

A sequential hypothesis testing, optimal stopping problem in underwater acoustic detection

Harilaos N. Psaraftis and Anastassios N. Perakis^{a)}

Massachusetts Institute of Technology, Cambridge, Massachusetts 02139

(Received 18 March 1983; accepted for publication 17 August 1983)

A sequential hypothesis testing, optimal stopping problem in underwater acoustic detection is formulated and solved using dynamic programming. The problem calls for deciding whether acoustic signals being received over long ranges in the ocean are due to a source or to ambient noise alone, so as to minimize the expected value of a specified cost function over a given time horizon. The cost function incorporates a constant cost per observation as well as terminal costs for false acceptance of either hypothesis. According to previous work by the authors, and without loss of generality, modeling the acoustic signals assumes a two-state discrete-time Markov process for each of the two hypotheses, the state of the process depending on whether the intensity of the signal at the receiver is above a specified threshold or not. The decision process presented is based on observations of the signal's "interarrival times," that is, the time intervals between two successive detection events. The algorithm is then extended into more than two alternative hypotheses (several "false" targets) and results using both simulated and experimental acoustic data for the two and three hypotheses cases are presented. Computational issues in implementing the algorithms as well as possible extensions of this work are finally discussed.

PACS numbers: 43.60.Gk, 43.30.Bp

INTRODUCTION

In previous work (Psaraftis *et al.*, 1981a,b), it was shown that under fully saturated phase-random multipath acoustic fluctuations (for a definition see Hamblen, 1977; Mikhailevsky and Dyer, 1978; Mikhailevsky, 1979, 1980), the underwater acoustic detection process has memory. It was demonstrated theoretically and also verified via a comparison with data, that the occurrence of a detection event at some instant of time strongly influences the probability of another detection event shortly thereafter. Detection events were defined as occurring whenever p , the root mean square pressure at the receiver, exceeded a specified threshold level. In Psaraftis *et al.* (1981b) it was shown that the detection process can be modeled as a discrete-time two-state Markov process (see Fig. 1) being in state U (for "up") for $p > p_0$ and in state D (for "down") for $p < p_0$. The transition probabilities a and b were evaluated as functions of p_0 , of some other statistical parameters of the signal and of the user-calibrated time increment ΔT , which separates two consecutive opportunities for a state transition. Probability mass functions (PMFs) of the time interval between two successive detections (or D-U transitions) and of the time interval between a D-U transition and the first U-D transition that follows, were derived in terms of a and b . The above time intervals will be referred to as "interarrival time" and "holding time," respectively.

The purpose of this paper is to develop an algorithm for deciding whether acoustic fluctuations being received over long ranges in the ocean are due to a source or to noise alone. Specifically, we assume that exactly one of two hypotheses is true: Hypothesis H_1 , according to which the fluctuations being received are coming from a source (plus some back-

ground noise), or hypothesis H_0 , according to which the fluctuations are due to noise alone. We also assume that the detection process is a discrete-time Markov one as outlined above, with known transition probabilities: (a_1, b_1) if H_1 is true and (a_0, b_0) if H_0 is true, and with the same time interval ΔT in both cases. Starting with an *a priori* probability that H_1 is true, we have to decide in at most N observations (stages) about which of H_1 and H_0 is true. At each but the last time stage, we may stop and declare one of H_1, H_0 as being true, or we may continue "observing" the process. In both cases we incur costs. In the former case, we incur costs in case of false acceptance, while in the latter case we incur observation costs. The objective at hand is to make the above decision so as to minimize its expected total costs.

The idea to use sequential algorithms in order to decide on the validity of a particular hypothesis chosen among a set of mutually exclusive hypotheses is not a new one. In Wald (1947), the Sequential Probability Ratio Test (SPRT) was first proposed as a procedure to test a simple hypothesis against a simple alternative. Modifications to the test have appeared in subsequent research, enabling it to deal with a wider class of problems in a satisfactory manner. Several authors have presented modifications to the decision thresholds of the SPRT so as to maintain a reasonable expected sample size when the same SPRT is used in testing a simple hypothesis against a composite alternative (Anderson, 1960).

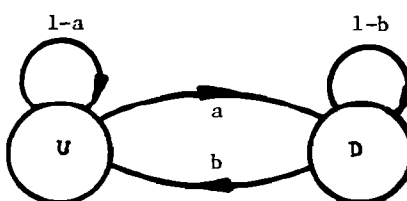


FIG. 1. Two-state Markov model.

^{a)} Present address: Dept. of Naval Architecture and Marine Engineering, The University of Michigan, Ann Arbor, MI 48109.

Resets were also introduced in Chien and Adams (1976) in order to detect a change from one hypothesis to another at some unknown time.

