A search game for optimizing information collection in UUV mission planning

John G. Baylog and Thomas A. Wettergren
Naval Undersea Warfare Center
1176 Howell Street
Newport, RI 02841
john.baylog@navy.mil
t.a.wettergren@ieee.org

***Abstract*—This paper addresses search planning to find hidden objects in undersea environments where the sensor detection process is subject to false alarms with a geographically varying likelihood. We develop a game-theoretic approach for maximizing the information flow that occurs as a multi-agent collaborative search is conducted over a bounded region. To accomplish this, we apply 1) a search channel formalism to the discrete search paradigm to formulate information measures as a function of searcher regional visitation and 2) a Receiver Operator Characteristic (ROC) analysis to map search strategies to ROC operating points which, in turn, define the properties of the search channel. This enables us to formulate a search game where the cost is driven by expended search effort and the payout is in terms of the information collected.**

We describe a new algorithmic capability for solving the search allocation problem of multiple unmanned undersea vehicles that can collaborate in their efforts to resolve false alarm outcomes. We discuss the properties of the information measure developed for a repeated look over the search channel and examine the impact of setting ROC operating points to their game equilibrium values. Doing so enables the regional visitation to be determined via a greedy scheduling solution. Results are presented to show the dominance of this approach over alternative ROC strategies and to demonstrate the capability to take advantage of known non-homogeneity in the search environment.

I. Introduction

The undersea environment is often non-homogeneous with spatial variability in search performance occurring due to oceanic, geologic, biologic and operational influences [1]. The planning of collaborative search operations in such environments remains a topic of contemporary research with advances addressing various aspects of the scheduling problem as in [2], [3], [4], [5]. A good search plan is one that maximizes the efficiency of the search, expediting the discovery of the intended search objects while minimizing the number of search agents required to do so. Often, the impact of false alarms are not considered in determining the overall effectiveness of the planning strategy. Yet false alarms compromise the efficacy of a search plan by disrupting search agent capabilities to identify undersea objects of interest. Detection thresholds can be set to minimize false alarm occurrences but doing so increases the corresponding probability of missed detection [6]. When classification of acoustic detections is uncertain, additional search effort must be applied beyond the established plan so to confirm or deny the nature of an acoustic contact.

This paper addresses search planning under false alarm conditions by applying a game theoretic approach to the optimization of search channels. Search channels were introduced in [7] and serve to develop bounds on achievable search performance. The search channel has information theoretic properties analogous to that of the communications channel. As in the communications channel, the objective is to maximize the information throughput of the channel [8].

The execution of a search plan can be thought of mathematically as a sequential game problem, in a very similar manner to the way that many parlor games have been studied over the past century [9]. The game theory solution seeks to find the equilibrium solution of the game rather than the optimal solution to a fixed objective of maximizing value-over-cost. Such an equilibrium solution is desirable in uncertain situations since it treats the uncertainty as a decision process of an unknown second player, and thus the solution is robust with respect to that player's decision. Such games have been extensively applied to many military operations [10] and are often referred to as a *two player game with complete information*. In the search theoretic context, such games have been studied as *search games* [11] by considering the second player to be the object of search. In the search context, the object of search can be positioned at various places according to the desire of the second player to stay hidden from the search. This notion of a second player captures the uncertainty that is inherent in the search process. When the search object is moving (again in an uncertain manner and/or direction), the problem is referred to as a *rendezvous game*. Both the search game and the rendezvous game are solved with similar computational mathematical procedures, and thus may be treated similarly in research studies.

In the search game, the allowable degrees of freedom afforded to each player must be articulated explicitly. We do this by devising games where channel characteristics are selected based upon a standard Receiver Operating Characteristic (ROC) analysis. ROC analyses has been applied extensively in undersea sonar applications [12]. ROC curves establish the relationship between the probability of detecting the actual objects of search (i.e., true positives) and the related probability of declaring nuisance detections such as clutter as valid (i.e., false positives). These relations can be determined experimentally for incompletely specified systems or theoretically when applying specific system models. Analytical development of ROC representations is well established and there are several good references [13], [14].

This paper is organized as follows. In section II we describe the discrete search paradigm, its related search channel realization and the formulation of search games from these formulations. In section III we develop the calculation of search channel information measures as they pertain to restricted and unrestricted game play. We also describe an algorithm for the simultaneous determination of ROC operating points and the searcher visitation schedules that correspond to them. lastly, in section IV we present results that confirm the properties of the search game equilibrium points as they relate to the optimization algorithm. Concluding remarks follow.

II. The Discrete Search Channel

The information flow that occurs as the undersea search domain is interrogated by a number of autonomous searchers is not unlike the message transfer of a communications channel. The message to be conveyed is the location of the objects of search. It is unknown which positions in the search space will be occupied by objects of interest but there may be knowledge of the relative likelihood of positions. In order to develop measures of the information flow through a search channel, first we must define the probability space over which the information measure is developed.

