Optimal Survivor Search with Multiple States

JOSEPH H. DISCENZA

United States Coast Guard, Springfield, Virginia

LAWRENCE D. STONE

Daniel H. Wagner, Associates, Paoli, Pennsylvania

(Received March 1979; accepted September 1980)

The optimal survivor search problem is to maximize the probability of finding a moving target alive. The target may be a person missing at sea or lost in a wilderness area. In a multistate target problem, the target is moving and changing states stochastically. Each state may imply different motion and detectability assumptions. We find necessary and sufficient conditions for optimal detection plans for multistate targets. Using these conditions, we extend the algoritnms of Brown and Stone for finding optimal detection search plans for moving targets to multistate target problems, and show that the optimal survivor search problem may be solved as a special case of the multistate target problem.

SEARCH AND RESCUE problems often involve a person missing at sea who might be adrift in a disabled vessel (State 1); the vessel may sink and the survivor may take to a life raft (State 2); the life raft may fail forcing the survivor into the water (State 3); and the person may die (State 4, absorbing). Each state has a different sweep width (detectability), and each reacts to wind with a different motion. A rescue mission succeeds only if the person is found alive. The optimal survivor search problem is to maximize the probability of finding the person alive.

This paper introduces the class of multistate search problems. In a multistate problem both the target's motion and its change of state are described by a stochastic process. There are usually different detectability and motion assumptions for each target state. There may be absorbing states in which the target may or may not be detected, but the reward for detecting the target in an absorbing state is negligibly small. Search effort cannot be transferred from one time to another, but it may be divided as finely over the search space as desired. The detection function is assumed to be concave.

We show that the optimal survivor search problem is a special case of a multistate target problem. The necessary and sufficient conditions of Stone (1979) are shown to apply to plans which maximize detection probability over an interval of time $[0, T]$ subject to constraints on the

Operations Research
Vol. 29, No. 2, March-April 1981

0030-364X/81/2902-0309 $01.25

amount of effort available at each time. Using these conditions, we extend the algorithms of Brown (1980) and Stone et al. (1978) in order to solve the optimal survivor search problem. As with the original algorithms, the extended one proceeds by solving a sequence of stationary target problems. However, because of the multistate nature of the target, the stationary problem becomes an uncertain sweep width problem which may be solved by the methods in Richardson and Belkin (1972) or Stone (1975). The algorithm yields a plan which maximizes the probability of detecting the target alive by time T, and may be easily extended to a wide class of multistate target problems.

Recently, the results of Stone (1979) have been further generalized by Stromquist and Stone (1980) to include a wide class of payoff functionals with search and other applications.

We present an example of an optimal survivor search plan obtained by our algorithm and show that some alternate methods of finding plans (commonly used in operational planning), can yield significantly lower success probabilities.

1. MULTISTATE TARGET PROBLEM

We describe the multistate target problem for a discrete time and space search. The corresponding description for continuous space and time is given in the Appendix.

The target's position and state are described by a stochastic process $(X, S) = \{(X_t, S_t), t = 0, \cdots, T\}$ where X_t gives the target position and S_t the target state at time t. Let $\mathbf{X}$ denote the space in which the target moves, in particular, $\mathbf{X}$ is a set $\{1, \cdots, J\}$ of J cells. This is also the space over which we can allocate effort. Effort cannot be allocated to subsets of the space $\mathbf{S} = \{1, \cdots, K\}$ of K possible target states. The probability laws of the process (X, S) are known to the search planner and there is a probability function p such that

$p(j, k, t)$ = probability of the target being in cell j and in state k at time t.

An allocation of search effort is specified by a function $\psi: \{1, \cdots, J\} \times \{0, \cdots, T\} \to [0, \infty)$, such that $\psi(j, t)$ gives the search effort applied to cell j at time t. The amount of search effort available at time t is $m(t)$, and allowable plans must satisfy

$$\textstyle\sum_{j=1}^{J} \psi(j, t) = m(t) \quad \text{for} \quad t = 0, \cdots, T. \qquad (1)$$

The class of search plans ψ satisfying (1) is denoted by $\Psi(m)$. Observe that effort cannot be transferred from one time to another but can be spread over space $\mathbf{X}$ as desired, i.e., effort is *infinitely divisible.*

Let ω denote a sample path of the process (X, S). We assume there is a detection function b:$[0, \infty) \to [0, 1]$ and a nonnegative sweep width function W which specifies the detectability of the target as a function of target location, state, and time. For a search plan ψ, the probability of detecting the target given that it follows the sample path ω is $b(\sum_{t=0}^{T} W(X_t(\omega), S_t(\omega), t)\psi(X_t(\omega), t))$. Letting E denote expectation over the sample paths and suppressing the variable ω, the total probability of detection by time T becomes $P_T[\psi] = E[b(\sum_{t=0}^{T} W(X_t, S_t, t)\psi(X_t, t))]$. If $\psi^* \in \Psi(m)$ satisfies

$$P_T[\psi^*] \geq P_T[\psi] \quad \text{for} \quad \psi \in \Psi(m), \tag{2}$$

then ψ^* is called *T-optimal within* $\Psi(m)$.