A usual assumption in problems of that nature, which we shall also adopt, is that only one of the hypotheses tested is true, for the entire duration of the observation interval. Attempts to include changes in hypotheses, such as shifts from one hypothesis to another at random points in time, usually cause severe computational difficulties in the solution; Chernoff and Zacks (1964) point out that in such cases only *ad hoc* and hence suboptimal procedures offer any practical solutions. In terms of our specific problem, the above assumption will mean that the source (target) under investigation either is or is not there, with no shifts allowed between these two states. The stationarity assumption of the phase random process is hence preserved in our hypothesis testing as well.

The remainder of the paper is organized as follows: Sec. I formulates and solves the two-hypotheses case via dynamic programming, using the probability of one of the hypotheses being true as the system's state variable. This probability is updated in a Bayesian manner based on observations on the interarrival time of the process. It is seen that the Markov assumption causes no loss of generality in our algorithms and that their generalization to any other process is straightforward if the PMF of the interarrival time of the process is given. This section also extends the above formulation for more than two hypotheses. Although the extension is straightforward, it is seen that computational issues are likely to make such an extension impractical for a number of hypotheses higher than three. Section II presents and discusses several numerical applications of the algorithms for the two and three hypotheses cases. Finally Sec. III discusses important issues regarding this work and suggests areas for further research.

I. PROBLEM FORMULATION AND SOLUTION

As mentioned in the previous section, the version of the problem we shall deal with calls for deciding between two hypotheses H_1 , the "signal-plus-noise" hypothesis, and H_0 , the "pure noise hypothesis." The following is assumed to be known about the decision process:

(1) Each hypothesis H_i ($i = 0, 1$) implies a discrete-time two-state Markov process for fluctuations arriving at the receiver, according to which the state of the process is either U or D ("up" or "down") with known transition probabilities a_i and b_i ($i = 0, 1$) from U to D and D to U, respectively. Opportunities for a state transition occur every ΔT units of time.

(2) The process starts with some user-supplied *a priori* probability $p_0(H_0)$ that H_0 is true. If such a probability is not available, we can always start with $P_0(H_0) = 0.5$. Based on all available information and at any subsequent point in time, we may decide to stop and declare that either H_0 or H_1 is true. In case of false acceptance of H_0 (H_1) we incur a known penalty of L_0 (L_1). Alternatively, we may decide to gather more information about the process in a way described below.

(3) Observing the process involves a known cost of C per observation. Each observation simply involves recording the process' "interarrival time," that is, the time interval between two consecutive detection events (or D-U transitions). In other words we assume that we will be using only the above variable to update our information about the system.

(4) Our objective is to declare one of the two hypotheses as being true in at most N observations, while minimizing the expected value of the total cost associated with that decision.

Before we proceed to the mathematical formulation some observations and clarifications are in order.

(1) As it will be seen later in this section, the Markov assumption about the process is not binding, and therefore causes no loss of generality. That assumption can be replaced by any other assumption about the detection process as long as the PMF of the interarrival times for each hypothesis of that process is an input and these interarrival times form a sequence of independent observations. In the Markov case, the above PMFs are derived from the Markov transition probabilities and the independence assumption holds.

(2) The false acceptance costs L_0 and L_1 , as well as the observation cost C are not constrained to be constants. We could generalize by assuming that these costs depend on when they are incurred, that is, on the stage of the process at which they are realized. The assumption that they are constant is for simplicity reasons and may be therefore easily dropped. Moreover, since we are minimizing the total expected costs, if the cost per observation is not just a function of time but a random variable whose expected value is a known function of time we will not have additional complications in the algorithm.

(3) The interpretation of the above costs is open to discussion. In a passive sonar surveillance situation, either a "monetary" or an "operational" interpretation can be given. For instance, C could be interpreted in terms of resources (sensors, CPU time, communication channels) that are used in order to make an additional observation. Similarly, the penalty costs L_0 and L_1 can be interpreted in terms of resources that would have to be used if a corrective action is necessary in case a false "verdict" is recognized. The issue of how to weigh these two kinds of costs is considered to be beyond the scope of this paper.

(4) The assumption that our only observation variable is the interarrival time of the process is, of course, simplistic and patently in contrast with the intricate signal processing and other information collecting mechanisms available in modern underwater surveillance. However we have made that assumption to provide a motivation on how the above important variable could affect the decision-making process in the version of the problem we described. A further discussion on the merits of that variable as opposed to others is carried out in Sec. III.

With the above clarifications, we now proceed to the formulation and solution of the problem.