A. Cell Visitation Model

In our search paradigm, the region to be searched is partitioned into a finite set of subregions or cells that are exclusive of one another and exhaustive over the space.

$$\mathcal{S} = \bigcup_{\ell} C_\ell,\ C_\ell \cap C_m = \emptyset,\ \ell \neq m \tag{1}$$

A search event space is constructed from the partition. Given the position $x \in \mathcal{S}$ of point objects, each cell is subject to a binary hypothesis on its occupancy state. For each cell $C_\ell, \ell = 1, \ldots, L$ let $\mathcal{H}_{1_\ell} : \exists x \mid x \in C_\ell$, denote the hypothesis that point objects have been placed within the cell. Let $\mathcal{H}_{0_\ell} : \forall x \mid x \notin C_\ell$ denote the null hypothesis that the cell is unoccupied. The number of objects within the cell, N_ℓ, can be arbitrarily specified. However, for ease of discussion, we presume that cell size is sufficiently small relative to the placement density such that $\frac{Pr(N_\ell \geq 1)}{Pr(N_\ell = 1)} \rightarrow 1$. We also presume that object placement is independent of any other search event with placement probability $P_{C_\ell} = Pr(x \in C_\ell)$.

The scheduling undertaken in developing a search plan specifies the search paths for each of the agents conducting the search. In lieu of specific search trajectories we develop cell visitation sequences that designate how much search effort is applied to each cell in the collection. This allows us to remain agnostic with regard to platform capabilities and constraints, yet articulate and exploit any natural variations that occur in performance due to placement, sensory and environmental factors. The underlying cell decomposition of the search space and an example uniform coverage search path is illustrated in figure 1.

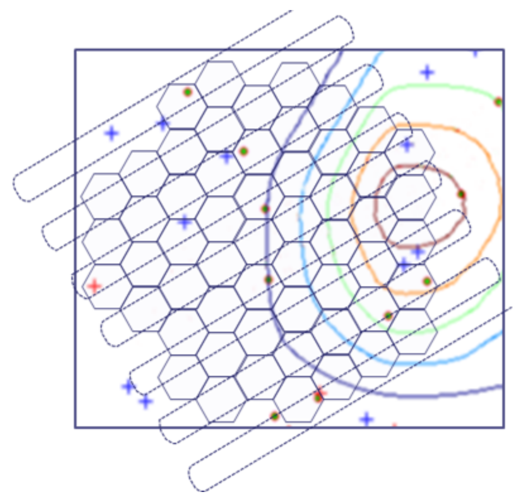

Fig. 1. Cell visitation sequence optimization

We employ a standard gauss-gauss channel model for energy detection within a cell [15]. If there are no objects of interest in the cell, a null hypothesis maps to a signal model whereby the energy detected by the sensor is comprised of noise alone, $\mathcal{H}_{0_\ell} : Z = W$. A valid detection implies that the cell is occupied and maps to a model comprised of acoustic signals plus noise, $\mathcal{H}_{1_\ell} : Z = S + W$. Both signal and noise within the cell are modeled as stationary zero mean white Gaussian processes.

Let γ_ℓ denote a detection threshold in the log likelihood ratio space with sufficient statistic $T(Z) = \sum_{i=1}^{N} z_i^2$ where $E\{z_i^2\} = \sigma^2$ under $\mathcal{H}_{0_\ell}$ and $E\{z_i^2\} = \sigma^2 + \sigma_s^2$ under $\mathcal{H}_{1_\ell}$. Signal-to-noise ratio $SNR_\ell = \sigma_s^2/\sigma^2$

is presumed to vary over the search region. Let $\breve{x}$ denote the position within the cell where the false alarm feature occurs. The probability of false alarm is calculated over the cell by the measuring the content of the spurious distribution above the detection threshold. Following Kay [14], this becomes

$$\begin{aligned} P_{F_\ell}(\gamma_\ell) &= \int_{C_\ell} Pr(d = 1 \mid \breve{x}; \gamma_\ell) f(\breve{x} | x \notin C_\ell) d\breve{x} \\ &\approx Q_{\chi^2_N}(\acute{\gamma_\ell}) \end{aligned} \tag{2}$$

where $\acute{\gamma_\ell} = \gamma_\ell / \sigma^2$ denotes a normalized detection threshold and

$$Q_{\chi^2_\nu}(x) = \int_x^\infty \frac{1}{2^{\frac{\nu}{2}} \Gamma(\frac{\nu}{2})} t^{\frac{\nu}{2}-1} \exp(-\frac{1}{2} t) dt \tag{3}$$

denotes the integral of the χ^2 distribution with ν degrees of freedom. The probability of a search pass detection within the cell becomes

$$\begin{aligned} P_{D_\ell}(\gamma_\ell) &= \int_{C_\ell} Pr(d \mid x; \gamma_\ell) f(x | x \in C_\ell) dx \\ &= Q_{\chi^2_N}\left(\frac{Q^{-1}_{\chi^2_N}(P_{F_\ell})}{SNR_\ell + 1} \right). \end{aligned} \tag{4}$$

Equation 4 has been formulated to allow a ROC curve analysis to be conducted without explicitly solving for detection threshold γ_ℓ. We simplify the calculation by presuming the N-dimensional message vector is comprised of *i.i.d.* components.

