

©Copyright 2016

Thomas Powers

Constrained Robust Submodular Sensor Selection with Application to Multistatic Sonar Arrays

Thomas Powers

A thesis submitted in partial fulfillment of the
requirements for the degree of

Master of Science in Electrical Engineering

University of Washington

2016

Committee:

Les Atlas

David W. Krout

Program Authorized to Offer Degree:
Electrical Engineering

University of Washington

Abstract

Constrained Robust Submodular Sensor Selection with Application to Multistatic Sonar Arrays

Thomas Powers

Chair of the Supervisory Committee:
Professor Les Atlas
Electrical Engineering

We develop a framework to select a subset of sensors from a field in which the sensors have an ingrained independence structure. Given an arbitrary independence pattern, we construct a graph that denotes pairwise independence between sensors, which means those sensors may operate simultaneously. The set of all fully-connected subgraphs (cliques) of this independence graph forms the independent sets of matroids over which we maximize the average and minimum of a set of submodular objective functions. The average case is submodular, so it can be approximated. The minimum case is both non-submodular and inapproximable. We propose a novel algorithm called MATSAT that exploits submodularity and, as a result, returns a near-optimal solution with approximation guarantees on a relaxed problem that are within a small factor of the average case scenario. We apply this framework to ping sequence optimization for active multistatic sonar arrays by maximizing sensor coverage for both average and minimum case scenarios and derive lower bounds for minimum probability of detection for a fractional number of targets. In these ping sequence optimization simulations, MATSAT exceeds the fractional lower bounds and reaches near-optimal performance, and submodular function optimization vastly outperforms traditional approaches and nearly achieves optimal performance.

TABLE OF CONTENTS

	Page
List of Figures	ii
Glossary	iii
Chapter 1: Introduction	1
Chapter 2: Submodularity and Independence Graphs	3
Chapter 3: Application to Ping Sequence Optimization	6
Chapter 4: Matroid Constrained SATURATE (MATSAT)	10
Chapter 5: Experimental Results	15
Chapter 6: Conclusion	22
Bibliography	23

LIST OF FIGURES

Figure Number	Page
1.1 General system diagram for sensor selection.	2
3.1 Relationship between the independence graph, coverage regions and interference regions for a four buoy arrangement.	7
4.1 Visualization of lower bound as the saturation level c and the fractional guarantee $\lfloor \gamma M \rfloor$ vary. The submodular guarantee is fixed at $\alpha = \frac{1}{2}$	14
5.1 Probability of detection (PD) results for nine and thirty two buoy Monte Carlo simulations comparing the proposed SFO-GREEDY method to a round-robin strategy for single buoy selection, and SFO-GREEDY to exhaustive search for multiple buoy selection.	17
5.2 Coverage pattern for the nine buoys in a three by three grid and target locations.	18
5.3 Interference regions and independence graph for a nine buoy experiment.	19
5.4 Minimum probability of detection results for an eight target Monte Carlo simulation comparing the proposed MATSAT (dashed blue line) and SFO-GREEDY (dotted magenta line) methods to exhaustive search (solid green line), and $\frac{1}{k+1}$, $k = 4$, submodular guarantee lower bound (solid red line) for buoy selection. The bottom subfigure (b) zooms in on subfigure (a) to emphasize the difference between MATSAT and SFO-GREEDY methods and exhaustive search.	21

GLOSSARY

SUBMODULARITY: a property of set functions that reflects the notion of diminishing returns.

CLIQUE: a completely connected set of nodes of a graph.

STABLE-SET GRAPH: a graph whose vertices are the maximal stable, or independent, sets of a graph.

ACKNOWLEDGMENTS

This work was done in collaboration with Dr. David W. Krout, Professor Les Atlas, and Professor Jeff Bilmes. This work was funded by the U.S. Office of Naval Research, Contract Number N00014-12-G-0078, Delivery Order #24, and the ARO grant number W911NF1510450. I also want to note the helpful discussions with Scott Wisdom and Rishabh Iyer of the University of Washington Department of Electrical Engineering.

DEDICATION

to my wife, Heather

Chapter 1

INTRODUCTION

Subset selection problems are important for many applications in areas such as wireless communications, environmental monitoring, and speech processing [1, 9, 10, 13]. This paper addresses general sensor selection, with a specific focus on sensor networks with interfering sensors, and we demonstrate improved performance on active multistatic sonar arrays [14]. As is the case with most optimization problems, it would be advantageous for the problem to be convex. However, framing sensor selection problems as convex has two main problems. First, sensor selection is inherently a discrete optimization problem since selecting a sensor is an absolute choice. One cannot partially choose a sensor, and convex optimization utilizes continuous variables. Second, convex optimization is unable to handle dependence constraints between variables in a way that is appropriate for our task. Specifically, there is no known way to enforce dependent values between variables in a convex framework such that if a pair of sensors interfere, the two sensors will not be present in the solution together.

