

The Impact of Measurement Passing in Sensor Network Measurement Selection

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

Abstract—This paper considers a set of sensors, which as a group are tasked with taking measurements of the environment and sending a small subset of the measurements to a centralized data fusion center, where the measurements will be used to estimate the overall state of the environment. The sensors’ goal is to send the most informative set of measurements so that the estimate is as accurate as possible. This problem is formulated as a submodular maximization problem, for which there exists a well-studied greedy algorithm, where each sensor sequentially chooses a set of measurements from its own local set, and communicates its decision to the future sensors in the sequence. In this work, sensors can additionally share measurements with one another, in order to augment the decision set of each sensor. We explore how this increase in communication can be exploited to improve the results of the nominal greedy algorithm. Specifically, we show that this measurement passing can improve the quality of the resulting measurement set by up to a factor of $n + 1$, where n is the number of sensors.

I. INTRODUCTION

Multiagent robotic systems have received increased attention, with research addressing a large variety of challenges posed by such systems, for instance collaborative motion planning [2], [3], [4] and allocation [5], [6], [7], sensor coverage [8], [9], and state estimation [10], [11]. One area of importance is the coordination of vehicles by sharing of information on a resource constrained communication network. This paper studies a system of heterogeneous sensors are collecting data at high rate. Each sensor is connected to an off-site data fusion center with limited bandwidth, and the group of sensors is tasked with sending a small but high-quality set of measurements to the data fusion center for processing. In addition, the sensors can coordinate which measurements to send via a local communication network. This problem has been studied in [12], where only a small subset of measurements was carefully selected to be sent to the data fusion center. However, greater performance at the fusion center can be observed if a small amount of communication is utilized to coordinate among sensors. Motivated by this observation, this paper seeks to

quantify the extent to which inter-sensor coordination, through limited, localized information exchange, can improve fusion performance over the standard setting.

In this paper it is assumed that the objective function the sensors seek to maximize is *monotone submodular*, meaning that more measurements cannot decrease the objective function, and that adding a measurement to a set exhibits a “diminishing returns” property. Monotone submodularity is a common assumption for applications which need to quantify information content in some way [13], [14], [15]. Therefore, the sensors’ global objective is formulated as a submodular maximization problem. Such problems have also been well-studied in the literature, with applications to sensor placement [14], outbreak detection in networks [16], maximizing and inferring influence in a social network [17], [18], document summarization [19], clustering [20], assigning satellites to targets [21], path planning for multiple robots [22], and leader selection and resource allocation in multiagent systems [23], [24].

Maximizing a submodular function is an NP-hard problem for many relevant subclasses of submodular functions and constraints [25]. 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 in [26]. Here it was shown that for certain classes of problems, (e.g., submodular maximization subject to a uniform matroid constraint) the solution provided by the greedy algorithm must be within a multiplicative factor of $1 - 1/e \approx 0.63$ of the optimal, and within $1/2$ of the optimal for the more general class of problems [27]. Since then, more sophisticated algorithms have been developed to show that the submodular maximization problem that can be solved efficiently within the $1 - 1/e$ guarantee [28], [29]. 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 number of evaluations of the objective function [30].

One line of research has studied distributed methods for solving submodular maximization; see, for instance [31], [32]. The work in [33] addresses a similar problem as this paper, where vehicles are communicating measurements to a centralized entity. The constraints, however, are different from the current setting in that a subset of vehicles is chosen to share *all* measurements rather than all vehicles sharing a

This paper was submitted for review on September 23, 2020; a preliminary version appears in [1]. This research was supported in part by AFRL contracts FA8750-18-C-0014 and FA95550-20-1-0054 and ONR grants N00014-20-1-2359 and N00014-20-1-2093.