Returning to the optimal survivor search problem, we see that by estimating the probability distribution of the target's initial location and state as well as the probabilities of transitions to other states (e.g., the survival time distribution of a person in the water) and using our knowledge of how the target moves in response to winds and currents in each state combined with the distribution of these winds and currents, we can construct the multistate process (X, S) which represents our knowledge about the location and state of the survivor. Then, by taking the function W to be equal to the sweep width for detecting the target in each of its possible states, and to be equal to zero if the target is dead, we have created a multistate problem in which $P_T[\psi]$ is the probability of detecting the target *alive* by time T. Thus, maximizing detection probability in this multistate problem is equivalent to maximizing the probability of detecting the target alive.

In the Appendix of this paper, necessary and sufficient conditions are given for a plan ψ^* to be T-optimal within $\Psi(m)$ for a multistate target problem. These conditions are used to construct the algorithm for finding optimal survivor search plans that is discusssed in the next section. The Appendix also presents a proof that this algorithm converges to the optimal plan.

2. ALGORITHM FOR FINDING OPTIMAL SURVIVOR SEARCH PLANS

The algorithm developed here is for an exponential detection function. That is, if the target is in cell j and state k at time t, and we apply $\psi(j, t)$ effort, then detection occurs with probability

$$1 - \exp[-W(j, k, t)\psi(j, t)]$$

and is independent of detection at any other time. For convenience of notation, all normalizing constants are put into W. For example, if effort is measured in units of search track length, then $W(j, k, t)$ is the sweep width divided by the area of cell j. In terms of the multistate target model given in the previous section, we have taken $b(z) = 1 - e^{-z}$ for $z \geq 0$, and

$$P_T[\psi] = E[1 - \exp(-\sum_{t=0}^{T} W(X_t, S_t, t)\psi(X_t, t))].$$

The algorithm described in this section requires the calculation of a posterior distribution on target location and state and the solution of a single stage optimization problem involving that posterior distribution. To describe the algorithm, we shall first discuss the calculation of the posterior distribution, then the solution of the single stage optimization, and finally show how these two are used in the algorithm.

Calculation of Posterior Distribution

For the algorithm described in this section, one must be able to calculate

$$\tilde{p}_\psi(j, k, t) = \Pr\left\{\begin{array}{l}\text{target is in cell } j \text{ and state } k \text{ at time } t \text{ and has} \\ \text{not been detected by the search effort at all} \\ \text{times } u < t \text{ and will not be detected at all times} \\ u > t \text{ using the allocation } \psi\end{array}\right\}.$$

The method of calculating $\tilde{p}_\psi$ will depend on the stochastic process used to represent target motion and state.

Monte Carlo Simulation. Suppose one represents the stochastic process (X, S) by a finite number, N, of sample paths obtained from a Monte Carlo stimulation of (X, S). Let the nth path be

$$\omega_n = ((x_0^n, s_0^n), (x_1^n, s_1^n) \cdots (x_T^n, s_T^n)),$$

and let

$$\alpha_n = \Pr\{\text{target follows path } \omega_n\} \quad \text{for} \quad n = 1, \cdots, N.$$

Then

$$\tilde{p}_\psi(j, k, t) = \sum_{\{n: x_t^n = j, s_t^n = k\}} \alpha_n \exp(-\sum_{u \neq t} W(x_u^n, s_u^n, u)\psi(x_u^n, u)).$$

For this method, it is not necessary that the target motion take place in discrete space. One could have the target moving in continuous space and have the cells determine only how finely effort can be allocated. Then the above summation would be over all points x_t^n which fall in the jth cell. This is how target location probability distributions are obtained in the U.S. Coast Guard CASP search planning system (see Richardson and Discenza [1980]).

Markov Chain. For the example of an optimal survivor search plan presented in the next section, we use a Markov chain model for (X, S) and a modification of Brown's *reach* and *survive* matrices to compute $\tilde{p}_\psi$. Specifically, let $\tau_t(j, k, i, l)$ be the probability of transitioning from cell j and state k at time t to cell i and state l at time $t + 1$. Recall that $p(j, k,$

0) is the initial probability of the target being in cell j and state k at time 0, and define the *reach* matrix R for an allocation ψ as follows:

$$R(j, k, 0, \psi) = p(j, k, 0) \quad \text{for} \quad j = 1, \cdots, J \quad \text{and} \quad k = 1, \cdots, K$$

and for $1 \leq t \leq T$,

$$R(j, k, t, \psi) = \textstyle\sum_{i=1}^{J} \sum_{l=1}^{K} R(i, l, t-1, \psi) \cdot \exp[-W(i, l, t-1)\psi(i, t-1)]\tau_{t-1}(i, l, j, k). \tag{3}$$

Similarly, define the *survive* matrix, S as follows:

$$S(j, k, T, \psi) = 1 \quad \text{for} \quad j = 1, \cdots, J \quad \text{and} \quad k = 1, \cdots, K$$

and for $0 \leq t \leq T - 1$,

$$S(j, k, t, \psi) = \textstyle\sum_{i=1}^{J} \sum_{l=1}^{K} \tau_t(j, k, i, l) \cdot \exp[-W(i, l, t+1)\psi(i, t+1)]S(i, l, t+1, \psi). \tag{4}$$

