

Execution Order Matters in Greedy Algorithms with Limited Information

Rohit Konda, David Grimsman, Jason R. Marden

Abstract—In this work, we study the multi-agent decision problem where agents try to coordinate to optimize a given system-level objective. While solving for the global optimum is intractable in many cases, the *greedy algorithm* is a well-studied and efficient way to provide good approximate solutions - notably for submodular optimization problems. Executing the greedy algorithm requires the agents to be ordered and execute a local optimization based on the solutions of the previous agents. However, in limited information settings, passing the solution from the previous agents may be nontrivial, as some agents may not be able to directly communicate with each other. Thus the communication time required to execute the greedy algorithm is closely tied to the order that the agents are given. In this work, we characterize interplay between the communication complexity and agent orderings by showing that the complexity using the best ordering is $O(n)$ and increases considerably to $O(n^2)$ when using the worst ordering. Motivated by this, we also propose an algorithm that can find an ordering and execute the greedy algorithm quickly, in a distributed fashion. We also show that such an execution of the greedy algorithm is advantageous over current methods for distributed submodular maximization.

I. INTRODUCTION

Many real-world problems are well-modeled as multiagent decision problems, including building energy management [1], stock trading [2], water resource allocation [3], [4], traffic light management [5], the power grid [6], and robot path planning [7]. In these scenarios, the set of n decision makers, or agents, coordinate to a joint decision that maximizes some objective function.

In general, finding the optimal decision set is computationally intractable, even for a centralized authority. Therefore, there exist a multitude of techniques for arriving at a joint decision, which may be an approximation of the optimal. For instance, consensus algorithms offer a way for agents to converge as a group toward a unified decision [8]–[10]. In other settings, a game-theoretic approach is advantageous, where agents arrive at a joint decision which is some form of equilibrium (e.g., Nash equilibrium [11], Wardrop equilibrium [12], Stackelberg equilibrium [13], etc.). Of course even finding such equilibria can be challenging [14], but there are subclasses of problems where this can be done efficiently [15].

Another common approach to multiagent decision problems is a greedy algorithm [16], [17]. A common theme among greedy algorithms is that at each iteration of the algorithm, a myopic choice is made: simply pick the best

immediate option, ignoring the effect on future iterations of the algorithm. As with the other algorithms mentioned above, greedy algorithms in general are not always guaranteed to find an optimal solution to a given problem, however, they are often easy to implement, execute quickly, and in some cases provide some degree of optimality.

This work focuses on the scenario where a greedy algorithm is used to solve a multiagent decision problem. In this setting, a greedy algorithm is implemented by first ordering the agents. Then each agent sequentially makes its decision by choosing the action that maximizes the objective function, based solely on the decisions of previous agents in the sequence. An underlying element of the greedy algorithm is that the agents are able to coordinate with each other via some network. In the best case, such a network would allow for each agent to communicate with all other agents directly. In many applications, however, this is not realistic; communication between agents i and j must pass through other agents in the network. If i and j are on opposite ends of the network, or if the network has highly-limited bandwidth, this communication may be delayed. In light of this, two questions arise:

- 1) Given the structure of the communication network, how does the ordering of the agents affect the time it takes to complete the greedy algorithm?
- 2) Can the agents coordinate among themselves to find the ordering that will cause the greedy algorithm to complete as fast as possible?

We address the first question by showing, given a network structure, that the greedy algorithm finishes in $O(n^2)$ time steps for the worst ordering and $O(n)$ time steps for the best ordering. We then address the second question by presenting a fully-distributed algorithm whereby agents can find a near-optimal ordering while simultaneously running the greedy algorithm.

Of particular import in this work are submodular maximization problems, which are prevalent in modeling many applications, for e.g., see [18]–[21]. A key feature that is shared among the objective functions in these various domains is a property of *diminishing returns*. For example, in outbreak detection in networks, the added benefit of placing an outbreak sensor on a node in a network is valuable when there are few other sensors in the network, and less valuable when there are already many other sensors present. Objectives that exhibit such properties are *submodular*.

