example, a platform set may be resolved into one or more individual platforms or platform subsets--or an emitter may be turned off and its identity may be in doubt when it is turned on again. Such transformations in the probable identity of an entity cannot be represented in the form of position or observation data for an individual entity.

To provide for such identity transformations within the simulation, it is necessary to give the new entity (or entities) a new name and to provide it with a new data file. The time-sequence relationships between such entities are represented through the use of predecessor or successor links which make it possible to identify predecessors or successors. In some cases, there can be ambiguity about which entity is the real predecessor or the real successor. Consequently, each such link contains a provision for a probability figure which reflects the estimated probability that the link points to a true predecessor or a true successor.

In addition to the predecessor and successor linkages, there is a potential need for cohort linkages to link between records that may be alternative representations of the same object. Although this concept may not be implemented in early versions of the system, a provision for it as a growth potential seems advisable.

A.5.2 RELATIONSHIP LINKAGES

The relationship linkages make it possible to trace the relationships and interactions that exist within the system at any point in historical time. Fig. A-1 shows the basic structure of relationships that are represented in the system.

The links reflect the system's current assessment of the interactions between sensors, signals, emitters, and platform sets. As noted, a single sensor (such as S_1 or S_3) can be simultaneously tracking several signals. Moreover the emissions of a single emitter can be simultaneously detected as signals by several different sensors (as illustrated for emitters E_2 and E_4). Although by definition each signal originates with a unique emitter (or source), the actual source can sometimes be ambiguous, so (as illustrated with K_6) a signal can be associated with some probability with more than one source. Similarly, as illustrated with E_3 in the center of the figure, an emitter may

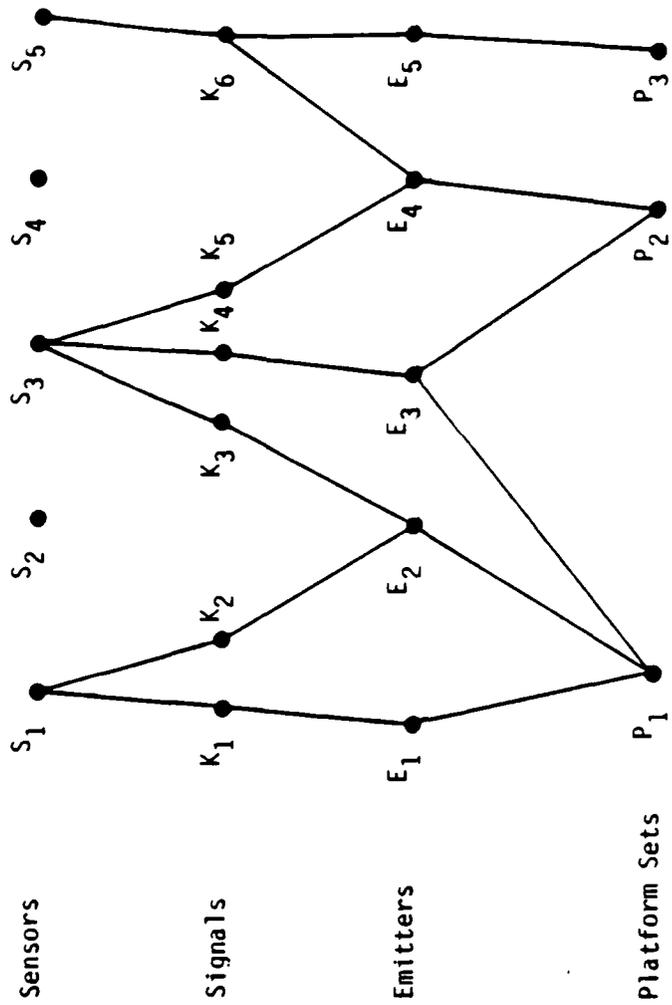


Figure A-1. Illustrating Relationship Linkages in Knowledge Representation

not be localized well enough to uniquely associate it with any specific platforms (or platform set) so it can be simultaneously associated with more than one platform set with some probability. Like all the other links in the system, the interaction links are two-directional links that allow processing to proceed both upward and downward in the linkage network.

In the initial implementation of the system, many of the potential linkages will not be activated; for example, the time-sequence relationships for sensors and signals may not be needed. However, the framework is designed to be general enough to have a great deal of growth potential.

A.5.3 DATA MANAGEMENT UTILITY FUNCTIONS

To facilitate updating operations within this knowledge structure, a number of utility subroutines will be provided that can accomplish routine functions such as: update the status of specific entities; add or delete entities; update relationships between entities; and subdivide or combine entities. The extent to which such capabilities will exist or will have to be developed will, of course, depend on the specific data base system that is selected for the project.

A.6 ILLUSTRATIVE DATA STRUCTURES

The detailed data structures that will be used in the final system will in all probability need to be modified to fit within the framework of whatever data base system is to be used for the model. The data structure described here, therefore, is not intended to be final or definitive. Its purpose is to clarify the functional operation of the system by providing an illustrative implementation in terms of simple data records and linkage pointers. For any specific data base management system, the detailed relationships might be different, but conversion of the illustrative data structure to achieve the same functional operation should not be difficult in any appropriate system.

The representation of historical knowledge in the illustrative data structure is contained in a series of lists that are related through linkage associations. Table A-1 illustrates the data content that will be required in the various types of data records. Since we are concerned at this time with a developmental system, computer efficiency in data storage has been sacrificed to some degree for developmental efficiency. This is reflected in part in the inclusion of reserved data space to deal with presently unanticipated data requirements. It is also reflected in the inclusion of a redundant data classification code which should serve as a useful diagnostic to detect any errors in list processing logic. The actual position of data words in the records shown in Table A-1 is not intended to be significant.

To facilitate the programming of the linkage structure, it is helpful to minimize the number of lists required. This is accomplished by combining the records shown in Table A-2 into a set of only three different lists, which are illustrated in Table A-3. Notice that each basic list uses exactly the same basic data processing structure. The first word contains a linkage and a code, which is followed by a list of N data words. From a data processing perspective, the length of this list is the only feature which varies from list to list, so the same list processing logic can be used for all lists. In each case, the end of a list is signaled by a linkage entry which is either zero or negative. (The negative data link in the entity linkage record is used to maintain this convention.)

A.7 DEDUCTIVE METHODOLOGY

A.7.1 GENERAL DEDUCTIVE APPROACH

The deductive logic utilized within the model differs from formal Bayesian inference in two ways. First, where Bayesian type methods are appropriate, they are almost always replaced by mathematical approximations which avoid the combinatorial

TABLE A-1

DATA RECORDS AND LINKAGE STRUCTURE

1. Entity Data Record
(To be defined individually for each entity type)
2. Entity Linkage Record
Word
 - 1 Code. Specifies entity type
(sensor, signal, emitter,
platform set, ship)
--Space reserved for
subcodes--
 - 2 Time origin Starting time for entity
 - 3 Time end. End time for entity
 - 4 Data pointer. Points to entity data
record (above)
 - 5 History pointer Points to localization
history
 - 6 Observation pointer Points to localization
observations
 - 7 Successors. Points to list of
successors
 - 8 Predecessors. Points to list of
predecessors
 - 9 Cohorts Points to list of cohorts
 - 10 Up pointer. Points to upward
interactions (Fig. A-1)
 - 11 Down pointer. Points to list of downward
interactions
 - 12 Reserved for future needs
3. Entity History Record
 - 1 Code. Entity type and link type
 - 2 Link. Links to next history
entry
 - 3 Time. Time of localization
 - 4 Localization. Points to localization
record
 - 5 Reserved for future needs

TABLE A-1 (Cont.)

