The Impact of Message Passing in Agent-Based Submodular Maximization

David Grimsman, Matthew R. Kirchner, João P. Hespanha and Jason R. Marden

***Abstract*— Submodular maximization problems are a relevant model set for many real-world applications. Since these problems are generally NP-Hard, many methods have been developed to approximate the optimal solution in polynomial time. One such approach uses an agent-based greedy algorithm, where the goal is for each agent to choose an action from its action set such that the union of all actions chosen is as high-valued as possible. Recent work has shown how the performance of the greedy algorithm degrades as the amount of information shared among the agents decreases, whereas this work addresses the scenario where agents are capable of sharing more information than allowed in the greedy algorithm. Specifically, we show how performance guarantees increase as agents are capable of passing messages, which can augment the allowable decision set for each agent. Under these circumstances, we show a near-optimal method for message passing, and how much such an algorithm could increase performance for any given problem instance.**

I. Introduction

Submodular maximization problems are relevant to many fields and applications, including sensor placement [1], outbreak detection in networks [2], maximizing and inferring influence in a social network [3], [4], document summarization [5], clustering [6], assigning satellites to targets [7], path planning for multiple robots [8], and leader selection and resource allocation in multiagent systems [9], [10]. An important similarity among these applications is the presence of an objective function which exhibits a "diminishing returns" property. For instance, consider a company choosing locations for its retail stores. If, for a given city, the company has no retail stores, and chooses to add a location in that city, the marginal gain in revenue from that store would be higher than if the company chose a similar city where it already had 100 retail stores. Objectives (such as revenue in this instance) satisfying this property are *submodular*.

While submodular minimization can be solved in polynomial time [11], submodular maximization is an NP-hard problem for certain subclasses of submodular function. Therefore, much effort has been devoted to finding and improving algorithms which approximate the optimal solution in polynomial time. A key result of this line of research is that algorithms exist which give strong guarantees as to how well the optimal solution can be approximated.

One such algorithm is the greedy algorithm, first proposed in the seminal work [12]. Here it was shown that for certain classes of constraints the solution provided by the greedy algorithm must be within $1 - 1/e \approx 0.63$ of the optimal, and within $1/2$ of the optimal for the more general case of constraints [13]. Since then, more sophisticated algorithms have been developed to show that there are many instances of the submodular maximixation problem that can be solved efficiently within the $1-1/e$ guarantee [14], [15]. It has also been shown that progress beyond this level of optimality is not possible using a polynomial-time algorithm, where the indicator step for the time complexity is the evaluation of the objective function [16].

In addition to the strong performance guarantees, a nice benefit to using the greedy algorithm to solve a submodular maximization problem is that it is easy to implement, even in distributed settings. One recent line of research has studied a distributed version of the greedy algorithm, which can be implemented using a set of n agents, each with its own action set [17]. In this algorithm, the agents are ordered and choose sequentially, each agent greedily choosing its best action relative to the actions which the previous agents have chosen. The solution provided is the set of all actions chosen. Like the standard greedy algorithm, it has been shown that this distributed greedy algorithm guarantees the solution is within 1/2 the optimal [17].

In this setting, recent literature has emerged which attempts to quantify how information impacts the performance of the algorithm, specifically as the information among the agents degrades. For instance, [18] and [19] describe how the 1/2-guarantee decreases as agents can only observe a subset of the actions chosen by previous agents. The work in [20] shows how an intelligent choice of which action to send to future agents can recover some of this loss in performance. Other work has also begun to explore how additional knowledge of the structure of the action sets can offset this loss [21], [22].

This paper addresses the impact of information in the other direction: how increasing the amount of actions that can be shared among the agents improves the performance. We introduce the concept of *message passing*, wherein agents not only choose an action as part of the algorithm solution, but also choose some actions to share with future agents. Future agents may choose these shared actions as part of the solution, thus message passing is a way to augment the action sets of future agents in the sequence to offset any agent that may not have access to valuable actions.

Message passing gives rise to two key questions: what policy should agents use to select actions to pass and how does it affect performance? We address the first question in Section II-C, where we propose an augmented greedy policy

D. Grimsman (davidgrimsman@ucsb.edu), M. R. Kirchner (kirchner@ucsb.edu), J. R. Marden (jmarden@ece.ucsb.edu), and J. P. Hespanha (hespanha@ece.ucsb.edu) are with the Department of Electrical and Computer Engineering, University of California, Santa Barbara, CA

This research was supported in part by AFRL contracts FA8750-18-C-0014 and FA95550-20-1-0054 and ONR grant N00014-20-1-2359.

that is near-optimal in the limit of a large number of agents or a large number of shared messages. The performance question is addressed both from a worst-case and a best-case perspective. It turns out that it is possible to find "bad" problem instances for which message sharing brings little benefit when the number of agents is large. Moreover, we prove in Theorem 1 that this is so *regardless of the message passing policy used*. On the flip side, there are also "good" problem instances for which message passing can improve performance significantly, by a multiplicative factor that can be as large as the number of agents (Theorem 2) and such performance gains are achieved by the proposed augmented greedy policy.