While such problems are NP-Hard in general, the property of submodularity can be exploited to show that certain algorithms can achieve near-optimal results. The seminal work in [22] shows that a centralized greedy algorithm can,

R. Konda and J. R. Marden are with the Department of Electrical and Computer Engineering at the University of California, Santa Barbara, CA and D. Grimsman is with the Department of Computer Science at Brigham Young University. This work is supported by AFOSR Grant #FA9550-20-1-0054 and AFOSR Grant #FA9550-21-1-0203.

in fact, provide a solution that is guaranteed to be within $1/2$ of the optimal solution. More sophisticated algorithms have pushed this guarantee from $1/2$ to $1 - 1/e \approx 0.63$ [23], [24]. Progress beyond this approximation frontier is not possible for polynomial time algorithms as it was also shown that no such algorithm can achieve higher guarantees than $1 - 1/e$, unless $P = NP$ [25].

Recently, work has emerged wherein submodular maximization problems are modeled as multiagent decision problems [26]–[29]. In Section V, we will show how our version of the greedy algorithm applies to this setting, and how this algorithm compares to existing techniques in terms of runtime and performance guarantees. It will be shown that the greedy algorithm will complete in fewer time steps than existing methods, while still maintaining $1/2$ -optimality in the resulting decision set.

In Section II, we introduce our model. In Section III, we present our main results on communication time guarantees versus different orderings. We empirically verify our theoretical results in Section IV and discuss the implications in submodular maximization problems in Section V. We conclude in Section VI. The relevant code is found at [30] and a full version of this paper with proofs is found at [31]. We sincerely thank Gilberto Diaz-Garcia for the helpful discussions.

II. MODEL

Consider a distributed optimization problem with n agents $\mathcal{I} = \{1, \dots, n\}$, where each agent is endowed with a decision or action set \mathcal{A}_i . We denote an action as $a_i \in \mathcal{A}_i$, and a joint action profile as $a \in \mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_n$. We assume that each agent i has the ability to “opt out” of participating in the decision process. This is modeled by having an action $a_i^\emptyset \in \mathcal{A}_i$, so that when agent i chooses action a_i^\emptyset , agent i is opting out. The quality of each joint action profile is evaluated with a global objective function $W(a) : \mathcal{A} \rightarrow \mathbb{R}_{\geq 0}$ that a system designer seeks to maximize. In other words, the goal of the system designer is to coordinate the agents to a joint action profile that satisfies

$$a^{\text{opt}} \in \arg \max_{a \in \mathcal{A}} W(a). \quad (1)$$

In general, solving the multi-agent decision problem in Eq. (1) is infeasible, due to computational, informational, communication constraints etc. Therefore, fast, distributed algorithms are employed to compute good approximate solutions. The *greedy algorithm* has cemented its place as a universal approach to arrive at approximate solutions in many application domains. In this algorithm, the set of agents is ordered (for instance, according to its index i) and then each agent sequentially solves the reduced optimization problem

$$\tilde{a}_i \in \arg \max_{a_i \in \mathcal{A}_i} W(\tilde{a}_1, \dots, \tilde{a}_{i-1}, a_i, a_{i+1}^\emptyset, \dots, a_n^\emptyset), \quad (2)$$

where each agent i chooses the best action \tilde{a}_i given that the previous agents in the sequence have also played their best action and the successive agents in the sequence have opted out. After each agent chooses according to Eq. (2), then

the algorithm is complete and the resulting set of decisions $(\tilde{a}_1, \dots, \tilde{a}_n)$ comprises the joint decision set \tilde{a} . The process completes in n time steps, where a time step is comprised of an agent making a decision and communicating that decision to future agents in the sequence.

However, the greedy algorithm makes a key assumption that agents have access to the decisions of the previous agents. In purely distributed systems, this assumption may be infeasible. There have been prior works that study the performance of the greedy algorithm with relaxed informational assumptions, in that agent i only knows the decisions of some strict subset of the previous agents $S \subset \{1, \dots, i-1\}$ [17], [32]. However, this work takes a different approach, where we assume that agents can make up for their informational deficiencies through a communication infrastructure. We model the communication constraints through an underlying graph structure $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where each vertex in \mathcal{V} corresponds to an agent in \mathcal{I} and each edge $(i, j) \in \mathcal{E} = \mathcal{V} \times \mathcal{V}$ implies that agents i and j can communicate with one another. The graph \mathcal{G} is assumed to be connected and undirected throughout this paper, unless explicitly stated. The set of agents that agent i can communicate with is agent i 's neighborhood \mathcal{N}_i .

