# Leader Selection Games under Link Noise Injection Attacks

Andrew Clark*, Linda Bushnell†, and Radha Poovendran*
*Network Security Lab (NSL), EE Dept., University of Washington, Seattle, WA, 98195, USA
†Networked Control Systems Lab, EE Dept., University of Washington, Seattle, WA, 98195, USA
Email: {awclark, lb2, rp3}@uw.edu

## ABSTRACT

In a leader-follower multi-agent system, the states of a set of leader agents are controlled directly by the system owner and used to influence the behavior of the remaining follower agents. When deployed in hostile environments, leader-follower systems may be disrupted by adversaries introducing noise in the communication links between agents through interference or false packet insertion, thus corrupting the states of the follower agents. In this paper, we study the problem of mitigating the effect of noise injection attacks by selecting leader agents. We address two cases within a supermodular game-theoretic framework. In the first case, a fixed set of leaders is chosen when the system is initialized. We model this case as a Stackelberg game, in which the system moves first by choosing leaders in order to minimize the worst-case error and the adversary responds by introducing noise. In the second case, the set of leaders varies over time. We study the second case as a simultaneous-move game between the system and an adversary. We show that the game formulations for both cases have equilibria that can be approximated up to a provable bound using supermodular optimization techniques. We illustrate our approach via simulations.

## Categories and Subject Descriptors

J.2 [**Computer Applications**]: Engineering

## General Terms

Theory

## Keywords

Submodular optimization, multi-agent systems, game theory

## 1. INTRODUCTION

Networked multi-agent systems (MAS) are prevalent in a variety of settings, such as formation control of unmanned vehicles [7]. In such systems, each agents receives inputs from its neighbors, performs computations to update its state, and then broadcast its updated state information to its neighbors. An important sub-class of MAS consists of leader-follower systems, in which a set of leader agents are controlled directly by the system owner and influence the remaining agents [8].

In a hostile environment, an adversary can disrupt the performance of a leader-follower system by injecting noise into the communication links between agents. The injected noise corrupts the inputs broadcast from leaders to followers, or between follower agents, causing agents to update their states based on incorrect information. These incorrectly updated states are then broadcast and used as inputs by other agents, causing noise-induced errors to propagate through the system. This noise injection attack can be performed, for example, by broadcasting an interfering signal in the vicinity of the communicating agents.

The effect of noise on a leader-follower system can be mitigated by selecting leader agents to minimize errors due to noise [4]. Leaders that are selected in order to minimize error due to benign, environmental noise, however, may leave the network vulnerable to an attack because an intelligent adversary can observe the set of leaders and inject link noise accordingly. There are two cases for leader selection in MAS in hostile environments. In the first case, a fixed set of leaders is used for the lifetime of the MAS, and hence must be chosen to minimize the worst-case error [8]. In the second case, the agents may be equipped with sensing hardware that allows them to monitor their environment, observe increased noise levels, and update the set of leaders accordingly [9]. Currently, however, there is no analytical approach for selecting leaders in either of these cases.

In this paper, we study the problem of leader selection to mitigate the effects of noise injection attacks. We develop our approach within a two-player game framework, in which the MAS owner selects a set of leaders in order to minimize the mean-square error in the agent states, while the adversary injects noise on a set of communication links in order to maximize this error. We make the following specific contributions:

- We study the problem of leader selection in MAS in the presence of an adversary mounting a link noise injection attack. We study two classes of the leader selection problem: (a) the problem of selecting a fixed set of leaders, and (b) the problem of adaptively choosing leaders in response to an attack.
- We model the selection of a fixed set of leaders as a


*HiCoNSŠ12,* April 17–18, 2012, Beijing, China.

supermodular Stackelberg game, leading to efficient algorithms for approximating the optimal leader set up to a provable bound. As an intermediate step, we prove that the limit of a sequence of supermodular functions and the integral of a collection of supermodular functions are supermodular.

- We formulate a repeated, simultaneous-move game modeling the interaction between the adversary and the MAS for the case where the leader set may change over time. We develop efficient algorithms for approximating a mixed-strategy Nash equilibrium for the game, and provide bounds on its optimality for each player.

- We evaluate the performance of MAS under both models via simulation study. We compare our leader selection methods to other approaches, including random and degree-based leader selection, and show that our scheme leads to lower overall mean-square error in the agent states under link noise injection attacks. We further observe that allowing the leader set to vary over time improves the resilience of the MAS to noise injection.

## 2. RELATED WORK

Current approaches to mitigating link noise injection attacks on leader-follower systems focus on securing the communication protocol used by the agents, or designing the agent state dynamics to be robust to noise. Protocol-based methods, such as frequency hopping, attempt to hide the communication channel used for inter-agent communication and thereby prevent the adversary from injecting noise into the channel [11]. From a control-theoretic standpoint, the state dynamics of the agents can be designed to be robust to noise. In [14], a convex optimization approach to deriving the agent state dynamics in order to minimize the mean-square error due to link noise was proposed. While both protocol-based and control-theoretic methods can be used to improve the resilience of a MAS with given leaders to link noise injection, neither of these methods specifies which leaders should be chosen.

Choosing leader agents to act as control inputs to MAS has been examined in [4, 8]. These approaches are based on MAS operating in the absence of adversaries, and hence may lead to a suboptimal leader set when an intelligent adversary attempts to disrupt the MAS.

The rest of this paper is organized as follows. In Section 3, the system model is presented, along with background on game theory and supermodular functions. In Section 4, we introduce a game-theoretic model for selection of a fixed set of leaders in the presence of adversaries. In Section 5, we formulate a game for the case where the set of leader agents changes in response to adversarial actions. Section 6 presents our simulation results. Section 7 concludes the paper.

## 3. SYSTEM MODEL AND PRELIMINARIES

In this section, the system and adversary models are presented. Needed background information on game theory and the theory of supermodular functions is also given.

### 3.1 System Model

An MAS consisting of $n$ agents, indexed in the set $V = \{1, \ldots, n\}$, is considered. Each agent $i$ is assumed to have a time-varying state $x_i(t) \in \mathbb{R}$. The state variable may represent heading or velocity (in the case where the MAS is a vehicle formation) or sensor measurements. Let $\mathbf{x}(t) \in \mathbb{R}^n$ denote the vector of agent states at time $t$.

