

Efficient, Optimal k -Leader Selection for Coherent, One-Dimensional Formations

Stacy Patterson¹, Neil McGlohon¹, and Kirill Dyagilev²

Abstract—We study the problem of optimal leader selection in consensus networks with noisy relative information. The objective is to identify the set of k leaders that minimizes the formation’s deviation from the desired trajectory established by the leaders. An optimal leader set can be found by an exhaustive search over all possible leader sets; however, this approach is not scalable to large networks. In recent years, several works have proposed approximation algorithms to the k -leader selection problem, yet the question of whether there exists an efficient, non-combinatorial method to identify the optimal leader set remains open. This work takes a first step towards answering this question. We show that, in one-dimensional weighted graphs, namely path graphs and ring graphs, the k -leader selection problem can be solved in polynomial time (in both k and the network size n). We give an $O(n^3)$ solution for optimal k -leader selection in path graphs and an $O(kn^3)$ solution for optimal k -leader selection in ring graphs.

I. INTRODUCTION

We explore the problem of leader selection in leader-follower consensus systems. Such systems arise in the context of vehicle formation control [1], distributed clock synchronization [2], and distributed localization in sensor networks [3], among others. In these systems, several agents act as *leaders* whose states serve as the reference trajectory for the entire system. The leaders may be controlled autonomously or by a system owner. The remaining agents are *followers*. Each follower updates its state based on noisy measurements of the states of its neighbors. The objective of the leader-follower system is for the entire formation to maintain a desired global state.

Since the follower agents’ measurements are corrupted by stochastic noise, the agents cannot maintain the formation exactly. However, the variance of the deviation of the agents’ states from the desired states is bounded [4]. This variance is related to the *coherence* of the formation, and is quantified by an H_2 norm of the leader-follower system [5], [6]. It has been shown that the coherence of such leader-follower consensus systems depends on which agents act as leaders [4], [5], and so, by judiciously choosing the leader set, one can minimize the total variance of the formation. The *k -leader selection problem* is precisely to select a set of at most k leaders for which the total variance of the deviation of the follower nodes is minimized.

The problem of finding a leader set that minimizes the total variance of the deviation has received attention in

recent years. The optimal leader set can be found by an exhaustive search over all subsets of agents of size at most k , but this solution is not tractable in large networks. In [7], the authors present a leader selection algorithm based on a convex-relaxation of the problem. They then give an efficient algorithm for this relaxed problem. While this algorithm offers no guarantees on the optimality of the chosen leader set, its good performance was demonstrated on example networks. In [8], the authors show that the total variance of the deviation from the desired trajectory is a super-modular set function [9]. This super-modularity property implies that one can use a greedy, polynomial-time algorithm to find a leader for which the total variance is within a provable bound of optimal. Another recent work [10] has shown a connection between the optimal leader set and the information centrality measures. They have used this connection to give efficient algorithms for finding the optimal single leader and optimal pair of leaders in several network topologies. Other works have explored optimal leader selection in leader-follower systems without stochastic disturbances [11] and in systems where both the leaders and followers are subject to stochastic disturbances and the leaders also have access to relative state information [7]. Despite this recent interest, the question of whether it is possible to find an optimal leader set using a non-combinatorial approach remains open.

In this work, we take a first step towards answering this question. We show that, in one-dimensional weighted graphs, namely path graphs and ring graphs, the k -leader selection problem can be solved in time that is polynomial in both k and the network size n . Our approach is to first transform the k -leader selection problem into the problem of finding a minimum-weight path in a weighted directed graph, a problem that can be solved polynomial time. We then use a slightly modified version of the well-known Bellman-Ford algorithm [12] to find this minimum-weight path and thus the optimal leader set. For path graphs, our algorithm finds the optimal leader set of size at most k in $O(n^3)$ time. In ring graphs, our algorithm finds the optimal leader set of size at most k in $O(kn^3)$ time.

Our approach was inspired by recent work on facility location on the real line [13]. In this work, the authors solve the facility location problem by reducing it to a minimum-weight path problem over a directed acyclic graph. We note that our graph construction and our path-finding algorithm are both different from those in [13].