The primary focus of this work is to examine the interplay between the communication graph \mathcal{G} and the order in which the greedy algorithm in Eq. (2) is solved under. For a given graph \mathcal{G} , the order $\pi : \mathcal{V} \rightarrow \mathcal{I}$ is defined by which label i given to each vertex v . Therefore, given \mathcal{G} , we would like to characterize the communication time guarantees of the worst order and the best order. To analyze the spectrum of possible guarantees with respect to different ordering methods, we define the following two quantities

$$T_{\min}(\mathcal{G}) = \min_{\pi} T(\mathcal{G}, \pi), \quad (3)$$

$$T_{\max}(\mathcal{G}) = \max_{\pi} T(\mathcal{G}, \pi), \quad (4)$$

where $T(\mathcal{G}, \pi)$ refers to the time it takes for the communication process to finish for a given graph \mathcal{G} and ordering π . We will use π_{best} and π_{worst} to refer to the orderings that are the solutions of Eq. (3) and Eq. (4) respectively. We remark that only in the full information setting, where \mathcal{G}_c is the complete graph, is the run-time for any order the same, with $T_{\max}(\mathcal{G}_c) = T_{\min}(\mathcal{G}_c) = n - 1$.

We can describe the k -hop communication, in which an agent i 's greedy action \tilde{a}_i is passed along to agents outside of its neighborhood \mathcal{N}_i , using the following graph-theoretic notation. A *walk* on the graph \mathcal{G} is a sequence of vertices $\gamma = (v_1, \dots, v_m)$, in which each successive pair $(v_j, v_{j+1}) \in \mathcal{E}$ for all $1 \leq j < m$. We denote the length of the walk as $|\gamma|$ being the number of vertices in the sequence. A *spanning walk* is a walk in which all vertices in the graph are visited and a *minimum spanning walk* is a spanning walk with shortest length. A *path* p is a walk in which all the vertices $\{v_j\}_{j \leq m}$ are all distinct. The expression of $T(\mathcal{G}, \pi)$ is given as

$$T(\mathcal{G}, \pi) = \sum_{i=1}^{n-1} \left(\min_{p_{i \rightarrow i+1}} |p_{i \rightarrow i+1}| - 1 \right), \quad (5)$$

where $p_{i \rightarrow i+1}$ is a path on the graph from the vertex (labeled with) i to the vertex $i + 1$. This expression is motivated by a natural communication process, where initially agent 1 computes its greedy action \tilde{a}_1 at time 0. Then agent 1 communicates \tilde{a}_1 to agent 2 through a k -hop walk through the graph, where each hop is assumed to take 1 time step. Then agent 2 computes \tilde{a}_2 given \tilde{a}_1 and passes both actions to agent 3 through another k -hop walk. Continuing this process, agent $n - 1$ passes $\{\tilde{a}_j\}_{j < n}$ to agent n and agent n computes \tilde{a}_n finishing the process. To isolate the run-time analysis with respect to only the communication time, we also assume that agents can solve for their greedy action \tilde{a}_i arbitrarily fast.

III. MAIN RESULTS ON COMMUNICATION TIME

A. Motivating Example

To make the communication process concrete, consider when the given communication graph is a line graph as shown in Figure 1. In this graph scenario, agent 1 initially computes its greedy response \tilde{a}_1 and passes the action it has played to agent 2 at $t = 0$. Then at $t = 1$, agent 2 (knowing \tilde{a}_1) can compute \tilde{a}_2 and passes both \tilde{a}_1 and \tilde{a}_2 along to agent 3. Continuing this to $t = 5$, agent 6 will have been passed the greedy actions of all previous agents 1 through 5 and play its greedy action \tilde{a}_6 , completing the greedy algorithm in Eq. (2). This will complete in $T(\mathcal{G}, \pi) = 5$ time steps, which is the best that one can hope for when implementing the greedy algorithm in a limited information setting.




Fig. 1: Example of a line graph, where we have labeled the vertices according to the best ordering π_{best} . In this example, the agents compute their greedy action and pass it down the line.

However, consider the following order in Figure 2. This situation can occur if the order π is improperly picked by the system operator. Under this ordering, agent 2 can only receive the greedy action of agent 1 through a 3-hop path through agents 6 and 4, since there is not a direct communication link between agent 1 and agent 2. Following this logic, the greedy algorithm will complete at time $T(\mathcal{G}, \pi) = 3 + 5 + 4 + 3 + 2 = 17$, which can be seen to be significantly higher than the previous well chosen order.