Let n be the stage variable of the process. Each stage corresponds to a potential observation. The process starts at $n = 0$ and ends at a stage n which can be no more than N . Each decision to continue observations or stop and declare

H_0 or H_1 true is made immediately after each observation. It is, in that respect, important to emphasize that the time interval between stages n and $n + 1$ is not constant, but is equal to the duration of the interarrival time that elapses in between. Of course, this time interval is an integer multiple of ΔT .

Let also x_n be the state variable of the process. x_n is the probability that H_0 is true at stage n , that is, after n interarrival time observations. Obviously, $x_0 = p_0(H_0)$.

We denote as $f_i(k)$ ($i = 0, 1$ and k positive integer) the PMF of the interarrival time of the process if hypothesis H_i is true. k is the interarrival time itself and is expressed in multiples of ΔT . It was shown by Psarftis *et al.* (1981b) (or can be easily derived) that for a two-state Markov process with transition probabilities a_i and b_i , $f_i(k)$ is of the following form:

$$f_i(k) = \begin{cases} ab \sum_{j=1}^{k-1} (1-a)^{j-1} (1-b)^{k-j-1}, & k = 2, 3, 4, \dots, \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

We note that our notation implies that $k = 2$ (the lowest possible value for k) in case of a D-U-D-U transition.

Our observation variable is denoted as z_n . z_n is the value the n th observed interarrival time, also expressed in multiples of ΔT .

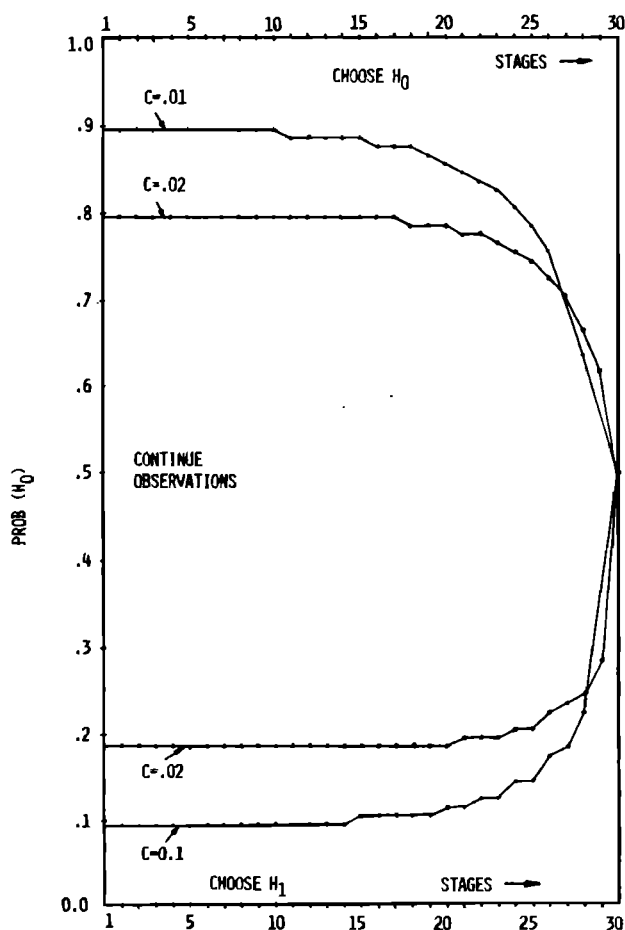


FIG. 2. Stages in a 2-hypotheses testing algorithm.

TABLE I. Summary of input data.

Figure	Horizon (stages, N)	Observation cost C	Transition probabilities						False acceptance costs			A priori probabilities				Case Data records			
			a_0	a_1	a_2	b_0	b_1	b_2	L_0	L_1	L_2	$p(H_0)$	$p_1(H_1)$	$p_2(H_2)$	H_0	H_1	H_2		
2	30	0.01	0.5002	0.5718	...	0.0637	0.1475	...	1	1	146	140
3	7	0.02	0.5002	0.5718	0.2523	0.0637	0.1475	0.2163	1	1	1	146	140	132
4	30	0.02	0.5002	0.5718	...	0.0637	0.1475	...	1	1	...	0.2 to 0.7	0.8 to 0.3	...	146	140
5	30	0.02	0.5002	0.5718	...	0.0637	0.1475	...	1	1	...	0.2 to 0.7	0.8 to 0.3	...	146	140
6	7	0.02	0.5002	0.5718	0.2523	0.0637	0.1475	0.2163	1	1	1	0.4	0.4	0.2	146	140	132
7	7	0.02	0.5002	0.5718	0.2523	0.0637	0.1475	0.2163	1	1	1	0.4	0.4	0.2	146	140	132