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

Result	Measure	Lower bound	Upper bound
Theorem 1	$\frac{\text{AUG. GREEDY}}{\text{NOMINAL GREEDY}}$	$\frac{1}{2 - \frac{(\min(m/k, 1))^{n-1}}{\sum_{i=0}^{n-1} (\min(m/k, 1))^i}}$	$2 + \min(n-1, m/k)$
Theorem 2	$\frac{\beta\text{-APPROX. AUG. GREEDY}}{\beta\text{-APPROX. NOM. GREEDY}}$	$1 + \frac{1}{\beta} - \frac{1}{\beta \frac{(\beta \min(m/k, 1))^{n-1}}{\sum_{i=0}^{n-1} (\beta \min(m/k, 1))^i}}$	$\frac{2}{\beta} + m/k, \text{ if } m \leq k$ $1 + \frac{1}{\beta} \left(1 + \min\left(n-1, \frac{m}{k}\right)\right), \text{ if } m > k$

Fig. 1: A brief summary of the results from this paper, where a sensor selects up to k measurements as its “action” and up to m measurements to share. We explore how well an augmented greedy algorithm, which includes measurement passing, performs. Theorem 1 provides a range as to how well the augmented greedy algorithm can perform versus the nominal greedy algorithm. Theorem 2 describes the scenario where each agent can only approximate the solution to its local problem within a factor of β .

subset of measurements. Nevertheless, it is shown that an auction-style method of communication among the agents results in a greedy solution and thus a $1 - 1/e$ guarantee. Of particular interest to this work is a distributed version of the greedy algorithm, where the collective decisions result from a sequential process where partial decisions are made using a defined selection rule [34]. Each sensor greedily chooses the measurements that provide the most added value relative to those chosen by the previous sensors in the sequence. Like the standard greedy algorithm, it has been shown that this distributed greedy algorithm guarantees the solution is within $1/2$ the optimal [34], [27].

In this setting, recent literature has emerged which explores relaxing the requirement that each sensor has the ability to compute its contribution relative to the decisions of *all* previous sensors. For instance, [35] describes how the $1/2$ -guarantee decreases as sensors can only access a subset of the measurements chosen by previous sensors. The work in [36] shows how a particular type of communication between the sensors can be employed to recover some of this loss in performance; for example, rather than communicate its own decision to future sensors, a sensor can choose to communicate a more valuable decision of another sensor. Work has also begun to explore how additional knowledge of the structure of the initial measurement distribution can offset this loss [37], [38].

This paper addresses the extent to which inter-sensor communications can improve the performance of distributed algorithms. Specifically, we consider a type of submodular maximization problem with n sensors, where each sensor has available a large set of measurements and needs to downselect this large set to a subset of k measurements to be sent to the fusion center. This transmission to the fusion center is overheard by the other sensors (or explicitly communicated to the other sensors). On top of this, we have measurement passing, by which a sensor can forward up to m of its own measurements that subsequent sensors may choose to pass to the fusion center instead (or in combination) to their own measurements. This effectively relaxes the standard constraint in the distributed setting that sensor i can only send measurements from its own local set. As an example, consider a scenario where

two flying vehicles are trying to identify the positions of a set of targets. Each vehicle captures many images, but can only send k of them to a central satellite, which uses the sent measurements to estimate the location of the targets. In this scenario, vehicle 1 can use the local communication network to share with vehicle 2 the k images which it sent to the satellite, and can additionally send m more. Vehicle 2 could then send to the satellite any combination of k images from its original set as well as these new communicated measurements. In the extreme case where vehicle 2 was not able to capture any “valuable” images by itself, it could still send to the fusion center the m that came from vehicle 1, thus offsetting a potentially poor system performance.

The focus on this paper is to study how the increased communication allowed by measurement passing can improve the value of the resulting set of measurements. Theorem 1 addresses this directly by giving bounds on how well a measurement passing policy can perform compared to the distributed greedy algorithm: it shows that there exist problem instances for high m where measurement passing can outperform the greedy algorithm by a multiplicative factor of $n + 1$. For smaller m , that factor reduces to $2 + m/k$. Theorem 1 also shows that any measurement passing algorithm can be outperformed by the greedy algorithm for carefully chosen problem instances, but always by a factor less than 2, and as low as $2 - 1/n$ for high m .