Submodular function optimization (SFO) provides a more intuitive framework for handling these two problems, since it inherently uses set functions and can be constrained to optimize over the independence sets of matroids. Matroids are a structure that generalize the notion of linear independence from vector spaces to set systems, and can be used to form constraints in SFO. They will be addressed further in Chapter 2.

The main focus of this paper is a novel optimization algorithm, MATSAT, that is related to SATURATE [13]. MATSAT maximizes the minimum of a set of submodular functions subject to matroid independence constraints rather than cardinality constraints, and as a result offers a new strategy for robust submodular optimization over constraint sets (such as matroids) that are not easily expanded as in the case of a cardinality constraints. We also




Figure 1.1: General system diagram for sensor selection.

investigate modeling sensor networks as graphs and then using the graph structure to form matroid independence sets for use as constraints in submodular function optimization. SFO can handle constraints that make problems nonconvex or non-polynomial (NP) hard and find polynomial time solutions that are provably near-optimal with performance guarantees [6]. Chapter 2 describes how to model sensor network interference patterns as an independence graph can be folded directly into SFO as a series of matroids [3, 4]. This new approach as well as our previously developed SFO-GREEDY approach is applied to scheduling active multistatic deep-water sonar arrays, or ping sequence optimization (PSO), in which we repeatedly optimize to select a subset of buoys that maximize a probabilistic coverage metric. We detail this application in Chapter 3. In order to demonstrate the specific advantages of MATSAT in the worst-case coverage case, we compare its performance to SFO-GREEDY and exhaustive search approaches in Chapter 5. For the average-case coverage scenario, we compare SFO-GREEDY to the standard round-robin sensor selection, greedy single buoy selection, and exhaustive search approaches in Chapter 5 as well.

Chapter 2

SUBMODULARITY AND INDEPENDENCE GRAPHS

The discrete nature of sensor selection makes optimization difficult. Typically, one represents the sensor nodes in an indicator vector with a selected sensor node as ones and unselected sensor nodes as zeros. These independence constraints make optimization problems nonconvex. One of the main contributions of this paper is modeling independence constraints on the sensor networks.

Submodularity is a property that describes set functions similar to how convexity describes functions in a continuous space. For ping sequence optimization, submodular functions can be used to find optimal subsets of buoys to achieve objectives like maximizing coverage of non-interfering buoys, or maximizing probability of target detection in a target tracking scenario. Rather than exhaustively searching over all combinations of subsets, submodular functions provide a fast and tractable framework to compute a solution [6,12,15].

Let the set of available objects, known as the ground set, be denoted as V . A submodular function f maps a set of objects denoted by a binary indicator vector of length V to a real number. The binary indicator vector is represented by the expression 2^V since the variable can take two values and is indexed by elements of set V . As mentioned previously, a value of 1 or 0 for the i^{th} element of the indicator vector denotes the inclusion or exclusion of the i^{th} element of the ground set V , and therefore any subset $A \subseteq V$ can be placed in one-to-one correspondence with incidence vectors.

A submodular function f is defined as one with the following property: for all $A, B \subseteq V$,

$$f(A) + f(B) \geq f(A \cup B) + f(A \cap B) \quad (2.1)$$

In light of the above equivalence between subsets and incidence vectors, a submodular function can also be seen to be one that operates on 0/1-vectors with entries elements indexed by

elements of V . Submodularity is sometimes viewed as a discrete analog to convexity [15], although it should be noted that submodularity and convexity are quite distinct for many reasons.

Submodularity can equivalently, and perhaps more intuitively, be expressed via the notion of diminishing returns, i.e., the incremental gain of the objective diminishes as the context grows. Defining the incremental gain of adding v to A as $f(v|A) = f(A \cup \{v\}) - f(A)$, then submodularity is defined as any function with $f(v|A) \geq f(v|B)$. for all $A \subseteq B \subset V$ and a $v \notin B$.

Submodularity is very closely tied to structures known as matroids, which generalize the notion of linear independence in vector spaces [4]. One can think of matroids as a generalization of matrices, which extend the definition of rank beyond column vectors of a matrix to more general independent subsets over a finite ground set. More importantly, SFO allows for matroid independence constraints to be placed on the problem, which means complicated variable dependence patterns can be encoded into the problem and polynomial time solutions can be obtained. Given a finite set V and a finite set of subsets $\mathcal{I} = \{I_1, I_2, \dots\}$, the pair (V, \mathcal{I}) is said to be a matroid when the family of sets \mathcal{I} satisfies the following three properties:

1. $\emptyset \in \mathcal{I}$
2. $I_1 \subseteq I_2 \in \mathcal{I}$
3. $I_1, I_2 \in \mathcal{I}, |I_1| < |I_2| \implies \exists v \in I_2 \setminus I_1 : I_1 \cup v \in \mathcal{I}$

This leads us to the independence graphs, where nodes on the graph represent sensors and edges denote pairwise independence between sensors. An edge between two sensors, in other words, means both sensors can be used at the same time. For this setup, the nodes in any fully connected subgraph (clique) is an allowable subset. The set of all cliques, also called the clique complex, from this independence graph, denoted by G , can form a partition matroid if the stable sets of the complimentary graph \bar{G} form a partition. This is not true in general for

the independence graphs generated from the interference patterns in a sensor field. However, the clique complex can be modeled as the intersection of k matroids [11]. Kashiwabara et al. give the following algorithm to construct the k matroids.