II. Model

Let S be a set of elements and $f : 2^S \to \mathbb{R}$ a scalar-valued function. We restrict f to have the following three properties:

- Normalized: $f(\emptyset) = 0$.
- Monotone: $f(A) \leq f(B)$ for $A \subseteq B \subseteq S$.
- Submodular: $f(A \cup \{s\}) - f(A) \geq f(B \cup \{s\}) - f(B)$ for all $A \subseteq B \subseteq S$ and $s \in S \setminus B$.

This paper focuses on a subclass of distributed submodular optimization problems. To that end, let $N = \{1, \ldots, n\}$ be a set of agents where each $i \in N$ is associated with a local set of elements $S_i \subseteq S$ that contains the various sets of elements from which agent i can choose. For a given set $A \subseteq S$, we use $(A)^l$ to mean the subsets of A of size at most l. We focus on the scenario where each agent i can select as most $k \geq 1$ elements from S_i, and hence we express this choice as $x_i \in (S_i)^k$. We denote a choice profile by $x = (x_1, \ldots, x_n) \in (S_1)^k \times \cdots (S_n)^k$, and will evaluate the objective f of this profile x as $f(x) = f(\cup_i x_i)$. The central goal is to solve

$$\mathrm{OPT}(f, (S_i)_i, k) = \max_{x \in (S_1)^k \times \cdots \times (S_n)^k} f(x). \qquad (1)$$

We denote $x^{\mathrm{opt}}(f, (S_i)_i, k)$ to mean a choice profile which maximizes (1).

A. The Greedy Algorithm

It is well known that characterizing an optimal allocation $x^{\mathrm{opt}}(f, (S_i)_i, k)$ is an NP-Hard problem in the number of agents for certain subclasses of submodular functions f. However, there are very simple algorithms that can attain near optimal behavior for this class of submodular optimization problems. One such algorithm, termed the greedy algorithm [13], proceeds according to the following rule: each agent $i \in N$ sequentially selects their choice $x_i \in (S_i)^k$ by greedily choosing the action which yields the greatest benefit to the objective f, i.e.,

$$x_i^{\mathrm{ng}} \in \arg\max_{x_i \in (S_i)^k} f\left(x_i, \{x_j^{\mathrm{ng}}\}_{j<i}\right). \qquad (2)$$

Note that while the greedy algorithm can be implemented in a distributed fashion, there is an implicit informational demand on the system as each agent $i \in N$ must be aware of the choices of all previous agents $j < i$, i.e., $\{x_1, \ldots, x_{i-1}\}$.

In addition to being easy to implement, the greedy algorithm is also high-performing: it is guaranteed to produce a solution x^{ng} which is within 1/2 of the optimal [13], i.e., $f(x^{\mathrm{ng}}) \geq (1/2)\dot{f}(x^{\mathrm{opt}})$. In fact, the greedy algorithm is shown to give the highest performance guarantees possible for any algorithm which runs in polynomial time for some classes of distributed submodular maximization problems [16].

B. A Motivating Example

Consider a scenario in which n flying vehicles carry on-board cameras that capture images of ground vehicles of interest and return their pixel coordinates. Each vehicle $i \in N$ has access to a large collection of pixel coordinate measurements taken by its own camera, which comprise the local element set S_i. However, each vehicle i needs to select a much smaller subset of these measurements (no more than k) to send to a centralized location for data fusion. The goal is to select the best set of k measurements that each vehicle should send to the centralized location so that an optimal estimate $\hat{\theta}$ of the ground vehicle's position θ can be recovered by fusing the measurements from all the vehicles. It was shown in [23] that an optimal estimator that achieves the Cramér–Rao lower bound results in an error covariance matrix that can be written as

$$\mathrm{E}\left[(\hat{\theta} - \theta)(\hat{\theta} - \theta)'\right] = \Big(Q_0 + \sum_{s \in S} Q_s\Big)^{-1}$$

where Q_0 is a symmetric positive definite matrix that encodes a-priori information about the position of the ground vehicle, $S = \cup_i S_i$ is the complete set of measurements, and Q_s is a symmetric positive semi-definite matrix that depends on the relative position and orientation of the flying vehicle's camera with respect to the ground vehicle for each measurement $s \in S$. In addition, it was shown that the function $f : 2^S \to \mathbb{R}$ defined for any $A \subseteq S$ as

$$f(A) = \log\det\Big(Q_0 + \sum_{s \in A} Q_s\Big) - \log\det Q_0 \qquad (3)$$

is normalized, monotone, and submodular [23]. It turns out that maximizing (3) corresponds to the so-called D-optimality, which essentially minimizes the volume of confidence intervals.