Theorem 2 addresses the practical issue that solving the “local” problem that each sensor must solve of (i) selecting the k “best” measurements to send to the fusion center and (ii) selecting the m “best” measurements to forward to other sensors, can be by themselves intractable problems. This is typically the case for the example scenario described above, when each of the flying vehicles has available a large collection of local measurements. Realistically, the optimal solution to the local problem must be approximated using some computationally-feasible algorithm. Assuming that agents can approximate the solution to the local problem within a factor of β , Theorem 2 shows how these local approximations impact the results of Theorem 1. Interestingly, these local approximations do not impact the resulting performance guarantees in a significant way. A summary of the theoretical results can be found in

Figure 1.

Finally, this paper provides a numerical example to show how measurement passing can improve performance. We show that on average, measurement passing helps the most with few agents (e.g., $n = 3$) and many targets (e.g., 10 targets). However, even with a large number of agents and fewer targets (e.g., $n = 10$ and 4 targets), message passing improves performance in the vast majority of the cases (over 99%).

II. MODEL

Consider a set of n sensors, where each sensor $i \in \{1, \dots, n\}$ has a local set of noisy measurements, denoted S_i . Each sensor can communicate up to k measurements to a fusion center, where the total nk measurements received will be used to estimate the state of the environment. The collective goal of the sensors is to send the most informative set of measurements to the fusion center so that the estimation is as close as possible to the actual state of the environment.

We model such an environment as a multiagent decision problem where the choice set of each agent $i \in N$ is $(S_i)^k := \{S' \subseteq S : |S'| \leq k\}$ for the space of possible measurements S and where the system-level objective is captured by a function $f : 2^S \rightarrow \mathbb{R}$. Thus each problem instance can be defined by the tuple $I = (f, S_1, \dots, S_n)$. In this work we assume that the objective function f has the following properties:

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

We will refer to such functions simply as *submodular*. In the classical setting, the sensors attempt to solve

$$\max_{x_i \in (S_i)^k} f(x_1 \cup \dots \cup x_n). \quad (1)$$

This problem is an instance of a submodular maximization problem subject to a partition matroid constraint, a known NP-Hard problem¹.

Example (Flying Vehicles [12]). *Consider the scenario where the sensors are cameras carried on board n flying vehicles that capture images of ground targets and return their pixel coordinates. Each vehicle $i \in \{1, \dots, n\}$ 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 satellite for data fusion. The goal of the vehicles is to select the best set of k measurements that each vehicle should send to the satellite so that an optimal estimate $\hat{\theta}$ of the targets' positions θ can be recovered by fusing the measurements received from all the vehicles.*

To facilitate this goal, one can employ the use of the Fisher Information Matrix $\text{FIM}(x)$ for a set of measurements $x \subseteq S$, which is defined as follows:

$$\text{FIM}(x) := Q_0 + \sum_{s \in x} Q_s, \quad (2)$$

¹Furthermore, it can also be shown that this particular subclass of submodular maximization is itself NP-Hard.

where

$$Q_0 := \frac{\partial \log p(\theta)}{\partial \theta} \cdot \frac{\partial \log p(\theta)^T}{\partial \theta}, \quad (3)$$

$$Q_s := \mathbb{E}_\theta \left[\frac{\partial \log p(s|\theta)}{\partial \theta} \cdot \frac{\partial \log p(s|\theta)^T}{\partial \theta} \middle| \theta \right], \quad (4)$$

and $p(\theta)$ is the *a-priori* probability density function of θ and $p(s|\theta)$ is the likelihood of measurement $s \in x$. The positive semidefinite matrices Q_0 and Q_s encode the prior information and the informative contribution of measurement s , respectively. The FIM is helpful in the current setting, given the Cramér-Rao lower bound (CRLB), which states that for an unbiased estimator,

$$\mathbb{E} \left[(\hat{\theta}(x) - \theta)(\hat{\theta}(x) - \theta)^T \right] \geq \text{FIM}(x)^{-1}, \quad (5)$$

where we use \geq in the sense that if $A \geq B$, then $A - B$ is positive semidefinite. According to (5), for any optimal estimator that achieves the CRLB, a set x of measurements that “minimizes” $\text{FIM}(x)^{-1}$ also minimizes the error covariance. A scalar metric that is commonly used to measure the information content of a set of measurements is the *D-optimality* [39], which in our context can be defined by

$$f(x) := \log \frac{\det(\text{FIM}(x))}{\det(\text{FIM}(\emptyset))}, \quad (6)$$

which has been shown to be submodular [14], [40].