The remainder of this paper is organized as follows. In Section II, we present the system model and formalize the k -leader selection problem. In Section III, we present our

¹S. Patterson and N. McGlohon are with the Department of Computer Science, Rensselaer Polytechnic Institute, Troy, NY 12180, USA, sep@cs.rpi.edu, mcglon@rpi.edu

²K. Dyagilev is with the Department of Computer Science, Johns Hopkins University, Baltimore, MD 21218, USA, kirilld@cs.jhu.edu

algorithms for finding the optimal leader set in path and ring graphs. Section IV gives computational examples comparing the optimal leader set to the leader set selected by the greedy algorithm. Finally, we conclude in Section V.

II. PROBLEM FORMULATION

In this section, we describe the dynamics of leader-follower consensus system and formally define the k -leader selection problem.

A. System Model

We consider a network of agents, modeled by an undirected, connected graph $G = (V, E)$, where the node set $V = \{1, 2, \dots, n\}$ represents the agents and the edge set E captures the communication structure of a network. We use the terms node and agent interchangeably. If $(i, j) \in E$, then nodes i and j can exchange information. We denote the neighbor set of a node i by $N(i)$, i.e., $N(i) \triangleq \{j \in V \mid (i, j) \in E\}$. In this work, we restrict our study to one-dimensional graphs, namely path graphs and ring graphs. For path graphs, we assume the nodes IDs are assigned in order along the path. For ring graphs, we assume that the node IDs are assigned in ascending order around the ring in a clockwise fashion, starting with an arbitrary node.

Each agent i has a state $x_i(t) \in \mathbb{R}$ which represents, for example, the position of agent i at time t . The objective is for each pair of neighbor agents i and j to maintain a pre-specified difference p_{ij} between their states,

$$x_i(t) - x_j(t) = p_{ij} \quad \text{for all } (i, j) \in E. \quad (1)$$

If $x_i(t)$ represents an agent's position, p_{ij} is the desired distance between agents i and j . We assume that each agent i knows p_{ij} for all $j \in N(i)$.

A subset $S \subset V$ of agents act as leaders. The states of these agents serve as reference states for the network. The state of each leader $s \in S$ remains fixed at its reference value \bar{x}_s . Let $\bar{\mathbf{x}}$ denote the vector of states that satisfy (1) when the leader states are fixed at their reference values. We assume that at least one such $\bar{\mathbf{x}}$ exists.

The remaining agents $v \in V \setminus S$ are followers. A follower updates its state based on noisy measurements of the differences between its state and the states of its neighbors. The dynamics of each follower agent is given by

$$\dot{x}_i(t) = - \sum_{j \in N(i)} W_{ij} (x_i(t) - x_j(t) - p_{ij} + \epsilon_{ij}(t)),$$

where W_{ij} is the weight for link (i, j) and $\epsilon_{ij}(t)$, $(i, j) \in E$, are zero-mean white noise processes with autocorrelation functions $\mathbf{E}[\epsilon_{ij}(t)\epsilon_{ij}(t + \tau)] = \nu_{ij}\delta(\tau)$. Here $\delta(\cdot)$ denotes the unit impulse function. Each ϵ_{ij} is independent, and ϵ_{ij} and ϵ_{ji} are identically distributed. As in [8], we select the edges weights as $W_{ij} = \frac{\nu_{ij}}{\Delta_i}$, where $\Delta_i = \sum_{j \in N(i)} (1/\nu_{ij})$. This edge weight policy corresponds to the best linear unbiased estimator of the leader agents' states when $x_j(t) = \bar{x}_j$ for all $j \in N(i)$ [4].

B. Performance Measure

Without the noise processes ϵ_{ij} , the agents' states would converge to $\bar{\mathbf{x}}$. With these noise processes, the agents' states deviate from $\bar{\mathbf{x}}$; however, the variances of these deviations are bounded in the mean-square sense. For a follower agent i , let r_i be the steady-state variance of the deviation from \bar{x}_i ,

$$r_i \triangleq \lim_{t \rightarrow \infty} \mathbf{E} [(x_i(t) - \bar{x}_i)^2].$$

The steady-state variances r_i , $i \in V \setminus S$ can be obtained from the weighted Laplacian matrix of the graph G , whose components are defined as

$$L_{ij} = \begin{cases} -\frac{1}{\nu_{ij}} & \text{if } (i, j) \in E \\ \Delta_i & \text{if } i = j \\ 0 & \text{otherwise.} \end{cases}$$

For a leader set S , let the follower-follower sub-matrix L_{ff} be the matrix L with the rows and columns corresponding nodes in S removed. It has been shown that [4],

$$r_i = \frac{1}{2} (L_{ff}^{-1})_{ii}. \quad (2)$$

We note that since G is connected, L_{ff} is positive definite [5], and therefore, (2) is well-defined.