Algorithm 1 Clique Complex to Matroids

- 1: Construct the stable-set graph $\mathcal{S}(G)$
 - 2: Find a k -coloring of $\mathcal{S}(G)$
 - 3: Group the nodes of each color to form the partitions of one of the k matroids
-


The stable-set graph of G , denoted $\mathcal{S}(G)$, is a graph whose vertices are the maximal stable, or independent, sets of G . An edge exists between vertices that share vertices of G . For the k -coloring problem, the goal is to assign each vertex one of k colors such that adjacent vertices are different colors. In terms of $\mathcal{S}(G)$, vertices of the same color do not share any vertices from G , and thus each vertex forms a partition. Kashiwabara et al. also prove that the clique complex of a graph of n vertices can be represented by the intersection of at most $n - 1$ partition matroids [11]. By turning the interference pattern of a sensor field into a set of matroid constraints, we can guarantee that two interfering sensors will not be chosen in the solution.

Chapter 3

APPLICATION TO PING SEQUENCE OPTIMIZATION

We apply this sensor selection framework to active sonar arrays, where each buoy has a co-located transmitter and receiver that operates monostatically. However, since SFO allows for multiple buoys to be selected, the array functions multistatically in that multiple receivers are operating simultaneously and at potentially overlapping regions. An example of a spatial buoy arrangement where some of the buoys interfere can be found in Fig. 3.1a. The four buoys are arranged in a diamond pattern with locations represented by black dots. In Fig. 3.1a, the blue rings denote the coverage regions for each buoy and the red rings denote the regions where another buoy will interfere with a given buoy. Coverage is defined by the probability of target detection for a buoy. If two interfering buoys transmit simultaneously, the direct path signal from the first will arrive at the second when the second buoy's reflections would arrive. Fig. 3.1a also shows the relationship between the coverage and interference regions for the buoys. In this arrangement, the buoys across from each other, i.e. the top and bottom pair and left and right pair, will interfere with each other, since the buoys in each pair are in the red interference region of the other buoy. However, any other pair of buoys can ping simultaneously [14].

In order to find out the maximum number of buoys that can ping simultaneously, the largest set of nodes is picked such that all the nodes in the set are connected to every node in the set. Note that self-loops are implied, since a buoy does not interfere with itself. The problem of finding the largest subset of fully connected nodes is a well known problem in computer science [2]. Exact methods for solving this problem run in exponential time, but for reasonable graph sizes (a hundred vertices), the algorithm runs fairly quickly. For example, if the graph meets certain conditions, i.e. if the graph is “planar” or “perfect,” finding the



(a) Four buoys (black dots) with coverage regions (blue rings) and interference regions (red rings).

(b) Independence graph for four buoys as nodes with edges that signify pairwise independence.

Figure 3.1: Relationship between the independence graph, coverage regions and interference regions for a four buoy arrangement.

largest clique can be solved in polynomial time [8]. For the arrangement in Fig. 3.1a, there is a four-way tie for largest clique, which are the adjacent pairs (top and left buoys, left and bottom buoys, bottom and right buoys, and right and top buoys). The independence graph for this arrangement is depicted in Fig. 3.1b. In a real scenario, the detection regions will not be perfect rings, so one of the pairs might have better coverage than the others. A more complicated interference pattern will emerge as the number of buoys is increased, which is demonstrated in Fig. 5.3b.

Our objective function is a variant of probabilistic coverage. It utilizes target state estimates to help determine which buoys are selected. Let V be the set of N buoys $b_i, i = 1 \dots N$. Let $B \subseteq V$ such that B is a clique of G , where G is the independence graph determined by the interference pattern of all the buoys b_i in V . Let the set of all sets of sensors that form cliques on the graph be a partition matroid \mathcal{I} . Coverage is a positive, non-decreasing objective, so the goal is to maximize the objective function. Two different problem formula-

tions can be used, which correspond to maximizing the average coverage over a set of target locations [17], Equation (3.1), and maximizing the worst-case coverage over a set of target locations, Equation (3.2). Then the optimal set of buoys for the average-case scenario is given by

$$B^* \in \operatorname{argmax}_{B \in \mathcal{I}} \frac{1}{M} \sum_{i=1}^M f_i(B). \quad (3.1)$$

The worst-case scenario is given by

$$B^* \in \operatorname{argmax}_{B \in \mathcal{I}} \min_i f_i(B). \quad (3.2)$$

In both cases, $i = 1, \dots, M$ corresponds to the predicted target locations and M is the number of targets, and the functions $f_i : 2^V \rightarrow \mathbb{R}$ are given by the equation

$$f_i(B) = 1 - \prod_{b_i \in B} (1 - P_{i,b_i}) \quad (3.3)$$

where P_{i,b_i} is the probability of detection of buoy b_i at location i determined by a table look-up for pre-computed probability of detection maps for each buoy.