C. The Greedy Algorithm with Message Passing

The focus of this work is on understanding the degree to which inter-agent communications can be exploited to improve the performance guarantees associated with the greedy algorithm. Accordingly, here we propose message passing, wherein each agent $i \in N$ is tasked with making both a selection and communication decision, denoted by x_i and z_i respectively. We focus on the situation where agent i can communicate up to $m \geq 0$ of its measurements to the forthcoming agents $j > i$, i.e., $z_i \in (S_i)^m$. Agent j can then select its choices either from among original set S_j or to also include some subset of the communicated measures $z_1 \cup \cdots \cup z_{j-1}$. Accordingly, we replace the decision-making

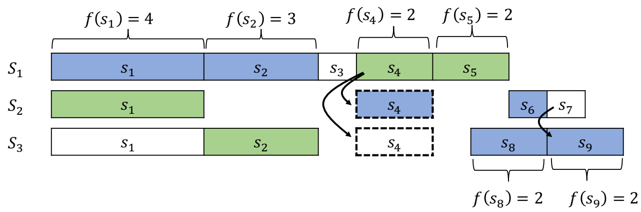

(a) An example problem, where $n = 3$, $k = 2$, $m = 1$. Each box represents an element of S, and row represents S_i, i.e., the elements in S to which agent i has access. The function f is represented by the width of each box, where the width of elements not specifically labeled in the diagram is 1. For $A \subseteq S$, $f(A)$ is the total amount of horizontal space covered by the elements in A. For instance, $f(\{s_6, s_8\}) = 2$ and $f(\{s_5, s_6, s_8\}) = 3$. Here we assume the use of some policy π for selection and message passing. The arrows indicate the message passing dictated by π, for instance $z_1^\pi = \{s_4\}$. The boxes with the dashed outline indicate that s_4 is not in S_2 or S_3, but is included as part of the agents' augmented action set, should they choose to use it. The boxes shaded in blue indicate the elements x_i^π chosen by π, and the boxes shaded in green are the optimal choices, where those differ.

Solution method	x_1	x_2	x_3	$f(x)$
Optimal	$\{s_4, s_5\}$	$\{s_1, s_6\}$	$\{s_2, s_9\}$	14
Independent	$\{s_1, s_2\}$	$\{s_1, s_6\}$	$\{s_1, s_2\}$	8
Greedy	$\{s_1, s_2\}$	$\{s_6, s_7\}$	$\{s_8, s_9\}$	11
Policy π	$\{s_1, s_2\}$	$\{s_4, s_6\}$	$\{s_8, s_9\}$	13

(b) A table representing the performance for 4 different solution methods. First, the optimal solution to (1) is given. Then, the independent solution is shown, which is the solution where each agent chooses independent of the other agents. The solution to the greedy algorithm assumes that the agents choose according to (2). Finally, last row assumes agents choose according to some policy π which employs message passing.

Fig. 1: An example problem illustrating the extended model from Section II-C.

rule of each agent $i \in N$ given in (2) with a new rule π_i which dictates how the agent selects its choice x_i and its message z_i in response to the previous selections and messages. In particular, we focus on rules π_i of the form

$$\{x_i^\pi, z_i^\pi\} = \pi_i(S_i, \{x_j^\pi\}_{j<i}, \{z_j^\pi\}_{j<i}), \tag{4}$$

where $x_i^\pi \in (S_i \cup z_1^\pi \cup \cdots \cup z_{i-1}^\pi)^k$ and $z_i^\pi \in (S_i)^m$. For k, m, and n, we denote $\Pi(k, m, n)$ as the set of admissible policy profiles $\pi = \{\pi_1, \ldots, \pi_n\}$. It is important to highlight that the performance of π is ultimately gauged by the performance of the resulting allocation $x^\pi = (x_1^\pi, \ldots, x_n^\pi)$, as the communicated messages $(z_1^\pi, \ldots, z_n^\pi)$ are merely employed to influence these decision making rules. See Figure 1 for an example which highlights this extended model.

The following algorithm highlights an opportunity for message passing to potentially improve the performance guarantees associated with the greedy algorithm through augmenting the agents' choice sets.