A subset $S \subseteq V$, consisting of the leader agents, receives state values directly from the MAS owner. Let $\mathbf{x}_f(t) \in \mathbb{R}^{n-|S|}$ and $\mathbf{x}_l(t) \in \mathbb{R}^{|S|}$ denote the vectors of follower and leader states, respectively. Assume, without loss of generality that the indices are chosen such that $\mathbf{x}(t) = [x_f(t) \quad x_l(t)]^T$. The leader agent states $\mathbf{x}_l(t)$ are input by the MAS owner, while the follower agents update their state values.

Each follower agent $i \in V \setminus S$ receives a relative state value $r_{ij}$ from each neighboring agent $j$, where $r_{ij} = x_i(t) - x_j(t) + \epsilon_{ij}(t)$. $\epsilon_{ij}(t)$ is a white noise process with mean 0 and variance $\nu_{ij}$. The set of neighbors of agent $i$ is denoted $N(i)$. If $j \in N(i)$, then we say a link $(i, j)$ exists. Let $E$ denote the link set. It is assumed that if $(i, j) \in E$, then $(j, i) \in E$ and $\nu_{ij} = \nu_{ji}$. The degree of agent $i$ is defined to be the number of neighbors of $i$, $|N(i)|$.

In order to estimate its correct state value relative to the leader set, agent $i \in V \setminus S$ updates its state according to a best linear unbiased estimator of $x_i$, defined by [2]

$$\dot{x}_i(t) = -D_i^{-1} \sum_{j \in N(i)} \nu_{ij}^{-1} r_{ij}(t) \tag{1}$$

where $D_i = \sum_{j \in N(i)} \nu_{ij}^{-1}$. It is assumed that each agent $i$ has a mechanism to estimate the noise characteristics of each link $(i, j)$ and choose the weights in (1) accordingly.

Define the $n \times n$ matrix $L$ by

$$L = (L_{ij}) = \begin{cases} -\nu_{ij}^{-1}, & (i,j) \in E \\ D_i, & i = j \\ 0, & \text{else} \end{cases} \tag{2}$$

$L$ can be written in the form

$$L = \left( \begin{array}{c|c} L_{ff} & L_{fl} \\ \hline L_{lf} & L_{ll} \end{array} \right) \tag{3}$$

so that the overall dynamics of the follower agents can be written in vector form as

$$\dot{\mathbf{x}}_f(t) = -D^{-1}(L_{ff} x_f(t) + L_{fl} x_l(t)) + W(t) \tag{4}$$

where $W(t)$ is a white process. The following theorem, first proved in [2], describes the mean-squared error of each follower agent's state when the dynamics of (1) are used.

THEOREM 1. *Let $\mathbf{x}^* \in \mathbb{R}^n$ denote the target state of the MAS, defined by $\mathbf{x}^* = \mathbf{x}_r + x_0 \mathbf{1}$, where $\mathbf{x}_r$ is a known constant and $x_0$ is an unknown reference point. Suppose that $\mathbf{x}_s(t) \equiv \mathbf{x}_s^*$ for all $s \in S$. Then for each $u \in V \setminus S$,*

$$\lim_{t \to \infty} \mathbf{E}[(\mathbf{x}_u(t) - \mathbf{x}_u^*)^2] = (L_{ff}^{-1})_{uu} \tag{5}$$

The total mean-square error due to noise is given by the function

$$\chi(S, \nu) = \sum_{u \in V \setminus S} (L_{ff}^{-1})_{uu}. \tag{6}$$

We will analyze two cases of the leader set $S$. In the first case, the leader set $S$ is fixed throughout the system lifetime [8]. In the second case, the leader set $S$ may change over time [9].

## 3.2 Adversary Model

The MAS is assumed to be deployed in the presence of an adversary who is capable of injecting noise on the links between agents by broadcasting an interfering signal in the vicinity of the agents. This noise injection leads to an overall error variance $\nu_{ij} = \nu_{ij}^0 + \nu_{ij}^A$, where $\nu_{ij}^0$ and $\nu_{ij}^A$ are the variances of the error on link $(i,j)$ due to ambient noise and adversarial noise, respectively.

The variance $\nu_{ij}^A$ is equal to the received strength of the interfering signal broadcast by the adversary. The received strength depends on the position of the receiver, denoted $y_j \in \mathbb{R}^3$, the position of the adversary, denoted $z \in \mathbb{R}^3$, the transmit power of the adversary for link $(i,j)$, denoted $P_{ij}$, and the path-loss constant of the propagation medium, denoted $\alpha$. We assume that, in order to avoid detection, the adversary does not choose its position to coincide with any agents, so that $z \neq y_j$ for all $j$. The resulting error variance is given by $\nu_{ij}^A = P_{ij}||y_j - z||_2^{-\alpha}$. It is assumed that the adversary has a constraint $P_A$ on the total power available, so that

$$\sum_{(i,j)\in E} P_{ij} = \sum_{(i,j)\in E} \nu_{ij}^A ||z - y_j||_2^{\alpha} \leq P_A \tag{7}$$

The adversary is assumed to know the network topology and the environmental noise characteristics $\nu_{ij}^0$ for all $(i,j) \in E$. Furthermore, since the leader agents do not follow the dynamics (1), the adversary can determine the leader set $S$ by eavesdropping on the agents' state values and observing which agents do not update their states according to (1).

## 3.3 Background – Game Theory

A game is defined by a set of players $\{\mathcal{P}_1, \ldots, \mathcal{P}_m\}$. Each player $\mathcal{P}_i$ has a set of strategies $\mathcal{S}_i$ and a utility function $U_i : \mathcal{S}_1 \times \cdots \times \mathcal{S}_m \to \mathbf{R}$. The utility function $U_i$ represents $\mathcal{P}_i$'s benefit from its action $s_i \in \mathcal{S}_i$ and the actions of the remaining players. The goal of each player $\mathcal{P}_i$ is to maximize its utility $U_i$.

For simplicity, we restrict ourselves to two-player games ($m = 2$). It is assumed that each player knows the strategy space and utility function of the other player. In addition, in a *Stackelberg* game, player $\mathcal{P}_2$ observes the strategy $s_1 \in \mathcal{S}_1$ chosen by $\mathcal{P}_1$ before choosing a strategy $s_2 \in \mathcal{S}_2$. In order to maximize its utility, $\mathcal{P}_2$ will therefore choose strategy $s_2^* \in \mathcal{S}_2$ that satisfies[1]

$$s_2^* \in \arg\max_{s_2 \in \mathcal{S}_2} U_2(s_1, s_2) \tag{8}$$

Let $s_2^*(s_1)$ denote $\mathcal{P}_2$'s optimal strategy when $\mathcal{P}_1$ chooses strategy $s_1$. $\mathcal{P}_1$ will therefore choose strategy $s_1^*$ that maximizes its utility given $\mathcal{P}_2$'s response:

$$s_1^* \in \arg\max_{s_1 \in \mathcal{S}_1} U_1(s_1, s_2^*(s_1)) \tag{9}$$

By contrast, in a *simultaneous-move game*, the two players choose their strategies at the same time, so that neither player observes the strategy of the other player. In this case, when the players are rational, they will choose their strategies according to a *Nash equilibrium*, defined as follows [1].