Our approach emphasizes tracking in that the objective prioritizes covering areas where known targets are located, but it provides good coverage as well. After the algorithm addresses coverage of the known targets, it adds as many non-interfering buoys as are available, and thus provides an effective simultaneous track and search framework.

The average-case scenario objective function given in Problem (3.1) is submodular. Since the values of P_{i,b_i} in Equation are probabilities, $f_i(B)$ is a concave over modular function, the average of the functions is also submodular. We use the greedy algorithm from the SFO toolbox to solve the above optimization problem [12]. This submodular objective function is monotonic non-decreasing and subject to a matroid constraint. Submodular maximization for functions of this form have been studied and have certain performance guarantees. For a monotonic non-decreasing objective function subject to a matroid constraint, the solution has a worst case performance bound of $\frac{1}{2}$ using the greedy algorithm, and the bound scales

with the number of matroids, i.e. for k matroid constraints the worst case performance bound is $\frac{1}{k+1}$ [5]. The greedy algorithm has a computational complexity of $O(nk)$, where n is the ground set size and k is the size of the solution, but can be significantly sped up [16]. We evaluate this submodular greedy approach in Chapter 5.

Unfortunately, Problem (3.2) is not submodular, since the pointwise minimum of a set of submodular functions is not submodular. Even worse, Problem (3.2) is NP-hard and inapproximable [13]. However, as we explain in the next chapter, an approximate solution can be obtained by relaxing the problem.

Chapter 4

MATROID CONSTRAINED SATURATE (MATSAT)

To optimize the Equation (3.2), we develop a novel algorithm MATSAT, which generalizes the SATURATE algorithm created by Krause et al. [13]. Krause et al. use SATURATE to optimize an objective function of the form

$$A^* = \operatorname{argmax}_{|A| \leq k} \min_i f_i(A) \quad (4.1)$$

where $f_i(A)$ is a set of monotone submodular functions. SATURATE solves this worst-case optimization problem by proposing an alternative formulation and relaxing the cardinality constraint from $|A| \leq k$ to $|A| \leq \alpha k$. As long as α is large enough, the solution \hat{A} from the SATURATE algorithm guarantees that

$$\min_i f_i(\hat{A}) \geq \operatorname{argmax}_{|A| \leq k} \min_i f_i(A) \quad \text{and} \quad |\hat{A}| \leq \alpha k.$$

Krause et al. claim that the only way to achieve a non-trivial guarantee is to relax the constraint, which limits both the types of constraints that can be applied to the problem as well as the values the objective functions can take, i.e. integral or rational valued objective functions. Matroid constraints, for instance, have no immediately obvious relaxation. One way to relax the matroid constraint, however, might be to expand the bases of the matroid. For our task, this leads to undesirable solutions. For example, for ping sequence optimization, the resulting solutions might contain interfering buoys. However, there is another way to achieve non-trivial guarantees, which is to relax the objective itself, leaving the constraints intact, and produce a fractional bound on the objective function, something that is made possible thanks to the use of submodularity. The proposed algorithm, MATSAT, uses this alternative approach to find a solution such that a fraction γ of the submodular functions

Algorithm 2 Greedy (f^c, c)

```

1:  $B \leftarrow \emptyset$ 
2: while  $\exists b \in V \setminus B$  s.t.  $B \cup \{b\} \in \mathcal{I}$  and  $f^c(b|B) > 0$  do
3:    $S \leftarrow \{b \in V \setminus B : B \cup \{b\} \in \mathcal{I}\}$ 
4:    $s^* \in \operatorname{argmax}_{s \in S} f^c(b|B)$ 
5:    $B \leftarrow B \cup \{s^*\}$ 
6: end while
7: return  $B$ 

```

are above a minimum value β . Moreover, the user can set particular values of β or γ , as long as β is less than the submodular guarantee α and $\gamma < 1$. The derivation for the lower bound is given below.