A. The Greedy Algorithm

Although (1) is intractable in general, there are simple algorithms that can attain near optimal behavior for this class of submodular optimization problems. One such algorithm, termed the *greedy algorithm* [27], proceeds according to the following rule: each sensor i sequentially selects its choice $x_i \in (S_i)^k$ by “greedily” choosing the action which yields the greatest immediate benefit to the objective f , i.e.,

$$x_i \in \arg \max_{\tilde{x} \in (S_i)^k} f(x_1 \cup \dots \cup x_{i-1} \cup \tilde{x}). \quad (7)$$

While this greedy algorithm can be implemented in a distributed fashion, there is an informational demand on the system, as each sensor i must be aware of the union of the choices of all previous sensors. Further, each sensor must also be able to compute the optimal choice as defined in (7).

While relatively simple, the greedy algorithm is also high-performing as it yields a solution $x = (x_1, \dots, x_n)$ which is within 1/2 of the optimal, i.e., $f(x_1 \cup \dots \cup x_n) \geq (1/2)\text{OPT}$, where OPT is the value of the solution to (1).

B. Measurement Sharing

This work explores the notion of *measurement sharing* as an extension of the greedy algorithm, where we relax the constraint that the k measurements sent to the fusion center by sensor i are a subset of S_i . Specifically, consider the case where sensor i can share up to $m \geq 0$ of the measurements in S_i to the forthcoming sensors $j > i$, and we denote by $z_i \in (S_i)^m$ those shared measurements. The subsequent sensors $j > i$ can then select their choices from among

their original set S_i , but also can include some of the shared measurements $z_1 \cup \dots \cup z_{j-1}$ from previous sensors in the sequence.

In this paper we consider the set of decision-making algorithms of the form $\pi = (\pi_1, \dots, \pi_n)$, where π_i is a rule employed by sensor i to select $k \geq 0$ measurements and “communicate” $m \geq 0$ measurements. Specifically, given the decisions of previous agents $x_1^\pi, \dots, x_{i-1}^\pi$ and measurements shared from previous agents $z_1^\pi, \dots, z_{i-1}^\pi$, π_i specifies both decision and communication of appropriate dimension, i.e.,

$$\{x_i^\pi, z_i^\pi\} \in \pi_i^{k,m}(S_i, x_{1:i-1}^\pi, z_{1:i-1}^\pi), \quad (8)$$

subject to the constraint that $x_i^\pi \in (S_i \cup z_1^\pi \cup \dots \cup z_{i-1}^\pi)^k$ and $z_i^\pi \in (S_i)^m$. This constraint ensures that each sensor i can only select elements either from its own set S_i or elements shared by previous sensors $j < i$. We denote $x^\pi(I)$ to be the resulting decision set for policy π on problem instance I , and abuse notation so that $f(x^\pi(I)) = f(x_1^\pi \cup \dots \cup x_n^\pi)$.

Measurement sharing effectively generalizes the nominal greedy algorithm in that it yields solutions which are not in the set $(S_1)^k \times \dots \times (S_n)^k$. In the same way that the nominal greedy algorithm approximates (1), a measurement passing policy π approximates

$$\max_{\substack{x_i \in (S_i)^m \\ x_i \in (S_i \cup z_1 \cup \dots \cup z_{i-1})^k}} f(x_1 \cup \dots \cup x_n) \quad (9)$$

When $m = 0$, (9) is equivalent to (1), therefore it is also at least NP-Hard in general.²

We note that the set of allowable policies is sequential in nature - of course, this need not be the case in general. For instance, one could consider a distributed auction as a type of policy that could select a set of nk measurements from $S_1 \cup \dots \cup S_n$. However, under these circumstances, it's not immediately clear how such a set of policies scales with the number of allowed measurements to be shared. Since the greedy algorithm has a well-known performance guarantee and the set of policies in $\Pi^{k,m}$ are well-defined as m scales, this setting allows one to focus on how increasing m affects performance. The study of more general policies will be a topic for future work.