B. Search Channel Messaging

Search channel information flow follows Shannon's model of the communications channel [16]. The intended message, $\mathbf{M} = M_1 M_2 \ldots M_L$, is comprised of the placement status over the collection of cells. The collaborative search is conducted so as to replicate this message in an estimate $\hat{\mathbf{M}}$. To accomplish this, a search plan is devised that specifies the number of visits that the searchers make over the respective cells. The planned visitations comprise an encoding scheme whereby the repeated look introduces a redundancy that mitigates uncertainty in search outcomes. The criterion imposed on the observed detection outcomes, $D_\ell = d_1 d_2 \ldots d_{n_\ell}$ represents error correction within the decoding framework. In this way, search planning embodies the block coding of the communications channel.

The information flow through the search channel can be maximized by judicious selection of the controllable parameters. In this paper, we are concerned primarily with the simultaneous optimization of the cell visitation sequence and the governing probabilities of detection and false alarm specified via ROC threshold setting. These are illustrated in figure 2 with the selected operating point depicted in figure 2a defining the channel properties in figure 2b.

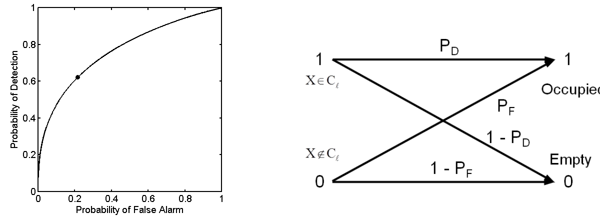

a) ROC threshold b) binary asymmetric channel

Fig. 2. Regulating the search channel

When the search is comprised of a single pass over a cell, $D_\ell = d_1$. From [7], the mutual information between cell search channel input and output takes the form

$$I(M_\ell; \hat{M_\ell}) = I(M_\ell; D_\ell) = H(D_\ell) - H(D_\ell \mid M_\ell) \tag{5}$$

where the placement conditional detection entropy is given by

$$\begin{aligned} H(D_\ell \mid M_\ell) &= P_{C_\ell} H(D_\ell \mid x \in C_\ell)) \\ &+ (1 - P_{C_\ell}) H(D_\ell \mid x \notin C_\ell). \end{aligned} \tag{6}$$

The conditional entropy components are given by

$$H(D_\ell \mid x \in C_\ell) = -P_{D_\ell} log_2 P_{D_\ell} - (1 - P_{D_\ell}) log_2 (1 - P_{D_\ell}) \tag{7}$$

$$H(D_\ell \mid x \notin C_\ell) = -P_{F_\ell} log_2 P_{F_\ell} - (1 - P_{F_\ell}) log_2 (1 - P_{F_\ell}) \tag{8}$$

and the marginal detection entropy is given by

$$H(D_\ell) = -P_{d_\ell} log_2 P_{d_\ell} - (1 - P_{d_\ell}) log_2 (1 - P_{d_\ell}) \tag{9}$$

where the marginal probability of detection takes the form $P_{d_\ell} = P_{C_\ell} P_{D_\ell} + (1 - P_{C_\ell}) P_{F_\ell}$.

C. Search Games

At each step within the search, the planner must decide which cells to search with which sensors. Each search pass comprises an independent Bernoulli trial in which, if the cell is occupied, will present a lottery as to whether the object will be detected or not. If the cell is unoccupied, a similar lottery will be presented as to whether a false alarm will occur on not. Together, the two comprise a compound Bernoulli trial that determines if a detection event is to occur.

The clearance game involves selecting among search alternatives for a bank of parallel cell search channels. The objective of the game is to find every object of interest using the allowable amount of search effort that can be applied. Object placement within any given cell is independent of the placement events that occurs in other cells. The search channel and its player action is depicted in figure 3.

Using a ROC analysis to map game play to search performance, the varying levels of game complexity can be summarized as follows.