By relaxing the problem, we can consider the following constrained optimization problem:

$$A^* = \operatorname{argmax}_{A \in \mathcal{I}} \min_i f_i(A) \quad (4.2)$$

where \mathcal{I} are the independent sets of a matroid. For a fixed value of c , which can be thought of as the saturation level, we can determine if $f_i(A) \geq c$ via submodular maximization of the following surrogate function:

$$f^c(B) = \frac{1}{M} \sum_{i=1}^M \min\{f_i(B), c\}. \quad (4.3)$$

$f^c(B)$ is a submodular function, because it is a non-negatively weighted sum of functions $\min\{f_i(B), c\}$ which are submodular [7]. At each iteration of MATSAT, we run a greedy algorithm described in Algorithm 2. Basically, the greedy approach selects the buoy which provides the best incremental gain in coverage and does not interfere with any of the other previously selected buoys.

MATSAT is outlined in Algorithm 3. Given monotone submodular functions (f_1, \dots, f_M) , approximation guarantee α for the matroid constrained submodular maximization problem, and tolerance threshold ϵ , we first set c_{min} and c_{max} to values that ensure the true optimal

Algorithm 3 MATSAT $(f_1, \dots, f_M, \alpha, \epsilon)$

```

1:  $c_{min} \leftarrow 0, c_{max} \leftarrow \min_i f_i(V)$ 
2: while  $(c_{max} - c_{min}) > \epsilon$  do
3:    $c \leftarrow (c_{max} + c_{min}) / 2$ 
4:    $f^c(B) \leftarrow \frac{1}{M} \sum_{i=1}^M \min\{f_i(B), c\}$ 
5:    $\hat{B} \leftarrow \text{Greedy}(f^c, c)$ 
6:   if  $f^c(\hat{B}) < \alpha c$  then
7:      $c_{max} \leftarrow c$ 
8:   else
9:      $c_{min} \leftarrow c, B_{best} \leftarrow \hat{B}$ 
10:  end if
11: end while
12: return  $B_{best}$ 

```

value lies in the interval. While performing a binary search over c , we test the value of the approximate solution $f^c(\hat{B})$ against the lower bound αc . If the approximate solution is less than the lower bound, we know that the true optimal is less than c , so we limit the search to the lower half of the interval. Likewise, if the lower bound is met, we store the solution (which, as we describe below, is fractionally good w.r.t. the current c) and then continue attempting to find a better one (higher c) by searching over the upper half of the interval. We stop when the range falls within the tolerance.

Theorem 1. *Given a value $\beta < \alpha$, MATSAT finds a solution \hat{B} that guarantees the following fraction γ of the M functions $\min\{f_i(\hat{B}), c\} \geq \beta c$:*

$$\gamma \geq \frac{\alpha - \beta}{1 - \beta}$$

where α is the approximation guarantee for matroid independence set constrained submodular maximization problem.

Proof $f^c(B) \geq c$ only if $\min_i f_i(B) \geq c$. When all $f_i(B) \geq c$, then $f^c(B) = c$. Likewise, when any $f_i(B) < c$ then the $f^c(B) < c$, since the maximum value of $f^c(B)$ is c . The greedy solution \hat{B} for maximizing a monotone submodular function subject to a matroid constraint is $f^c(\hat{B}) \geq \alpha f^c(B^*)$, where α depends on the algorithm chosen. If line 6 of Algorithm 3 is true, then $f^c(\hat{B}) < \alpha c$ which implies that $\min_i f_i(B^*) < c$. Line 6 being true also implies that c is too large, so $c_{max} \leftarrow c$. At line 12 of Algorithm 3, $(c_{max} - c_{min}) \leq \epsilon$ and the true optimal value $\min_i f_i(B^*)$ is in the interval $[c_{min}, c_{max}]$. The submodular approximation guarantee ensures that $f^c(\hat{B}) \geq \alpha f^c(B^*)$ and $f^c(\hat{B}) \geq \alpha c$. Given a value for β , the terms of $f^c(\hat{B})$ can be split into two groups, one that have value less than βc and the other greater than or equal to βc :

$$f^c(\hat{B}) = \frac{1}{M} \sum_{i: \min\{f_i(B), c\} < \beta c} \min\{f_i(B), c\} + \frac{1}{M} \sum_{i: \min\{f_i(B), c\} \geq \beta c} \min\{f_i(B), c\}.$$

Let γ be the fraction of terms that meet the βc threshold. Then, the two summation terms become $f^c(\hat{B}) = (1 - \gamma) \beta c + \gamma c \geq \alpha c$. Rearranged, the expression becomes $\gamma \geq \frac{\alpha - \beta}{1 - \beta}$. A visualization of the lower bound over values of the saturation level c and fractional number of targets $\lceil \gamma M \rceil$ is shown in Fig. 4.1. In this case, the submodular guarantee for a monotone submodular function constrained by a single matroid via the greedy algorithm is $\alpha = 1/2$.




Figure 4.1: Visualization of lower bound as the saturation level c and the fractional guarantee $\lfloor \gamma M \rfloor$ vary. The submodular guarantee is fixed at $\alpha = \frac{1}{2}$.