Definition 1 (Augmented Greedy Algorithm): In the augmented greedy algorithm each agent $i \in N$ is associated with a selection rule π_i as in (4) of the form

$$x_i^{\mathrm{ag}} \in \underset{x_i \subseteq (S_i \cup z_1^{\mathrm{ag}} \cup \cdots \cup z_{i-1}^{\mathrm{ag}})^k}{\arg\max} f(x_i, \{x_j^{\mathrm{ag}}\}_{j<i}), \tag{5a}$$

$$z_i^{\mathrm{ag}} = z_i^k \cup z_i^{m-k}, \tag{5b}$$

$$z_i^k \in \underset{z_i \in (S_i)^{\min(k,m)}}{\arg\max} f(z_i, \{x_j^{\mathrm{ag}}\}_{j\le i}) \tag{5c}$$

$$z_i^{m-k} \in (S_i)^{\max(m-k,0)} \tag{5d}$$

The communication depicted above entails that when $m \le k$ each agent forwards the best m measurements that it is unable to select to the remaining agents. When $m > k$, the best k measurements are chosen, and the remaining $m - k$ measurements can be chosen arbitrarily. Then, each remaining agent can choose whether or not its selection should include these augmented choices. Note that we will require a policy to be deterministic, so the rules in (5a)–(5d) do not constitute a specific policy, since (5d) doesn't specify a rule for finding z_i^{m-k}, and also since the $\arg\max$ in (5a) and (5c) may not be unique. Therefore we refer to a policy which has the form of (5a)–(5d) as an augmented greedy policy.

III. A Worst-Case Analysis

In this section we explore whether message passing can increase worst-case guarantees beyond the 1/2 guarantee of the nominal greedy algorithm. We show that an augmented algorithm is a near-optimal algorithm in this setting, but also show that the benefits in terms of worst-case analysis decrease as the number of agents increases. Additionally, we show that even an optimal policy does not generally increase performance by much.

A. Performance Measure

Given a submodular, monotone, normalized function f with element set S, we measure the performance of a policy $\pi \in \Pi(k, m, n)$ and across all policies, respectively, as

$$\gamma_\pi(f, S, k) := \min_{\{S_i \subseteq S\}_{i\in N}} \frac{f(x^\pi)}{\mathrm{OPT}(f, (S_i)_i, k)}, \tag{6}$$

$$\gamma(\Pi(k, m, n), f, S) := \max_{\pi \in \Pi(k,m,n)} \gamma_\pi(f, S, k), \tag{7}$$

where the minimization in (6) captures the notion that we are interested in how well π performs regardless of how the base set S is distributed among the different agents.

This section focuses on worst-case guarantees over all criteria f, which are characterized by

$$\gamma_\pi = \inf_{f:2^S \to \mathbb{R}} \gamma_\pi(f, S, k), \tag{8}$$

$$\gamma = \inf_{f:2^S \to \mathbb{R}} \gamma(\Pi(k, m, n), f, S, k), \tag{9}$$

where the (worst-case) $\inf$ are taken over all possible criteria f that are submodular, monotone, and normalized. When we only consider decision-rules aligned with the greedy algorithm, as in (2), the bound given in (8) is $1/2$ (see [13], [17]).

B. Main Result

Theorem 1: For any $k \ge 1, m \ge 0$ and $n \ge 2$,

$$\gamma \le \frac{1}{2 - \frac{\min((n-1)m/k, 1)}{n-1+\min((n-1)m/k,1)}}, \tag{10}$$

i.e., no policy can guarantee a performance higher than the above expression. Furthermore, when π is an augmented greedy policy, then

$$\gamma \ge \gamma_\pi \ge \frac{1}{2 - \frac{(\min(m/k,1))^{n-1}}{\sum_{i=0}^{n-1}(\min(m/k,1))^i}}. \tag{11}$$

We will give the formal proof for the theorem below, with a brief description here. The statement in (10) is shown by presenting a canonical example, for which no policy can guarantee a performance above the expression in (10). The lower bound in (11) is proven by showing that if an augmented greedy policy is implemented, the system performance cannot be below the expression in (11).

Assuming that one could design a policy π^* such that γ meets the upper bound in (10), we see that in general the guaranteed performance does not increase much above the 1/2 guaranteed by the standard greedy algorithm given by (2), especially for large n. However, in cases where n is small, one could see a nontrivial increase in guaranteed performance: for instance, when $n = 2$ and $m/k = 1$, then $\gamma = 2/3$.

Another observation about Theorem 1 is that the bounds in (10) and (11) are equal when $m/k \geq 1$ - a range for m where increasing m does not affect γ. This implies that increasing m higher than k does not offer any benefit in terms of worst-case performance guarantees. Therefore, when considering constraints on how much information agents may share with one another, it may not benefit the system to increase capacity beyond that bound.