[1]When more than one strategy $s_2^*$ satisfies (8), we assume that $\mathcal{P}_2$ chooses the strategy satisfying (8) that minimizes $U_1$.

DEFINITION 1. *A pair of strategies $(s_1^*, s_2^*) \in \mathcal{S}_1 \times \mathcal{S}_2$ is a* Nash equilibrium *if and only if*

$$s_1^* = \arg\max_{s_1 \in \mathcal{S}_1} U_1(s_1, s_2^*) \tag{10}$$

$$s_2^* = \arg\max_{s_2 \in \mathcal{S}_2} U_2(s_1^*, s_2) \tag{11}$$

In words, if $(s_1^*, s_2^*)$ is a Nash equilibrium, then $\mathcal{P}_i$ cannot change its strategy $s_i^*$ without decreasing his utility. Any strategy $s_1^*$ (resp. $s_2^*$) satisfying (10) (resp. (11)) is a *best response* for player $\mathcal{P}_1$ (resp. $\mathcal{P}_2$).

Each player may attempt to improve its performance by randomizing over a set of strategies. This concept is defined as follows.

DEFINITION 2. *A mixed strategy for player $\mathcal{P}_i$ consists of a set of ordered pairs $\{(s_i^{(1)}, p_1), \ldots, (s_i^{(r)}, p_r)\}$, where $s_i^{(j)} \in \mathcal{S}_i$ for all $j$ and the $p_j$'s are nonnegative real numbers that sum to* 1. *Under this strategy, player $\mathcal{P}_i$ chooses strategy $s_i^{(j)}$ with probability $p_j$. A* mixed strategy equilibrium *is a Nash equilibrium in which one or more players uses a mixed strategy.*

## 3.4 Background – Supermodular Functions

Let $V$ be a finite set, and let $2^V$ denote the set of subsets of $V$. A supermodular function is defined as follows.

DEFINITION 3. *A function $f : 2^V \to \mathbf{R}$ is defined to be supermodular if, for any subsets $S, T \subseteq V$ with $S \subseteq T$ and $v \in V \setminus T$,*

$$f(S) - f(S \cup \{v\}) \geq f(T) - f(T \cup \{v\}) \tag{12}$$

Definition 3 can be interpreted as a diminishing returns property of $f$ as a function of $S$. The following two lemmas give methods for constructing supermodular functions [10].

LEMMA 1. *If $f_1(S), \ldots, f_r(S)$ is a set of supermodular functions of $S$ and $\alpha_1, \ldots, \alpha_r$ are nonnegative constants, then $f(S) = \sum_{i=1}^r \alpha_i f_i(S)$ is a supermodular function.*

LEMMA 2. *If $f(S)$ is a supermodular function and $r \geq 0$ is constant, then the function $\hat{f}(S) = \max\{f(S), r\}$ is supermodular.*

The following theorem concerns limits of sequences of supermodular functions, and to the best of our knowledge has not appeared elsewhere in the literature.

THEOREM 2. *Suppose $\{f_k(S)\}_{k=1}^{\infty}$ is a collection of supermodular functions in $S$, and suppose that there exists a function $f : 2^V \to \mathbb{R}$ such that, for every $S \subseteq V$ and every $\epsilon > 0$, there exists $K$ such that $k > K$ implies*

$$|f_k(S) - f(S)| < \epsilon. \tag{13}$$

*Then $f(S)$ is supermodular.*

A proof of Theorem 2 can be found in the appendix.

# 4. PROBLEM FORMULATION – STATIC GAME

In this section, we consider the problem of choosing a fixed set of leaders in the presence of an adversary injecting noise.

## 4.1 Game Definition

In this setting, the MAS owner first chooses a set of up to $k$ leaders. The adversary then observes the set of leaders and chooses a set of error variances $\nu_{ij}^A$ satisfying the adversary's power constraint, given by (7). The goal of the adversary is to choose the vector $\nu^A$ such that the total system error $\chi(S, \nu^0 + \nu^A)$ of (6) is maximized, while the MAS owner's goal is to choose a set of leaders $S$ such that the total error in the worst case is minimized.

We formulate this problem as a Stackelberg game, in which the first player, $\mathcal{P}_1$, is the MAS owner and the second player, $\mathcal{P}_2$, is the adversary. The strategy space $\mathcal{S}_1$ of the MAS owner is given by the set of possible leader sets, $\mathcal{S}_1 = \{S \subseteq V : |S| \leq k\}$. The adversary's strategy space, $\mathcal{S}_2$, consists of the set of feasible error variances (7).

The MAS utility is given by $U_{MAS}(S, \nu^A) = -\chi(S, \nu^0 + \nu^A)$, while the adversary's utility is given by $U_A(S, \nu^A) = \chi(S, \nu^0 + \nu^A)$, so that the MAS owner's goal is to minimize the total error variance, while the adversary's goal is to maximize it. In what follows, we explore the optimal pure strategies for each player in detail, leaving the analysis of possible mixed strategies as future work.

## 4.2 Solution Algorithms for Adversary

For a given leader set, the adversary's goal is to choose the error variances $\nu^A$ such that the total system error $\chi(S, \nu)$ is maximized. Hence the adversary's optimal strategy is given by the solution to the optimization problem

$$\begin{array}{ll} \text{maximize} & \chi(S, \nu^0 + \nu^A) \\ \nu_{ij}^A & \\ \text{s.t.} & \nu_{ij}^A \geq 0 \quad \forall (i,j) \in E \\ & \sum_{(i,j)\in E} \nu_{ij}^A ||z - y_j||_2^\alpha \leq P_A \end{array} \tag{14}$$

The following theorem leads to efficient solution algorithms for (14).

THEOREM 3. *The function $\chi(S, \nu)$ is a concave function of $\{\nu_{ij} : (i,j) \in E\}$.*

The proof of this theorem can be found in the appendix.

COROLLARY 1. *Problem (14) is a concave optimization problem.*

PROOF. The proof follows from Theorem 3 and the fact that the constraints of (14) are convex in $\nu_{ij}^A$. □

As a result, the optimal strategy for the adversary can be computed in polynomial time using interior point algorithms [3].