From the above definitions of R and S we can see that $R(j, k, t, \psi)$ is the probability that the target reaches cell j in state k at time t without being detected by the effort at times $0, 1, \cdots, t-1$ and that $S(j, k, t, \psi)$ is the probability that if the target starts in cell j and state k at time t, it will not be detected by the search effort at times $t+1$, $t+2, \cdots, T$. Thus

$$\tilde{p}_\psi(j, k, t) = R(j, k, t, \psi)S(j, k, t, \psi).$$

Single Stage Optimization

Having computed $\tilde{p}_\psi(\cdot, \cdot, t)$ for some t, we perform an optimization (which is equivalent to finding an optimal allocation for maximizing the probability of detecting a stationary target with uncertain sweep width; see Richardson and Belkin [1972], or Stone [1975]). In particular, let

$$G(t) = \textstyle\sum_{j=1}^{J} \sum_{k=1}^{K} \tilde{p}_\psi(j, k, t).$$

Then

$$q(j, k) \equiv \tilde{p}_\psi(j, k, t)/G(t) = \Pr\left\{\begin{matrix}\text{target in cell } j \text{ and state } k \\ \text{at time } t\end{matrix} \;\middle|\; \begin{matrix}\text{failure to detect at all times} \\ \text{other than } t.\end{matrix}\right\}.$$

Let us consider the problem of allocating $m(t)$ effort to maximize the probability of detecting a stationary target with location and state distribution given by q. If we place effort z in cell j and the target is in cell j and state k, then the probability of detecting the target is $1 - \exp[-W(j, k, t)z]$. The probability of the target being in cell j and being detected with effort z is

$$\textstyle\sum_{k=1}^{K} q(j, k)(1 - e^{-W(j,k,t)z}) \quad \text{for} \quad z \geq 0. \tag{5}$$

In order to put this in the form of a standard stationary target problem as given in Stone (1975), we let

$$\bar{q}(j) = \textstyle\sum_{k=1}^{K} q(j, k)$$

$$\bar{q}(k|j) = q(j, k)/\bar{q}(j) = \Pr\{\text{target in state } k \,|\, \text{target in cell } j\},$$

and define

$$\beta(j, z) = \textstyle\sum_{k=1}^{K} \bar{q}(k|j)(1 - e^{-W(j,k,t)z}) \quad \text{for} \quad z \geq 0, j = 1, \cdots, J. \tag{6}$$

If we allocate $f(j)$ effort to cell j for $j = 1, \cdots, J$, then the probability of detecting the target is

$$Q[f] \equiv \textstyle\sum_{j=1}^{J} \bar{q}(j)\beta(j, f(j)). \tag{7}$$

Let $\beta'(j, \cdot)$ denote the derivative of $\beta(j, \cdot)$. Observe that $\beta(j, \cdot)$ is concave for $j = 1, \cdots, J$. By Corollary 2.2.6 of Stone (1975) and the comment following it, f^* is an optimal allocation of $m(t)$ effort if and only if there exists a $\lambda(t) \geq 0$ such that

$$\begin{aligned} \bar{q}(j)\beta'(j, f^*)) &= \lambda(t) \quad \text{if} \quad f^*(j) > 0 \\ &\leq \lambda(t) \quad \text{if} \quad f^*(j) = 0 \quad \text{for} \quad j = 1, \cdots, J. \end{aligned} \tag{8}$$

To find f^*, we choose a λ_0 and solve (numerically) for the $f(j)$ which satisfies (8) for $j = 1, \cdots, J$. Observe that $f(j)$ is monotone decreasing as a function of λ for $j = 1, \cdots, J$. Calculate $m_0 = \sum_{j=1}^{J} f(j)$. If $m_0 < m(t)$ choose $\lambda_1 < \lambda_0$, or if $m_0 > m(t)$ choose $\lambda_1 > \lambda_0$. Then repeat the above procedure to find new values of $f(j)$ for $j = 1, \cdots, J$. By performing a binary search in λ we solve (numerically) for $\lambda(t)$ and hence f^*. Define a transformation Ξ_t which operates on a search allocation ψ, by finding f^* to solve the above problem and replacing ψ by the new allocation

$$\Xi_t\psi = \begin{cases} \psi(\cdot, u) & \text{for} \quad u \neq t \\ f^* & \text{for} \quad u = t. \end{cases} \tag{9}$$

Algorithm

We now describe the algorithm for finding optimal survivor search plans. The stopping rule for the algorithm involves use of Washburn's (1981) upper bound which for each search plan ψ gives a number $B(\psi)$ such that $P_T[\psi^*] - P_T[\psi] \leq B[\psi]$ where ψ^* is the T-optimal plan. Washburn's bound is described in the Appendix.