We also see that the upper and lower bounds are equal depending on how many agents are in the system. For instance, when $n = 2$, the lower and upper bounds are equal. This implies that an augmented greedy policy is optimal in this setting. Likewise, as $n \to \infty$, we see that the bounds are increasingly tighter, showing that for high n, any augmented greedy policy is a near-optimal policy. This provides motivation for further studying the augmented greedy algorithm in the next section.

C. Proof for Theorem 1

To show (10), we introduce a problem instance $(f, (S_i)_i, k)$ such that for any π, $f(x^\pi)/f(x^{\text{opt}})$ is lower than the bound in (10).

Fix π and assume first that $(n-1)m \leq k$. Suppose that f and $(S_i)_i$ are as represented in Figure 2a. Here the format of example is the same as in Figure 1. We assume that all of the small rectangles are of width 1, and that the large rectangle s^{big} is of length $k(n-1)+1$. Essentially, for agent $i \in \{1, \dots, n-1\}$, all elements in S_i are identical according to f, since none of these agents is aware of the elements in S_n.

Therefore, without loss of generality, assume that π_i selects the blue elements for x_i^π, $i = 1, \dots, n-1$. Likewise, assume without loss of generality that π_i selects the orange elements for z_i^π, $i = 1, \dots, n-1$. Notice that $|S_n| = m(n-1) \leq k$, so $x_n^\pi = S_n$ is feasible, and an optimal choice regardless of the previous agents' decisions. This implies that $f(x^\pi) = k(n-1) + m(n-1)$.

A set of optimal selections, where different from blue, are shaded in green, showing that $f(x^{\text{opt}}) = 2k(n-1) + m(n-1)$. It follows that for this problem instance,

$$\frac{f(x^\pi)}{f(x^{\text{opt}})} = \frac{k(n-1) + m(n-1)}{2k(n-1) + m(n-1)}$$

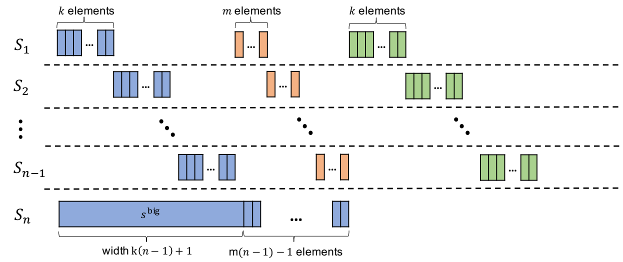

(a) An example for proving (10) when $m(n-1) \leq k$. The key is that $|S_n| \leq k$, and the orange elements offer no value beyond S_n. Each small box has width 1 and s^{big} has width $k(n-1)+1$. Thus $f(x^\pi) = k(n-1) + m(n-1)$ and $f(x^{\text{opt}}) = 2k(n-1) + m(n-1)$.

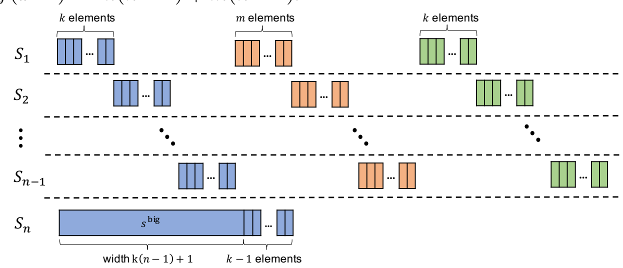

(b) An example for proving (10) when $m(n-1) \geq k$. Unlike the example above, here S_n consists of exactly k elements, thus $f(x^\pi) = k(n-1) + k$ and $f(x^{\text{opt}}) = 2k(n-1) + k$.

Fig. 2: Problem instances used in the proof for Theorem 1. This representation is similar to that in Figure 1, in that each row represents some S_i and $f(A)$ is the amount of horizontal space covered by the boxes in $A \subseteq S$. Blue squares represent elements chosen by policy π, orange are those elements which are passed as messages using policy π, and green are the optimal choices x^{opt}, when those differ from the algorithm. We have omitted the dashed boxes, with the understanding that an orange box in S_i is available to agents $j > i$.

$$= \frac{1}{2 - \frac{\min((n-1)m/k, 1)}{n-1+\min((n-1)m/k, 1)}}. \tag{12}$$

In the case where $m(n-1) \geq k$, we leverage the example in Figure 2b. Here S_i are the same as in Figure 2a, for $i = 1, \dots, n-1$, implying again without loss of generality that $x_1^\pi, \dots, x_{n-1}^\pi$ are the respective blue elements. In this case, however, $|S_n| = k$, but note that $f(x_1^\pi, \dots, x_{n-1}^\pi, x_n) = kn$ for any $x_n \in (S_n)^k$, thus $f(x^\pi) = kn$. The green elements are the optimal selection, where different from the blue, showing that $f(x^{\text{opt}}) = 2k(n-1) + k$. In this case we see that

$$\frac{f(x^\pi)}{f(x^{\text{opt}})} = \frac{kn}{2k(n-1) + k} = \frac{n}{2(n-1) + 1}. \tag{13}$$

Equations (12) and (13) establish (10) for all cases.