## 4.3 Solution Algorithms for the MAS Owner

Since the adversary will choose the noise injection strategy that maximizes the error due to link noise, the goal of the MAS owner is to select leaders such that this worst-case error is minimized. The optimal strategy is therefore given as the solution to the optimization problem

$$\begin{array}{ll} \text{minimize} & \max_{\nu^A} \chi(S, \nu^0 + \nu^A) \\ S & \\ \text{s.t.} & |S| \leq k \end{array} \tag{15}$$

Define $\chi(S) \triangleq \max_{\nu^A} \chi(S, \nu)$. Problem (15) involves optimizing over $\binom{n}{k}$ possible leader sets, which is infeasible when $n$ or $k$ is large. Moreover, functions of the form $g(S) = \max_i f_i(S)$, where $f_i(S)$ is supermodular, are not supermodular in general. We instead introduce an equivalent, supermodular formulation and derive a solution algorithm for (15) as a result.

As a preliminary, define $F_\zeta(S)$ by

$$F_\zeta(S) \triangleq \frac{1}{|\mathcal{S}_2|} \int_{\mathcal{S}_2} \max\{\chi(S, \nu^0 + \nu^A), \zeta\}\, d\nu^A \tag{16}$$

where $|\cdot|$ denotes the Lebesgue measure of a set. The function $F_\zeta(S)$ is supermodular as a function of $S$ (see Appendix, Lemma 6). An algorithm for solving (15), based on the supermodularity of $F_\zeta(S)$, is as follows.

First, select parameters $\beta \geq 1$ and $\delta > 0$. The algorithm finds a set $S$ satisfying $\chi(S) \leq \chi^*$, where $\chi^*$ is the optimal value of (15), and $|S| \leq \beta k$.

Define $\zeta_{min}^0 = 0$ and $\zeta_{max}^0 = \chi(\{1\})$, the error $\chi$ corresponding to leader set $S = \{1\}$. At the $j$-th iteration, let $\zeta^j = \frac{\zeta_{max}^{j-1} + \zeta_{min}^{j-1}}{2}$. The goal of the $j$-th iteration is to determine if there is a set $S^j$ satisfying $\chi(S^j) \leq \zeta^j$ with $|S| \leq \beta k$. This is accomplished by solving the optimization problem

$$\begin{array}{ll} \text{minimize} & |S| \\ \text{s.t.} & F_\zeta(S) \leq \zeta^j \end{array} \tag{17}$$

noting that $\chi(S) \leq \zeta^j$ if and only if $F_\zeta(S) \leq \zeta^j$.

Problem (17) is solved as follows. Initialize $S^j = \emptyset$. At each iteration, choose $v^*$ as

$$v^* = \arg\max \{F_\zeta(S^j) - F_\zeta(S^j \cup \{v\}) : v \in V \setminus S\} \tag{18}$$

Set $S^j = S^j \cup \{v\}$. The process continues until $F_\zeta(S^j) \leq \zeta^j$.

If $S^j$ satisfies $|S^j| \leq \beta k$, then set $\zeta_{max}^j = \zeta^j$ and $\zeta_{min}^j = \zeta_{min}^{j-1}$. Otherwise, set $\zeta_{min}^j = \zeta^j$ and $\zeta_{max}^j = \zeta_{max}^{j-1}$. The algorithm terminates when $\zeta_{max}^j - \zeta_{min}^j < \delta$ and returns the set $S^j$. A pseudocode description of this approach is given as algorithm **Fixed-$k$**.

**Fixed-$k$**: Algorithm for selecting a fixed set of up to $k$ leaders under worst-case noise injection attack

**Input:** Maximum number of leaders, $k$
Topology $G = (V, E)$, error variances $\nu^0(i,j)$
Parameters $\beta$ and $\delta$
Node positions $y_j$, $j \in V$, adversary position $p$
**Output:** Set of leaders $S$

$\zeta_{min} \leftarrow 0$, $\zeta_{max} \leftarrow \max_{\nu^A} \chi(\{1\}, \nu^0 + \nu^A)$ 1
$j \leftarrow 0$ 2
**while** $\zeta_{max} - \zeta_{min} \geq \delta$ 3
  $\zeta \leftarrow (\zeta_{min} + \zeta_{max})/2$ 4
  $S^j \leftarrow \emptyset$ 5
  **while** $F_\zeta(S^j) \geq \zeta$ 6
    $v^* \leftarrow \arg\max_{v \in V \setminus S^j} \left\{F_\zeta(S^j) - F_\zeta(S^j \cup \{v\})\right\}$ 7
    $S^j \leftarrow S^j \cup \{v^*\}$ 8
  **end while** 9
  **if** $|S^j| > \beta k$ 10
    $\zeta_{min} \leftarrow \zeta$ 11
  **else** 12
    $\zeta_{max} \leftarrow \zeta$ 13
  **end if**; $j \leftarrow j + 1$ 14
**end while** 15
$S \leftarrow S^j$, return $S$ 16

THEOREM 4. *If $\beta$ satisfies*

$$\beta \geq 1 + \log\left\{\frac{\max_{v\in V} F_\zeta(\{v\})}{\zeta^*}\right\} \tag{19}$$

*then **Fixed**-$k$ returns a set $S$ satisfying $\chi(S) \leq \chi(S^*)$ and $|S| \leq \beta k$.*

The proof can be found in the appendix. Corollary 1 implies that the adversary can compute the best-response to the chosen leader set $S$, while Theorem 4 proves that the strategy chosen by the MAS is within a provable bound of the optimum. These strategies, when taken together, therefore form an approximate Stackelberg equilibrium.

# 5. PROBLEM FORMULATION – REPEATED GAME

We now consider the case where the set of leaders can change over time.

## 5.1 Game Definition

Under this model, the MAS owner periodically updates the leader set $S$ in order to minimize the overall system error $\chi(S,\nu)$, based on the observed noise characteristics $\nu$. The adversary, upon observing a change in $S$, chooses a new noise injection strategy in order to maximize the error experienced by the MAS. This leads to a repeated game model for the interaction between the MAS and the adversary.

Formally, the MAS owner is the first player, $\mathcal{P}_1$, while the adversary is the second player, $\mathcal{P}_2$. At the $t$-th iteration of the game, the MAS owner selects a leader set $S_t$ satisfying $|S_t| \leq k$. The adversary is unaware of the leader set for a time $T$, and hence chooses a vector of link error variances $\nu_t^A \in \mathcal{S}_2$ without knowledge of $S_t$, where $\mathcal{S}_2$ is defined as in Section 4. After time $T$ elapses, the adversary discovers $S_t$ and chooses a new vector of link error variances, $\tilde{\nu}_t^A \in \mathcal{S}_2$. An additional time $T'$ elapses until the next iteration.