We measure the performance of the leader-follower system for a given leader set S by the total steady-state variance of the deviation from $\bar{\mathbf{x}}$,

$$R(S) \triangleq \sum_{i \in V \setminus S} r_i = \frac{1}{2} \mathbf{tr} (L_{ff}^{-1}) \quad (3)$$

A formation with a small $R(S)$ exhibits good coherence, i.e., the formation closely resembles a rigid formation.

C. The k -Leader Selection Problem

A natural question that arises is how to identify a leader set S of a certain size that minimizes the total steady-state variance in (3). This question is formalized as the k -leader selection problem.

Problem Statement 1: The k -leader selection problem is

$$\begin{aligned} & \text{minimize} && R(S) \\ & \text{subject to} && |S| \leq k \end{aligned} \quad (4)$$

A naïve solution to this problem is to construct all subsets of V of size at most k , evaluate $R(S)$ for each leader set, and choose the set S^* for which $R(S)$ is minimized. The computational complexity of this solution is combinatorial since the number of leader sets that would need to be evaluated is $\sum_{i=1}^k \binom{n}{i}$. In the next section, we show that for one-dimensional graphs, the optimal leader set can be found with time complexity that is polynomial in both n and k .

III. LEADER SELECTION ALGORITHM

In this section, we present our leader selection algorithms for path and ring graphs. Our approach for both graph types is to first reduce the leader selection problem to the problem of finding a minimum-weight path in a digraph. We then solve this minimum-weight path problem using a modified version of the Bellman-Ford algorithm [12].




Fig. 1: Digraph generated from four node path graph.

A. Leader Selection for a Path Graph

We first note that, in a path graph with k leader nodes, the matrix L_{ff} is block diagonal with at most $k + 1$ blocks. Further, each block is tridiagonal. For example, consider a path graph with leader set $S = \{\ell_1, \ell_2, \dots, \ell_k\}$, where $\ell_i < \ell_{i+1}$ for $i = 1, \dots, (k - 1)$. The matrix L_{ff} is formed from the weighted Laplacian L by removing the rows and columns corresponding to the nodes in S . L_{ff} can be written as,

$$L_{ff} = \begin{bmatrix} L_{[1, \ell_1]} & \mathbf{0} & \cdots & \cdots & \mathbf{0} \\ \mathbf{0} & L_{(\ell_1, \ell_2)} & \mathbf{0} & \cdots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ \mathbf{0} & \ddots & \ddots & L_{(\ell_{k-1}, \ell_k)} & \mathbf{0} \\ \mathbf{0} & \cdots & \cdots & \mathbf{0} & L_{(\ell_k, n]} \end{bmatrix},$$

where the matrices $L_{[1, \ell_1]}$, $L_{(\ell_i, \ell_{i+1})}$, $i = 1, \dots, (k - 1)$, and $L_{(\ell_k, n]}$ are defined as follows. $L_{[1, \ell_1]}$ is the sub-matrix of L consisting of the rows and columns 1 through $\ell_1 - 1$ of L . This matrix models the interactions between the follower nodes, $i = 1, \dots, (\ell_1 - 1)$. Note that node ℓ_1 is the only leader that affects the states of these nodes. $L_{(\ell_i, \ell_{i+1})}$ is the sub-matrix of L consisting of the rows and columns indexed from $\ell_i + 1$ to $\ell_{i+1} - 1$, inclusive. This matrix models the interactions of the nodes between leader node ℓ_i and leader node ℓ_{i+1} ; these follower nodes are influenced by both leaders. Finally, $L_{(\ell_k, n]}$ consists of the rows and columns $i = (\ell_k + 1), \dots, n$ of L , inclusive, and models the interactions of the follower nodes $i = (\ell_k + 1), \dots, n$. Leader node ℓ_k is the only leader that influences the states of these nodes. If there are leaders $u, v \in S$ such that $v = u + 1$, then the corresponding sub-matrix of L will be of size 0.