1. Let ϵ be a small positive number.
2. Let $\psi_{-1}(j, t) = 0$ for $j = 1, \cdots, J$, and $t = 0, \cdots, T$.
3. Let $i = 0$.
4. Perform Step 5 for $t = 0, \cdots, T$.
5. Set $\psi_{(T+1)i+t} = \Xi_t\psi_{(T+1)i+t-1}$.

6. If $B[\psi_{(T+1)(i+1)-1}] < \epsilon$, stop: the answer is $\psi_{(T+1)(i+1)-1}$.
7. Otherwise, increase i by 1 and return to Step 4.

The number ϵ is the tolerance set for the problem solution. The first pass through Step 5 produces the "myopic" plan. This plan simply chooses $\psi(\cdot, t)$ to maximize the probability of detecting the target alive during time t given the target has not been detected prior to time t. Generally, the myopic plan will not be T-optimal.

In the Appendix we show that the above algorithm produces a sequence of plans whose probability of success approaches that of the T-optimal plan.

3. EXAMPLE AND COMPARISONS

In this section we use an example to compare the optimal survivor search plan obtained by the algorithm described above to plans obtained by some other techniques commonly in use in operational search planning. For simplicity, we shall assume that the sweep width W is a function only of the state and write $W(k)$ in place of $W(j, k, t)$.

Each of the other techniques replaces the multistate target problem with a proxy single state problem and finds the myopic allocation for the proxy problem. As in Brown, and Stone (1975), the myopic or incrementally optimal plan is the one which allocates the effort available at time t in a fashion which maximizes the detection probability at that time given failure to detect before time t. The myopic plan is obtained by the first iteration through $t = 0, \cdots, T$ of the algorithm described in Section 2.

To calculate the proxy single-state plans, a sequence of single-stage or stationary target optimizations are performed for the posterior distribution $\tilde{p}_\psi$ given in Section 2. Although ψ may differ, the same method is used to calculate $\tilde{p}_\psi$ in all these plans, so we can describe the proxy problem by describing the stationary search optimization that is performed at each time period in place of the correct multistate optimization.

To describe the single-stage optimization, we fix t and write $\tilde{p}_\psi(j, k)$ in place of $\tilde{p}_\psi(j, k, t)$ for the probability that the target is in cell j and state k.

The first approximation is the simplest: Choose one state, and ignore the others. We will call this the *selective* method and let k^* denote the state selected. Set

$$r(j) = \tilde{p}_\psi(j, k^*) \quad \text{for} \quad j = 1, 2, \cdots, J. \tag{10}$$

Search is then allocated to be optimal for the (defective) distribution r and sweep width $W(k^*)$.

The second approximation, which we call the *additive* method, computes a single probability distribution in space which is the sum of the state probabilities over states k for which $W(k) > 0$, ignoring the differences in sweep width. Let $\mathbf{K}^+ = \{k{:}1 \le k \le K \text{ and } W(k) > 0\}$. The additive method finds

$$q(j) = \textstyle\sum_{k\in\mathbf{K}^+} \tilde{p}_\psi(j, k) \quad \text{for} \quad j = 1, 2, \cdots, J. \tag{11}$$

For the sweep width one uses the average $\bar{W}$ where

$$\bar{W} = \textstyle\sum_{k\in\mathbf{K}^+} W(k) \sum_{j=1}^{J} p(j, k, 0) / \sum_{k\in\mathbf{K}^+} \sum_{j=1}^{J} p(j, k, 0), \tag{12}$$

and plans the search according to the (defective) distribution q and sweep width $\bar{W}$. Note that $\bar{W}$ is the expectation of W with respect to the prior distribution on states and does not change as the optimization proceeds.

The additive distribution represents the actual location distribution because the probability mass from all states is included, whereas the selective distribution omits the probability mass from all states except one. For display purposes, r and q should be rescaled to add to one, but the calculation of optimal effort distributions depends only on the ratios of the probabilities in different regions, and for this purpose scaling is unnecessary.

A third method, called *weighting*, is not in operational use, but we propose it as an approximate method. This method finds

$$c(j) = \textstyle\sum_{k=1}^{K} W(k)\tilde{p}_\psi(j, k) \quad \text{for} \quad j = 1, \cdots, J, \tag{13}$$

and then rescales c to be a probability distribution. Again $\bar{W}$ as defined in (12) is used for the sweep width in finding f_c, the optimal allocation for this single-state problem. In addition we use f_c as the initial guess for the solution of the single-stage multistate optimization described in Section 2.

For our example we use a 10-cell problem having 3 states and 6 time periods during which available effort $m(t)$ increases. This situation may represent a survivor either in a vessel (State 1), on a raft (State 2), or drowned (State 3), where targets in both states are moving toward a region of peril. We take

$$W(1) = 5, \quad W(2) = 1, \quad W(3) = 0,$$

$$p(1, 1, 0) = 0.4, \quad p(4, 2, 0) = 0.6,$$

$$m(0) = m(1) = 0.2, \quad m(2) = m(3) = 0.3, \quad m(4) = m(5) = 0.5.$$

Motion from cell to cell and state change are modeled by independent Markov chains. Motion is permitted only from cell j to j or $j + 1$. A target in State 1 moves from cell j to $j + 1$ with probability 0.1, and a target in State 2 moves from j to $j + 1$ with probability 0.6. A target in any cell in State 1 changes to State 2 with probability 0.1, but a target in cell j and State 2 changes to State 3 with probability $j/10$, making the higher-numbered cells more "dangerous" to State 2 targets.