The penalty of the MAS at the $t$-th iteration is given by the average system error experienced, so that

$$U_{MAS}(S_t,\nu_t^A,\tilde{\nu}_t^A) = -\left(\frac{T}{T+T'}\chi(S,\nu^0+\nu_t^A) + \frac{T'}{T+T'}\chi(S,\nu^0+\tilde{\nu}_t^A)\right) \tag{20}$$

Similarly, the utility of the adversary is equal to the average system error:

$$U_A(S_t,\nu_t^A,\nu_t^{A\prime}) = \frac{T}{T+T'}\chi(S,\nu^0+\nu_t^A) + \frac{T'}{T+T'}\chi(S,\nu^0+\tilde{\nu}_t^A) \tag{21}$$

In what follows, it is assumed that it takes more time for the adversary to determine the leader set than for the MAS to detect the increase in error due to the noise injection attack, so that $T \gg T'$. Since the adversary and MAS are not aware of each other's strategies during this interval, this is a repeated simultaneous game with $U_{MAS}(S_t,\nu_t^A) \approx -\chi(S,\nu^0+\nu_t^A)$ and $U_A(S_t,\nu_t^A) \approx \chi(S,\nu^0+\nu_t^A)$. We first study the best-response behavior of each player. We then analyze the equilibria of the game based on the best-response behavior.

## 5.2 Best-Response Strategies

We first analyze the best-response strategy for the adversary at each iteration $t$. For a given choice of $S_t$, the adversary's best response to $S_t$ is given by

$$\begin{array}{ll}\text{maximize} & \chi(S_t,\nu^0+\nu_t^A)\\ \text{s.t.} & \nu_t^A \in \mathcal{S}_2\end{array} \tag{22}$$

The function $\chi(S_t,\nu^0+\nu_t^A)$ is a convex function of $\nu_t^A$ by Theorem 3. Hence the adversary's optimal noise allocation $\nu_t^A$ at each iteration can be obtained by solving (22) via convex optimization.

The MAS's problem of choosing the optimal leader set to minimize noise in response to $\nu_t^A$ is formulated as

$$\begin{array}{ll}\text{minimize} & \chi(S_t,\nu^0+\nu_t^A)\\ \text{s.t.} & |S_t| \leq k\end{array} \tag{23}$$

Problem (23) can be solved by supermodular optimization, as shown by the following theorem, proved in [4].

THEOREM 5. *For fixed $\nu$, $\chi(S,\nu)$ is a supermodular function of $S$.*

While minimizing a supermodular function is NP-hard in general, a greedy algorithm can be used to approximate the optimal leader set $S_t$ [10]. In the algorithm, the leader set $S_t = \emptyset$ initially. At each iteration, the agent $v^*$ satisfying

$$v^* = \arg\max\Big\{v \in V \setminus S_t : \chi(S_t,\nu^0+\nu_t^A) - \chi(S_t\cup\{v\},\nu^0+\nu_t^A)\Big\} \tag{24}$$

is added to the leader set. The algorithm terminates after $k$ iterations. A pseudocode description as algorithm **BestResponse**-$k$.

**BestResponse**-$k$: Algorithm for selecting up to $k$ leaders in response to an adversary's strategy

**Input:** Maximum number of leaders, $k$
Graph topology $G = (V, E)$
Link error variances $\nu_t = \nu_t^0 + \nu_t^A$
**Output:** Leader set $S_t$
$S_t \leftarrow \emptyset$, $l \leftarrow 0$ 1
**while** $l < k$ 2
  $v^* \leftarrow \arg\max\{v \in V \setminus S_t : \chi(S_t,\nu_t) - \chi(S_t\cup\{v\},\nu_t)\}$ 3
  $S_t \leftarrow S_t \cup \{v^*\}$, $l \leftarrow l+1$ 4
**end while** 5
**return** $S_t$ 6

THEOREM 6. *Let $S_t^*$ be the optimal solution to (23). Then the set $S$ returned by **BestResponse**-$k$ satisfies*

$$\chi(S_t,\nu_t^0+\nu_t^A) \leq \left(1-\frac{1}{e}\right)\chi(S_t^*,\nu_t^0+\nu_t^A) + \frac{1}{e}\chi_{max} \tag{25}$$

*where $\chi_{max} \triangleq \max_i \sum_{j\in V}\chi(i,j)$.*

PROOF. The proof follows from Proposition 4.1 of [10] and the supermodularity of the error $\chi$. □

## 5.3 Equilibrium Analysis

As discussed in Section 3.3, the MAS owner and the adversary will maximize their utilities by playing a Nash equilibrium strategy at each iteration $t$. In general, determining

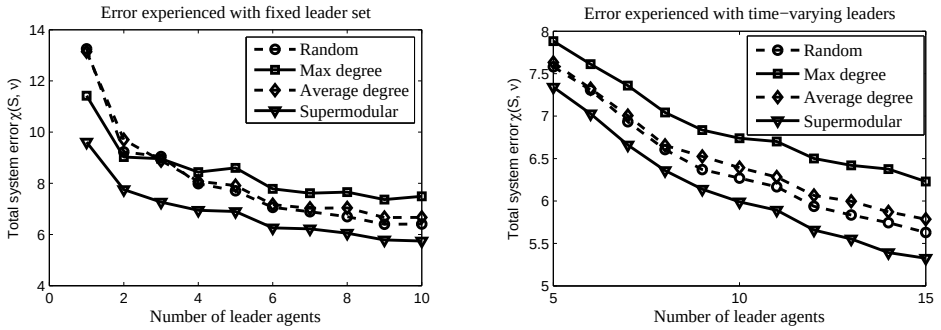


**Figure 1: (a) Simulation of selection of a fixed set of leaders in the presence of adversarial noise using random, degree-based, and supermodular algorithms. While all schemes experience a decrease in performance, the supermodular selection approach provides the most robustness to noise. (b) Simulation of the approximate Nash equilibrium that arises from updating the leader set over time. All schemes outperform the case of fixed leader selection, while the supermodular optimization approach outperforms both random and degree-based approaches.**

Nash equilibria of two-player games is PPAD-complete [12, Ch. 4.1]. Moreover, in this case, the MAS's best-response strategy is NP-hard to compute. Instead, the following algorithm, first proposed for general two-player games [5], can be used to efficiently compute an approximate mixed-strategy equilibrium. We note that there may be other equilibria, which we will characterize in future work.