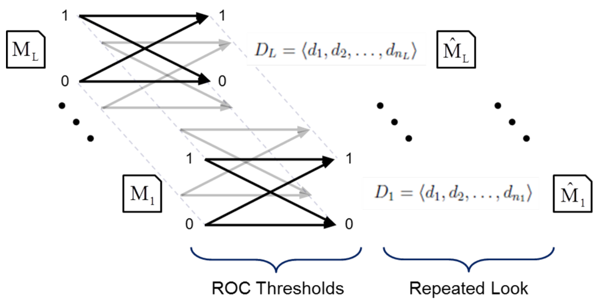

Fig. 3. The clearance channel search game

- Cell visitation: apply a single ROC with a fixed operating point over all the cells and determine the best cell visitation sequences.
- Scalar threshold: apply a distinct ROC over each cell and solve jointly for cell visitation sequences and a single false alarm probability applied to all cells.
- Vector threshold: apply a distinct ROC over each cell and solve jointly for cell visitation sequences and the set of false alarm probabilities, one for each cell in the collection.
- Matrix thresholds: apply a distinct ROC over each cell and solve jointly for cell visitation sequences and the set of false alarm probabilities, one for each cell search pass conducted.
- Matrix thresholds with searcher diversity: apply multiple ROCs per cell and solve jointly for the optimal searcher-cell visitation sequences and the set of false alarm probabilities, one for each searcher cell pass combination.

Let S denote the set of all possible search strategies with $s_i \in S$ denoting a specific strategy. A *Nash* equilibrium point is any strategy, $s_i^* \in S$ such that no other strategy available to the search planner will produce a greater information payout while all other players keep their strategies unchanged.

III. Information Payout

The information gain that a search plan provides will, of course, depend upon how much search effort the planner is willing to extend in order to achieve a successful search. In undersea search applications, the available search effort is typically finite. Hence, we devise two categories of search games.

- Fixed search effort: Every search plan applies the same number of cell search passes and no information discounting is necessary.
- Variable search effort: The number of allowable cell search passes may change from search plan to search plan. Information discounting is applied beyond a minimal search effort to model the diminishing return of untimely information gain.

For this paper, we address the fixed search effort clearance game only.

A. Waiting Time Information Measures

When the same probabilities are applied for each search pass, the complexity of the search game calculations can be reduced by developing the information measure as a function of waiting time probabilities. Waiting time denotes the number of trials that must be attempted before the search event occurs. Let $(W_r = n)$ denote the waiting time event that the r^{th} detection outcome occurs during search pass n. The probability of this event is given by the negative binomial distribution [17].

$$Pr(W_r = n) = \begin{pmatrix} n-1 \\ r-1 \end{pmatrix} p^r (1-p)^{n-r}. \tag{10}$$

Let $P_{\delta_\ell}(n, r) = Pr(W_r = n)$ denote the waiting time probability in cell C_ℓ. If the outcome of each Bernoulli trial is independent of the other outcomes, then the cumulative probability of detection (CPD) in cell C_ℓ, denoted as $P_{\Sigma_\ell}(n, r) = Pr(W_r \leq n)$, can be calculated directly as the sum of the search pass probabilities,

$$P_\Sigma(n, r) = \sum_{i=1}^{n} P_\delta(i, r). \tag{11}$$

As indicated in figure 2, separate waiting time events occur according to which state the placement lottery produces. False alarms cannot occur if the cell is occupied and missed detections cannot occur if the cell is unoccupied. Let $P^d_{\Sigma_\ell}(n, r)$ denote the cumulative probability of making r valid detections during the search of cell C_ℓ with

$$P^d_{\Sigma_\ell}(n, r) = Pr(W_r \leq n \mid x \in C_\ell) = \sum_{i=1}^{n} P^d_{\delta_\ell}(i, r). \tag{12}$$

Let $P^f_{\Sigma_\ell}(n, r)$ denote the cumulative probability of incurring r false alarms during the same cell search with

$$P^f_{\Sigma_\ell}(n, r) = Pr(W_r \leq n \mid x \notin C_\ell) = \sum_{i=1}^{n} P^f_{\delta_\ell}(i, r). \tag{13}$$

Let $P_{\Sigma_\ell}(n, r)$ denote the marginal probability that r detections of indeterminate origin will occur during an n pass search of the cell. Then, with P_{C_ℓ} denoting the placement probability, the marginal waiting time probability becomes

$$P_{\Sigma_\ell}(n, r) = P_{C_\ell} P^d_{\Sigma_\ell}(n, r) + (1 - P_{C_\ell}) P^f_{\Sigma_\ell}(n, r). \tag{14}$$

Within each cell search channel, the repeated look an an n pass search yields n waiting time probabilities for each placement state. With these and the marginal probabilities derived from them, the cell information can be calculated

$$
\begin{aligned}
I(M_\ell; \hat{M}_\ell) &= I(M_\ell; W_{1_\ell}) \\
&+ \sum_{r=2}^{n} Pr(W_{r_\ell - 1} \leq n) I(M_\ell; W_{r_\ell}). \quad (15)
\end{aligned}
$$

The set of waiting time probabilities comprise a sufficient statistic that can be mapped one-to-one to the sufficient statistic comprised of the set of detection number probabilities, $\{\{Pr(N_r)\}_{r=0}^{n}$ where $N_r = \sum_{i=0}^{n} d_i$ as developed in [18]. The mutual information for the bank of cell search channels reduces to the sum of the cell information measures.

$$
I(\mathbf{M}; \hat{\mathbf{M}}) = \sum_{\ell=1}^{L} I(M_\ell; \hat{M}_\ell) \quad (16)
$$

Following Shannon's convention for measuring information [16], we apply the logarithm where the base is determined by the size of the input alphabet. So, we use log_2 for the cell channels and normalize equation 16 by dividing by $log_2(2^l) = L$, the number of cells in the collection.