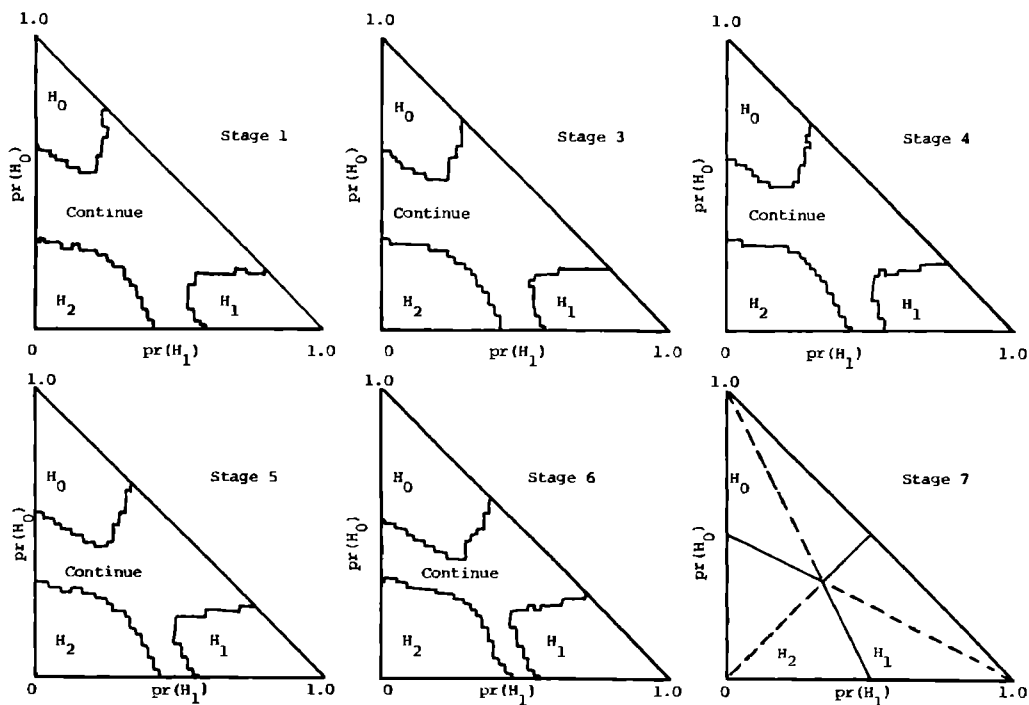


FIG. 3. Stages in a 3-hypotheses testing algorithm.

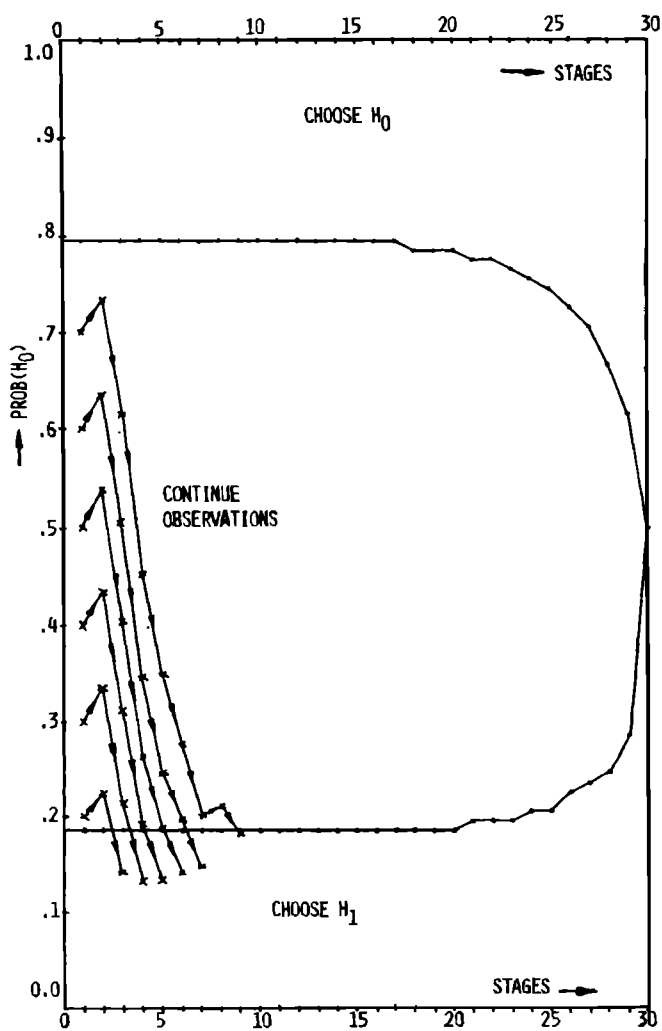


FIG. 4. Real data, 2-hypotheses case, for various $p_0(H_0)$'s.