The MAS first chooses a set of leaders $S$. The adversary computes $\hat{\nu}^A$ by solving the best-response problem of (22) based on the leader set $S$. The MAS approximates the best-response $S'$ to $\hat{\nu}^A$ using **BestResponse**-$k$.

The MAS's strategy is to choose leader set $S$ with probability 1/2 and to choose $S'$ with probability 1/2, corresponding to a mixed strategy $\{(S, 1/2), (S', 1/2)\}$. The adversary's strategy is to choose link error variances $\hat{\nu}^A$ with probability 1. A pseudocode description of the algorithm for approximating a Nash equilibrium is given as algorithm **Approx-NE**.

| **Approx-NE**: Algorithm for computing an approximate Nash equilibrium | |
|---|---|
| **Input:** Maximum number of leaders, $k$ | |
| Graph topology $G = (V, E)$ and variances $\nu^0$ | |
| Adversary's position $y$ and power constraint $P$ | |
| **Output:** Noise error variances $\hat{\nu}^A$ for adversary | |
| Mixed strategy $\{(S^{(1)}, p_1), (S^{(2)}, p_2)\}$ for MAS | |
| $S \leftarrow$ **BestResponse**-$k(\nu^0, G, k)$ | 1 |
| $\hat{\nu}^A \leftarrow$ **NoiseInjection**$(\nu^0, G, y, P)$ | 2 |
| $S' \leftarrow$ **BestResponse**-$k(\nu^0 + \hat{\nu}^A, G, k)$ | 3 |
| **return** $\{(S, 1/2), (S', 1/2)\}$, $\hat{\nu}^A$ | 4 |

The following theorem gives a bound on the approximation error of **Approx-NE**.

THEOREM 7. *Let $\hat{U}_{MAS}$ be the utility of the MAS under the strategy defined above, and let $\hat{U}_A$ be the utility of the adversary. Let $U^*_{MAS}$ be the best-response utility of the MAS to the adversary's strategy, and let $U^*_A$ be the best-response utility of the adversary. Then*

$$\hat{U}_{MAS} \geq \frac{1}{2}\left[\left(1 - \frac{1}{e}\right) U^*_{MAS} - \frac{1}{e}\chi_{max}\right] - \frac{1}{2}\chi_{max} \quad (26)$$

$$\hat{U}_A \geq \frac{1}{2} U^*_A \quad (27)$$

PROOF. Let $S^*$ be the MAS's best response to the adversary's strategy $\hat{\nu}^A$, and let $U_{MAS}(S^*)$ be the resulting utility of the MAS. Let $\mathbf{S}$ be a random variable corresponding to the leader set under the mixed strategy returned by **Approx-NE**. Then under this mixed strategy,

$$\begin{aligned}\mathbf{E}(U_{MAS}(\mathbf{S}, \hat{\nu}^A)) &= U_{MAS}(S, \hat{\nu}^A) Pr(\mathbf{S} = S) \\ &\quad + U_{MAS}(S', \hat{\nu}^A) Pr(\mathbf{S} = S') \\ &\geq -\frac{1}{2}\chi_{max} + \frac{1}{2}\left[\left(1 - \frac{1}{e}\right) U_{MAS}(S^*) \right. \\ &\quad \left. - \frac{1}{e}\chi_{max}\right] \end{aligned} \quad (28)$$

where (28) follows from Theorem 6 and the fact that $\chi_{max}$ is an upper bound on the error experienced by the MAS. This proves (26).

Suppose that the adversary's best response to $\{(S, \frac{1}{2}), (S', \frac{1}{2})\}$ is given by $\nu^{A*}$. Then the adversary's payoff from the strategy $\hat{\nu}^A$ is given by

$$\begin{aligned}\mathbf{E}(U_A(\mathbf{S}, \hat{\nu}^A)) &= U_A(S, \hat{\nu}^A) Pr(\mathbf{S} = S) \\ &\quad + U_A(S', \hat{\nu}^A) Pr(\mathbf{S} = S') \\ &= \frac{1}{2} U_A(S, \hat{\nu}^A) + \frac{1}{2} U_A(S', \hat{\nu}^A) \\ &\geq \frac{1}{2} U_A(S, \hat{\nu}^{A*}) \end{aligned} \quad (29)$$

The last inequality follows from the fact that $U_A \geq 0$ and $U_A(S, \hat{\nu}^A) \geq U_A(S, \hat{\nu}^*)$, since $\hat{\nu}^A$ is by definition the best response to leader set $S$. □

## 6. SIMULATION STUDY

The performance of leader-follower systems in the presence of adversaries, including the case where the leader set is fixed as well as the case where the leader set varies over time, was analyzed using Matlab™. Simulations were conducted assuming a set of 100 agents, deployed uniformly at random over a 1000m x1000m square area, with each agent's

radio range set to 300m. It was assumed that the environmental noise had variance proportional to the distance between agents. An adversary positioned at random within the deployment area, and with power budget equal to $10^6$ was simulated. The path-loss parameter was set to $\alpha = 2$. Each plotted data point represents an average of 30 trials.

For both cases, the proposed leader selection algorithms were compared with three alternative heuristics: random leader selection, selection of the highest-degree nodes as leaders, and selection of the nodes with average degree as leaders.

*Case 1 – Fixed leader set:* The performance of leader-follower systems under noise injection attack when the leader set is fixed is illustrated in Figure 1(a). The algorithm **Fixed-$k$** was used to select leaders, with $\beta = 1$ and $\delta = 1$. The supermodular optimization approach **Fixed-$k$** outperforms the degree-based and random selection heuristics, achieving an error of 5, compared to 7 for random selection and 8 for maximum degree-based selection. Furthermore, we observe that the random selection approach achieves comparable performance to the average degree-based selection, and outperforms selection of high-degree nodes as leaders.

*Case 2 – Time-varying leader set:* Figure 1(b) shows the total error variance when the leader set is allowed to vary over time, based on the approximate Nash equilibrium computed using **Approx-NE**. By allowing the leader set to vary in response to attack, all of the schemes considered achieve better performance than the case of fixed leaders. The supermodular optimization approach outperforms the other three heuristics, achieving a total error variance of 5.25 when 15 leaders are chosen, compared to an error variance of roughly 6 for the random and average degree heuristics, and an error variance of 6.5 for the maximum degree heuristic. Furthermore, random selection of leaders outperforms selecting high-degree agents to act as leaders.