Fig. 2: Example of a line graph, but instead we consider the adversarial ordering π_{worst} in which the vertices are labeled intermittently. In this instance, the k -hop communication path must bounce back and forth between agents to complete the greedy algorithm.

Extending this argument to n agents, under a line graph, the greedy algorithm under the best ordering π_{best} will

complete in $T(\mathcal{G}, \pi_{\text{best}}) = n - 1$ steps and under the worst ordering π_{worst} will complete in $T(\mathcal{G}, \pi_{\text{worst}}) = \lfloor n^2/2 \rfloor - 1$ steps, where $\lfloor a \rfloor$ is the floor function. Therefore, there may be a significant gap in the communication complexity that results from choosing different orderings. We analyze the possible gap by characterizing the quantities $T_{\text{max}}(\mathcal{G})$ and $T_{\text{min}}(\mathcal{G})$ in this paper.

B. Communication Run-time Characterizations

We outline the main theorem of the paper below, where the worst case communication time over any graph structure is given for the best and worst orderings. The corresponding graph structures and orders that attain the worst-case communication time are displayed in Figure 2 and Figure 3.




Fig. 3: A 7 node star graph with agent 2 in the center. We note that the communication time $T(\mathcal{G}, \pi)$ on this graph using any ordering must be greater than $2 \cdot 7 - 4 = 10$. For $n \geq 3$ agents, the n -node star graph is the worst case graph for the best ordering π_{best} .

Theorem 1. Let $n \geq 3$ be the number of agents. The maximum communication time required to complete the greedy algorithm in Eq. (2) for any undirected, connected communication graph \mathcal{G} with the best and worst orderings is equal to

$$\max_{\mathcal{G}} T_{\text{min}}(\mathcal{G}) = 2n - 4 \quad (6)$$

$$\max_{\mathcal{G}} T_{\text{max}}(\mathcal{G}) = \lfloor n^2/2 \rfloor - 1, \quad (7)$$

where $T_{\text{min}}(\mathcal{G})$ is defined in Eq. (3) and $T_{\text{max}}(\mathcal{G})$ is defined in Eq. (4) and $\lfloor a \rfloor$ is the largest integer that is below a .

C. Distributed Orderings that are Near-Optimal

According to Theorem 1, there is a significant complexity gap between using the best ordering π_{best} and worst ordering π_{worst} for the communication time. However, finding the best order π_{best} in general is not practical either due to computational restrictions or lack of information about the graph. So we would like to be able compute orderings that get as close to the run-time with π_{best} as possible in a feasible manner. Therefore in this section, we construct an algorithm that can quickly compute a good ordering in conjunction with executing the greedy algorithm. An outline of the proposed algorithm is in Algorithm 1 with a fully distributed implementation found in [30]. The proposed design in essence computes a spanning walk on the graph that is close to the length of the minimum spanning walk through a variant of a depth-first search algorithm.

The distributed implementation of Algorithm 1 to compute an approximate solution to Eq. (1) offers significant benefits over other distributed approaches. The communication scheme is simple, which allows for linear-time guarantees. This also means that the message complexity is low, where the bulk of the message is comprised of the previous agent's actions and the communication is robust to time delays. Lastly, since the base of Algorithm 1 is the greedy algorithm, we also inherit the corresponding performance guarantees. To be able to run Algorithm 1, we assume that each agent (vertex) can store and access the following variables.

- $v.\text{actions} = \emptyset$ is the set of greedy actions that v knows.
- $v.\text{order} = \emptyset$ is the index in \mathcal{I} that v is labeled with.
- $v.\text{parent} = \emptyset$ is v 's parent in the depth first search.
- $v.\text{neighborhood}$ is the neighborhood set of v .

We also assume that a seed v_{seed} is given as the starting point of the Algorithm 1. The communication time of Algorithm 1 is equivalent to the total number of calls to MESSAGE, where the vertex v messages either a vertex that hasn't been visited or its parent $v.\text{parent}$. We keep track of the communication time through the variable t . The communication time guarantees of Algorithm 1 is given below.

Proposition 1. *Let n be the number of agents and $T_{\text{alg}}(\mathcal{G}, v_{\text{seed}})$ be the output of Algorithm 1 given a communication graph \mathcal{G} and a seed vertex $v_{\text{seed}} \in \mathcal{V}$. The maximum communication time for any undirected, connected \mathcal{G} and seed v_{seed} is*

$$\max_{\mathcal{G}, v_{\text{seed}}} T_{\text{alg}}(\mathcal{G}, v_{\text{seed}}) = 2n - 2 \quad (8)$$

Thus the communication guarantees of Algorithm 1 is only off by a constant of 2 from the optimal communication guarantee of $2n - 4$ from the best ordering π_{best} . We remark that this difference can be further reduced if the termination condition is changed from ' $v.\text{parent}$ is not empty' to ' $v.\text{order} = n$ ', where n is the number of agents.