The corresponding error $R(S)$ is given by

$$R(S) = \frac{1}{2} \text{tr}(L_{[1, \ell_1]}^{-1}) + \frac{1}{2} \sum_{i=1}^{k-1} \text{tr}(L_{(\ell_i, \ell_{i+1})}^{-1}) + \frac{1}{2} \text{tr}(L_{(\ell_k, n]}^{-1}) \quad (5)$$

Note that, if any sub-matrix of L is of size 0, then the trace of the inverse of this sub-matrix is 0. Similarly, $\text{tr}(L_{[1, 1]}^{-1}) = 0$ and $\text{tr}(L_{(n, n]}^{-1}) = 0$. By decomposing $R(S)$ in this manner, we observe that, for a given leader set S , the total steady-variance can be computed by finding the trace of the inverse of the $k + 1$ blocks.

We now construct a weighted digraph $\bar{G} = (\bar{V}, \bar{E}, W)$ based on L as follows. The set of nodes in the graph consists of the nodes in V and an additional source node s and target

Algorithm 1 Algorithm for finding the optimal solution to the k -leader selection problem in a path graph.

Input: $G = (V, E)$, edge weights $\frac{1}{\nu_{ij}}$, k
Output: Set of leader nodes S , error $R(S)$

$L \leftarrow$ Laplacian of G
 $\bar{G} \leftarrow$ digraph constructed from L
 $(P, w) \leftarrow \text{MODIFIEDBELLMANFORD}(\bar{G}, s, t, k + 1)$

Construct leader set from min. weight path:

$S \leftarrow \emptyset$
for $v \in P$ **do**
 if $v \neq s$ and $v \neq t$ **then**
 $S \leftarrow S \cup \{v\}$
 end if
end for
 $R(S) \leftarrow w$
return $(S, R(S))$

node t , i.e., $\bar{V} = V \cup \{s, t\}$, as shown in Figure 1. The edge set \bar{E} contains edges from s to every node $v \in V$. The weight of edge (s, v) is $\text{tr}(L_{[1, v]}^{-1})$. This edge weight is the total steady-state variance for nodes $1, \dots, (v - 1)$ when node v is a leader node and there are no other leader nodes u with $u < v$. \bar{E} also contains edges from from each node $u \in V$ to each node $v \in V$ with $u < v$. The weight of edge (u, v) is $\text{tr}(L_{(u, v)}^{-1})$. This weight is the total variance of the nodes $i = (u + 1), \dots, (v - 1)$ when nodes u and v are leaders and there are no other leader nodes w with $u < w < v$. Finally, \bar{E} contains edges from every node $v \in V$ to node t . The weight of edge (v, t) is $\text{tr}(L_{(v, n]}^{-1})$. This weight is the total variance of the nodes $i = v + 1, \dots, n$ when v is a leader and there are no other leader nodes u with $u > v$. Note that the weights of edges $(s, 1)$, (n, t) , and $(u, u + 1)$, $u = 1 \dots (n - 1)$ are 0.

We observe that the weight of a path of length $k + 1$ from s to t in \bar{G} is equivalent to $R(S)$ in (5), where each vertex on the path, excepting s and t , is an element of S . Thus, to find a leader set S with $|S| \leq k$ that minimizes (5), one seeks a minimum-weight path from s to t in \bar{G} that has at most $k + 1$ edges. The optimal leaders are the nodes along this path between s and t . To find this minimum-weight path, we use a slightly modified implementation of the Bellman-Ford algorithm [12]. The Bellman-Ford algorithm is an iterative algorithm that finds the minimum-weight paths (of any length) from a source node to every other node in the graph. While there are more efficient algorithms that solve this same problem, Bellman-Ford offers the benefit that, in each iteration m , the algorithm finds the minimum-weight paths of m edges. Therefore, we can execute the Bellman-Ford algorithm for $k + 1$ iterations to find the minimum-weight path of at most $k + 1$ edges. We have made slight modifications to this algorithm to compute not only the weight of the path but the list of nodes traversed in this path. Our modified version of Bellman-Ford is detailed in the technical report [14].