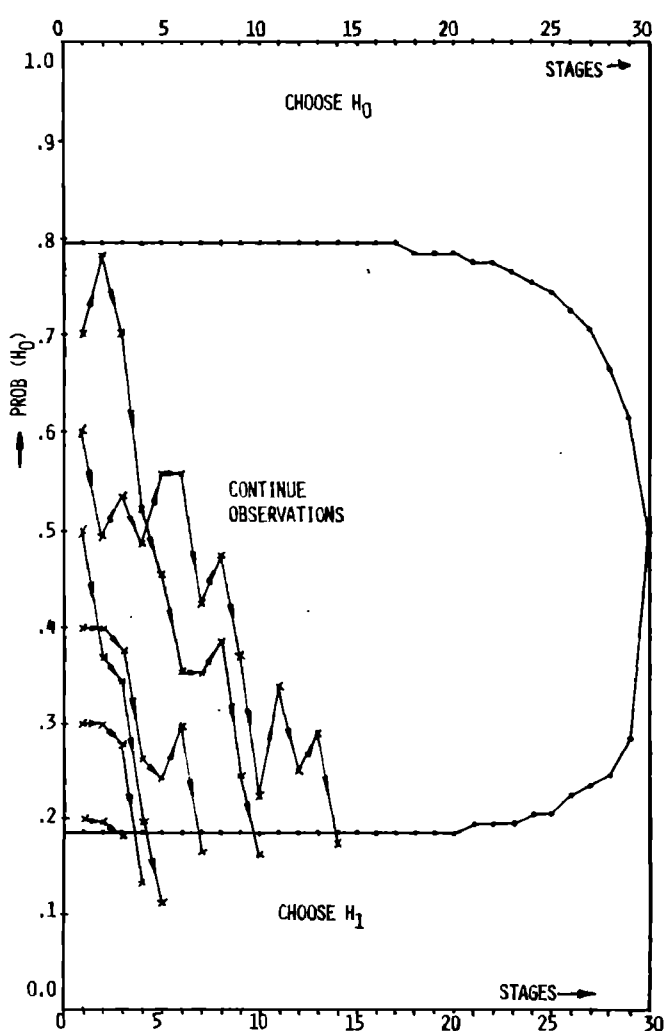


FIG. 5. Simulated data, 2-hypotheses case, for various $p_0(H_0)$'s.

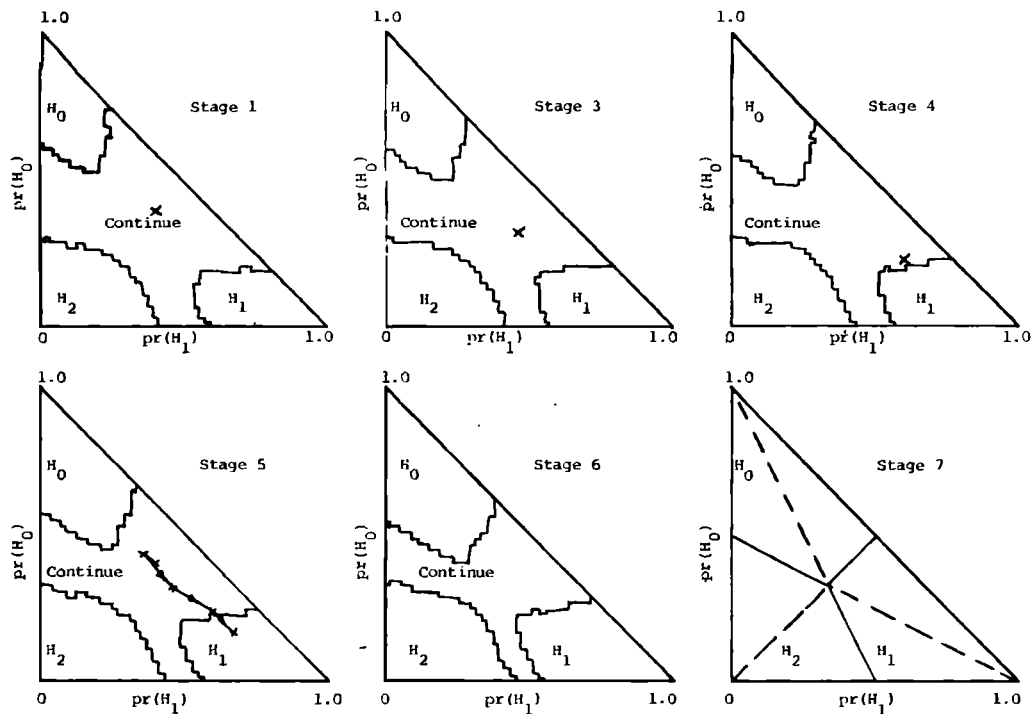


FIG. 6. Real data, 3-hypotheses.