B. Unconstrained Game Play

The search game complexity increases when detection and false alarm probabilities are allowed to vary over search pass. Such action may be warranted if a definitive information source can be applied to resolve conflicting search outcomes, as in [7]. For this case, waiting time probabilities cannot be calculated using equation 10. However, a probability recursion can be developed directly from the Bernoulli trial event associated with the search pass. For example, the rth occurrence yields the event equivalence

$$
\begin{aligned}
(W_r \leq n) \quad \equiv \quad & (W_r \leq n-1) \vee \\
& (W_{r-1} \leq n-1) \wedge (d_n = 1) \quad (17)
\end{aligned}
$$

Manipulating the corresponding event probabilities yields the following recursive formulation for the calculation. For the first occurrence, $r = 1$,

$$
P_\delta(n, r) = p(n) \left[1 - P_\Sigma(n-1, r)\right] \quad (18)
$$

and for subsequent occurrence, $r > 1$,

$$
P_\delta(n, r) = p(n) \left[P_\Sigma(n-1, r-1) - P_\Sigma(n-1, r)\right]. \quad (19)
$$

As above, recursions are conditioned on placement state. Hence, $p(n) = P_{D_\ell}(n)$ when conditioning on the cell being occupied and $p(n) = P_{F_\ell}(n)$ when the cell is presumed unoccupied. Cumulative probabilities follow from equations 12 and 13.

The cumulative probabilities of valid detection and false alarm remain useful measures of search performance. However, for the general unconstrained game, they are no longer sufficient statistics and cannot be used to calculate the channel information. To see this, consider the posterior probability of cell placement given the detection outcome sequence $D_\ell = \langle d_1, d_2, \ldots, d_{n_\ell} \rangle$ following an n pass search of the cell.

$$
Pr(x \in C_\ell \mid D_\ell) = \frac{Pr(D_\ell(n_\ell), x \in C_\ell)}{Pr(D_\ell, x \in C_\ell) + Pr(D_\ell, x \notin C_\ell)} \quad (20)
$$

where, given that the repeated looks are conditionally independent,

$$
Pr(D_\ell, x \in C_\ell) = \left[\prod_{i=1}^{n_\ell} Pr(d_i \mid x \in C_\ell)\right] P_{C_\ell} \quad (21)
$$

and

$$
Pr(D_\ell, x \notin C_\ell) = \left[\prod_{i=1}^{n_\ell} Pr(d_i \mid x \notin C_\ell)\right] (1 - P_{C_\ell}) \quad (22)
$$

The information through the search channel for the repeated look becomes

$$
I(\mathbf{M}; \hat{\mathbf{M}}) = \sum_{\ell=1}^{L} \left[H(M_\ell) - H(M_\ell \mid D_\ell)\right] \quad (23)
$$

where the entropy of the placement posterior is calculated as

$$
\begin{aligned}
H(M_\ell \mid D_\ell) = -\sum_{\{D_\ell\}} Pr(D_\ell) \cdot \\
\left[\sum_{\{M_\ell\}} Pr(M_\ell \mid D_\ell) log_2(Pr(M_\ell \mid D_\ell(n_\ell)))\right] \quad (24)
\end{aligned}
$$

When probabilities are allowed to vary over search passes, each becomes an independent variable in equation 21 and 22. The degrees of freedom cannot be reduced. Hence, equations 23 and 24 must be applied directly in the information calculation.

C. Visitation Sequence Optimization

To optimize cell visitation sequences, it is useful to calculate how much information will be gained as additional search passes are planned for the cell. Denoting the preceding detection outcome sequence as $D_\ell(n_\ell - 1) = \langle d_1, d_2, \ldots, d_{n_\ell - 1} \rangle$, then from equation 23, the information gain in cell C_ℓ is given by

$$
\begin{aligned}
\Delta I(M_\ell; D_\ell) &= I(M_\ell; D_\ell(n_\ell)) - I(M_\ell; D_\ell(n_\ell - 1)) \\
&= I(M_\ell; d_n) - I(d_n; D_\ell(n_\ell - 1)) \quad (25)
\end{aligned}
$$

$I(M_\ell; d_n)$ represents the information gained during search pass n independent from the preceding search passes. $I(d_n; D_\ell(n_\ell - 1))$ indicates the mutual information between the current search pass and all the preceding passes. Equation 25 indicates how much new information is brought to bear by adding one more pass to the cell search.