Chapter 5

EXPERIMENTAL RESULTS

In this chapter, we design two sets of experiments. In the first set, we compare the performance of the proposed ping sequence optimization algorithms, SFO-GREEDY and MATSAT, to the standard round-robin approach, optimized single buoy selection, and exhaustive search in three Monte Carlo simulations. For the first experiment, we have nine buoys in a grid pattern, with each buoy 60 km away from its neighbors. The independence graph for the buoys can be found in Fig. 5.3b, and the interference pattern in Fig. 5.3a. For the second experiment, we have thirty two buoys in a ring. For all of the experiments, we assume the buoys have a probability of detection of $P = 0.8$ in the coverage region and $P = 0$ everywhere else. We assume here that there is no sensor drift during the experiment. Two targets with random initial location, constrained to be within the buoy array’s detection area, and constant velocities are present for each trial. The first experiment consists of ten thousand trials with each trial lasting thirty time-steps or until a target moved out of the array’s detection area, and the second experiment consists of one thousand trials with each trial lasting sixty-four time-steps or until a target moved out of the array’s detection area.

The results in Table 5.1, Fig. 5.1a, and Fig. 5.1b demonstrate the practical utility of the proposed SFO-GREEDY method. In the nine buoy experiment, a round-robin approach to sensor selection detected the targets about 7% of the time on average, whereas the proposed SFO-GREEDY approach detected the targets over 30% of the time. In addition, the greedy algorithm performed nearly as well as an exhaustive search over all sets of non-interfering buoys. The results also demonstrate the advantage of using multiple buoys. By allowing simultaneous pinging, we gained 50% better coverage over optimized single buoy selection. The low overall probability of detection in Fig. 5.1a can be attributed to the fact that nine

buoys have gaps in their overall coverage, so many simulated targets were impossible to detect given the arrangement. However, in no trial did the round-robin approach beat the proposed algorithm. The worst case bound for the greedy algorithm provides a nice floor for worst-case behavior, but in practice the algorithm is as good as exhaustively searching over exponential growing sets of non-interfering buoys. Likewise, in the thirty two buoy experiment, SFO-GREEDY significantly outperformed the conventional round-robin method and almost matched exhaustive search for average probability of detection.


Table 5.1

Method	Single Buoy		Multiple Buoys	
	Round-robin	SFO-GREEDY (Proposed)	Exhaustive search	SFO-GREEDY (Proposed)
Mean PD, 9 Buoy Exp	0.068	0.210	0.320	0.315
Mean PD, 32 Buoy Exp	0.060	0.505	0.781	0.779


Probability of detection (PD) results for Monte Carlo simulation.

In the second set of experiments, we compare the performance of the proposed ping sequence optimization algorithm MATSAT to the previously proposed SFO-GREEDY and exhaustive search. MATSAT and exhaustive search are optimized with respect to Equation (3.2), while SFO-GREEDY is optimized with respect to Equation (3.1). If SFO-GREEDY were to be optimized with respect to Equation (3.2), the algorithm would always select the empty set, except in the degenerate case where there is a buoy that can detect every single target.

For the third experiment, there are nine buoys in a grid pattern spaced 128 km away from



(a) Probability of detection (PD) results for nine buoy Monte Carlo simulation.



(b) Probability of detection (PD) results for thirty two buoy Monte Carlo simulation.

Figure 5.1: Probability of detection (PD) results for nine and thirty two buoy Monte Carlo simulations comparing the proposed SFO-GREEDY method to a round-robin strategy for single buoy selection, and SFO-GREEDY to exhaustive search for multiple buoy selection.

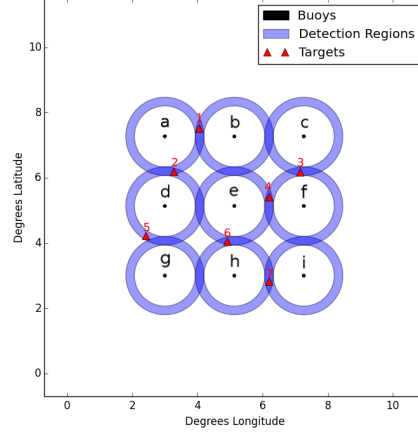



Figure 5.2: Coverage pattern for the nine buoys in a three by three grid and target locations.

their neighbors and seven targets. The nine coverage regions and seven target locations are represented by blue rings the red triangles, respectively, in Figure 5.2. In this experiment, we allow the targets to have unequal probabilities of detection. Ordered from top to bottom and left to right, the targets have probability of detection of $P = \{0.6, 0.6, 0.9, 0.2, 0.6, 0.9, 0.9\}$ in each buoy's coverage region and $P = 0$ everywhere else. The independence graph for the buoys can be found in Fig. 5.3b, and the interference pattern in Fig. 5.3a. By applying Kashiwabara's method for creating a set of matroid constraints to the interference graph in Fig. 5.3b, the clique complex can be represented as the intersection of four partition matroids, whose partitions are given as $\{ad, be, cf, hi, g\}$, $\{ab, de, fi, gh, c\}$, $\{bc, ef, dg, a, h, i\}$, and $\{eh, a, b, c, d, e, f, i\}$.