Table I presents the resulting search plans from six different methods: (a) myopic selective, State 1; (b) myopic selective, State 2; (c) myopic additive; (d) myopic weighted; (e) myopic multistate (first pass through the T-optimal algorithm); and (f) T-optimal.

A comparison of the results shows the T-optimal plan to be strikingly different from the myopic ones. For example, plans (a), (d), and (e) concentrate on the easy-to-find, State 1 targets for at least two periods of

TABLE I

COMPARISON OF EFFORT ALLOCATIONS OBTAINED FROM THE MYOPIC PLANS AND THE T-OPTIMAL PLAN

(a) Myopic selective, $k^* = 1$, $P_T = 0.38$

t	Cell 1	2	3	4	5	6
0	0.20					
1	0.20					
2	0.24	0.06				
3	0.14	0.16				
4	0.20	0.22	0.08			
5	0.15	0.17	0.17			

(b) Myopic selective, $k^* = 2$, $P_T = 0.40$

t	Cell 1	2	3	4	5	6
0				0.20		
1					0.20	
2					0.30	
3	0.14	0.16				
4	0.03	0.45	0.02			
5	0.01	0.31	0.18			

(c) Myopic additive, $P_T = 0.46$

t	Cell 1	2	3	4	5	6
0	0.02			0.18		
1	0.20					
2	0.20				0.10	
3	0.06	0.21				0.02
4	0.18	0.21	0.11			
5	0.07	0.16	0.21	0.05		

(d) Myopic weighted, $P_T = 0.38$

t	Cell 1	2	3	4	5	6
0	0.20					
1	0.20					
2	0.28				0.02	
3	0.02	0.26				0.02
4	0.26	0.16	0.03			0.04
5	0.05	0.18	0.27			

(e) Myopic multistate, $P_T = 0.39$

t	Cell 1	2	3	4	5	6
0	0.20					
1	0.20					
2	0.17				0.13	
3	0.13	0.17				
4	0.16	0.20	0.06			0.08
5	0.13	0.17	0.20			

(f) T-optimal, $P_T = 0.47$

t	Cell 1	2	3	4	5	6
0				0.20		
1	0.11				0.09	
2	0.21				0.09	
3	0.13	0.17				
4	0.19	0.21	0.10			
5	0.14	0.18	0.18			

search before searching for the State 2 targets. The T-optimal plan, however, immediately concentrates on the State 2 targets in the more perilous region, and devotes a significant amount of effort there for three periods. The T-optimal plan exhibits almost a 20% improvement in detection probability over the myopic multistate plan, plan (d), and at least some improvement over each of the others. By coincidence, the myopic additive plan and the T-optimal plan are very similar in this situation, producing almost identical detection probabilities. However if one changed the above example by setting $W(1) = 1$ and $W(2) = 5$, then the myopic additive plan would produce a detection probability of 0.62 compared to 0.69 for the T-optimal plan.

The algorithm converges quickly, reaching $B[\psi_{(T+1)(i+1)-1}] = 0.001$ in three iterations and 0.0004 in four iterations.

APPENDIX

In this appendix we state necessary and sufficient conditions for a search plan to be T-optimal within $\Psi(m)$ for a multistate target search. In addition, we demonstrate how the algorithm described in Section 2 converges to the T-optimal plan and how Washburn's upper bound is computed.

Multistate Target Problem in Continuous Space and Time

The target's position and state are described by a stochastic process $(X, S) = \{(X_t, S_t), 0 \leq t \leq T\}$ where X_t gives the target position and S_t the target state at time t. Let $\mathbf{X}$ denote the space in which the target moves, in particular, $\mathbf{X}$ is a subset of Euclidean n-space. This is also the space over which we can allocate effort. Effort cannot be allocated to subsets of the space $\mathbf{S}$ of possible target states. The probability laws of the process (X, S) are known to the search planner and there is a density function p such that

$p(x, s, t)$ = probability density of the target being at point x and in state s at time t.

An allocation of search effort is specified by a Borel function $\psi:\mathbf{X} \times [0, T] \rightarrow [0, \infty)$, such that $\psi(x, t)$ gives the rate at which effort density is applied to the point x at time t. The rate at which search effort is available is specified by the function $m:[0, T] \rightarrow (0, \infty)$ and allowable plans must satisfy

$$\int_{\mathbf{X}} \psi(x, t)dx = m(t) \quad \text{for} \quad 0 \leq t \leq T.$$

The class of search plans ψ satisfying the above equality is denoted by $\Psi(m)$. Again, effort cannot be transferred from one time to another but can be spread as finely over space as desired.