DATA RECORDS AND LINKAGE STRUCTURE

4. Entity Observation Record
 - 1 Code. Entity type and link type
 - 2 Link. Links to next observation entry
 - 3 Time. Time of observation
 - 4 Observation Points to observation
 - 5 Reserved for future needs
5. Genealogy Record (Successor, Predecessor, Cohort)
 - 1 Code. Entity type and link type--(predecessor, successor, cohort)
 - 2 Link. Links to next record in list
 - 3 Pointer Points to entity of interest
 - 4 Probability Probability of genealogy
 - 5 Reserved for future needs
6. Interaction Record
 - 1 Code. Entity type and link type (up or down)
 - 2 Link. Links to next record in list
 - 3 Pointer Points to entity of interest
 - 4 Probability Estimated probability of interaction
 - 5 Reserved for future needs
7. Localization Data
 - 1 $X, Y, \delta_x, \delta_y, \delta_{xy}, Z, \delta_z$. . . Derivative pointer
 - 2 $\dot{X}, \dot{Y}, \dot{\delta}_x, \dot{\delta}_y, \dot{\delta}_{xy}, \dot{Z}, \dot{\delta}_z$. . . Derivative pointer
 - 3 $\ddot{X}, \ddot{Y}, \ddot{\delta}_x, \ddot{\delta}_y, \ddot{\delta}_{xy}, \ddot{Z}, \ddot{\delta}_z$. . . Derivative pointer

Note: A zero pointer implies no higher order derivative data is available.

TABLE A-2

LINKAGE ASSOCIATIONS

A. Time-Sequence Links for Entities

(Sensors, signals, emitters, platform sets, ships)

1. History Files
 - Localization history
 - Observation list
2. Genealogy Lists
 - Predecessors
 - Successors
 - Cohorts

B. Relationship Links

1. Interactions

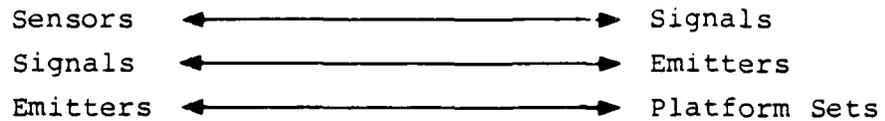


TABLE A-3

LIST STRUCTURES

A. Standardized Data Structure (8 words)

Serves as: Entity data record and localization data

Word

1a Link. Continuation link for longer data

1b Code. Data-type code

2-8 Data

B. Entity Linkage Record

Serves as: Sensors, signals, emitters, and platform sets

Word

1a Link data Points to associate data

1b Code. Entity type

2a Time origin Starting time for entity

2b Time end. End time for entity

3a Observations. Points to observation list

3b History Points to history list

4a Uplink. Points to up associate list

4b Dnlink. Points to down associate list

5a Predecessors. Points to predecessor list

5b Successors. Points to successor list

6a Cohorts Points to abort list

6b Reserved

C. Linkage and History Lists (3 words)

Serves for the following linkage lists: Interactions, clues, ship classifications and genealogy using following format:

TABLE A-3 (Cont.)

LIST STRUCTURES

- 1a Link Links to next list entry
- 1b Code
- 2a Pointer Points to associate
entity
- 2b Probability Associate probability
- 3 Reserved

Serves for observation and history lists using the following format:

- 1a Link Links to next list entry
- 1b Code
- 2a Pointer Points to observation
- 2b Time Or localization record
- 3 Reserved

complexity of the traditional methods of Bayesian inference. Second, to avoid excessive dependence on a priori estimates of enemy strategy, the deductive processes are designed like common sense deductions to be conservative. Instead of attempting to assess the relative probability of alternative possibilities, the deductive processes are designed to distinguish between areas where conclusions can be drawn with considerable confidence and other areas where we have to conclude that we simply do not know.

To implement this basic deductive concept, the model is designed so that it can start from a basic hypothesis that the location of all platforms is unknown. This baseline hypothesis is then abandoned only when there is positive evidence to the contrary.

Within the model this approach is implemented by starting with an a priori assumption that all platforms are included in a single large platform set which is labeled "unknown" and is assumed to be located in some unobserved or unobservable part of the ocean.

When specific signals are detected coming from observed parts of the ocean, it of course becomes most unlikely that all platforms are in the unobserved platform set "unknown." Consequently, on the basis of positive evidence that is inconsistent with the previous hypothesis, new platform sets are introduced. This basic conceptual approach to the deductive process is used in many different applications within the model. In its generalized form it can be viewed as a principle of "positive separation." No new platform sets are introduced unless signals are received that give positive evidence that some ships are separated from the existing platform sets. Similarly, no new emitters are introduced unless there is positive evidence of signals that cannot be explained by the emitters already postulated. This principle of positive separation provides an attractive way of approaching a conservative deductive process

because it does not rely on any a priori knowledge of enemy tactics.

A.7.2 DEDUCTIVE PROCESSING STEPS

The deductive processing logic can be viewed as a combination of two basic processes:

1. Tracking. A tracking algorithm is used to track signals, emitters, and platform sets. Some research will be required to select an appropriate algorithm, since the need for retrospective historical tracking and for conservatism to minimize vulnerability to deception impose design requirements that are somewhat different than in conventional tracking applications.
2. Correlation methods. A correlation procedure is used to decide when multiple tracks can be associated with a common platform source. This procedure is used to generate higher level constructs--such as emitters associated with signals; or platform sets associated with emitters. It is also used to decide when a platform set can be resolved into subsets or individual platforms. The end product of these two steps is a list of platform sets and a list of emitters which (with some probability) are assumed to be associated with each platform set.

The following sections provide an introduction to each of these deductive components.

A.7.3 THE TRACKING ALGORITHM

The use of a standardized representation of the signal observations and the localization history for each of the entities in the simulation should make it possible to use a single tracking algorithm for most of the entities that need to be tracked. The key entities which require tracking include signals, emitters, and platform sets. For each such entity,

there is a time-ordered list of observations which provide the raw data from which the localization history must be calculated. In order to convert these observations into an estimated localization history, two things are required: a suitable tracking algorithm and a representation of the laws of motion that characterize the dynamical behavior of the entity that is being tracked.

The laws of motion appropriate for an ocean surveillance system should probably include at least two separate sets of dynamics, one for ships and one for aircraft. In both cases, the motion is characterized by a tendency to travel a fairly straight course but with the capability to maneuver sharply at unpredictable intervals. The aircraft dynamics of course permits motion in three dimensions, whereas the ships are limited to the two-dimensional surface of the ocean.

The development of a suitable tracking subroutine will have to be viewed as a specific research task because of a number of special characteristics of the present tracking problem. Whereas most tracking filters are concerned only with providing an estimate of the current localization of each track, the present tracker must also provide an efficient retrospective estimate of position and position uncertainty as a function of time. This historical position data for each object tracked is required in order to permit an accurate assessment of the association of signals with specific emitters and platforms. Thus each time an additional observation becomes available, the tracking algorithm will have to take the observation into account, not only as it affects the estimate of current position, but also as it affects the estimate of previous historical positions.

Based on experience in tracking such objects we are confident that satisfactory performance can be obtained using some variant of a Kalman filter with variable noise. However, we believe that other approaches should be explored since it may be

possible to provide comparable or even better performance in this context with a simpler form of tracking algorithm.