The results for experiment one can be found in Table 5.2. MATSAT performs as well as the exhaustive search, which is to say it achieved the optimal solution both in terms of worst case and average coverage objectives. SFO-GREEDY, however, chooses the buoys $\{a, f, h\}$ instead of $\{b, d, f, h\}$, which misses one of the targets and results in suboptimal worst case and average coverage.

The fourth experiment tests the fractional worst-case performance of MATSAT against



(a) Interference pattern for the nine buoys in a three by three grid.

(b) Independence graph for the nine buoys in a three by three grid.

Figure 5.3: Interference regions and independence graph for a nine buoy experiment.

SFO-GREEDY and exhaustive search. In this case, fractional performance means the worst-case probability of detection of a fraction γ of the targets. For this experiment, exhaustive search represents the optimal worst-case probability of detection over all of the targets. This experiment has thirty two buoys arranged in a ring, eight targets, and runs for two hundred trials. The rest of the experimental parameters are the same as the second experiment.

The results of the fourth experiment can be found in Fig. 5.4. Both MATSAT and SFO-GREEDY have significantly higher fractional worst-case probability of detection than the lower bound across all fractions of the eight targets $\lfloor \gamma M \rfloor$. Even when the lower-bound provides a trivial guarantee, $\lfloor \gamma M \rfloor \geq 3$, both MATSAT and SFO-GREEDY match the optimal performance. In fact, the performance of MATSAT and SFO-GREEDY are practically indistinguishable from exhaustive search for any fraction of the targets. Fig. 5.4b zooms in on Fig. 5.4a to show the minute differences between the approximate solutions of the proposed algorithms and the true optimal. The only differences are at $\lfloor \gamma M \rfloor < 3$, where MATSAT and

Method	MATSAT	SFO-GREEDY	Exhaustive search
Min PD	0.200	0.000	0.200
Mean PD	0.671	0.586	0.671
Buoys Selected	$\{b, d, f, h\}$	$\{a, f, h\}$	$\{b, d, f, h\}$

Table 5.2: Probability of detection (PD) results for seven targets and nine buoys.

SFO-GREEDY have slightly lower minimum probabilities of detection. While Fig. 5.4b also shows that MATSAT and SFO-GREEDY have equivalent minimum probabilities of detection, MATSAT provides a better theoretical guarantee than SFO-GREEDY. The lower bound is shown by the red line in Fig. 5.4a. The bound reflects the submodular guarantee of the forward greedy algorithm used in the experimental code, which is $\frac{1}{k+1}$ [5]. Since the clique complex can be represented as the intersection of four matroids, the guarantee is 0.2.