For each cell with a fixed operating point on the corresponding ROC curve, the information gain must be a monotone decreasing function. This is due to the fact that probability measures are inherently submodular functions. Given the branching property, this must also be true of the information measure. When the operating point is not fixed, the decreasing monotonicity of the incremental information gain is not ensured. However, when the change in operating point is set such as to transition from the $(n-1)th$ pass equilibrium point to the nth pass equilibrium point, the decreasing monotonicity property of the incremental information gain is preserved. The proof of this is straightforward but is beyond the scope of this paper.

The consequence is that a greedy algorithm can be applied to determine which cells are the best to search. The greedy algorithm represents a local search over the collection of cells for the cell with the highest information gain. In this paper, we assume all search agents have the same sensory capabilities. Hence, the greedy sort will be indifferent to searcher order. A depiction of the greedy cell sort is provided in figure 4 below. The test vector is created and sorted. For K search agents, the first K cells in the sorted test vector are assigned to the search agents as no two searchers are allowed to search the same cell at the same time. Upon assignment, the test vector is recreated and the process continues until all the allowable search effort has been assigned.

Pass	C1	C2	C3	C4	C5	C6	C7	C8	C9
1	0.022	0.044	0.036	0.044	0.084	0.044	0.048	0.044	0.022
2	0.018	0.026	0.014	0.026	0.049	0.026	0.019	0.026	0.018
3	0.014	0.016	0.006	0.016	0.028	0.016	0.008	0.016	0.014
4	0.011	0.010	0.002	0.010	0.016	0.010	0.003	0.010	0.011
5	0.009	0.006	0.001	0.006	0.010	0.006	0.001	0.006	0.009
6	0.007	0.003	0.000	0.003	0.006	0.003	0.000	0.003	0.007
Test	0.022	0.016	0.014	0.016	0.016	0.026	0.019	0.026	0.022

Fig. 4. Greedy search over incremental information gain

Our algorithm for maximizing the information flow through the search channel can be summarized as follows. Let N denote the maximum number of search passes allotted to any given cell.

1) For $n = 1, \ldots, N$ and $\ell = 1, \ldots, L$, determine the optimal value for the probability of false alarm, P_{F_ℓ} and its corresponding information gain.

$$I^*(n, \ell) = \max_{P_F} I(M_\ell; D_\ell(n_\ell))$$

2) Develop the payout matrix by calculating the information gain from search pass to search pass

$$\pi(n, \ell) = I^*(n, \ell) - I^*(n-1, \ell)$$

3) Run the greedy algorithm for determining the cell visitation sequences based on the payout matrix

IV. Numerical Results

A set of results are presented to demonstrate the efficacy of the information measure in maximizing search performance. In particular, we demonstrate how an optimized strategy serves to simplify the numerical calculations necessary to determine equilibria. This is shown in a simplified ROC analysis where the strategy space can be viewed directly and in an extended example exhibiting ROC spatial variability.

A. Elementary Search Game Example

To illustrate the properties of the equilibrium point, we construct an elementary search game comprising a search over a single cell. The search is restricted to two search passes. Player action consists of choosing two threshold settings from a nominal ROC curve where equations 2 and 4 are applied with $SNR = 1.65$. The strategy space consists of the cross product of the two allowable threshold ranges, $S = P_{F_1} \times P_{F_2}$. The information measure is developed over the strategy space using equation 23. This is depicted in figure 5 below. Optimization occurs through an exhaustive search of the strategy space.

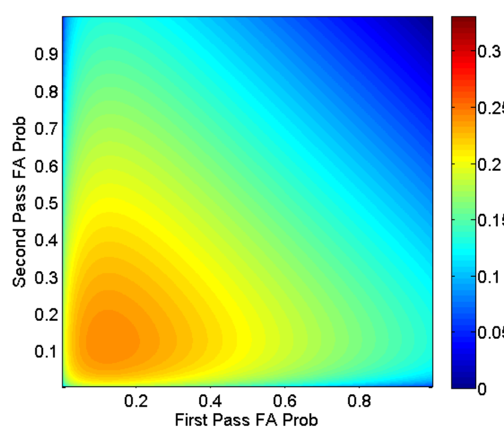

Fig. 5. Mutual information with two degrees of freedom

The symmetry of the information measure over the threshold space is evident. Optimization yields a single equilibrium point with $P_{F_1} = P_{F_2} = 0.14$ and

$$\max_{\{P_{F_1}, P_{F_2}\}} I(M; d_1, d_2) = 0.242. \qquad (26)$$

The commonality of the two thresholds at the equilibrium point means that the information maximum can be calculated as above or by using the waiting time sufficient statistics. This is illustrated below in figure 6 where the information measure is presented as a function of the probability of false alarm threshold. The first pass information values are plotted as the solid black line. For each first pass $P_F = P_{F_1}$, an information maximizing value for P_{F_2} is determined and the resultant information value is depicted as the solid blue line. This curve corresponds to a second pass optimization strategy given a specified first pass threshold setting. The dashed curve corresponds to the information achieved using a single threshold and calculated with the waiting time probabilities.