For efficiency it would be desirable if the algorithm could be designed so that a new observation would result in an automatic updating of the localization history which extends only as far backward (and forward) in time as is required for reasonable accuracy. Obviously, an observation at a particular point in time will have the greatest effect on position estimates at about that time. Estimates of position at much earlier or later times are likely to be only very slightly affected because these estimates are controlled primarily by other observations taken at the earlier or later times.

Although the design details of the tracking algorithm remain to be developed, they are not required for the design of the rest of the system. From the point of view of the rest of the deductive processes, we can simply assume the existence of a suitable tracking algorithm that can be used to update the localization history for any object when a new observation becomes available.

A.7.4 CORRELATION METHODOLOGY

The basic purpose of the correlation methodology is to provide an association between signals and emitters and between emitters and platform sets. In the relationship between signals and emitters, the correlation methodology is used to decide which signals are to be attributed to the same emitter. At the level of emitters and platform sets, the correlation methodology determines which emitters are to be attributed to the same platform set, and it is also used to decide when a platform set should be decomposed into subsets or individual platforms.

Since in many cases the associations may be ambiguous, the methodology is designed so that it can provide a probabilistic estimate of the required associations. When there are several

reasonably probable associations separate probabilistic links are provided to each of the alternative associations. However, when the association probability falls below a certain critical threshold (probably around 1% or less), the corresponding probability links are dropped from the data base to avoid an undue burden of data storage and computation.

A.7.4.1 The Design Principle

The basic correlation methodology follows a common design principle in each of the cases where it is applied. A measure of merit for each of the plausible associations is computed, and corresponding probabilities of association are estimated that are proportional to the various measures of merit. The measures of merit are designed so that they provide a rough estimate of the probability that a separation as large, or larger (than the apparent separation between sources) would be observed if they were actually a single source. The observed "separation" of course can be in a multidimensional space which might include factors such as frequency, waveform, and repetition rate, as well as estimated position in time and space. For simplicity, in cases where the measure of merit reflects such a multidimensional correlation, the measures of merit can be calculated as a product of simple measures of merit--appropriate to the various dimensions of the correlation. (To avoid generating combined measures of merit that are inevitably smaller than any of the simple measures, the individual measures may be multiplied by a factor of two--so that a product of measures of merit corresponding to the typical 50% probability would yield a correspondingly typical combined measure of merit).

The foregoing methodology provides a conceptually simple and logically consistent way of estimating the probability of association with alternative existing emitters or platform sets. It leaves, however, the problem of how to estimate the probability that is not associated with any existing entity--so that a new emitter or a new platform set must be created.

In a formal Bayesian procedure, one would have to ask how likely it is that a different emitter (or platform set) would be located within the observed separation from an existing emitter (or platform set). As noted earlier, such estimates can be extremely sensitive to one's interpretation of enemy tactics. The probability could be very high if the opponent deliberately chooses to locate emitters and platforms very close together in frequency or geographic space. The probability could be low if he typically avoids such collocation. To avoid having deductive processes that are unduly sensitive to estimated tactics, the design makes use of a principle of positive separation. New emitters and platform sets are not created unless the data is quite inconsistent with the hypothesis that the observations are associated with an existing source. To implement this principle, the correlation system operates as if the a priori probability of a new entity is low. For example, if the a priori probability of a new platform set is taken as 1%, then the probability assigned to a new platform set will tend to be small so long as the platform set is within the 1% probability range of existing platforms. The specific a priori probability used can be viewed as a probability threshold for the creation of a new entity. In most cases, system performance should not be very sensitive to the specific threshold selected, and experimental tests can be conducted to determine what range of values gives the most satisfactory performance. Of course, once a new entity has been created, it will tend to increase rapidly in probability if confirming evidence continues to support its existence.

It is worth noting that for entities that have been tracked for a considerable period of time, the appropriate measure of separation will include a product of measures of merit corresponding to different periods of tracking. Thus, if substantial positive separation is observed at any time during the tracking period, a new entity would be created and maintained.

A.7.4.2 Typical Processing Steps

To illustrate the operation of the correlation methodology, it may be useful to discuss the processing steps required for the analysis of a new signal. When a signal is observed by a sensor, it is checked against historical signals observed by the same sensor to determine whether it can be interpreted as a continuation of a previously observed signal. If it can, it is used to update the localization history for that signal. If it cannot be associated with a previous signal, it is entered as a new signal. In an unusual case where it can be reasonably associated with more than one previous signal, a new signal is entered and probabilistic predecessor links are provided to the prior signals.

After the signal records are updated, the next step is to update the emitter records. For a new signal, a search is made of known emitters to determine whether that signal can be associated with any existing emitter. If not, a new emitter is entered. If correlation is found with one or more existing emitters, appropriate probability links are entered. If the calculated association probability is very high (for example, above 90%) with any specific emitter, the signal can be viewed as uniquely associated with the emitter and can be used to aid in the localization of that emitter. As signals are updated, they are systematically used to update the localization of emitters with which they are uniquely associated. In addition, the association probability is updated. If the association probability falls below a critical level, the unique association can be broken and alternative associations either with a new entity or with another existing entity will be defined.

After the emitter records have been updated, the next step is to update the associated platform sets. Essentially, the same set of steps are repeated--first, to associate a new emitter with existing platform sets if possible; second, to update the

platform set localizations where appropriate; and finally, to create new platform sets when the existing ones are not adequate to reflect the available data.

The end product of the tracking and correlation analysis is a time-dependent set of entities with appropriate time-dependent location data and time-dependent association linkages that reflect all of the correlation and tracking results up to any specific simulated time. The output information that is pertinent to the analysis of platform set identity, however, is limited simply to the set of emitters that is estimated with some probability to be associated with each platform set. The system is designed so that this information can be utilized at any time by the ship identity inference system to provide an up-to-date assessment of ship identity.

ANNEX E
CONCEPTUAL DESIGN FOR THE SHIP IDENTITY INFERENCE SYSTEM

B.1 INTRODUCTION

The main paper of this report described a new approach for automatic fusion of multisensor ocean surveillance data. The basic methodology involves two basic components:

1. An historical correlator which provides a retrospective association of the clues with specific current platform sets, and
2. A ship identity inference system, which provides an assessment of what can be known about the identity of ships in each platform set on the basis of the available clues.

This paper describes the design concept for the second system (specifically the ship identity inference system).* The historical background and the intuitive rationale for the approach are developed in Sections B.2 and B.3 of this Annex. Sections B.4 and B.5 discuss the mathematical formation and solution methodology, and the last two sections discuss some remaining theoretical problems.

B.1.1 STRUCTURE OF THE PROBLEM

The basic input data required for the ship identity analysis consists of the tabulation of emitters associated with each platform set that is provided by the historical correlator. This list of emitters includes not only active emitters observed such as radars, communication systems, and noisy propellers, but also it includes passive emitters such as a ship hull that

*DSA's original formulation of this ship identity algorithm, which was reported in DSA Report No. 276, used a slightly different mathematical form based on information theory. This annex supersedes that earlier version.

reflects optical and radar signals or a warm hull that radiates infrared. The totality of such emitters identified by the historical correlator for all platform sets constitutes the set of clues that are available to be analyzed by the ship identity inference system.

The basic purpose of the identity inference system is to decide what can be known about the identity of the ships in each platform set when this totality of available clues is analyzed in the context of available intelligence and background information such as electronic order of battle data. Because of the size of the problem and the importance of coordinated tactics in determining which emitters will be turned on, it is not computationally feasible to approach the problem rigorously using the formal Bayesian methods of mathematical inference.