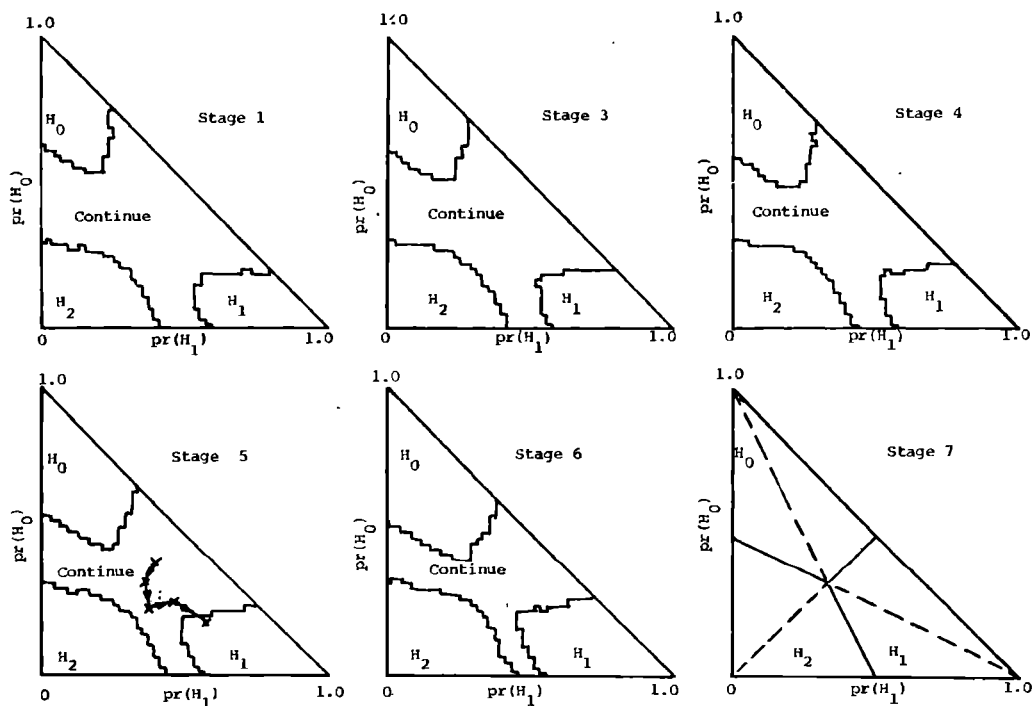


FIG. 7. Simulated data, 3-hypotheses.

Finally, let us define as $J_n(x_n)$, the minimum achievable additional expected cost to terminate the decision process given that after n observations the probability that H_0 is true is x_n . Following Bertsekas (1976) it is now straightforward to solve the problem by dynamic programming. The relevant relationships are

$$J_N(x_N) = \min[(1 - x_N)L_0, x_N L_1], \quad (2)$$

$$J_n(x_n) = \min \begin{cases} (1 - x_n)L_0 \\ x_n L_1 \\ C + E_{z_{n+1}}[J_{n+1}(x_{n+1})] \end{cases} \quad (n = 0, 1, \dots, N-1), \quad (3)$$

with

$$x_{n+1} = \frac{x_n f_0(z_{n+1})}{x_n f_0(z_{n+1}) + (1 - x_n) f_1(z_{n+1})}, \quad (4)$$

and

$$x_0 = p_0(H_0), \quad (5)$$

and the expectation taken over z_{n+1} which obeys the PMF given by

$$\begin{aligned} \text{prob}(z_{n+1}) &= x_n f_0(z_{n+1}) + (1 - x_n) f_1(z_{n+1}), \\ \text{for } z_{n+1} &= 2, 3, \dots \end{aligned} \quad (6)$$

The numerical implementation requires a discretization of the state space.

The extension of the above algorithm to more than two hypotheses is also straightforward. For $M > 2$ hypotheses H_0, H_1, \dots, H_{M-1} , the state variable becomes a $(M-1)$ -dimensional vector each of its elements denoting the current probability that the corresponding hypotheses is true. Perhaps the most significant difference between the $M = 2$ and $M > 2$ cases concerns the computational effort of the algorithm. If we discretize the $[0, 1]$ interval of each probability into s values, then storage for the algorithm grows as Ns^{M-1} and running time as NMs^{M-1} . It is clear that this growth becomes explosive even for relatively small values of M (e.g., 4 or 5), unless the degree of discretization also becomes very small. But even in the most crude discretization scheme—in which $s = 2$, “high probability” or “low probability”—the growth still remains an exponential function of M .