D. Directed Communication Graphs

In this section, we consider the communication time guarantees with respect to the more general class of connected, directed graphs using different orderings. Under the class of undirected graphs, there is a significant gap in the communication time guarantees for the best π_{best} and the worst π_{worst} orderings. Not surprisingly, when we relax to the optimization problem $\max_{\mathcal{G}_{\text{dir}}} T_{\text{max}}(\mathcal{G}_{\text{dir}})$ over the class of directed graphs, the worst case guarantees also increase. However, when considering the optimization problem for the best ordering over directed graphs $\max_{\mathcal{G}_{\text{dir}}} T_{\text{min}}(\mathcal{G}_{\text{dir}})$, we have that the worst case guarantees are also of quadratic order. Therefore in directed graphs, the gap between the performance guarantees under different orderings is relatively small. This is elaborated on in the following lemma.

Lemma 1. *Let $n \geq 4$ be the number of agents. The maximum communication time required to complete the greedy algorithm in Eq. (2) for any directed, connected communication*

Algorithm 1 Distributed Near-Optimal Ordering

Require: graph \mathcal{G} and a vertex $v_{\text{seed}} \in \mathcal{V}$

Output: time t

initialize the time $t \leftarrow 0$

INIT($v_{\text{seed}}, \emptyset, \emptyset, 1$)

return t

procedure INIT($v, v_{\text{par}}, \alpha, i$)

label v as done

update $v.\text{order} \leftarrow i$ and $v.\text{parent} \leftarrow v_{\text{par}}$

let v compute \tilde{a}_i from Eq. (2) given actions α

update $v.\text{actions} \leftarrow \alpha \cup \{\tilde{a}_i\}$

MESSAGE($v, v.\text{actions}$)

end procedure

procedure MESSAGE(v, α)

update $v.\text{actions} \leftarrow \alpha$

if exists w in $v.\text{neighborhood}$ not labeled done **then**

increment $t \leftarrow t + 1$

INIT($w, v, \alpha, |\alpha|+1$)

else if $v.\text{parent}$ is not empty **then**

increment $t \leftarrow t + 1$

MESSAGE($v.\text{parent}, \alpha$)

end if

end procedure

graph \mathcal{G}_{dir} with the best and worst ordering is

$$\max_{\mathcal{G}_{\text{dir}}} T_{\text{min}}(\mathcal{G}_{\text{dir}}) \geq \left\lfloor \frac{n}{2} \right\rfloor \cdot \left\lceil \frac{n}{2} \right\rceil, \quad (9)$$

$$\max_{\mathcal{G}_{\text{dir}}} T_{\text{max}}(\mathcal{G}_{\text{dir}}) = (n - 1)^2, \quad (10)$$

where $\lfloor a \rfloor$ is largest integer smaller than a and $\lceil a \rceil$ is the smallest integer larger than a .

IV. SIMULATIONS

We analyze our theoretical results for the communication time guarantees empirically through a simulation, presented in Figure 4 and Figure 5. The code used to generate the simulations can be found in [30]. We use the model of Erdos-Renyi networks [33], where each possible undirected pair of edges (i, j) has a probability $P \leq 1$ of existing, to generate a sample set of possible graph structures. For Figure 4, we sample 200 instances of Erdos-Renyi networks with 6 nodes and a probability parameter of $P = .3$. For each graph, we calculate the communication time for the greedy algorithm using the best ordering π_{best} , the ordering given by Algorithm 1, and a randomly assigned ordering. For Figure 5, we sample 300 instances of Erdos-Renyi networks with 40 nodes and a probability parameter of $P = .05$. In Figure 5, we calculate the communication time for only the ordering given by Algorithm 1 and a randomly assigned ordering.