We now show the lower bound on γ in (11) by assuming π takes on the form in (5a)–(5d), and showing that for any $(f, (S_i)_i, k)$, $f(x^\pi)/f(x^{\text{opt}})$ cannot be lower than the expression in (11). We first define the function Δ:

$$\Delta(A|B) := f(A \cup B) - f(B), \tag{14}$$

for $A, B \subseteq S$. This can be thought of as the marginal contribution of A given B according to f. Likewise, define

$$\Delta^l(A|B) := \max_{\tilde{A} \in (A)^l} \Delta(\tilde{A}|B). \tag{15}$$

This can be thought as the value of the highest-valued subset of A with size at most l, with respect to B. Note that if $l \geq |A|$, then $\Delta^l(A|B) = \Delta(A|B)$. It should also be clear that

$$\begin{aligned}\Delta^l(A|B) \geq & \frac{\min(l,|A|)}{|A|}\Delta(A|B) \\ = & \min(l/|A|,1)\Delta(A|B). \end{aligned} \tag{16}$$

We also appeal to the following result:

Lemma 1: Assume a policy π is applied to problem instance $(f,(S_i)_i,k)$ such that for $\alpha \geq 1$

$$\alpha\Delta^k(z_i^\pi|x_{1:i}^\pi) \geq \Delta(x_i^{\text{opt}}|x_{1:i}^\pi). \tag{17a}$$

Then

$$\frac{f(x^\pi)}{f(x^{\text{opt}})} \geq \frac{1}{2-\frac{1}{\sum_{i=0}^{n-1}\alpha^i}} \tag{18}$$

We omit the proof here to conserve space [1]. In the current setting it suffices to show a valid value for α in order to show the lower bound on γ_π (11). It follows that

$$\Delta^k(z_i^{\text{ag}}|x_{1:i}^{\text{ag}}) \geq \Delta(z_i^k|x_{1:i}^{\text{ag}}) \geq \Delta(x_i^{\text{opt}}|x_{1:i}^{\text{apx}}) \tag{19a}$$

$$\geq \Delta^m(x_i^{\text{opt}}|x_{1:i}^{\text{apx}}) \geq \min(m/k,1)\Delta(x_i^{\text{opt}}|x_{1:i}^{\text{apx}}) \tag{19b}$$

where (19a) is true by (5c) (recall that z_i^k is defined in (5c)), and (19b) is true by (16). Thus we can set $\alpha = 1/\min(m/k,1)$. Then by Lemma 1,

$$\begin{aligned}\frac{f(x^{\text{apx}})}{f(x^{\text{opt}})} \geq & \frac{1}{2-\frac{1}{\sum_{i=0}^{n-1}(1/\min(m/k,1))^i}} \\ = & \frac{1}{2-\frac{(\min(m/k,1))^{n-1}}{\sum_{i=0}^{n-1}(\min(m/k,1))^i}}.\end{aligned}$$

■

IV. A Best-Case Analysis

In this section we consider an optimistic approach to understanding how an increase in message passing affects the performance of the system. We assume the use of an augmented greedy policy, which Theorem 1 shows is near-optimal, and study its potential effects. In particular, since the nominal greedy algorithm in (2) is merely the augmented greedy algorithm when $m = 0$, comparing solutions of the two on an instance-by-instance basis can give insight into the potential benefits of message passing. While the previous section focused on worst-case scenarios, here we ask the question: how much *could* action sharing increase performance for any individual problem instance? The following result gives insight on the answer to this question:

Theorem 2: For any $n \geq 2$, any $m \geq 0$, and any problem instance $(f,(S_i)_i,k)$, the following holds:

$$2+\min\left(\frac{m}{k},n-1\right) \geq \frac{f(x^{\text{ag}})}{f(x^{\text{ng}})} \geq \frac{1}{2-\frac{(\min(m/k,1))^{n-1}}{\sum_{i=0}^{n-1}(\min(m/k,1))^i}}, \tag{20}$$