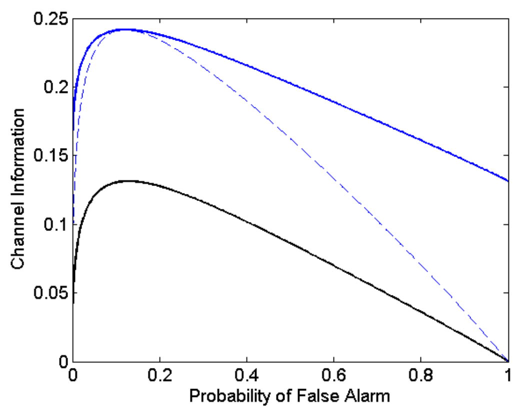

Fig. 6. Properties of the equilibrium point

Clearly, the second pass optimization strategy dominates the single threshold strategy everywhere except at the equilibrium point. To determine the information at the equilibrium point, one need merely iterate along the waiting time information curve until the maximum is attained. However, if the equilibrium point is disallowed, then optimization must be performed over both threshold using the general formulation for search channel information lest the suboptimal solution become obscured.

B. Information Collection Example

To introduce spatial variability into the search channel, we have created a scenario with a $L = 100$ cell region and we vary the SNR randomly across the cells by drawing from a uniform distribution with $0.3 \leq SNR \leq 3.0$. This is an artificial scenario intended only to demonstrate the properties of the algorithm discussed in III-C. Four identical search agents are presumed.

The variation in SNR creates a corresponding variation in ROC curves across the cells. A nominal search strategy is applied whereby search effort is applied uniformly across the search region with each cell receiving $N = 4$ search passes. The probability of false alarm is also set to a constant value over the space. The value applied maximizes the achieved information for the nominal ROC curve with $SNR = 1.65$. This search plan produces the information depicted by the solid black line in figure 7 for the collection of cells. The cells are ordered in order of increasing SNR and this is reflected in the corresponding information values. Also shown is impact of a random variation in the ROC thresholds for each of the 4 passes. Those resulting cell strategies that are dominated by the nominal search plan are indicated by red dots. Those cell strategies that yield information greater than the nominal plan are indicated by blue stars. These appear in cells not in proximity to the nominal ROC curve.

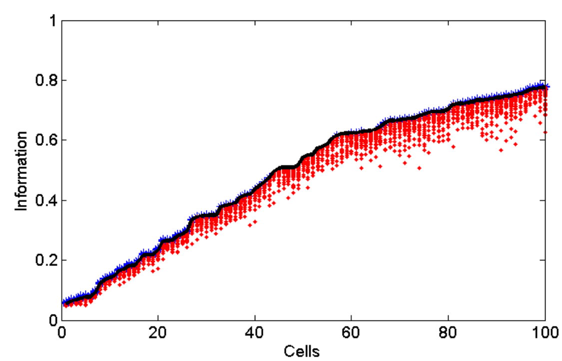

Fig. 7. Nominal uniform coverage search information

Next, we apply the algorithm discussed in section III-C to determine the information gain that is possible with optimization. Figure 8 illustrates the monotone nature of the payout matrix. Each line in the plot represents the cell information gain that is achieved as additional passes are added. All payouts are non-negative and monotone decreasing. Note, however, that the cells that offer the greatest information gain varies with search pass number. The algorithm modifies the search plan. Figure 9 illustrates the resulting search passes devised for each cell relative to the nominal plan. The total search effort is fixed at $4L = 400$ cell visits. The impact of the greedy algorithm is to shift search effort away from the low SNR cell toward those providing the greatest information gain.

The information achieved by the optimized search plan is shown in figure 10. The optimized strategy information is indicated by the solid black line with random variations also depicted. Note, however, that all random strategies are dominated by the optimized strategy. With greedy optimization, the total channel information payout increases from that of the nominal plan, 0.489 in normalized units, to a value of to 0.52 for a modest information gain.