Here we give an example policy, which will be analyzed in this paper.

Definition 1 (Augmented Greedy Policy). A policy $\pi \in \Pi^{k,m}$ is an augmented greedy policy if each sensor $i \in \{1, \dots, n\}$ is associated with a selection rule π_i of the form


$$x_i^\pi \in \arg \max_{\tilde{x} \subseteq (S_i \cup z_1^\pi \cup \dots \cup z_{i-1}^\pi)^k} f(x_1^\pi \cup \dots \cup x_{i-1}^\pi \cup \tilde{x}) \quad (10a)$$

$$z_i^\pi = z_i^k \cup z_i^{m-k}, \text{ where} \quad (10b)$$

$$z_i^k \in \arg \max_{\tilde{z} \in (S_i)^{\min(m,k)}} f(x_1^\pi \cup \dots \cup x_i^\pi \cup \tilde{z}) \quad (10c)$$

$$z_i^{m-k} \in \arg \max_{\tilde{z} \in (S_i)^{\max(m-k,0)}} f(x_1^\pi \cup \dots \cup x_i^\pi \cup z_i^k \cup \tilde{z}) \quad (10d)$$

²Like (1), it can be shown that (9) is a submodular maximization problem subject to a matroid constraint, for which there exist generic centralized algorithms that approximate the solution [27], [28]. However, this particular subclass of matroid, which is sequentially constructed, has not been explicitly studied as a submodular maximization constraint, as far as we know.



(a) An example problem, where $n = 3$, $k = 2$, $m = 1$. Each box represents an element of S , and each row represents S_i for each sensor, i.e., the local measurements to which the sensor 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, if the agents choose the elements shaded in blue, then $f(x^\pi)$ is represented by the combined lengths of the blue line segments at the bottom of the image. Clearly f is submodular. In this example π is an augmented greedy policy. The arrows indicate the measurement 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 sensors' augmented decision set, should they choose to use it.

Solution method	x_1	x_2	x_3	$f(x)$
Greedy	$\{s_1, s_2\}$	$\{s_6, s_7\}$	$\{s_8, s_9\}$	11
Augmented Greedy (π)	$\{s_1, s_2\}$	$\{s_4, s_6\}$	$\{s_8, s_9\}$	13

(b) A table representing the performance for 3 different solution methods. First, the greedy algorithm, where sensors choose according to (7), is shown. The second row assumes sensors choose according to an augmented greedy policy.

Fig. 2: An example problem illustrating measurement sharing introduced in Section II-B.

Here each sensor greedily selects the k best measurements for x_i based on what measurements have previously been selected. When $m \leq k$, the messages in z_i are the “next best” m measurements in S_i . When $m > k$, z_i is selected in two stages: first, the “next best” k measurements in S_i are selected, then the “next best” $m-k$ after that.³ We note that the rules in (10) are not deterministic: the $\arg \max$ may be multivalued. Thus there are many augmented greedy policies that satisfy (10) in conjunction with some tiebreaking rule. We similarly use the term *nominal greedy policies* in conjunction with the decision rule in (7). See Figure 2 for an example problem instance where an augmented greedy policy is used.

III. COMPARISON TO GREEDY

In this section, we compare the nominal greedy policies to augmented greedy policies. Since nominal greedy policies are equivalent to the subset of augmented greedy policies where $m = 0$, this comparison illustrates how measurement passing can improve system performance. We note that in many realistic settings, solving the local optimization problems (7) and (10) may not be feasible. Here we assume the vehicles can solve these problem precisely in order to highlight the advantage of measurement passing. In Section IV we relax this assumption.

³It can be shown that a policy that does not use this two-stage selection rule performs no better than an augmented greedy policy. In particular, as $m \rightarrow \infty$, the performance guarantees decrease to the point that message passing gives no advantage.