The pseudocode for our k -leader selection algorithm is given in Algorithm 1. Our algorithm returns the optimal set

of leaders of size at most k . A leader set may have cardinality $h < k$ if the inclusion of more than h leaders does not decrease $R(S)$.

B. Leader Selection for a Ring Graph

In a ring graph with k leaders, the leader-follower system can be decomposed into k independent subsystems (with some possibly consisting of zero nodes). Each of these subsystems corresponds to a segment of the graph where two leader nodes form the boundaries of this segment. As before, we can define a weighted digraph where the weight of edge (u, v) is the total steady-state variance of the nodes on the segment between u and v when both u and v are leaders. The total steady state-variance of the leader-follower system is then given by the weight of a path through this digraph. We now present the details of our k -leader selection algorithm for ring graphs. Pseudocode is given in Algorithm 2.

To find the optimal leader set of size at most k , we first select one node i as a leader. We then translate the problem finding the remaining $k-1$ leaders into a problem of finding a minimum weight path of at most k edges over a weighted digraph. The digraph construction is described below. To ensure that our algorithm finds the optimal leader set, the algorithm performs this translation and path-finding for each possible initial leader i , $i = 1, \dots, n$. The optimal leader set is the set with the minimum weight path among these n minimum weight paths (one for each initial leader i).

For a given initial leader i , we construct its weighted digraph $\overline{G}_i = (\overline{V}_i, \overline{E}_i, \overline{W}_i)$ as follows. The vertex set of \overline{V}_i contains a source node s_i , a target node t_i , and the vertices in V excepting i , i.e.,

$$\overline{V}_i = \{s_i, t_i\} \cup (V \setminus \{i\}).$$

The edge set \overline{E}_i contains directed edges from s_i to every node $v \in (V \setminus \{i\})$. The weight of edge (s_i, v) is the total steady-state variance for the nodes between i and v in the clockwise direction on the ring graph when both i and v are leaders and there are no other leader nodes after i and before v in the clockwise direction. \overline{E}_i also contains directed edges from every node $v \in (V \setminus \{i\})$ to t_i . The weight of edge (v, t_i) is the total steady-state variance of the nodes between v and i in the clockwise direction on the ring graph when both i and v are leaders and there are no other leader nodes after v and before i in the clockwise direction. The weight of edge (u, v) is the total variance of the nodes between u and v in the clockwise direction on the ring graph when both u and v are leaders and there are no other leader nodes following u and preceding v .

To compute the weight for edge (u, v) , where u precedes v on the ring in the clockwise direction, we first construct the matrix M from L by shifting the rows and columns of L so that node u corresponds to the first row and column of M . The index (row and column of M) corresponding to a node u after this shift is 1, and the index corresponding to node v is $v - u + 1$ modulo n . The weight of edge (u, v) is

$$w_{uv} = \mathbf{tr} \left(M_{(1, v-u+1 \pmod{n})}^{-1} \right), \quad (6)$$

Algorithm 2 Algorithm for finding the optimal solution to the k -leader selection problem in a ring graph.

Input: $G = (V, E)$, edge weights $\frac{1}{\nu_{ij}}$, k
Output: Set of leader nodes S , error $R(S)$

```

 $L \leftarrow$  weighted Laplacian for  $G$ 
 $minWeight \leftarrow \infty$ 
 $minP \leftarrow \perp$ 
for  $i = 1 \dots n$  do
   $\overline{G}_i \leftarrow$  digraph constructed from  $L$  for first leader  $i$ 
   $(P, w) \leftarrow$  MODIFIEDBELLMANFORD( $\overline{G}_i, s_i, t_i, k - 1$ )
  if  $w < minWeight$  then
     $minWeight \leftarrow w$ 
     $minP \leftarrow P$ 
  end if
end for

Construct leader set from min. weight path:
 $S \leftarrow \emptyset$ 
for  $v \in minP$  do
  if  $v \neq s$  and  $v \neq t$  then
     $S \leftarrow S \cup \{v\}$ 
  end if
end for
 $R(S) \leftarrow minWeight$ 
return  $(S, R(S))$ 