Let E denote expectation over the sample paths of (X, S). Then in analogy to the discrete case, the total probability of detection by time T when using plan ψ is

$$P_T[\psi] = E\left[b\left(\int_0^T W(X_t, S_t, t)\psi(X_t, t)dt\right)\right].$$

If $\psi^* \in \Psi(m)$ satisfies $P_T[\psi^*] \geq P_T[\psi]$ for $\psi \in \Psi(m)$, then ψ^* is called T-optimal as before.

The multistate model and the theorems given below have the expected analogs in the case of mixtures of discrete and continuous space and time.

Necessary and Sufficient Conditions

To guarantee the existence of the kernel of the Gateaux differential defined below, we restrict the class of search plans to the following normed linear space F. Let F be the space of real-valued Borel functions defined on $\mathbf{X} \times [0, T]$ such that

$$\|f\| \equiv \int_0^T \text{ess sup}_{x \in \mathbf{X}} |f(x, t)| dt < \infty.$$

Let $F^+ = \{f \in F: f \geq 0\}$. We require that the density function $p(x, s, \cdot)$ be Borel measurable for all $(x, s) \in \mathbf{X} \times \mathbf{S}$. We also need to assume some topological structure on $\mathbf{S}$ and require that W be measurable with respect to the Borel subsets of $\mathbf{X} \times \mathbf{S}$. Having done this, we can assume the process (X, S) has Borel sample paths and be guaranteed that $\int_0^T W(X_u, S_u, u)\psi(X_u, u)du$ is well-defined. Let b' denote the derivative of b, E_{xst} indicate expectation conditioned on $(X_t, S_t) = (x, s)$, and define

$$D_T(\psi, x, t) = \int_{\mathbf{S}} E_{xst}\left[b'\left(\int_0^T W(X_u, S_u, u)\psi(X_u, u)\, du\right)\right] \cdot W(x, s, t)p(x, s, t)ds. \tag{A-1}$$

THEOREM 1. *Suppose that b is concave with a bounded nonnegative derivative b' and that W is bounded. Assume the sample paths of (X, S) are Borel measurable and that D_T is well defined for (x, t) such that $\int_{\mathbf{S}} p(x, s, t)ds > 0$. Then ψ^* is T-optimal within $\Psi(m)$ if and only if $\psi^* \in \Psi(m)$ and there exist $\lambda:[0, T] \to [0, \infty)$ such that*

$$\begin{aligned} D_T(\psi^*, x, t) &= \lambda(t) \quad \text{if } \psi^*(x, t) > 0 \\ &\leq \lambda(t) \quad \text{if } \psi^*(x, t) = 0 \text{ for a.e. } (x, t) \in \mathbf{X} \times [0, T]. \end{aligned} \tag{A-2}$$

Proof. The proof proceeds by verifying that D_T is the kernel of the Gateaux differential of P_T. The remainder of the proof then follows from Theorem 1 of Stromquist and Stone.

For $\psi \in F^+$, let $C(\psi)$ be the cone of directions h such that $\psi + \theta h \in F^+$ for all sufficiently small nonnegative values of θ. For $\psi \in F^+$ and $h \in C(\psi)$, the Gateaux differential of P_T evaluated at ψ in the direction h is by definition

$$P_T'[\psi, h] = \lim_{\epsilon\to 0+} E\left\{(1/\epsilon)\left[b\left(\int_0^T W(X_t, S_t, t)[\psi(X_t, t) + \epsilon h(X_t, t)]dt\right)\right.\right.$$
$$\left.\left. - b\left(\int_0^T W(X_t, S_t, t)\psi(X_t, t)dt\right)\right]\right\}.$$

Let M_1 and M_2 be bounds for b' and W, respectively. Then the above integrand is bounded by $M_1M_2\|h\|$, and we may invoke the dominated convergence theorem to obtain

$$\begin{aligned}
&P_T'[\psi, h] \\
&= E\left\{b'\left(\int_0^T W(X_t, S_t, t)\psi(X_t, t)dt\right)\int_0^T W(X_t, S_t, t)h(X_t, t)dt\right\} \\
&= E\left\{\int_0^T W(X_t, S_t, t)b'\left(\int_0^T W(X_u, S_u, u)\psi(X_u, u)du\right)h(X_t, t)dt\right\} \\
&= \int_0^T E\left[W(X_t, S_t, t)b'\left(\int_0^T W(X_u, S_u, u)\psi(X_u, u)du\right)h(X_t, t)\right]dt \\
&= \int_0^T \int_{\mathbf{X}}\int_{\mathbf{S}} E_{xst}\left[b'\left(\int_0^T W(X_u, S_u, u)\psi(X_u, u)du\right)\right] \\
&\quad \cdot W(x, s, t)p(x, s, t)ds\,h(x, t)dxdt \\
&= \int_0^T \int_{\mathbf{X}} D_T(\psi, x, t)h(x, t)dxdt.
\end{aligned}$$

Thus, D_T is the kernel of the Gateaux differential and by Theorem 1 of Stromquist and Stone, our result is proved.