A. Direct Comparison to Greedy

The first comparison we make between the two classes of policies is direct: we compare the ratio between the two for any given problem instance. Given k, m , we denote $\Pi^{k,m}$ to be the set of all admissible policies, i.e., all policies that satisfy (8). We also denote $\Pi_{\text{ag}}^{k,m}$ to be the set of augmented greedy policies and denote $\Pi_{\text{ng}}^{k,m}$ to be the set of nominal greedy policies that satisfy (7). Note that $\Pi_{\text{ag}}^{k,0} = \Pi_{\text{ng}}^{k,0}$.

Theorem 1. *Consider the measurement selection problem with n sensors. Then for any $k \geq 1, m \geq 0$ the best-case gain in performance and worst-case loss in performance associated with an optimal measurement passing policy within $\Pi^{k,m}$ satisfies*

$$\begin{aligned} \max_{\pi \in \Pi^{k,m}} \max_{I \in \mathcal{I}, \rho \in \Pi_{\text{ng}}^{k,m}} \frac{f(x^\pi(I))}{f(x^\rho(I))} &\leq 2 + \min(m/k, n-1), \quad (11) \\ \max_{\pi \in \Pi^{k,m}} \min_{I \in \mathcal{I}, \rho \in \Pi_{\text{ng}}^{k,m}} \frac{f(x^\pi(I))}{f(x^\rho(I))} &\leq \frac{1}{2 - \frac{\min((n-1)m/k, 1)}{n-1 + \min((n-1)m/k, 1)}}, \quad (12) \end{aligned}$$


where \mathcal{I} is the set of all problem instances. When restricting attention to augmented greedy policies, the best-case gain in performance and worst-case loss in performance associated with any $\pi \in \Pi_{\text{ag}}^{k,m}$ satisfies

$$\begin{aligned} \max_{I \in \mathcal{I}, \rho \in \Pi_{\text{ng}}^{k,m}} \frac{f(x^\pi(I))}{f(x^\rho(I))} &\geq 2 + \min(m/k, n-1-1/k), \quad (13) \\ \min_{I \in \mathcal{I}, \rho \in \Pi_{\text{ng}}^{k,m}} \frac{f(x^\pi(I))}{f(x^\rho(I))} &\geq \frac{1}{2 - \frac{(\min(m/k, 1))^{n-1}}{\sum_{i=0}^{n-1} (\min(m/k, 1))^i}} \quad (14) \end{aligned}$$


where the bound in (13) becomes an equality of the form (11) when $m \leq nk - k - 1$ and the bound in (14) becomes an equality of the form (12) when $m \geq k$.

The theorem proof is given in the next subsection. The bounds given in Theorem 1 represent a range of possible values for $f(x^\pi(I))/f(x^\rho(I))$ when $\rho \in \Pi_{\text{ng}}^{k,m}$, $\pi \in \Pi_{\text{ag}}^{k,m}$ and for any problem instance I . If $m = 0$, i.e., π is equivalent to a nominal greedy policy, then $2 \geq f(x^\pi(I))/f(x^\rho(I)) \geq 1/2$, since there exist problem instances where there are (at least) two possible outcomes for the greedy algorithm: the solution to (1), and the other a worst-case outcome, which has 1/2 the value of the first outcome. Therefore, one would hope that measurement passing can increase this upper bound above 2 and lower bound above 1/2. Theorem 1 shows that this is the case.

Equation (11) gives the upper bound on $f(x^\pi(I))/f(x^\rho(I))$ for any π . As one might expect, this upper bound increases with m : the more measurement passing is permitted, the higher the possible performance increase. However, when $m > k(n-1)$, the upper bound remains constant: increasing m above this value no longer increases potential improvement. Equation (13) shows that any augmented greedy policy is optimal in this sense when $m \leq nk - k - 1$. Furthermore, the expressions in (13) and (11) are always within an additive factor of $1/k$, so any augmented greedy policy is also at least near-optimal in this sense.



(a) The bounds from Theorem 1 for different values of n . For any augmented greedy policy π , the shaded blue region represents all possible values of $f(x^\pi(I))/f(x^\rho(I))$ for the corresponding value of m/k when $n = 3$. For instance, the middle of the 3 solid lines indicates that when $m = k$, $3 \geq f(x^\pi(I))/f(x^\rho(I)) \geq 0.6$.