[1]The full proof is found at https://arxiv.org/abs/2004.03050

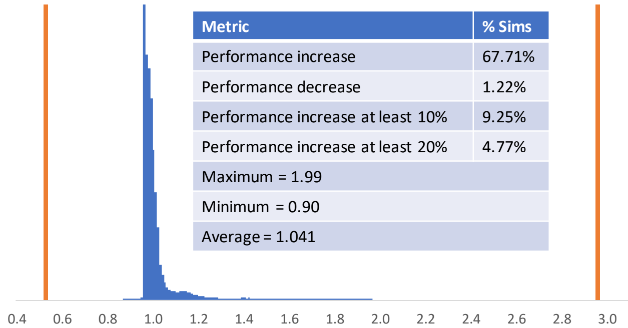

Fig. 3: Histograms of the relative performance for 1 million random simulations of the vehicle camera estimation problem, where $k = m = 2$ and $n = 3$. About $1/3$ of samples are at 1 (bar cropped to improve visualization), which is interpreted as augmented greedy performing the same as nominal greedy. Roughly 2/3 of examples showed improvement, whereas only about 1% showed a decrease in performance. The orange bars indicate the theoretical bounds provided in Theorem 2.

where x^{ag} is the solution to an augmented greedy algorithm and satisfies (5a) for all i, and x^{ng} is the solution to a nominal greedy algorithm and satisfies (2) for all i.

For space considerations, we present this result without its proof [2]. The focuses mainly on the upper bound in (20), relying, in a similar fashion to Theorem 1, on the properties of π^{ag} given in (5a)–(5d), and the submodularity and monotonicity of f. The lower bound, somewhat trivially, corresponds to the lower bound for γ_π in Theorem 1.

While the result in Theorem 1 was a somewhat negative result (in general, increasing m does not provide much higher performance guarantees), here we see a strong upside to message passing. For any given problem instance $(f,(S_i)_i,k)$, one could increase $f(x)$ by a factor of $n+1$. And, though this will not be the case for every problem instance, we will see in the next section that one can expect $f(x)$ to increase by some amount.

Theorem 2 also gives insight into how much increasing m might help any given problem instance. Whereas previously we saw that increasing m above k offered no further guarantees, here we see potential benefits to increasing m all the way up to $k(n-1)$. We also see the potential drawback: on any given instance, $f(x)$ could decrease by almost a factor of 1/2, although we will see in the next section that such scenarios aren't necessarily common.

V. Numerical Example

In this section, we present the results for instances of the state estimation problem in Section II-B, where $n = 3$ flying vehicles move along curved paths, each carrying a side-looking camera with a 90-degree field of view, with focal length 50 pixels and measurement noise in the image plane with standard deviation $\sigma = 1$ pixel. In each problem instance, the start/stop positions and curvature of each vehicle path are uniformly randomly selected. Vehicle i collects 10 images along its path, which are processed into $|S_i| = 10$

[2]The full proof is found at https://arxiv.org/abs/2004.03050

measurement estimates of 3 targets (which are uniformly randomly placed) by sending $k = 2$ measurement estimates to the fusion center. Each measurement estimate $s \in S$ has an associated matrix Q_s which carries information about the quality of that estimate, based on the vehicle location and camera angle. Details of how to construct the objective function f given in (3), along with the matrices Q_0 and Q_s, are found in [23].

Figure 3 compares the resulting decision set of the augmented greedy algorithm (5a)–(5d), where $m = 2$, and that of the nominal greedy algorithm (2). The orange bars show the Theorem 2 bounds for these problem instances:

$$0.5714 \leq \frac{f(x^{\mathrm{ag}})}{f(x^{\mathrm{ng}})} \leq 3. \tag{21}$$

While one could construct problem instances that meet either bound, our 1 million randomly-generated simulations did not yield a value of $f(x^{\mathrm{ag}})/f(x^{\mathrm{ng}})$ lower than 0.9 or higher than 2. The results show, however, that about 2/3 of the problem instances showed a performance increase with the augmented greedy algorithm. Only about 1% of instances showed a decrease in performance. Therefore, while the theoretical bounds were not seen in simulation, we see that message passing improves performance on average.

VI. Conclusion

In this paper we have shown how message passing affects the performance guarantees of a group of agents using the greedy algorithm. We showed that when $m > k$, we receive no additional guarantees. We also showed that a simple augmented greedy policy gives near-optimal performance guarantees. Using such a policy, this paper explored how much performance could increase for any problem instance, and showed by simulation that these results are relevant in a real-world application.

Future work will continue to explore message passing, first by attempting to create tighter bounds on γ. We also have some preliminary results related to settings where (5a) and (5c) cannot be computed directly, only approximated. Another direction could be to apply these results to situations where agents can only see the actions of a subset of previous agents, and again ask questions about what actions should be shared and selected.