For the convenience of the reader, we state the discrete counterpart to Theorem 1. We assume that the search space **X**, state space **S**, and time are discrete. Specifically, $\mathbf{X} = \{1, \cdots, J\}$ and $\mathbf{S} = \{1, \cdots, K\}$. This eliminates the need for the Borel measurability assumptions made for Theorem 1.

The function p becomes

$$p(j, k, t) = \Pr\{(X_t, S_t) = (j, k)\}$$

and D_T becomes

$$D_T(\psi, j, t) = \sum_{k=1}^{K} \mathrm{E}_{jkt}[b'(\textstyle\sum_{u=0}^{T} W(X_u, S_u, u)\psi(X_u, u))]W(j, k, t)p(j, k, t). \tag{A-3}$$

The above theorem becomes:

THEOREM 2. *Suppose that b is concave with bounded nonnegative derivative b′ and that W is bounded. Assume that D_T is well-defined for (j, t) such that* $\sum_{k=1}^{K} p(j, k, t) > 0$. *Then* ψ^* *is T-optimal within* $\Psi(m)$ *if and only if there exists a T + 1 vector* $(\lambda(0), \cdots, \lambda(T))$ *with nonnegative components such that for* $t = 0, 1, \cdots, T$,

$$D_T(\psi^*, j, t) = \lambda(t) \quad \text{if} \quad \psi^*(j, t) > 0 \tag{A-4}$$
$$\leq \lambda(t) \quad \text{if} \quad \psi^*(j, t) = 0 \quad \text{for} \quad j = 1, \cdots, J.$$

Convergence of the Algorithm

We now show that the algorithm described in Section 2 converges to the optimal plan. The convergence proof follows that given by Brown.

There are a finite number of cells, a finite number of time periods, and the maximum effort that can be placed in any cell at any time is $\max_{0 \leq t \leq T} m(t) < \infty$. The sequence $\{\psi_{(T+1)i}: i = 0, 1, \cdots\}$ may be thought of as a sequence of points in a compact subset of $J \times (T + 1)$dimensional Euclidean space, and so there is a subsequence $\{\psi_{(T+1)i_n}: n = 1, 2, \cdots\}$ which converges to an allocation $\psi^* \in \Psi(m)$. Since the sequence of detection probabilities $\{P_T[\psi_{(T+1)i}]: i = 0, 1, 2, \cdots\}$ is monotone increasing, it has a limit, and by the continuity of P_T it follows that

$$\lim_{i\to\infty} P_T[\psi_{(T+1)i}] = \lim_{n\to\infty} P_T[\psi_{(T+1)i_n}] = P_T[\psi^*].$$

Let Ξ be the composition $\Xi = \Xi_T, \Xi_{T-1}, \cdots, \Xi_0$. Since Ξ_t is continuous for each t, it follows that Ξ is continuous. Thus,

$$P_T[\psi^*] = \lim_{n\to\infty} P_T[\psi_{(T+1)i_n}] = \lim_{n\to\infty} P_T[\psi_{(T+1)(i_n+1)}]$$
$$= \lim_{n\to\infty} P_T[\Xi \psi_{(T+1)i_n}]$$
$$= P_T[\Xi \psi^*].$$

Since the detection function β defined in Section 2 is strictly concave, it follows that if $P_T[\Xi_t \psi] = P_T[\psi]$, then $\Xi_t \psi = \psi$. Thus, $\Xi_t \psi^* = \psi^*$ for $t = 0, \cdots, T$. This means that for $t = 0, 1, \cdots, T$, there exists a $\lambda(t) \geq 0$ such that

$$\bar{q}(j)\beta'(j, \psi^*(j, t)) = \lambda(t) \quad \text{if} \quad \psi^*(j, t) > 0 \tag{A-5}$$
$$\leq \lambda(t) \quad \text{if} \quad \psi^*(j, t) = 0 \quad \text{for} \quad j = 1, \cdots, J.$$

But conditions (A-5) are equivalent to the necessary and sufficient conditions (A-4) for a T-optimal plan given in Theorem 2. To see this, let E_{jkt} denote expectation conditioned on $(X_t, S_t) = (j, k)$ and recall that we are taking $b(z) = 1 - e^{-z}$ so that D_T in (A-3) becomes

$$D_T(\psi, j, t) = \sum_{k=1}^{K} E_{jkt}[\exp[-\sum_{u=0}^{T} W(X_u, S_u, u)\psi(X_u, u)]]$$
$$\cdot W(j, k, t) p(j, k, t)$$

$$= \sum_{k=1}^{K} E_{jkt}[\exp[-\sum_{u\neq t} W(X_u, S_u, u)\psi(X_u, u)]] \cdot p(j, k, t) W(j, k, t) e^{-W(j,k,t)\psi(j,t)}.$$