B.1.2 RELATIONSHIP TO THE ELECTRONIC WARFARE DECISION AID (EWAR)

Several years ago, in connection with the ONR Operational Decision Aids program, DSA developed an electronic warfare decision aid (EWAR)^{1,2} which utilized emitter clues to provide an estimate of ship identity in a computational context much like the present problem. The EWAR system used a heuristic approximation to Bayesian inference, which had been developed to avoid the combinatorial calculations of the formal Bayesian approach. Originally it was hoped that this algorithm could be generalized for application to the counter surveillance problem. It was obvious, however, that the EWAR algorithm could not be used directly because it was limited to certain restrictive scenario assumptions that were not acceptable in the more general ocean surveillance problem. In particular, a satisfactory

¹D. F. Noble, An Emissions Control Decision Aid, DSA Report No. 66, July, 1978.

²D. F. Noble, Development and Evaluation of an Emissions Control Decision Aid, DSA Report No. 246, July, 1980.

counter surveillance algorithm would have to be able to deal with the following complexities that were not considered in the basic EWAR scenario.

- Unknown number of ships in a platform set. In the EWAR scenario it was assumed that radar blips could be resolved to individual ships, so it could be assumed that each platform set contained one and only one ship. The EWAR problem, therefore, was simply to decide what single ships are most compatible with the observed signals from the platform set. The possibility of explaining the signals in terms of two or more ships did not have to be addressed.
- Probabilistic association of emitters. In the EWAR algorithm it was assumed that emitters could always be uniquely associated with a single platform set. Thus, it was not necessary to consider situations in which an emitter could be associated with two or more alternative platform sets. Obviously, such probabilistic association of emitters introduces an additional complication in the deductive process.
- Problem of redundant or unreliable information. When the reality of multiple information processing centers is considered, new analytical problems arise because the information received from other analysis centers could be either redundant or unreliable. For example, processed analytical reports received from another center might incorporate information from sensors that are also being analyzed locally.

In an effort to generalize the EWAR algorithm a number of revised heuristic concepts were developed and tested, but it soon became clear that the old algorithm could not be satisfactorily generalized and that some fundamentally different approach would be required.

B.1.3 SOLUTION METHODOLOGY

The new methodology which is used in the ship identity inference system uses an extension of a principle of mathematical inference which is known as minimum cross-entropy. Although the resulting minimum cross-entropy algorithm is far more formal and systematic than the deductive methods that are actually used by human intelligence analysts, it appears to operate on the same basic deductive principles. The discovery that the required analytical processes could be simulated using a single well structured mathematical algorithm was one of the major research findings which made the overall design concept feasible.

To evaluate the performance of this new cross-entropy approach, a small demonstration version of the algorithm was developed and tested. The performance of the algorithm proved to be excellent. For a wide variety of small problems (including all those addressible by EWAR) it produced results that were numerically indistinguishable from the results that would be produced by Bayesian methods. Deviations from the Bayesian results seemed to occur only in cases where the validity of the assumptions required by the Bayesian analysis were in doubt, and in those cases the algorithm (much like a human intelligence analyst) produced results that were more conservative or less conclusive than the formal Bayesian results.

B.1.4 STATUS OF DESIGN EFFORT

This annex outlines a basic design concept based on the use of the cross-entropy algorithm. A considerable amount of theoretical and experimental work remains to be done to develop a full understanding of the relationship of this cross-entropy approach both to the formal Bayesian methods and to the intuitive methods used by intelligence analysts. In addition, a considerable amount of development effort will be needed to extend the basic design so that it can deal with all of the complexities of ocean surveillance intelligence problems.

Nevertheless, the research to date has established that the approach is computationally feasible, that it produces results which have an excellent correspondence with human probability estimates, and that it provides a basic foundation which appears to be rather easily expanded to include almost all of the required intelligence processes.

B.2 QUALITATIVE OVERVIEW OF DEDUCTIVE METHODOLOGY

B.2.1 DIFFICULTIES WITH BAYESIAN FORMULATIONS

The initial efforts to produce a new algorithm began with the familiar Bayesian formulations of the problem. However, because of the way tactics tend to be correlated among multiple ships and multiple emitters the rigorous Bayesian approach faces a number of apparently insurmountable problems.

If one were to try to approach the problem by Bayesian methods, one would have to explicitly consider every possible combination of ships in every platform set to determine how well each such combination would explain the observed emissions. The resulting likelihoods would then be multiplied by the a priori probabilities, and the resulting posteriori probabilities would then be normalized to 1.0.

There are at least three serious problems with this approach. First, the required computation is astronomical. For example, if we had just fifty ships and fifty platform sets we would have to deal with fifty factorial (or approximately 3×10^{64}) combinations. Second, there is no logical way of assigning reasonable a priori probabilities to the combinations. The assumption that all combinations are equally likely is certainly very wrong since most of the possibilities would not correspond to any conceivable reasonable tactic. Third, even if we could select suitable a priori probabilities, there would be no way to calculate the probability of specific patterns of

electronic emissions, since the choice of which emitters to turn on is an integral part of definition of a specific tactic. These general problems seem so fundamental that there does not appear to be any way of circumventing the difficulty by restructuring the Bayesian formulation.

B.2.2 THE COMMON-SENSE APPROACH TO DEDUCTIVE INFERENCE

When human intelligence analysts are faced with such deductive problems, they use analytical methods that are very different from the combinatorial Bayesian analysis that is now the standard theoretical approach.

The human analyst uses a kind of puzzle solving logic. First, he asks what types of ships at each location could "explain" the observed signals. Second, he considers the problems of global consistency. For example, if a specific ship is required at one location to explain the observations, then one must avoid using the same ship to explain observations at another location.

One other important characteristic of this logic is that it deliberately avoids reaching any significant conclusions that are not required by the available information. For example, if a specific radar such as an SPS-52 is observed, the analyst does not conclude that some specific ship (for example, the Kitty Hawk) is present. Such a conclusion would not be justified by the limited available information. Instead, the analyst will conclude that it could be any of the ships equipped with the SPS-52. Obviously, either of the two conclusions would be consistent with (and would "explain") the observed evidence, but only the more conservative conclusion is considered to be "justified."

If we are to simulate this human deductive process we will need to formalize the reasons why the second conclusion is viewed as a more "logical" or better "justified" deduction. One of the

most important distinguishing features of the second conclusion is that it is less specific than the first. In effect, the preferred deduction appears to be the one that is as nonspecific as possible--consistent with the available evidence.

This suggests a simple approach for the simulation of human deductive logic which formalizes this intuitive concept. A more formal statement of the basic deductive principle is as follows: "the legitimate conclusions that can be drawn from any set of evidence are those that are as nonspecific as possible consistent with the available evidence." The deductive logic used in the ship identity inference system is based on a mathematical formalization of this common-sense concept.

The logic used by human analysts appears to have one other very important feature which is often overlooked in mathematical formulations. The human analyst does not attempt to consider or assess the likelihood of a combinatorial set of possibilities. Evidently, his deductive logic is applied directly to the probability distributions for the individual ships. Thus, in a technical sense it appears that he works directly with the marginal probability distributions, without explicitly considering the joint probabilities.