References

[1] A. Krause, C. Guestrin, A. Gupta, and J. Kleinberg, "Near-optimal sensor placements: Maximizing information while minimizing communication cost," in *Proceedings of the International Conference on Information Processing in Sensor Networks*. ACM, 2006, pp. 2–10.

[2] J. Leskovec, A. Krause, C. Guestrin, C. Faloutsos, C. Faloutsos, J. VanBriesen, and N. Glance, "Cost-effective outbreak detection in networks," in *Proceedings of the ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2007, pp. 420–429.

[3] D. Kempe, J. Kleinberg, and É. Tardos, "Maximizing the spread of influence through a social network," in *Proceedings of the ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2003, pp. 137–146.

[4] M. Gomez-Rodriguez, J. Leskovec, and A. Krause, "Inferring networks of diffusion and influence," *ACM Transactions on Knowledge Discovery from Data (TKDD)*, vol. 5, no. 4, p. 21, 2012.

[5] H. Lin and J. Bilmes, "A class of submodular functions for document summarization," in *Proceedings of the Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*. Association for Computational Linguistics, 2011, pp. 510–520.

[6] B. Mirzasoleiman, A. Karbasi, R. Sarkar, and A. Krause, "Distributed submodular maximization: Identifying representative elements in massive data," in *Advances in Neural Information Processing Systems*, 2013, pp. 2049–2057.

[7] G. Qu, D. Brown, and N. Li, "Distributed greedy algorithm for multi-agent task assignment problem with submodular utility functions," *Automatica*, vol. 105, pp. 206–215, 2019.

[8] A. Singh, A. Krause, C. Guestrin, W. J. Kaiser, and M. A. Batalin, "Efficient planning of informative paths for multiple robots." in *International Joint Conferences on Artificial Intelligence*, vol. 7, 2007, pp. 2204–2211.

[9] A. Clark and R. Poovendran, "A submodular optimization framework for leader selection in linear multi-agent systems," in *Conference on Decision and Control and European Control Conference*. IEEE, 2011, pp. 3614–3621.

[10] J. R. Marden, "The role of information in distributed resource allocation," *IEEE TCNS*, vol. 4, no. 3, pp. 654–664, Sept 2017.

[11] M. Grötschel, L. Lovász, and A. Schrijver, "The ellipsoid method and its consequences in combinatorial optimization," *Combinatorica*, vol. 1, no. 2, pp. 169–197, 1981.

[12] 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.

[13] M. L. Fisher, G. L. Nemhauser, and L. A. Wolsey, "An analysis of approximations for maximizing submodular set functions II," in *Polyhedral combinatorics*. Springer, 1978, pp. 73–87.

[14] G. Calinescu, C. Chekuri, M. Pál, and J. Vondrák, "Maximizing a submodular set function subject to a matroid constraint," in *International Conference on Integer Programming and Combinatorial Optimization*. Springer, 2007, pp. 182–196.

[15] M. Gairing, "Covering games: Approximation through non-cooperation," in *International Workshop on Internet and Network Economics*. Springer, 2009, pp. 184–195.

[16] U. Feige, "A threshold of ln n for approximating set cover," *Journal of the ACM (JACM)*, vol. 45, no. 4, pp. 634–652, 1998.

[17] B. Gharesifard and S. L. Smith, "Distributed submodular maximization with limited information," *Trans. on Control of Network Systems*, vol. 5, no. 4, pp. 1635–1645, 2017.

[18] D. Grimsman, M. S. Ali, J. P. Hespanha, and J. R. Marden, "The impact of information in greedy submodular maximization," *Trans. on Control of Network Systems*, 2018.

[19] D. Grimsman, J. H. Seaton, P. N. Brown, and J. R. Marden, "The cost of denied observation in multiagent submodular optimization," in *IEEE Conference on Decision and Control (CDC)*. IEEE, 2020.

[20] D. Grimsman, J. P. Hespanha, and J. R. Marden, "Strategic information sharing in greedy submodular maximization," in *2018 IEEE Conference on Decision and Control (CDC)*. IEEE, 2018, pp. 2722–2727.

[21] M. Corah and N. Michael, "Distributed submodular maximization on partition matroids for planning on large sensor networks," in *Conference on Decision and Control*. IEEE, 2018, pp. 6792–6799.

[22] ——, "Distributed matroid-constrained submodular maximization for multi-robot exploration: Theory and practice," *Autonomous Robots*, vol. 43, no. 2, pp. 485–501, 2019.

[23] M. R. Kirchner, J. P. Hespanha, and D. Garagic, "Heterogeneous measurement selection for vehicle tracking using submodular optimization," in *Proc. of the 2020 IEEE Aerospace Conf.*, 2020, pp. 1–10.