(b) A zoomed-in view of the plot in (a) that shows the lower bounds in (12) and (14).

Fig. 3: The results of Theorem 1, illustrated for the case where $k = 8$ and for $n = 3, 4, 5, 6$. The solid lines indicate the bounds for an optimal message passing policy as given in (11) and (12) respectively, and the dashed lines indicate the proven lower bounds for any augmented greedy algorithm as given in (13) and (14). The plot in (b) is a close-up of the various values for in part (a).

Equation (12) shows that for any policy $\pi \in \Pi_{\text{ag}}^{k,m}$, one can always carefully construct problem instances where a nominal greedy policy (complete with tie-breaking rule) will perform better. In fact, no policy can guarantee that this worst-case performance loss is higher than the expression in (12). Again, we see that increasing m increases this lower bound, although here one sees no additional increase when $m > k$. Different from the upper bound is that the expression in (12) decreases to 1/2 as $n \rightarrow \infty$, regardless of the value of m . Theorem 1 states that any augmented greedy policy is optimal in this sense when $m \geq k$, it is also optimal when $n = 2$ or as $n \rightarrow \infty$.

Figure 3 gives an illustration of Theorem 1 for the case where $k = 8$ and for $n = 3, 4, 5, 6$. The solid colored lines

indicate the ‘‘optimal’’ bounds given in (11)–(12), and the dashed lines indicate the bounds for any augmented greedy algorithm as given in (13)–(14). For instance, the shaded blue region represents, for $n = 3$, the possible values of $f(x^\pi(I))/f(x^\rho(I))$ for any $\pi \in \Pi_{\text{ag}}^{k,m}$, $\pi_{\text{ng}}^{k,m}$, and $I \in \mathcal{I}$. The lowest solid black line indicates what is described above: that when $m = 0$, the values of the ratio range between 2 and 1/2. The middle solid black line represents the range of values when $m = k$; note that in this region any augmented greedy policy is optimal in that no other policy can provide a higher upper bound or a higher lower bound. Finally, the highest black line represents the range of values when $m = 4k$, where the lower bound is still optimal, but the upper bound is only guaranteed to be near-optimal.

While the performance increase that results from measurement passing is notable, there is some tradeoff with runtime. Note that the nominal greedy policy rule in (7) does not prescribe how to solve the local optimization problem, which is intractable in general. A full implementation of the nominal greedy algorithm will require $O\left(n \cdot \binom{\max_i |S_i|}{k}\right)$ number of calls to f . An augmented greedy policy, by comparison, will require $O\left(n \cdot \binom{mn + \max_i |S_i|}{\max(k, m-k)}\right)$. We address this intractability in Section IV.

Finally, we note that in the examples which serve to prove the bounds in (12) and (13), there is some reliance on overlap among the local measurement sets S_1, \dots, S_n . Clearly, when $S_i = S_j$ for all i, j , then all sensors have access to the same information, thus measurement passing is futile. However, the extent to which this is the case is a topic of future work.

B. Proof for Theorem 1

We begin with two lemmas that, that, given two policies $\pi, \rho \in \Pi^{k,m}$, show marginal contributions for x_i^π, z_i^π , and x_i^ρ affects $f(x^\pi)/f(x^\rho)$. We define marginal contribution of $A \subseteq S$ with respect to $B \subseteq S$ according to f as

$$\Delta(A|B) := f(A \cup B) - f(B). \quad (15)$$

We abuse notation so that $\Delta(A, B|C, D) = \Delta(A \cup B|C \cup D)$. Also, denote $x_{a:b} = \cup_{a \leq i \leq b} x_i$, and likewise for $z_{a:b}$.