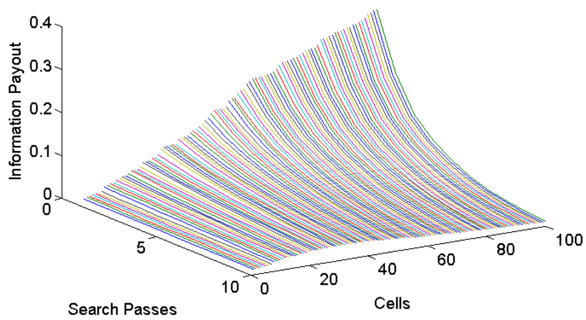

Fig. 8. Monotone nature of the payout matrix

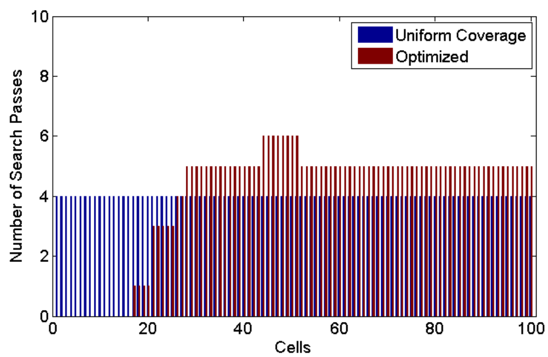

Fig. 9. Greedy cell visitation optimization

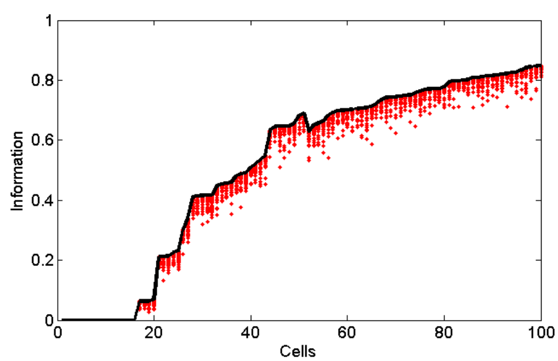

Fig. 10. Optimized information payout

V. Conclusion

In this paper, we have introduced the notion that search games can be constructed from search channel realizations with the game payout developed in terms of the information flow that is created through the channel. We have presented the clearance game directed toward finding an unknown number of search objects over a partitioned search space. We have demonstrated how a ROC analysis serves to map arbitrary game play into a framework suitable for general optimization. In the process, we have uncovered some essential properties of the game equilibrium points. Future work includes a more detailed investigation into the properties of the information measures including the impact of having available multiple ROCs to choose from in generating information.

Acknowledgment

This work was supported by internal (219) funding support at the Naval Undersea Warfare Center.

References

[1] A. R. Washburn. *Search and Detection*. INFORMS, 4th edition, 2002.

[2] J. Song, S. Gupta, and J. Hare. Game-theoretic cooperative coverage using autonomous vehicles. In *Proc. OCEANS 2014*, pages 1–6, Sept 2014.

[3] L. Paull, M. Seto, and J.J. Leonard. Decentralized cooperative trajectory estimation for autonomous underwater vehicles. In *Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Conference on*, pages 184–191, Sept 2014.

[4] N. Abreu and A. Matos. Minehunting mission planning for autonomous underwater systems using evolutionary algorithms. *Unmanned Systems*, 02(04):323–349, 2014.

[5] Clement Petres, Y. Pailhas, P. Patron, Y. Petillot, J. Evans, and D. Lane. Path planning for autonomous underwater vehicles. *Robotics, IEEE Transactions on*, 23(2):331–341, April 2007.

[6] J.G. Baylog and T.A. Wettergren. Multiple pass collaborative search in the presence of false alarms. In *SPIE Defense + Security*, April 2015.

[7] J.G. Baylog and T.A. Wettergren. On-line determination of the potential benefit of path adaptation in undersea search. *IEEE Transactions on Oceanic Engineering*, Vol. 39, 2014.

[8] A. J. Viterbi and J. K. Omura. *Principles of Digital Communication and Coding*. Dover, 2009.

[9] R. B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1991.

[10] R. Isaacs. *Differential Games: A Mathematical Theory with Applications to Warfare and Pursuit, Control and Optimization*. John Wiley and Sons, Inc., 1965.

[11] S. Alpern and S. Gal. *The Theory of Search Games and Rendezvous*. Kluwer Academic Publishers, Norwell, MA, 2003.

[12] R. Urick. *Principles of Underwater Sound*. Mcgraw-Hill, 1983.

[13] H. L. Van Trees. *Detection, Estimation, and Modulation Theory, Part 1*. Wiley-Interscience, 1968.

[14] S. Kay. *Fundamentals of Statistical Signal Processing, Volume II: Detection Theory*. Prentice Hall, 1998.

[15] A. Pezeshki, L.L. Scharf, J.K. Thomas, and B.D. Van Veen. Canonical coordinates are the right coordinates for low-rank gauss-gauss detection and estimation. *Signal Processing, IEEE Transactions on*, 54(12):4817–4820, Dec 2006.

[16] C. E. Shannon. A mathematical theory of communication. *Bell System Tech. J.*, 27:379–423, 623–656, 1948.

[17] G. Grimmett and D. Stirzaker. *Probability and Random Processes*. Oxford Unioversity Press, 2001.

[18] R.O. Duda and P.E. Hart. *Pattern Classification and Scene Analysis*. Wiley, 1973.