II. A NUMERICAL ILLUSTRATION

The DP algorithms described in the previous section were programmed at the MIT Joint Computer Facility's VAX/VMS system. Both the $M = 2$ and the $M = 3$ hypotheses cases were run. In both cases the probability state spaces were discretized; 101 values for $M = 2$ and 51×51 values for $M = 3$. The results of those runs appear in Figs. 2–7. Table I summarizes the values of the inputs to those runs. The Markov transition probabilities were calculated for a specific phase random acoustic environment and detection thresh-

old, as described in Psaraftis *et al.* (1981b). In all runs, the time interval ΔT of the Markov model was chosen to be equal to 0.4 s, being calibrated in a way also described in Psaraftis *et al.* (1981b). The PMFs of each Markov model's interarrival times can be obtained by substituting the values of the Markov transition probabilities (Table I) in Eq. (1). A discussion of the numerical results follows:

Figure 2 shows the decision regions for each stage in a two-hypotheses case, for two different observation costs and for $N = 30$. Everything else being equal, as it could be expected, the “continue-observations” decision region gets wider if the observation cost drops. It can also be seen that there is a gradual shrinking of the above region as one approaches the end of the decision horizon N , at which point the region in question actually vanishes.

Figure 3 is the counterpart of Fig. 2 for the three-hypotheses case and for $N = 7$. The acceptance and “continue observations” regions form now areas in the $[p(H_0), p(H_1)]$ plane, separated by distinct boundaries. As before, the “continue-observations” region shrinks as n approaches N . Clearly, two consecutive states of the decision process (for instance, two adjacent points on the abscissa of Fig. 2) are not necessarily separated by a fixed time interval, but correspond to instants when detection is made. The time interval between two consecutive stages is the interarrival time of the process which can take on any value. Hence, the horizontal axis of Fig. 2 cannot be considered as being to scale with any fixed time interval.

Figures 4 and 7 represent actual realizations of the decision process for both real and simulated data. All runs corresponding to those figures require the user to specify an *a priori* probability that each hypothesis is true. The “forward” part of the algorithms is then executed, taking into account observations of the interarrival time of the process. In the runs discussed here, these observations were sampled either from real acoustic records (Figs. 4 and 6) or from data obtained via a Monte Carlo simulation on the interarrival time PMFs (Figs. 5 and 7).

In Fig. 4 we test our algorithm with real data sampled from hypothesis H_0 (the actual true hypothesis), starting with six different values of the user-supplied *a priori* probability that H_0 is true, all of them inside the “continue observations” range at $n = 1$, and with $C = 0.02$. We can see that even when $p(H_0)$ is as low as 0.2, observations on the systems interarrival times are such that in all cases the algorithm identifies the correct hypothesis within at most nine observations, much shorter than the allowed horizon of 30.

Figure 5 is similar to Fig. 4 but data are now sampled by Monte Carlo simulation. The same comments as before are in order, but the algorithm here takes longer to identify the correct hypothesis. Figures 6 and 7 display similar tests for the three-hypotheses case, where the actual true hypothesis is H_1 . Figure 6 refers to real data while Fig. 7 to simulated data. In both cases the algorithm concludes that H_1 is true after five observations (in a horizon of seven).

Of course, the figures presented above by no means suggest that the algorithm will always identify the true hypothesis. Its ability to do so would be reduced if the algorithm is “forced” to come to a conclusion within a few observations,

because of, say, a high observation cost that would make the "continue observations" region narrower.

III. DISCUSSION: CONCLUSIONS AND POSSIBLE EXTENSIONS

The algorithms presented in this paper were developed to take advantage of recent work in the area of ocean acoustic detection (Psaraftis *et al.*, 1981a,b), work which demonstrated that one can predict the timing of detection events in the case of a phase random multipath process. The random variable which was chosen to reflect the timing of such events is the "interarrival time" of the process, that is, the time between two consecutive detections. In that respect, the algorithms developed use the above random variable as the "observation variable" for their Bayesian probability updates. One can say that the choice of that random variable as opposed to another has been dictated by methodological arguments rather than its unambiguous superiority for use as a statistic. This is a topic for discussion which is considered outside the scope of this paper. Nevertheless, one can justify the use of such a statistic in situations in passive underwater acoustic surveillance where it is particularly important to have a grasp of when a particular target will be "lost," or be detected again, and use such information to make conclusions about the nature of that target. Another random variable of the process which could be very well used for that purpose is the target's "holding time," that is, the time interval during which the signal is above the detection threshold, or, equivalently, the time during which the Markov process is in its "up" state. PMFs and various other properties of that variable have been derived also by Psaraftis *et al.* (1981a,b) and could be used, either alone, or in combination with the "interarrival time" PMFs, in algorithms similar to ours. The necessary modifications of the algorithm involve mainly the Bayesian update of the state variable [Eq. (4)] and are considered straightforward. It should be emphasized again that the phase randomness assumption or the use of a Markov model for the detection process causes no loss of generality for the algorithms developed. Any other process could be assumed, provided the PMF of the process' interarrival time could be derived. This PMF could then be readily used in the algorithms.