# 7. CONCLUSIONS

In this paper, improving the resilience of leader-follower multi-agent systems to noise injection attacks through leader selection was studied. Two leader selection problems were considered: first, the problem of choosing a fixed set of leaders to maximize robustness to a noise injection attack, and second, the problem of choosing a set of leaders that varies over time in response to attacks. Both cases were analyzed within a supermodular game framework. The first case was studied as a Stackelberg game between a MAS owner and an adversary. It was shown that the adversary's optimum strategy can be computed for a given leader set, while the MAS's best choice of leader set can be approximated up to a provable bound. In the second case, a simultaneous game framework was developed and an algorithm for efficiently approximating a mixed-strategy equilibrium was presented. Both cases were analyzed through simulation study, which demonstrated that choosing a varying set of leaders provides better robustness to noise injection than a fixed set, and that for both cases the supermodular optimization approach outperformed other leader selection algorithms.

In future work, we will study leader selection algorithms that improve on the bounds given in Sections 4 and 5. Moreover, we note that an adversary may employ additional techniques in order to disrupt system performance, including removing links from the MAS altogether through denial-of-service attack. We will further study leader selection methods for mitigating these different classes of attack.

# 8. REFERENCES


[1] T. Alpcan and T. Basar. *Network Security: A Decision and Game-Theoretic Approach*. Cambridge University Press, 2010.

[2] P. Barooah and J. Hespanha. Graph effective resistance and distributed control: Spectral properties and applications. In *45th IEEE Conference on Decision and Control*, pages 3479–3485. IEEE, 2006.

[3] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.

[4] A. Clark and R. Poovendran. A submodular optimization framework for leader selection in linear multi-agent systems. *Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference (CDC-ECC)*, 2011.

[5] C. Daskalakis, A. Mehta, and C. Papadimitriou. A note on approximate nash equilibria. *Internet and Network Economics*, pages 297–306, 2006.

[6] P. Doyle and J. Snell. Random walks and electric networks. *Carus mathematical monographs*, 22, 2000.

[7] J. Lawton, R. Beard, and B. Young. A decentralized approach to formation maneuvers. *IEEE Transactions on Robotics and Automation*, 19(6):933–941, 2003.

[8] B. Liu, T. Chu, L. Wang, and G. Xie. Controllability of a leader–follower dynamic network with switching topology. *IEEE Transactions on Automatic Control*, 53(4):1009–1013, 2008.

[9] M. Mesbahi and F. Hadaegh. Graphs, matrix inequalities, and switching for the formation flying control of multiple spacecraft. In *Proceedings of the 1999 American Control Conference*, volume 6, pages 4148–4152, 1999.

[10] G. Nemhauser, L. Wolsey, and M. Fisher. An analysis of approximations for maximizing submodular set functions. *Mathematical Programming*, 14(1):265–294, 1978.

[11] J. Proakis. Spread spectrum signals for digital communications. *Encyclopedia of Telecommunications*, 2001.

[12] Y. Shoham and K. Leyton-Brown. *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*. Cambridge University Press, 2009.

[13] L. Wolsey. An analysis of the greedy algorithm for the submodular set covering problem. *Combinatorica*, 2(4):385–393, 1982.

[14] L. Xiao, S. Boyd, and S. Kim. Distributed average consensus with least-mean-square deviation. *Journal of Parallel and Distributed Computing*, 67(1):33 – 46, 2007.


# APPENDIX

In this appendix, proofs of Theorems 2, 3, and 4 are given.

PROOF PROOF OF THEOREM 2. Let $\epsilon > 0$. For each $S \subseteq V$, let

$$K(S) \triangleq \min\{K : |f_k(S) - f(S)| < \epsilon/4 \ \forall k \geq K\}. \quad (30)$$

Further, let $K = \max_S K(S)$. Then for any $S \subseteq T$ and

$v \notin T$, and any $k > K$,

$$f(S) - f(S \cup \{v\}) > f_k(S) - f_k(S \cup \{v\}) - \epsilon/2 \quad (31)$$
$$\geq f_k(T) - f_k(T \cup \{v\}) - \epsilon/2 \quad (32)$$
$$> f(T) - f(T \cup \{v\}) - \epsilon \quad (33)$$

where (31) and (33) follow from the definition of $K$ and (32) follows from the supermodularity of $f_k$. Eq. (33) implies that $f(S) - f(S \cup \{v\}) \geq f(T) - f(T \cup \{v\})$, and hence $f$ is supermodular as a function of $S$. □

As a first step towards proving Theorem 3, the following intermediate lemmas are needed.

LEMMA 3. *Let $v$ and $J$ denote vectors in $\mathbf{R}^n$ such that $Lv = J$, $v_i = 0$ for all $i \in S$, $J_u = 0$, and $J_i = 0$ for all $i \in V \setminus (S + u)$.*

PROOF. The equation $Lv = J$ can be written in long form as

$$\begin{pmatrix} L_{ff} & L_{fl} \\ L_{lf} & L_{ll} \end{pmatrix} \begin{pmatrix} v_f \\ v_l \end{pmatrix} = \begin{pmatrix} J_f \\ J_l \end{pmatrix} \quad (34)$$

Substituting $v_l = 0$ yields $L_{ff} v_f = J_f$, which in turn implies that $v_f = L_{ff}^{-1} J_f$. Let $e_u$ denote the vector with a 1 in the $u$-th position and 0s elsewhere. Then multiplying both sides by $e_u^T$ yields

$$v_u = (L_{ff}^{-1})_{u1} J_1 + \cdots + (L_{ff}^{-1})_{u(n-|S|)} J_{n-|S|} \quad (35)$$

Since $J_i = 0$ for $i \neq u$, (35) reduces to $(L_{ff}^{-1})_{uu} J_u = (L_{ff}^{-1})_{uu} = v_u$, as desired. □

In order to prove the next lemma, the following definition is required.

DEFINITION 4. *A function $\mu : E \to \mathbf{R}$ is a* flow *if the following conditions hold for any $(i, j) \in E$. First, $\mu_{ij} = -\mu_{ji}$. Second, $\sum_{j \in N(i)} \mu_{ij} = 0$ for $i \in V \setminus (S + u)$. $\mu$ is defined to be a* unit flow *if the additional condition $\sum_{j \in N(u)} \mu_{uj} = 1$ holds.*

The following lemma gives a property of unit flows.