```

where $M_{(j,k)}$ is the sub-matrix of M consisting of the rows and columns between $j+1$ and $k-1$, inclusive. The weights of edges (s_i, v) , $v \in (V \setminus \{i\})$ can be computed by replacing u with i in the above procedure. The weights of edges (u, t_i) , $u \in (V \setminus \{i\})$ can be computed by replacing v with i in the above procedure. Note that each of the sub-matrices used in this procedure is a tri-diagonal matrix. Also, as with the path graph, the weight of an edge (u, v) is 0 if nodes u and v are adjacent in G .

Once the graph \overline{G}_i is constructed, the modified Bellman-Ford algorithm is used to find the minimum-weight path from s_i to t_i of at most k edges. The weight of this path is the minimal $R(S)$ when $i \in S$ and $|S| \leq k$. The leader set S for this error consists of node i and the nodes along this path, excepting s_i and t_i . By finding the optimal leader that contains i , for each $i \in V$, the algorithm is able to identify the initial leader i that minimizes $R(S)$ and thus the optimal solution to the k -leader selection problem.

C. Algorithm Analysis

We now analyze the computational complexity of our leader selection algorithms.

Theorem 1: For a path graph G with n nodes, the k -leader selection algorithm identifies the leader set S , with $|S| \leq k$, that minimizes $R(S)$ in $O(n^3)$ time.

Proof: For the path graph, our algorithm consists of two phases, each of which is performed once. The first phase is the construction of the digraph $\overline{G} = (\overline{V}, \overline{E})$. The edge set \overline{E} consists of n edges with s as their source (one to each $v \in V$), n edges with t as their sink (one from each $v \in V$), and one edge from each $u \in V$ to each $v \in V$ with $u < v$. Thus $|\overline{E}| \in O(n^2)$. To find each edge weight, we must find the diagonal entries of the inverse of a tridiagonal matrix of

size at most $(n - 1) \times (n - 1)$. These diagonal entries can be found in $O(n)$ operations [15]. Therefore the digraph \bar{G} can be constructed in $O(n^3)$ operations.

The second phase of the algorithm is to find the minimum-weight path of at most $k + 1$ edges from s to t in \bar{G} . For a graph with m edges, the Bellman-Ford algorithm finds the minimum-weight-path of length at most h edges in $O(hm)$ operations [12]. Therefore the second phase of our algorithm has complexity $O(kn^2)$.

Combining the two phases of the algorithm we arrive a computational complexity of $O(n^3)$. ■

Theorem 2: For a ring graph G with n nodes, the k -leader selection algorithm identifies the leader set S , with $|S| \leq k$, that minimizes $R(S)$ in $O(kn^3)$ time.

Proof: In the ring algorithm, n weighted digraphs \bar{G}_i , $i = 1, \dots, n$, are constructed, and a shortest-path algorithm is executed on each digraph. To construct these digraphs, first, the weight of each edge (u, v) , $u, v \in V, u \neq v$ is computed according to (6). For each computation, the shift operation can be performed in $O(n)$ operations to obtain a tridiagonal matrix. The trace of the inverse of the sub-matrix can also be found in $O(n)$. Thus, the weight of all pairs (u, v) can be computed in $O(n^3)$. These edge weights are used in every digraph \bar{G}_i and can be looked up in constant time (for example, by storing them in an $n \times n$ matrix).

To construct a digraph where edge weights can be computed in constant time requires $O(|V| + |E|)$ operations. Therefore, each \bar{G}_i can be constructed in $O(n^2)$. There are n such digraphs, so the construction of all digraphs requires $O(n^3)$ time.

Finally, for each digraph, the Bellman-Ford algorithm finds the minimum-weight path of at most k edges in $O(kn^2)$ time. Thus, to find the minimum-weight path over these n minimum-weight paths requires $O(kn^3)$ time.