This approach, of course, provides an enormous saving in the requirements for information storage and calculation. For example, in the case of 50 ships and 50 platform sets, there are only 50 possibilities that must be considered for each ship. If we sum this requirement over the 50 ships we find a total of 2500 possibilities that need to be explicitly considered. This is obviously a tremendous saving relative to 3×10^{64} combinatorial possibilities. As a practical matter it makes the difference between a feasible and a totally infeasible computational problem.

B.2.3 A MATHEMATICAL INTERPRETATION

The first step in developing a mathematical formalization of the common-sense deductive process is to choose a way of representing the "knowledge" that is to be maintained and updated. In keeping with our understanding of the human analytical methods, we have selected a representation of knowledge that includes only the marginal information. This allows us to represent our knowledge of ship identity in the form of a single probability matrix.

Specifically, the estimated probability that a specific ship i is associated with a specific platform set j is represented by the probability P_{ij} . Since each ship must be somewhere, we know that for each ship i , the summation of P_{ij} must be equal to one, when the summation is carried over all possible ship locations or platform sets. If we are to achieve the potential computational efficiency of this approach we will need to define a mathematical approach that allows the deductive calculations themselves to be performed using this same compact representation of the knowledge. In particular, we must avoid algorithms which would require us to explicitly consider the detailed joint probabilities for all combinations of ships.

The second step in developing the algorithm is to formalize mathematically what we mean when we say a deduction "explains," or is "consistent with", the evidence. In our mathematical formulation, this concept is implemented by defining a set of constraints on the derived probability distributions P_{ij} that must be met if the deductions are to be considered consistent with the evidence.

Fortunately, the information provided by the sensor observations can be conveniently represented in the form of such a set of constraints on the probability distribution P_{ij} . For example, if signals from a particular emitter type such as an SPS-52 have been observed coming from a particular platform set j

then we can feel quite sure that the platform set includes at least one SPS-52. This imposes a constraint on the allowable distributions, which requires that the summation of P_{ij} for platform set j over all ships i equipped with an SPS-52 must be at least equal to 1.0 (otherwise the distribution is not consistent with and cannot explain the observed signal from platform set j). Similarly, if radar observations have told us that the platform set includes exactly three ships, and no ships have been detected leaving or entering the group, then we can be quite sure that the platform set still includes three ship hulls and that the summation of P_{ij} over all ships must be equal to three. Thus, the sensor observations provide information which allows us to specify various constraints on the probability distribution P_{ij} .

The third step in developing a mathematical analogue of the human deductive logic is to decide how to represent the requirement that a deduction should be as "nonspecific" as possible consistent with the available evidence. Typically, there are many possible distributions P_{ij} that can satisfy the required constraints. Indeed, for our previous example any probability distribution P_{ij} will satisfy the constraints so long as it includes an expected value of three for the number of ships in platform set j , and an expected value of at least one ship equipped with an SPS-52 radar.

Thus, we need some criterion for selecting from among all the possible distributions which meet the constraints, a single preferred distribution which constitutes the "proper" deduction. The intelligence analyst typically looks for the least specific deduction which will "explain" the evidence. In order to provide a computer simulation of such a deductive process we will need a formal mathematical criterion corresponding to our intuitive concept of "least specific."

In the present system design we use an extension of the principle of minimum cross-entropy to provide the required mathematical definition of "least specific." In principle, the mathematical process operates as follows: first, we consider the set of all possible probability distributions P_{ij} that meet the specified constraints. Second, we select from this set the single unique probability distribution P_{ij} for which the calculated cross-entropy is lowest. By our formal definitions, we have then found the least specific set of assumptions concerning ship identity that is capable of "explaining" the observations.

We can now summarize the basic axioms of the approach as follows:

1. Knowledge is represented only in the form of marginal probability distributions P_{ij} .
2. A valid deduction is one which satisfies two requirements: (a) it is consistent with the evidence (i.e., it satisfies a set of constraints derived from the observations); (b) it is as nonspecific as possible (i.e., it minimizes total cross-entropy subject to meeting the constraints).
3. The total cross-entropy H is defined as the summation of the cross-entropy h_i associated with the marginal probability distribution P_{ij} for each ship i . Specifically:

$$H = \sum_i h_i \quad (2-1)$$

where

$$h_i = \sum_j P_{ij} \ln P_{ij}/P_{ij}^0 \quad (2-2)$$

and P_{ij}^0 represents the assumed a priori probability distribution for each ship among the various platform sets j .

B.2.4 MATHEMATICAL BACKGROUND

Obviously, the selection of cross-entropy as a way of formalizing the intuitive concept of "nonspecific" is somewhat arbitrary when it is interpreted only within the context of the common-sense deductive axioms. For example, any nonlinear function such as:

$$\sum P_{ij}^2 \quad \text{or} \quad \sum P_{ij}^2 \ln P_{ij} \quad (2-3)$$

which would penalize large values of P_{ij} could be selected as a definition of "nonspecific." The particular choice of the cross-entropy function was of course dictated by other theoretical considerations.

The cross-entropy methodology has an appealing relationship to information theory which suggests that it is a particularly appropriate form for the purpose. Moreover, in a wide variety of situations it can be shown that the cross-entropy formulation will produce results identical to Bayesian inference. Finally, there is an established literature which testifies to the general success of the cross-entropy methodology in producing appropriate mathematical inferences.

It is worth noting, however, that the version of cross-entropy which is being used here differs from what is usually used in formal scientific applications in that the entropy principle is being applied directly to the marginal distributions. In scientific applications the entropy concept

has typically been applied to complete and independent states, just as with the Bayesian method.

In the field of operations research it appears that the cross-entropy principle has been used for a number of applications where it has been applied directly to the marginal distributions, but the formal mathematical justification for this approach is unclear.

A recent article by John Shore and Rodney Johnson, "Axiomatic Derivation of the Principle of Maximum Entropy and the Principle of Minimum Cross-Entropy,"¹ appears at first sight to provide a theoretical foundation for the approach. However, an examination of the article reveals that their derivations assume a complete and independent representation of states. Consequently, it appears that at least some of their results are not applicable to the direct use of cross-entropy for marginal distributions.

Thus, considerable fundamental work is needed to develop a satisfactory theoretical understanding of the method.

B.2.5 RELATIONSHIP TO BAYESIAN INFERENCE

Our experimental tests of ship identification algorithms, have shown that for a wide variety of problems the marginal cross-entropy method produces results which are numerically indistinguishable from the Bayesian results. Nevertheless, we have identified certain classes of problems where the two methods produce quite different results. In these cases, the predictions of the cross-entropy algorithm seem to be in better agreement with military judgment than the more familiar Bayesian method. As yet, the precise theoretical nature of the difference has not

¹J. E. Shore and R. W. Johnson, "Axiomatic Derivation of the Principle of Maximum Entropy and Minimum Cross-Entropy," IEEE Transactions on Information Theory, January, 1981.

been established. However, a comparison of the assumptions required and the results produced by the alternative methods suggests that the primary difference is in the degree of conservatism, as opposed to the deductive efficiency provided by the two methods.

In order to use the Bayesian methodology, it is necessary to specify very detailed assumptions concerning the underlying a priori probability distributions (including the degree of correlation between different emitters and different ships). Moreover, it is necessary to formulate a complete and independent set of alternative hypotheses about the state of the system to be analyzed. The Bayesian method of inference treats all such detailed assumptions as immutable facts, and it produces the most efficient possible deductions--given that all of the specified assumptions are precisely true.

Because the Bayesian methodology treats its assumptions as absolutes, the resulting deductions are typically very sensitive to the assumptions. If the assumptions are inaccurate or unjustified then the methodology can produce apparently definitive conclusions that may be entirely unwarranted when the uncertainty in the assumptions is taken into account.