We observe in Figure 4 that indeed the best ordering π_{best} achieves the lowest distribution of communication times, centered closely to $n = 6$. The distribution of communication times of the ordering given in Algorithm 1 is noticeably close to the one of π_{best} , with indeed no run-times over $2n = 12$. The distribution of the communications using random orderings does perform the worst with the largest spread. In Figure

5, we compare communication times from the ordering from Algorithm 1 directly with the random ordering, as computing the best ordering π_{best} is infeasible for large n . We see the same trends in Figure 4 reflected in a more extreme fashion. The distribution of communication times using the ordering of Algorithm 1 is still upper bounded by $2n = 80$. However, the distribution times of communication times using the random ordering is now centered much higher with a larger spread as well. Therefore, we see significant benefits from using the ordering from Algorithm 1 rather than the naive approach of using random ordering.

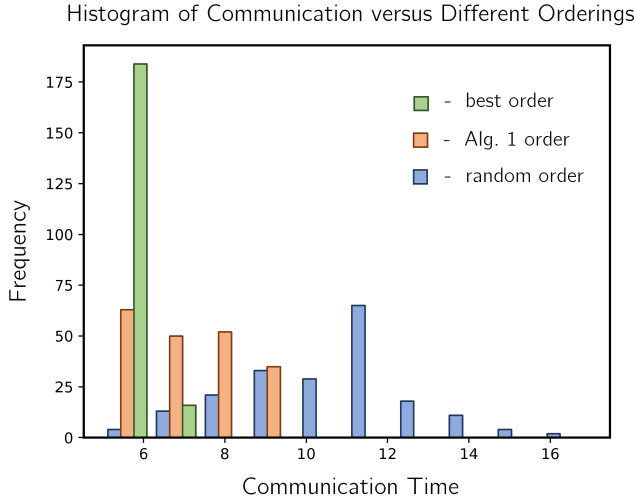


Fig. 4: We show the distribution over of communication times needed to complete the greedy algorithm for 200 instances of random graphs generated by a Erdos-Renyi process with respect to the random ordering, the best ordering π_{best} , and the ordering from Algorithm 1. For the graph parameters $n = 6$ number of agents and $P = .3$ probability of an edge existing, we see that π_{best} gives slightly lower average communication times than the ordering from Algorithm 1, but both offer significant improvements over the random ordering.

V. SUBMODULAR MAXIMIZATION

In this section, we discuss submodular maximization problems, which can be modeled as multiagent decision problems. Consider a base set of elements E , and let $a_i \subseteq E$, $\mathcal{A}_i \subseteq 2^E$, and $a_i^\emptyset = \emptyset$. The objective function takes the form $W(a) = f(\cup_{a_i \in a} a_i)$, where $f : 2^E \rightarrow \mathbb{R}$ has the following properties for any $A \subseteq B \subseteq E$:

- 1) *Submodular*: $f(A \cup \{x\}) - f(A) \geq f(B \cup \{x\}) - f(B)$ for all $x \in E \setminus B$
- 2) *Monotone*: $f(A) \leq f(B)$
- 3) *Normalized*: $f(\emptyset) = 0$

In this setting, it has been shown that the greedy algorithm that is implemented in Algorithm 1 guarantees that $W(\tilde{a}) \geq (1/2)W(a^{\text{opt}})$, where \tilde{a} is defined in Eq. (2).




Fig. 5: We compare the communication times under the random ordering to the times under the ordering algorithm proposed in Algorithm 1. For the graph parameters $n = 40$ number of agents and $P = .05$ probability of an edge existing, we see a marked decrease in the communication time from the random ordering to using the proposed algorithm over a set of 300 randomly generated graphs.

A. Comparison with Other Distributed Algorithms

As mentioned previously, much work has been done to develop other algorithms to solve submodular maximization. For instance, [28] presents a similar distributed algorithm, using a multilinear extension, and a distributed pipage rounding technique. At each time step, each agent performs a calculation for each action based on a sample of K actions drawn from a probability distribution. After T time steps, the performance guarantee is $(1 - 1/e)(1 - (2d(\mathcal{G})n + n/2 + 1)(n/T))$ with probability at least $1 - 2nTe^{-K/(8T^2)}$. Thus, for high T and $K = O(T^2)$, there is a high probability that the algorithm gives the $1 - 1/e$ guarantee. Using this information, the algorithm could provide a $1/2$ guarantee only for $T \geq 4.78(2d(\mathcal{G})n^2 + n^2/2 + 1)$, and only with high probability when $K = O(T^2)$.