Note that $\exp[-\sum_{u\neq t} W(X_u, S_u, u)\psi(X_u, S_u, u)]$ is the probability of failing to detect the target with the search effort placed at all times other than t. Since E_{jkt} is expectation conditioned on $(X_t, S_t) = (j, k)$ and $p(j, k, t) = \Pr\{(X_t, S_t) = (j, k)\}$, it follows that

$$\tilde{p}_\psi(j, k, t) = E_{jkt}[\exp[-\sum_{u\neq t} W(X_u, S_u, u)\psi(X_u, u)]] p(j, k, t)$$

and that

$$\begin{aligned} D_T(\psi, j, t) &= \sum_{k=1}^{K} \tilde{p}_\psi(j, k, t) W(j, k, t) e^{-W(j,k,t)\psi(j,t)} \\ &= \sum_{k=1}^{K} \bar{q}(j)\bar{q}(k|j) W(j, k, t) e^{-W(j,k,t)\psi(j,t)} \\ &= \bar{q}(j)\beta'(j, \psi(j, t)) \quad \text{for} \quad j = 1, 2, \cdots, J, t = 0, 1, \cdots, T. \end{aligned}$$

Thus, conditions (A-5) are equivalent to (A-4), and ψ^* satisfies the necessary and sufficient conditions of Theorem 2.

We have shown that the algorithm produces a sequence of plans whose probabilities of detection approach that of a T-optimal plan. Furthermore, there is a subsequence of these plans which converges to a T-optimal plan.

Washburn's Upper Bound

Washburn's upper bound may be applied to the algorithm as follows. Let ψ and $\tilde{\psi}$ be members of $\Psi(m)$. The concavity of P_T implies that

$$\begin{aligned} P_T[\tilde{\psi}] - P_T[\psi] &\leq P_T'[\psi, \tilde{\psi} - \psi] \\ &= \sum_{t=0}^{T} \sum_{j=1}^{J} [D_T(\psi, j, t)\tilde{\psi}(j, t) - D_T(\psi, j, t)\psi(j, t)]. \end{aligned}$$

For $t = 0, 1, \cdots, T$, let

$$\bar{\lambda}(t) = \max\{D_T(\psi, j, t): 1 \leq j \leq J\}$$

$$\underline{\lambda}(t) = \min\{D_T(\psi, j, t): 1 \leq j \leq J \quad \text{and} \quad \psi(j, t) > 0\}.$$

Then

$$P_T[\tilde{\psi}] - P_T[\psi] \leq \sum_{t=0}^{T} (\bar{\lambda}(t) - \underline{\lambda}(t))m(t) \equiv B(\psi). \qquad \text{(A-6)}$$

This is the bound used to provide the stopping rule for the algorithm. Since $\tilde{\psi}$ is any plan in $\Psi(m)$ and $B(\psi)$ does not depend on $\tilde{\psi}$, it is clear that $B(\psi)$ bounds the difference between $P_T[\psi]$ and the probability of detection $P_T[\psi^*]$ from the T-optimal plan ψ^*. From Theorem 2, we have that $B(\psi^*) = 0$, but there is no guarantee that $B(\psi_{(T+1)i}) \to 0$ as $i \to \infty$ in the algorithm. However, experience shows that the bound approaches 0 very quickly.

Washburn's bound has been generalized by Stromquist and Stone, and Pursiheimo (1980).

ACKNOWLEDGMENTS

The idea of modifying the existing algorithms that find optimal plans for moving targets in order to solve the optimal survivor search and the multistate target problems arose out of discussions between J. H. Discenza and H. R. Richardson at the NATO Advanced Research Institute on Search Theory and Applications in March 1979. S. S. Brown and W. R. Stromquist helped to extend the existing algorithms and necessary and sufficient conditions to multistate target problems.

This work received support from the Office of Naval Research, Contract No. N00014-76-C-0696.

The opinions in this paper are those of the authors and do not necessarily reflect those of the Commandant or the United States Coast Guard.

REFERENCES

BROWN, S. S. 1980. Optimal Search for a Moving Target in Discrete Time and Space. *Opns. Res.* **28,** 1275–1289.

PURSIHEIMO, U. 1980. Conjugate Duality in the Optimization of the Search for a Target with Generalized Conditionally Deterministic Motion. *J. Optim. Theory Appl.* **31,** 261–275.

RICHARDSON, H. R., AND B. BELKIN. 1972. Optimal Search with Uncertain Sweep Width. *Opns. Res.* **20,** 764–784.

RICHARDSON, H. R., AND J. H. DISCENZA. 1980. The United States Coast Guard Computer-Assisted Search Planning System (CASP). *Naval Res. Logist. Quart.* **27,** 659–680.

STONE, L. D. 1975. *Theory of Optimal Search.* Academic Press, New York.

STONE, L. D. 1979. Necessary and Sufficient Conditions for Optimal Search Plans for Moving Targets. *Math. Opns. Res.* **4,** 431–440.

STONE, L. D., S. S. BROWN, R. P. BUEMI AND C. R. HOPKINS. 1978. Numerical Optimization of Search for a Moving Target. Daniel H. Wagner, Associates, Paoli, Pa., Report to Office of Naval Research.

STROMQUIST, W. R. AND L. D. STONE. 1980. Constrained Optimization of Functionals with Search Theory Applications. *Math. Opns. Res.* **6** (Nov.).

WASHBURN, A. R. 1981. An Upper Bound Useful in Optimizing Search for a Moving Target. *Operations Research* (to appear).