LEMMA 4. *Let $w : V \to \mathbf{R}$ be any function on $V$ and let $\mu$ be a flow. Then for any $a, b \in V$,*

$$(w_a - w_b) \sum_{j \in N(a)} \mu_{aj} = \frac{1}{2} \sum_{(i,j) \in E} (w_i - w_j) \mu_{ij} \quad (36)$$

The proof of this lemma can be found in [6, Ch 1]. One final lemma is needed before the proof of Theorem 3.

LEMMA 5. *$(L_{ff}^{-1})_{uu}$ is equivalent to*

$$(L_{ff}^{-1})_{uu} = \min \left\{ \sum_{(i,j) \in E} \nu_{ij} \mu_{ij}^2 : \ \mu_{ij} \text{ is a unit flow} \right\}$$

PROOF. Let $v$ be as in the statement of Lemma 3, and define $\lambda_{ij} : E \to \mathbf{R}$ by $\lambda_{ij} = \nu_{ij}^{-1}(v_i - v_j)$. First, we show that $\lambda_{ij}$ is a unit flow. $\lambda_{ij} = -\lambda_{ji}$ follows from the definition. Further, $\sum_{j \in N(i)} \nu_{ij}^{-1}(v_i - v_j) = J_i$, since $v$ is defined by $Lv = J$. Thus $\sum_{j \in N(i)} \lambda_{ij} = 0$ if $i \in V \setminus (S + u)$ and $\sum_{j \in N(i)} \lambda_{ij} = 1$ if $i = u$, satisfying the second and third requirements for a unit flow.

We next show that

$$\sum_{(i,j) \in E} \nu_{ij} \lambda_{ij}^2 = (L_{ff}^{-1})_{uu} \quad (37)$$

From Lemma 4,

$$(v_u - v_s) \sum_{j \in N(u)} \lambda_{uj} = \frac{1}{2} \sum_{(i,j) \in E} (v_i - v_j) \lambda_{ij} \quad (38)$$

for any $s \in S$. By definition of $\lambda$, $(v_i - v_j) = \nu_{ij} \lambda_{ij}$, and so the right hand side of (38) is equivalent to $\frac{1}{2} \sum_{(i,j) \in E} \nu_{ij} \lambda_{ij}^2$. Meanwhile, since $\lambda$ is a unit flow, the left-hand side of (38) is equal to $v_u - v_s$. By definition, $v_s = 0$, and by Lemma 3, $v_u = (L_{ff}^{-1})_{uu}$.

Now, let $\mu_{ij}$ be another unit flow, and note that $\mu_{ij} = \lambda_{ij} + \phi_{ij}$, where $\phi_{ij}$ is a flow with $\sum_{j \in N(u)} \phi_{uj} = 0$. Then

$$\begin{aligned} \sum_{(i,j) \in E} \nu_{ij} \mu_{ij}^2 &= \sum_{(i,j) \in E} \nu_{ij} (\lambda_{ij} + \phi_{ij})^2 \\ &= \sum_{(i,j)} \nu_{ij} \lambda_{ij}^2 + 2 \sum_{(i,j)} \phi_{ij} (v_i - v_j) \\ &\quad + \sum_{(i,j) \in E} \nu_{ij} \phi_{ij}^2 \quad (39) \\ &= \sum_{i,j} \nu_{ij} \lambda_{ij}^2 + 4 v_u \sum_{j \in N(u)} \phi_{uj} \\ &\quad + \sum_{i,j} \nu_{ij} \phi_{ij}^2 \quad (40) \\ &= \sum_{i,j} \nu_{ij} \lambda_{ij}^2 + \sum_{i,j} \nu_{ij} \phi_{ij}^2 \quad (41) \\ &\geq \sum_{i,j} \nu_{ij} \lambda_{ij}^2 \quad (42) \end{aligned}$$

where (39) follows by expanding the previous equation and by definition of $\lambda$, (40) follows from Lemma 4 and (41) follows from the fact that $\sum_{j \in N(u)} \phi_{uj} = 0$. □

PROOF PROOF OF THEOREM 3. First, note that

$$\sum_{(i,j) \in E} \nu_{ij} \mu_{ij}^2 \quad (43)$$

is linear as a function of $\nu_{ij}$. By Lemma 5, $(L_f^{-1})_{uu}$ is a pointwise minimum of linear functions and is therefore concave. □

The following lemma is needed to prove Theorem 4.

LEMMA 6. *The function $F_\zeta(S)$ defined in (16) is supermodular as a function of $S$.*

PROOF. By the theory of Riemann integration, $F_\zeta(S)$ can be approximated by a sequence of functions

$$F_\zeta^k(S) = \sum_{l=0}^{L_k} \max \{\chi(S, R(P_l^k)), \zeta\} |C_l^k| \quad (44)$$

where $|\cdot|$ denotes the Lebesgue measure of a set, the sets $C_l^k$ satisfy $|C_l^k| = \delta_k$, with $\delta_k \to 0$ as $k \to \infty$, $P_l \in C_l^k$, and $\mathcal{S}_2 \subseteq \bigcup_{l=0}^{L_k} C_l^k$. Now, since $\chi(S, R)$ is supermodular for fixed $R$ (Theorem 5), $\max \{\chi(S, R), \zeta\}$ is supermodular as a function of $S$ (Lemma 2). $F_\zeta^k(S)$ is therefore a nonnegative weighted

sum of supermodular functions, and thus is supermodular by Lemma 1. This implies that $F_\zeta(S)$ is the limit of a sequence of supermodular functions, and so $F_\zeta(S)$ is supermodular by Theorem 2. □

LEMMA 7. *For fixed $\zeta$, for the optimization problem*

$$\begin{array}{cc} \textit{minimize} & |S| \\ \textit{s.t.} & F_\zeta(S) \leq \zeta \end{array} \tag{45}$$

*the greedy algorithm of lines 6-8 in **Fixed**-$k$ returns a set $S$ with*

$$\frac{|S|}{|S^*|} \leq 1 + \log_e \left\{ \frac{\max_v F_\zeta(\{v\})}{\zeta} \right\}, \tag{46}$$

*where $S^*$ is the optimum solution to (45).*

PROOF. The proof follows from Theorem 1 of [13] and the supermodularity of $F_\zeta(S)$. □

PROOF PROOF OF THEOREM 4. From Lemma 7, solving a problem of the form (45) with $\zeta = \zeta^*$ will return a set $S$ with $|S| \geq \beta k$ and $\beta$ is as in the statement of Theorem 4. Furthermore, the algorithm is guaranteed to reach $\zeta = \zeta^*$, because $\zeta^j$ is strictly decreasing as long as $|S^j| \leq \beta k$, which will hold for all $\zeta^j > \zeta^*$. □