The cross-entropy deductive method, in contrast, does not require the user to supply any more detail in the input assumptions than he intends to utilize in the results. As a consequence, it is not surprising that the calculated deductions seem to be more conservative, because there are fewer a priori assumptions. The cross-entropy deductions, of course, are still dependent on the assumed a priori distributions, but apparently because the number of assumptions is less the degree of sensitivity is far less.

In a tactical context (where a wide variety of deceptive tactics can be expected) a conservative method of deduction, that is relatively insensitive to dubious a priori assumptions seems likely to be far more satisfactory than the traditional Bayesian approach. This expectation seems to be confirmed by our observation that the cross-entropy inferences in our test cases appear to have a better correspondence with military intuition.

As noted earlier, the cross-entropy approach has a decisive advantage over the traditional Bayesian methods with regard to computation time. Because the typical Bayesian calculations scale factorially with the number of ships, the computational burden can reach astronomical levels even for modest problem sizes. The cross-entropy calculation in contrast remains computationally feasible even for very large problems.

Finally, the new formulation of the deductive process as an optimization problem provides a far more flexible theoretical structure that can be easily adapted to account for complexities and subtleties in the problems that are very difficult to include within the Bayesian method.

For all of these reasons we believe that the cross-entropy methodology provides an ideal approach for simulating the complex deductive processes involved in an ocean surveillance network.

B.3 HISTORY OF CROSS-ENTROPY AS A METHOD OF MATHEMATICAL INFERENCE

The historical origin¹ of cross-entropy began in the field of thermodynamics, where entropy was originally defined as a measure of the degree of disorder in a system. The higher the entropy, the more the disorder. Consider a probability

¹For a more detailed review see article by E. T. Jaynes, "Where Do We Stand on Maximum Entropy" in The Maximum Entropy Formalism, Levine & Tribus Editors, The MIT Press, 1979.

distribution for an arbitrary physical system which can have a large but finite number of possible states. If we represent the probability that the system is in any specific state j by a probability P_j then the entropy S , for the distribution is defined as:

$$S = \sum_j P_j \ln P_j \quad (3-1)$$

In 1948, C. E. Shannon published his paper "A Mathematical Theory of Communications"¹ in which entropy was also recognized as a quantitative measure of information content. The higher the entropy, the lower the information content. Thus, the selection of a probability distribution with the highest entropy content would also select the one with the lowest information content. To a first approximation this is what the present algorithm does. That is, for any specific set of evidence, the algorithm selects the least specific probability distribution P_{ij} (i.e., the one with the lowest marginal cross-entropy) which is consistent with the evidence.

The use of the principle of maximum entropy as a method of mathematical inference was originally suggested by E. T. Jaynes² in 1957. Jaynes' original formulation of the principle of maximum entropy using Eq. 3-1 still seems quite satisfactory in situations where there is no ambiguity about how one should define a "system state." However, in many practical applications the "system state" is more conveniently viewed as a continuous variable, and the choice of how the state should be quantized into the discrete states required by the standard definition of

¹C. E. Shannon, "A Mathematical Theory of Communication", Bell Syst. Tech. J., vol. 27, pp. 379-423, 623-656, 1948.

²E. T. Jaynes, "Information Theory and Statistical Mechanics I", Phys. Rev., vol. 106, pp. 620-630, 1957.

entropy used in Eq. 3-1 seems totally arbitrary. Moreover, when this traditional entropy measure is used in a deductive algorithm, the deductions are not independent of the choice of a coordinate system. Obviously, any valid deductive procedure should produce results that are invariant to the choice of a coordinate system. This theoretical deficiency in Jaynes' original entropy method, however, is rather easily corrected by using the "cross-entropy" measure rather than the simple thermodynamic "entropy."

Cross-entropy is defined simply as follows:

$$h = \sum_j P_j \ln (P_j/P_j^0) \quad (3-2)$$

where P_j^0 represents the a priori probability that the system will be found within the specific state j . In this formulation, the calculated value h of the cross-entropy is not changed when a quantized state is subdivided, because the ratio of the posteriori to a priori probability is independent of cell size.

This refinement of the approach, which is now identified as cross-entropy, was originally suggested by S. Kullback¹ in 1959, and was first labeled as cross-entropy by I.J. Good² in 1963. The use of the cross-entropy formulation provides a unique measure of the degree of order of a probability distribution P_j (relative to a known a priori distribution P_j^0) which is independent of the definition of cell size.

¹S. Kullback, Information Theory and Statistics. New York: Wiley, 1959.

²I. J. Good, "Maximum Entropy for Hypothesis Formulation, Especially for Multidimensional Contingency Tables", Annals Math. Stat., vol. 34, pp. 911-934, 1963.

It is this formulation of entropy that is used in the present ship identity inference algorithm. Specifically, an optimization method is used to minimize cross-entropy subject to the specified constraints. In this mathematical formulation, the minimization of cross-entropy guarantees that the inferred ship probability distribution will be as nonspecific as possible within the constraints defined by the observed signals.

In a very recent paper¹ John Shore and Rodney Johnson were able to derive the principle of minimum cross-entropy from more fundamental mathematical axioms. Specifically, they were able to show that among all possible deductive methods that are based on the optimization of a function, the cross-entropy approach is the only one that can avoid logically inconsistent results. Their proof, however, does not establish the existence of any "valid" method of deduction based on function optimization. Moreover, their derivation appears to assume complete and independent states so the relevance of their derivation to the use of cross-entropy on marginal distributions where the states may not be independent is unclear.

The practical value and usefulness of the cross-entropy methodology, both in detailed and marginal form, is not in doubt. The basic principle has been successfully applied to a wide variety of high technology problems that could not be easily addressed by conventional Bayesian methods. These include problems in thermodynamics, electromagnetic frequency analysis (where the observations are limited in time, space, or bandwidth), and problems in earthquake analysis and seismic detection of nuclear detonation. It has also been applied in both marginal and detailed form to a wide variety of decision

¹ Shore and Johnson, Ibid.

problems such as queuing theory, stock market analysis, production line decision-making and many others.

The correspondence which is developed in this paper between common-sense intelligence methods and the marginal form of cross-entropy suggests that the approach has a reasonable axiomatic foundation as a simulation of common-sense deductive processes. What remains unclear is the accuracy of the correspondence to human judgment and the theoretical relationship to more formal methods of mathematical inference.

B.4 MATHEMATICAL FORMULATION OF INFERENCE PROBLEM

B.4.1 INITIAL SIMPLIFICATIONS IN PROBLEM FORMULATION

When the ship identity inference problem is examined in its full detail it includes a number of complicating factors that do not fit neatly into a simple formulation, even using the marginal cross-entropy approach.

Some of the most important complications concern clues which have some elements of uncertainty, such that it seems inappropriate to classify them as constraints, but where it is also obvious that they must somehow be considered in the total system of evidence that determines the solution.

Some examples of this type of problem include:

1. Cases where the classification of an emitter type is uncertain.
2. Cases where the amplitude of the signal can be an indicator of the nature of the emitter, as in the case of radar returns.

3. Cases where the signal is well enough localized to strongly suggest a specific platform set as the point of origin, but where it is not well enough localized to exclude alternative explanations.

Although a number of heuristic approaches have been identified for dealing with these practical complications, it is not yet clear what approach will be most satisfactory. The development and justification of specific procedures for including these practical considerations will require future research that goes beyond the scope of this report.