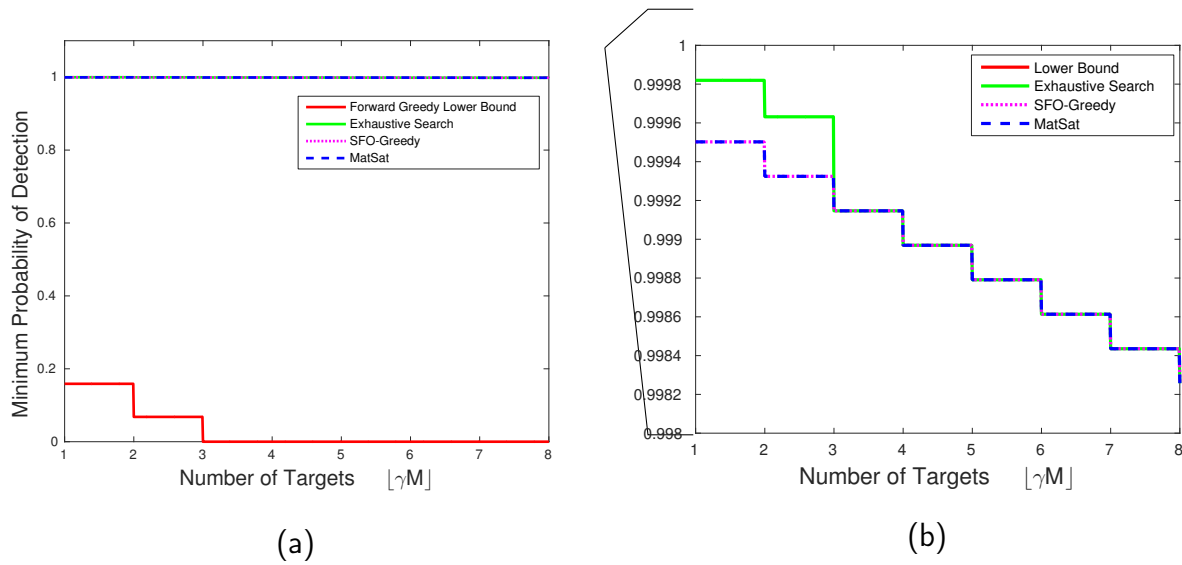


Figure 5.4: Minimum probability of detection results for an eight target Monte Carlo simulation comparing the proposed MATSAT (dashed blue line) and SFO-GREEDY (dotted magenta line) methods to exhaustive search (solid green line), and $\frac{1}{k+1}$, $k = 4$, submodular guarantee lower bound (solid red line) for buoy selection. The bottom subfigure (b) zooms in on subfigure (a) to emphasize the difference between MATSAT and SFO-GREEDY methods and exhaustive search.

Chapter 6

CONCLUSION

In this paper, we propose a new optimization algorithm MATSAT to solve a relaxed worst-case subset selection problem subject to a matroid constraint and derive lower bounds on the average performance as well as fractional worst-case performance. In applying independence graphs to a sensor selection problem, we demonstrate the utility of submodular function optimization (SFO) to the problem domain. Specifically for ping sequence optimization (PSO), SFO allows us go beyond the standard approach for buoy selection by allowing for simultaneous pinging. By posing PSO as a submodular optimization problem or as a discrete problem with submodular structure, we are able to derive near-optimal solutions for combinatorial problems. We can select buoys based on target state that significantly improve the probability of detecting targets over a standard approach and achieve equivalent performance to an optimal exhaustive search approach. Moreover, our approach allows for simultaneous search and track objectives within the system.

BIBLIOGRAPHY

- [1] I.F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci. A survey on sensor networks. *Communications magazine, IEEE*, 40(8):102–114, 2002.
- [2] C. Bron and J. Kerbosch. Algorithm 457: finding all cliques of an undirected graph. *Communications of the ACM*, 16(9):575–577, 1973.
- [3] G. Calinescu, C. Chekuri, M. Pál, and J. Vondrák. Maximizing a submodular set function subject to a matroid constraint. In *Integer programming and combinatorial optimization*, pages 182–196. Springer, 2007.
- [4] J. Edmonds. Matroids and the greedy algorithm. *Mathematical Programming*, 1(1):127–136, 1971.
- [5] M.L. Fisher, G.L. Nemhauser, and L.A. Wolsey. An analysis of approximations for maximizing submodular set functions-ii. In M.L. Balinski and A.J. Hoffman, editors, *Polyhedral Combinatorics*, volume 8 of *Mathematical Programming Studies*, pages 73–87. Springer Berlin Heidelberg, 1978.
- [6] S. Fujishige. *Submodular functions and optimization*, volume 58. Elsevier, 2005.
- [7] Toshihiro Fujito. Approximation algorithms for submodular set cover with applications. *IEICE Transactions on Information and Systems*, 83(3):480–487, 2000.
- [8] M.R. Garey, D.S. Johnson, and L. Stockmeyer. Some simplified np-complete graph problems. *Theoretical computer science*, 1(3):237–267, 1976.
- [9] R.K. Iyer and J.A. Bilmes. Submodular optimization with submodular cover and submodular knapsack constraints. In *Advances in Neural Information Processing Systems 26*, pages 2436–2444. Curran Associates, Inc., 2013.
- [10] S. Joshi and S. Boyd. Sensor selection via convex optimization. *Signal Processing, IEEE Transactions on*, 57(2):451–462, Feb 2009.
- [11] K. Kashiwabara, Y. Okamoto, and T. Uno. Matroid representation of clique complexes. *Discrete Applied Mathematics*, 155(15):1910–1929, 2007.

- [12] A. Krause. SFO: A toolbox for submodular function optimization. *J. Mach. Learn. Res.*, 11:1141–1144, March 2010.
- [13] A. Krause, B. McMahan, Guestrin C., and Gupta A. Robust submodular observation selection. *Journal of Machine Learning Research (JMLR)*, 9:2761–2801, December 2008.
- [14] D.W. Krout and T. Powers. Sensor management for multistatics. In *Information Fusion (FUSION), 2014 17th International Conference on*, pages 1–6, July 2014.
- [15] L. Lovász. Submodular functions and convexity. In Achim Bachem, Bernhard Korte, and Martin Grötschel, editors, *Mathematical Programming The State of the Art*, pages 235–257. Springer Berlin Heidelberg, 1983.
- [16] Michel Minoux. *Accelerated greedy algorithms for maximizing submodular set functions*, pages 234–243. Springer Berlin Heidelberg, Berlin, Heidelberg, 1978.
- [17] T. Powers, D.W. Krout, and L. Atlas. Sensor selection from independence graphs using submodularity. In *Information Fusion (Fusion), 2015 18th International Conference on*, pages 333–337, July 2015.