Combining all steps of the algorithm, we obtain a running time of $O(kn^3)$. ■


IV. COMPUTATIONAL EXAMPLES

In this section, we explore the results of our k -leader selection algorithms on several example graphs. For comparison, we also show results from the greedy leader selection algorithm presented in [8]. The greedy algorithm consists of at most k iterations. In each iteration, a node s is selected such that $s = \arg \min_{s \in V \setminus S} R(S \cup \{s\})$, and this node is added to the leader set. The greedy algorithm generates a leader set S of size at most k such that,


$$R(S) \leq \left(1 - \left(\frac{k-1}{k}\right)^k\right) R^* + \frac{1}{e} R_{\max},$$

where $R_{\max} = \max_{i \in V} R(i)$. As far as we are aware, the greedy algorithm is the only previously proposed algorithm that gives provable bounds on the optimality of the leader set. A complete description of the greedy algorithm and its analysis can be found in [8].

We have implemented both algorithms in Matlab. For each graph, we select the edge weights ν_{ij} uniformly at random




(a) Total variance for different leader set sizes k .




(b) Total variance of leader set S of size k produced by greedy algorithm, relative to total variance for optimal leader set S^* of size k .

Fig. 2: A comparison of the our optimal leader selection algorithm with the greedy leader selection algorithm on a 400 node path graph.



(a) Leaders selected by greedy algorithm.




(b) Optimal leaders.


Fig. 3: A comparison of the greedy leader selection and optimal leader selection for $k = 2$ and uniform edge weights in an 11 node path graph. The leaders selected by each algorithm are shown in grey.

from the interval $(0, 1)$. We have experimented with other edge weight selection policies, and results are similar to those shown below.

The results for a 400 node path graph are shown Figure 2. Figure 2a gives the total variance $R(S)$ for the optimal leader set, as found by our algorithm, and for the leader set found by the greedy algorithm for different leader set sizes k . Figure 2b shows total variance $R(S)$ of the leader set selected by the greedy algorithm relative to $R(S^*)$, where S^* is the optimal leader set. For $k = 1$, the optimal leader ℓ_1 is the weighted median of the path graph (see [16]), and both algorithms select this leader. An interesting observation is that for $k = 2$, the greedy algorithm demonstrates its worst relative performance. A reason for this can be observed in the example in Figure 3, where we show the optimal leaders and those selected by the greedy algorithm for $k = 2$ in an 11 node path graph with all edge weights equal to one. The locations of the optimal two leaders are symmetric. The greedy algorithm selects the best single leader, the node in the center of the path, in the first iteration. It selects a node nearer to the edge of the graph in the second iteration. The center node is a poor choice for $k = 2$, resulting in a



(a) Total variance for different leader set sizes k .



(b) Total variance of leader set S of size k produced by greedy algorithm, relative to total variance for optimal leader set S^* of size k .

Fig. 4: A comparison of the optimal leader selection algorithm with the greedy leader selection algorithm on a 400 node ring graph.

significantly larger total variance that the optimal.

We note that, overall, the greedy algorithm yields leader sets whose performance is fairly close to optimal. As the number of leaders increases, the total variance increases for both leader selection algorithms. The relative error of the greedy algorithm does not appear to vanish as k increases.

The results for a 400 node ring graph are shown in Figure 4. As before, Figure 4a gives the total variance $R(S)$ for optimal leader set and the leader set selected by the greedy algorithm, and Figure 4b shows total variance $R(S)$ of the leader set selected by the greedy algorithm relative to $R(S^*)$, where S^* is the optimal leader set as identified by our algorithm. For all k greater than 1, the greedy algorithm selects a sub-optimal leader set. The most interesting results are for smaller values of k . The greedy algorithm exhibits worse performance when k is odd than when k is even. We believe this is due to a similar phenomenon as observed for $k = 2$ in the path graph. The greedy algorithm makes poor choices when k is odd and can compensate for previous poor choices in subsequent iterations when k is even.

For larger values of k , the performance of the greedy algorithm appears to stabilize at around 1.05 times the optimal $R(S^*)$. While the greedy algorithm performs slightly worse for the ring graph than it does for the path graph, overall, the greedy solution yields good approximations for larger values of k .