In the following mathematical development, each of these issues is either simplified out of the problem statement, or treated by an approximation that is obviously imperfect. This approach allows us to develop a simple and mathematically consistent formulation (which is of course both incomplete and imperfect in its correspondence with the actual sea surveillance inference problem). The final sections of the paper discuss these limitations of the simplified theoretical formulation as a way of defining requirements for future research.

B.4.2 MATHEMATICAL FORMULATION OF SIMPLIFIED PROBLEM

The simplified ship identity inference problem can be formulated mathematically as follows. For each ship i we define a probability distribution P_{ij} corresponding to the estimated probability that the ship i is located within any specific platform set j . In addition we define an a priori estimate P_{ij}^0 of these same probabilities. We can then define the total cross-entropy for the distribution as a summation of the cross-entropy over all ships considered. Specifically, this leads to the following mathematical representation of the marginal cross-entropy, H , which is to be minimized subject to certain constraints that remain to be specified. Specifically:

$$H = \sum_i \sum_j P_{ij} \ln (P_{ij}/P_{ij}^0) \quad (4-1)$$

If the distribution P_{ij} is to be compatible with the observed signals, then it must provide at least one emitter of an appropriate type and location to explain each observed signal l . We assume that each such signal has been classified as originating with an emitter of type k with a confidence S_{kl} . Moreover, we assume that electronic order of battle information on all ships of interest is available, so that the number of emitters of type k on each ship i is known and can be represented in an array a_{ik} . We also assume that the localization information for each signal l can be summarized in the form of a set of likelihood estimates X_{jl}^0 that define an a priori estimate of likelihood that the observed signal could have originated at each of the platform sets j . Finally, for each signal l we define a set of parameters X_{jl} which will reflect the estimated posteriori probability that each signal l actually originates from each platform set j . Table B-1 summarizes the foregoing definitions.

Based on the foregoing problem definition we can now formulate a set of constraint relationships which must be satisfied by the final probability distribution P_{ij} . If we are to explain each signal then the summation of X_{jl} and X_{jl}^0 over all platform sets j must be equal to 1.0. Specifically:

$$\sum_j X_{jl}^0 = 1.0 \quad (\text{for all } l) \quad (4-2a)$$

$$\sum_j X_{jl} = 1.0 \quad (\text{for all } l) \quad (4-2b)$$

TABLE B-1

SUMMARY OF DEFINITIONS

P_{ij}^o, P_{ij}

a priori and posteriori probability that ship i is located at platform set j

S_{kl}

confidence that signal l originates with an emitter of type k

X_{jl}^o, X_{jl}

a priori and posteriori probability that signal l originates with platform set j

a_{ik}

number of emitters of type k located on ship i

Moreover, the number of emitters of each type k at each platform set j must be sufficient to account for the estimated values of X_{j1} . In particular:

$$\sum_i P_{ij} a_{ik} \geq \sum_l X_{jk} S_{kl} \quad (4-3)$$

The equation is written with an inequality because an emitter that is present may be turned off, and even if it is turned on it may not be detected as an observed signal l .

One of the important issues that is not yet satisfactorily resolved is how the posteriori values X_{j1} really should be influenced by the a priori likelihood estimates X_{j1}^0 . To avoid unnecessary complications in the present derivation we will simply assume the following very simple relationship:

$$X_{j1} = X_{j1}^0 \quad (\text{a simplifying assumption}) \quad (4-4)$$

In a limiting case when all values of X_{j1} are either zero or 1.0 this assumption is obviously a good one. However, in the case of intermediate values of X_{j1} it must be viewed as a very rough approximation. The implications of this simplification will be discussed after the basic derivation is complete. Using this simplification the basic constraint Eq. 4-3 can be rewritten as follows. For all j, k

$$\sum_i P_{ij} a_{ik} \geq \sum_l X_{j1}^0 S_{kl} \quad (4-5)$$

In order to provide a practical computational algorithm we need to develop an efficient way of solving for values of P_{ij}

that will minimize the cross-entropy H subject to meeting the constraints specified in Eq. 4-5 above.

B.5 SOLUTION METHODOLOGY

Although in principle one can think in terms of a large and comprehensive set of ship locations i and solve the problem in this way, it is more convenient in most practical problems to consider two separate categories of locations: first a small number of specifically identified platform sets, j , for which there is some form of concrete evidence (usually in the form of signals received); and second, another very large group of locations for which there is no specific evidence. We will identify this second large set of locations as a single generalized platform set which we will label as the unknown set u .

Using this revised notation we can rewrite the total cross-entropy Eq. 4-1 as follows:

$$H = \sum_i \left\{ P_{iu} \ln (P_{iu}/P_{iu}^0) + \sum_j P_{ij} \ln (P_{ij}/P_{ij}^0) \right\} \quad (5-1)$$

where we define P_{iu} and P_{ij} as follows:

$$P_{iu} = 1.0 - \sum_j P_{ij} \quad (5-2)$$

and

$$P_{iu}^0 = 1.0 - \sum_j P_{ij}^0 \quad (5-3)$$

In order to generate probability distributions that will minimize cross-entropy subject to specified constraints, it is

important to know how the cross-entropy H as defined in Eqs. 5-1, 5-2, and 5-3 above change when changes are made in the individual probabilities P_{ij} . To determine this dependence we can differentiate H with respect to the individual probabilities P_{ij} . Carrying out the differentiation we obtain:

$$\frac{\partial H}{\partial P_{ij}} = \ln \frac{P_{ij}}{P_{ij}^0} - \ln \frac{P_{iu}}{P_{iu}^0} \quad (5-4)$$

where P_{iu} and P_{iu}^0 are defined as in Eqs. 5-2 and 5-3.

To minimize the cross-entropy H subject to the specified constraints,

$$\sum_i P_{ij} a_{ik} \geq \sum_l X_{jl}^0 S_{kl} \quad (5-5)$$

we form a Lagrangian function L in which each of the above constraint relationships is multiplied by a separate Lagrange multiplier λ_{jk} . Specifically, the Lagrangian takes the form:

$$L = H(P_{ij}) + \sum_{j,k} \lambda_{jk} \left(\sum_i P_{ij} a_{ik} - \sum_l X_{jl}^0 S_{kl} \right) \quad (5-6)$$

For any specified set of Lagrange multipliers λ_{jk} the minimum of this Lagrangian will occur where the derivative relative to all of the probabilities P_{ij} is zero. Differentiating Eq. 5-6 relative to P_{ij} we find:

$$\frac{\partial L}{\partial P_{ij}} = \ln \frac{P_{ij}}{P_{ij}^0} - \ln \frac{P_{iu}}{P_{iu}^0} + \sum_k \lambda_{jk} a_{ik} \quad (5-7)$$

which yields the following conditions for a minimum of the Lagrangian:

$$\frac{P_{ij} P_{iu}^0}{P_{ij}^0 P_{iu}} = \exp \left(\sum_k \lambda_{jk} a_{ik} \right) \quad (5-8)$$

and this can be rewritten:

$$\frac{P_{ij} P_{iu}^0}{P_{ij}^0 P_{iu}} = \prod_k \exp (\lambda_{jk} a_{ik}) \quad (5-9)$$