The paper [29] describes a Jacobi-style algorithm, where at each time step agent i creates a strategy profile, i.e., a probability distribution across each of its actions. Then, it chooses K of those values to share with its neighbors to propagate through the network. It was shown that the resulting decision set approaches being within $1/2$ the optimal as the number of iterations increases. It is only shown in the paper that the probability of achieving the $1/2$ guarantee is $1 - O(1/T)$ rather than an explicit time expectation. However, the examples in the paper suggest that it may take $T \geq n^2$ or more time steps to realize this.

In another example, [27] presents the Constraint-Distributed Continuous Greedy (CDCG), a consensus-style algorithm, in which agent i shares an m -vector with all its neighbors at each time step, where m is the number

of actions available to i . It is shown that the resulting decision set approaches being within $1 - 1/e$ of the optimal as the number of iterations T increases. The error in the performance guarantee vanishes at a rate of $O(n^{5/2}/T)$, and therefore, it may require $T \geq n^{5/2}$ time steps in order to reach an acceptable error.

In each of the three methods listed above, each time step requires each agent do perform some calculation for each of its actions. The time requirement for each to reach an acceptable solution is expected to be greater than $2n - 2$, which is the number of time steps it takes to complete Algorithm 1. This suggests that there is a tradeoff between performance guarantees and time complexity: Algorithm 1 achieves the $1/2$ guarantee quickly, but other algorithms converge to a solution within $1 - 1/e$, but more slowly.

VI. CONCLUSION

In this work, we analyze the greedy algorithm in a multi-agent context. More specifically, when the agents have limited information about the other agents but can communicate with a subset of the other agents, we characterize the effects of using different agent orderings on the communication time over the network. First, we have shown that under the best ordering over any graph, the communication time is $2n - 4$, where n is the number of agents. When considering the worst ordering, we also show that this bound increases to $\lceil n^2/2 \rceil - 1$ holding for any communication network. We also provide an algorithm that can be used in a distributed fashion to obtain a communication time guarantee of $2n - 2$, which is a constant factor away from the best communication guarantee. Furthermore, we verify these results computationally in Section IV and describe the implications in the context of submodular optimization problems in Section V. Future work is comprised of extending this work to analyze the interplay between agent orderings, the communication time guarantees, and possible performance guarantees.