V. CONCLUSION

We have investigated the problem of optimal k -leader selection in noisy leader-follower consensus systems. A naïve solution to the leader selection problem has combinatorial complexity; however, it is unknown whether the leader

selection problem is NP-Hard, in general, or if efficient polynomial-time solutions can be found. In this work, we have taken a step towards addressing this open question. We have shown that, in one-dimensional weighted graphs, namely path graphs and ring graphs, the k -leader selection problem can be solved in polynomial time (in both k and the network size n). Further, we have given an $O(n^3)$ solution for optimal k -leader selection in path graphs and an $O(kn^3)$ solution for optimal k -leader selection in ring graphs.

It is our intuition that more efficient algorithms for leader-selection in one-dimensional networks can be found. We plan to explore this in future work. Further, we plan to extend our approach to optimal k -leader selection in other network topologies. Finally, we will explore using similar techniques for leader selection in other dynamics, for example, networks where leaders are also subject to stochastic disturbances.

REFERENCES

- [1] W. Ren, R. W. Beard, and T. W. McLain, "Coordination variables and consensus building in multiple vehicle systems," *Cooperative Control*, vol. 309, pp. 171–188, 2005.
- [2] J. Elson, R. M. Karp, C. H. Papadimitriou, and S. Shenker, "Global synchronization in sensor networks," in *Latin American Theoretical Informatics*, 2004, pp. 609–624.
- [3] P. Barooah and J. Hespanha, "Error scaling laws for linear optimal estimation from relative measurements," *IEEE Trans. Inf. Theory*, vol. 55, no. 12, pp. 5661–5673, Dec 2009.
- [4] —, "Graph effective resistance and distributed control: Spectral properties and applications," in *Proc. 45th IEEE Conf. on Decision and Control*, Dec 2006, pp. 3479–3485.
- [5] S. Patterson and B. Bamieh, "Leader selection for optimal network coherence," in *Proc. 49th IEEE Conf. on Decision and Control*, 2010, pp. 2692–2697.
- [6] B. Bamieh, M. R. Jovanovic, P. Mitra, and S. Patterson, "Coherence in large-scale networks: Dimension-dependent limitations of local feedback," *IEEE Trans. Autom. Contr.*, vol. 57, no. 9, pp. 2235–2249, May 2012.
- [7] F. Lin, M. Fardad, and M. Jovanovic, "Algorithms for leader selection in stochastically forced consensus networks," *IEEE Trans. Autom. Contr.*, vol. 59, no. 7, pp. 1789–1802, Jul 2014.
- [8] A. Clark, L. Bushnell, and R. Poovendran, "A supermodular optimization framework for leader selection under link noise in linear multi-agent systems," *IEEE Trans. Autom. Contr.*, vol. 59, no. 2, pp. 283–296, Feb 2014.
- [9] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher, "An analysis of approximations for maximizing submodular set functions," *Mathematical Programming*, vol. 14, no. 1, pp. 265–294, 1978.
- [10] K. Fitch and N. Leonard, "Information centrality and optimal leader selection in noisy networks," in *Proc. 52nd IEEE Conf. on Decision and Control*, Dec 2013, pp. 7510–7515.
- [11] A. Clark, B. Alomair, L. Bushnell, and R. Poovendran, "Minimizing convergence error in multi-agent systems via leader selection: A supermodular optimization approach," *IEEE Trans. Autom. Contr.*, vol. 59, no. 6, pp. 1480–1494, Jun 2014.
- [12] T. H. Cormen, C. E. Leiserson, R. L. Rivest, C. Stein, *et al.*, *Introduction to algorithms*, 3rd ed. MIT Press, 2010.
- [13] D. Z. Chen and H. Wang, "New algorithms for facility location problems on the real line," *Algorithmica*, vol. 69, no. 2, pp. 370–383, Jun 2014.
- [14] S. Patterson, N. McGlohon, and K. Dyagilev, "Efficient, optimal k -leader selection for coherent, one-dimensional formations," *eprint arXiv:1412.6595*, 2014.
- [15] G. B. Rybicki and D. G. Hummer, "An accelerated lambda iteration method for multilevel radiative transfer. i-non-overlapping lines with background continuum," *Astronomy and Astrophysics*, vol. 245, pp. 171–181, 1991.
- [16] S. Patterson, "In-network leader selection for acyclic graphs," *eprint arXiv:1410.6533*, 2014.