Fortunately this set of equations can be solved directly to yield specific values of P_{ij} for any specified choice of the multipliers λ_{jk} . For convenience in the solution we define a new set of parameters Z_{ij} :

$$Z_{ij} = \prod_k \exp (\lambda_{jk} a_{ik}) \quad (5-10)$$

We can now rewrite Eq. 5-9 as follows:

$$P_{ij} = P_{ij}^0 Z_{ij} \frac{P_{iu}}{P_{iu}^0} \quad (5-11)$$

In this equation all parameters except P_{iu} are assumed to be known. But from Eq. 5-2 we know that P_{iu} can be expressed as a function of the summation of the P_{ij} over j . This sum can be calculated by summing Eq. 5-11 over j . Carrying out the summation we obtain:

$$\sum_j P_{ij} = \left(\frac{P_{iu}}{P_{iu}^0} \right) \sum_j P_{ij}^0 Z_{ij} \quad (5-12)$$

But from Eq. 5-2 we know:

$$\sum_j P_{ij} = 1.0 - P_{iu} \quad (5-13)$$

Substituting into Eq. 5-12 we obtain:

$$1.0 - P_{iu} = \left(\frac{P_{iu}}{P_{iu}^0} \right) \sum_j P_{ij}^0 Z_{ij} \quad (5-14)$$

Now dividing by P_{iu} on both sides of the equation we obtain:

$$\frac{1.0}{P_{iu}} - 1.0 = \frac{1.0}{P_{iu}^0} \sum_j P_{ij}^0 Z_{ij} \quad (5-15)$$

Now solving for P_{iu} we obtain:

$$P_{iu} = \frac{P_{iu}^0}{P_{iu}^0 + \sum_j P_{ij}^0 Z_{ij}} \quad (5-16)$$

And finally substituting this value for P_{iu} into Eq. 5-11 we obtain:

$$P_{ij} = \frac{P_{ij}^0 Z_{ij}}{P_{iu}^0 + \sum_j P_{ij}^0 Z_{ij}} \quad (5-17)$$

To put this result in a more simple and symmetrical form we can define:

$$z_{iu} = 1.0 \quad (5-18)$$

and then extend the summation over platform sets j to include the unknown platform set u . When this is done we obtain a final simple form for P_{ij} as follows:

$$P_{ij} = \frac{P_{ij}^0 z_{ij}}{\sum_{j=1,u} P_{ij}^0 z_{ij}} \quad (5-19)$$

where

$$z_{ij} = \prod_k (w_{ijk}) \quad (5-20)$$

and

$$w_{ijk} = \exp (\lambda_{jk} a_{ik}) \quad (5-21)$$

In order to obtain a valid solution it is only necessary to find the lowest possible set of nonnegative Lagrange multipliers λ_{jk} such that the constraints defined in Eq. 5-5 will be satisfied. This reduces the problem to a standard problem in lagrange multiplier optimization for which a wide variety of iterative methods are available that can be used to converge to the proper values of the lagrange multipliers.

The Eqs. 5-19, 5-20, and 5-21 have a rather simple intuitive interpretation. It appears that there is a rather fundamental quantity $P_{ij} z_{ij}$ which can be viewed as an

unnormalized probability. To calculate the actual probabilities P_{ij} the first step is to calculate the unnormalized probability $P_{ij}Z_{ij}$ and then normalize the result by dividing the summation of the unnormalized probabilities. The unnormalized probabilities are obtained simply by multiplying the a priori probabilities by a series of multiplicative factors $W_{ijk} \geq 1.0$ (as shown in Eq. 5-20). The presence of the lagrange multipliers λ_{jk} in these factors simply indicates that each of the multiplicative factors is to be independently adjusted to whatever level is required to meet the corresponding constraint. Thus, the basic solution methodology involves a relatively straightforward search for the lowest multiplicative factors $W_{ijk} \geq 1.0$ which will satisfy the specified constraints.

The test program that was developed for the project (which used rather standard interactive methods to converge to the appropriate set of multipliers) was found to produce satisfactory solutions for small problems with a negligible amount of computer time.

B.6 LAGRANGE MULTIPLIERS AS AN INDICATOR OF SIGNAL IMPORTANCE

In information warfare applications it can be very important to know which signals are most important in giving away information. Similarly, it is important to be able to recognize quickly, new information which is in conflict with previous observations or previous interpretations of the data.

The lagrange multipliers in the present formulation provide a very sensitive indicator of the importance or significance of the individual signals. If a new signal is completely consistent with previous data, and thus provides no new information, then the lagrange multiplier for the signal will be zero. If the signal provides new information which substantially reduces the ambiguity in the interpretation of the situation, then the lagrange multiplier for the signal will be quite high. On the other hand, if a signal provides information which is logically

inconsistent with previous evidence, then the lagrange multipliers both for the new signal and for the conflicting previous signals will rise toward a maximum or limiting value which is determined by the confidence level assigned to the signals. Thus, when a lagrange multiplier approaches this confidence limit it can provide a warning that the signal is incompatible with previous data.

The availability of such a quantitative indicator of the importance of each signal is an important by-product of the present formulation which seems likely to enhance the usefulness of the algorithm, particularly in information warfare applications.

B.7 REMAINING ISSUES FOR THEORETICAL ANALYSIS

The simplifications used in the foregoing derivation suggest a number of issues that will require additional theoretical and analytical research. This section discusses some of these specific problems that will have to be addressed to provide the diversity of capabilities required for the analysis of practical sea surveillance and counter sea surveillance problems.

B.7.1 AMBIGUOUS LOCALIZATION OF SIGNAL ORIGIN

The first and most obvious problem concerns our simplifying assumption which treats the a priori probability distribution for the point of origin of a signal as a rigid constraint on the posteriori distributions. It seems clear that it would be better if this could be treated not as a rigid constraint, but rather as a somewhat flexible preference. For example, if the observed signal corresponds to an emitter that is already known (from other evidence) to be present in a particular platform set, it seems unreasonable to impose a constraint that requires a predetermined fraction of the emitter to be accounted for in the probability distributions for other platform sets that theoretically are possible points of origin for the signal. Several alternative approaches have been identified that could

provide a reasonable way of representing this information as a "soft" rather than a rigid constraint. However, considerable work remains to be done to select a preferred approach and to develop a satisfactory theoretical justification for the choice.

B.7.2 AMBIGUOUS IDENTIFICATION OF EMITTER OF ORIGIN

A second very important issue concerns the treatment of suggestive, but not definitive, evidence concerning either the unique identification of individual emitters or the identification of a specific emitter type.

The foregoing mathematical formulation uses "confidence of emitter identification," S_{k1} , as if it is a rigid constraint on the solution. It seems clear that this is not a satisfactory approach. For example, suppose we have a 75% "confidence" that a specific signal was produced by a specific emitter unique to the carrier Kitty Hawk. But, in addition, we have a 95% confidence that another signal coming from another part of the ocean is also associated with an emitter that is unique to the Kitty Hawk. Obviously, it is inappropriate to impose both constraints on the posteriori solution. A consistent and systematic way is needed, therefore, to balance such suggestive evidence concerning the specific emitter (or emitter type) of origin. The present prototype version of the ship identity inference system uses a simple ad hoc method of resolving conflicts between such incompatible clues, but a systematic method is needed for which an appropriate theoretical rationale can be provided.

B.7.3 SUGGESTIVE INFORMATION DERIVED FROM PASSIVE SIGNAL AMPLITUDE

The foregoing mathematical treatment has implicitly assumed active emitters that can be turned off and on as desired. Consequently, the absence of a signal has been treated as if it contains little useful information. Although this assumption is correct for most active emitters, it certainly is not correct for passive emissions such as radar returns, infrared, and optical signals.