REFERENCES

- [1] P. Zhao, S. Suryanarayanan, and M. G. Simoes, "An energy management system for building structures using a multi-agent decision-making control methodology," *IEEE Transactions on Industry Applications*, vol. 49, no. 1, pp. 322–330, 2013.
- [2] Y. Luo, K. Liu, and D. N. Davis, "A multi-agent decision support system for stock trading," *IEEE Network*, vol. 16, pp. 20–27, jan 2002.
- [3] M. T. Kahil, A. Dinar, and J. Albiac, "Cooperative water management and ecosystem protection under scarcity and drought in arid and semiarid regions," *Water Resources and Economics*, vol. 13, pp. 60–74, 2016.
- [4] K. Madani, "Game theory and water resources," *Journal of Hydrology*, vol. 381, no. 3-4, pp. 225–238, 2010.
- [5] L. Cruz-Piris, D. Rivera, S. Fernandez, and I. Marsa-Maestre, "Optimized sensor network and multi-agent decision support for smart traffic light management," *Sensors (Switzerland)*, vol. 18, no. 2, 2018.
- [6] H. Bai and B. Zhao, "A survey on application of swarm intelligence computation to electric power system," *Proceedings of the World Congress on Intelligent Control and Automation (WCICA)*, vol. 2, no. 60421002, pp. 7587–7591, 2006.
- [7] N. SinghPal and S. Sharma, "Robot Path Planning using Swarm Intelligence: A Survey," *International Journal of Computer Applications*, vol. 83, no. 12, pp. 5–12, 2013.
- [8] G. T. Nguyen and K. Kim, "A survey about consensus algorithms used in Blockchain," *Journal of Information Processing Systems*, vol. 14, no. 1, pp. 101–128, 2018.
- [9] S. Kar and J. M. Moura, "Distributed consensus algorithms in sensor networks: Quantized data and random link failures," *IEEE Transactions on Signal Processing*, vol. 58, no. 3 PART 1, pp. 1383–1400, 2010.
- [10] D. Ongaro and J. Ousterhout, "In search of an understandable consensus algorithm," *Proceedings of the 2014 USENIX Annual Technical Conference, USENIX ATC 2014*, pp. 305–319, 2019.
- [11] J. F. Nash *et al.*, "Equilibrium points in n-person games," *Proceedings of the national academy of sciences*, vol. 36, no. 1, pp. 48–49, 1950.
- [12] J. G. Wardrop, "Road paper. some theoretical aspects of road traffic research.," *Proceedings of the institution of civil engineers*, vol. 1, no. 3, pp. 325–362, 1952.
- [13] H. Von Stackelberg, *Market structure and equilibrium*. Springer Science & Business Media, 2010.
- [14] I. Gilboa and E. Zemel, "Nash and correlated equilibria: Some complexity considerations," *Games and Economic Behavior*, vol. 1, no. 1, pp. 80–93, 1989.
- [15] D. Monderer and L. S. Shapley, "Potential games," *Games and economic behavior*, vol. 14, no. 1, pp. 124–143, 1996.
- [16] B. Marzouki, O. Belkahlil Driss, and K. Ghédira, "Multi Agent model based on Chemical Reaction Optimization with Greedy algorithm for Flexible Job shop Scheduling Problem," *Procedia Computer Science*, vol. 112, pp. 81–90, 2017.
- [17] B. Ghareisifard and S. L. Smith, "On distributed submodular maximization with limited information," *American Control Conference*, vol. 2016-July, pp. 1048–1053, 2016.
- [18] A. Krause, A. Singh, and C. Guestrin, "Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies," *Journal of Machine Learning Research*, vol. 9, pp. 235–284, 2008.
- [19] A. Singh, W. Kaiser, M. Batalin, A. Krause, and C. Guestrin, "Efficient planning of informative paths for multiple robots," in *International Joint Conference on Artificial Intelligence*, pp. 2204–2211, 2007.
- [20] A. Clark and R. Poovendran, "A submodular optimization framework for leader selection in linear multi-agent systems," in *IEEE Conference on Decision and Control*, pp. 3614–3621, IEEE, 2011.
- [21] J. Leskovec, A. Krause, C. Guestrin, C. Faloutsos, J. Vanbriesen, and N. Glance, "Cost-effective outbreak detection in networks," in *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 420–429, 2007.
- [22] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher, "An analysis of approximations for maximizing submodular set functions—i," *Mathematical programming*, vol. 14, no. 1, pp. 265–294, 1978.
- [23] G. Calinescu, C. Chekuri, M. Pál, and J. Vondrák, "Maximizing a monotone submodular function subject to a matroid constraint," *SIAM Journal on Computing*, vol. 40, no. 6, pp. 1740–1766, 2011.
- [24] Y. Filmus and J. Ward, "The power of local search: maximum coverage over a matroid," in *Symposium on Theoretical Aspects of Computer Science*, pp. 601–612, LIPIcs, 2012.
- [25] U. Feige, "A threshold of $\ln n$ for approximating set cover," *Journal of the ACM*, vol. 45, no. 4, pp. 634–652, 1998.
- [26] B. Mirzasoleiman, A. Karbasi, R. Sarkar, and A. Krause, "Distributed submodular maximization," *Journal of Machine Learning Research*, vol. 17, no. 1, pp. 8330–8373, 2016.
- [27] A. Robey, A. Adibi, B. Schlotfeldt, G. J. Pappas, and H. Hassani, "Optimal algorithms for submodular maximization with distributed constraints," *arXiv*, vol. xxx, pp. 1–20, 2019.
- [28] N. Reza zadeh and S. S. Kia, "Distributed strategy selection: A submodular set function maximization approach," *arXiv preprint arXiv:2107.14371*, 2021.
- [29] B. Du, K. Qian, C. Claudel, and D. Sun, "Jacobi-style iteration for distributed submodular maximization," *arXiv preprint arXiv:2010.14082*, 2020.
- [30] R. Konda, "Kondaacc2021." <https://github.com/rohit-konda/KondaACC2021.git>, 2021.
- [31] R. Konda, D. Grimsman, and J. Marden, "Execution order matters in greedy algorithms with limited information," *arXiv preprint arXiv:2111.09154*, 2021.
- [32] D. Grimsman, M. S. Ali, J. P. Hespanha, and J. R. Marden, "Impact of information in greedy submodular maximization," in *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pp. 2900–2905, IEEE, 2017.
- [33] P. Erdos, A. Rényi, *et al.*, "On the evolution of random graphs," *Publ. Math. Inst. Hung. Acad. Sci.*, vol. 5, no. 1, pp. 17–60, 1960.