Results from various computer runs have shown that the *a priori* probabilities of each hypothesis being true, as well as the interarrival time PMFs do play an important role both in the final "verdict" of the algorithm and in the time it takes to reach such a conclusion. The system is quite sensitive to values of the observations in the sense that even when the user-specified *a priori* probabilities are strongly in favor of the wrong hypothesis, a few observations are likely to make the algorithm get on the right track.

Computationally, running time and storage will depend on the fineness of discretization of the probability state space in a polynomial manner, and on the number of alternative hypotheses in an exponential manner. More than three hypotheses are likely to create serious computational difficulties, unless a very crude discretization, or some other heuristic rule is adopted.

An extension of our algorithms to the infinite horizon case would necessitate either the discounting of future costs, or a redefinition of the objective in terms of minimization of a long term average of some cost function. The latter is considered a more sensible direction.

Another direction for possible extension involves the issue of stationarity. We have already mentioned that the observations process is stationary. To take $M = 2$ as an example, stationarity means that, for the entire extent of our observations, either H_0 or H_1 is true. In Chien and Adams (1976) a refinement is made for having a "degradation" occurring during the observations, the real state of the process switches after some observation from H_0 to H_1 (or vice versa) with no further switches allowed. This can be used for our purposes in problems in which the target is held for a rather long time and then goes away, that is, the observations process can be thought of as approximately stationary up to the point at which the switch is made. This refinement results in a modified SPRT, the so-called Sequential Reset Detection Test (SRDT), which is formulated in Walker (1980). This minimizes for a given expected time to a false alarm, the expected time to detection of a shift from H_0 to H_1 , following a long persistence of H_0 . The (approximate) results converge to the true ones as the expected time to false alarm approaches infinity. In the case of passive sonar surveillance, the assumptions of this model are justified, since targets over long ranges in the ocean can be thought of as being at a certain "location" for a relatively long period of time before they move away.

ACKNOWLEDGMENTS

Work for this paper has been supported by the Office of Naval Research (Code 431) under contract N00014-79-C-0238. We would like to thank the Associate Editor and the reviewer for their comments.

- Anderson, T. W. (1960). "A Modification of the Sequential Probability Ratio Test to Reduce Sample Size," *Ann. Math. Stat.* 31, 165-197.
- Bertsekas, D. P. (1976). *Dynamic Programming and Stochastic Control* (Academic, New York).
- Chernoff, H., and Zacks, X. (1964). "Estimating the Current Mean of a Normal Distribution which is Subjected to Changes in Time," *Ann. Math. Stat.* 35 (3), 999-1018.
- Chien, T. T., and Adams, M. B. (1976). "A Sequential Failure Detection Technique and its Applications," *IEEE Trans. Autom. Control* AC-21, 750-757.
- Hamblen, W. R. (1977). "A Phase Random Multipath Model for Acoustic Signal Fluctuation in the Ocean and its Comparison with Data," Ph.D. thesis, Dept. of Ocean Engineering, M.I.T.
- Mikhalevsky, P. N., and Dyer, Ira (1978). *J. Acoust. Soc. Am.* 63, 732-738.
- Mikhalevsky, P. N. (1979). *J. Acoust. Soc. Am.* 66, 751-762.
- Mikhalevsky, P. N. (1980). *J. Acoust. Soc. Am.* 67, 812-815.
- Psaraftis, H. N., Perakis, A. N., and Mikhalevsky, P. N. (1981a). "New Models on the Ocean Acoustic Detection Process," *J. Acoust. Soc. Am.* 69, 1724-1734.
- Psaraftis, H. N., Perakis, A. N., and Mikhalevsky, P. N. (1981b). "Memory Detection Models for Phase Random Acoustic Fluctuations," in *Proceedings, International Conference on Communication (IEEE)*, Denver, CO.
- Wald, A. (1974). *Sequential Analysis* (Wiley, New York).
- Walker, B. K. (1980). "A Semi-Markov Approach to Quantifying Fault-Tolerant System Performance," Sc.D. thesis, M.I.T., Dept. of Aeronautics and Astronautics.