Lemma 1. Assume that policy $\rho \in \Pi^{k,m}$ is applied to instance $I \in \mathcal{I}$ and that there exists $\alpha \geq 1$ such that

$$\alpha \Delta(x_i^\rho | x_{1:i-1}^\rho) \geq \max_{\tilde{x} \in (S_i)^k} \Delta(\tilde{x} | x_{1:i-1}^\rho), \quad \forall i \quad (16)$$

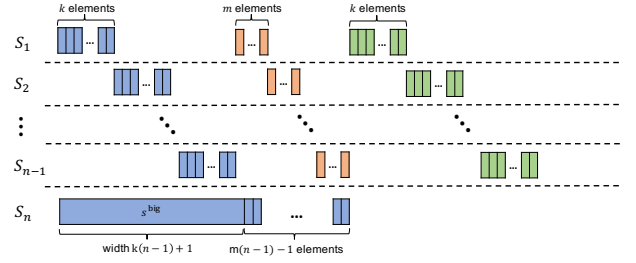
Then for any $\pi \in \Pi^{k,m}$,

$$\frac{f(x^\pi(I))}{f(x^\rho(I))} \leq \begin{cases} 2\alpha + \frac{m}{k}, & \text{if } m \leq k \\ 1 + \alpha \left(1 + \min\left(\frac{m}{k}, n-1\right)\right), & \text{if } m > k. \end{cases} \quad (17)$$

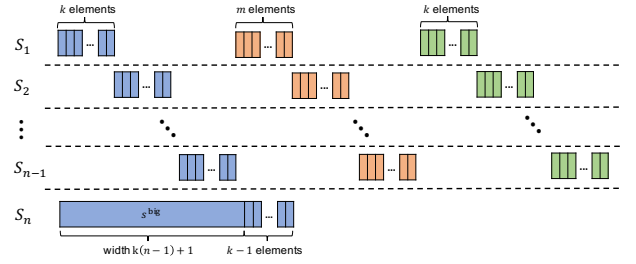
Lemma 2. Assume that policy $\pi \in \Pi^{k,m}$ is applied to instance $I \in \mathcal{I}$ and that there exist $\alpha_1, \alpha_2 \geq 1$ such that

$$\alpha_1 \Delta(x_i^\pi | x_{1:i-1}^\pi) \geq \max_{\tilde{x} \in (S_i)^k} \Delta(\tilde{x} | x_{1:i-1}^\pi) \quad (18a)$$


$$\alpha_2 \cdot \max_{\tilde{z} \in (z_i^\pi)^k} \Delta(\tilde{z} | x_{1:i}^\pi) \geq \max_{\tilde{x} \in (S_i)^k} \Delta(\tilde{x} | x_{1:i-1}^\pi). \quad (18b)$$



(a) An example for proving (12) when $m(n-1) \leq k$. The key is that $|S_n| \leq k$, and the orange measurements offer no value beyond S_n . Each small box has width 1 and s^{big} has width $k(n-1)+1$. Thus $f(x^\pi(I)) = k(n-1) + m(n-1)$ and $f(x^\rho(I)) = 2k(n-1) + m(n-1)$.



(b) An example for proving (12) when $m(n-1) \geq k$. Unlike the example above, here S_n consists of exactly k measurements, thus $f(x^\pi(I)) = k(n-1) + k$ and $f(x^\rho(I)) = 2k(n-1) + k$.



(c) An example for proving (13). Here the greedy algorithm chooses the green measurements which are ‘‘covered’’ by s^{big} , thus $f(x^\pi(I)) = k$ and $f(x^\rho(I)) = 2k + \min(m, k(n-1) - 1)$.

Fig. 4: Problem instances used in the proof for Theorem 1. This representation is similar to that in Figure 2, in that row i represents S_i and $f(A)$ is the amount of horizontal space covered by the boxes in $A \subseteq S$. Blue squares represent measurements chosen by π , orange are those measurements which are passed using π , and green are the nominal greedy choices $x^\rho(I)$, when those differ from $x^\pi(I)$. We have omitted the dashed boxes, with the understanding that an orange box in S_i is available to sensors $j > i$.

Then for any $\rho \in \Pi^{k,m}$ such that $x_i^\rho \subseteq S_i$ for all i ,

$$\frac{f(x^\pi(I))}{f(x^\rho(I))} \geq \frac{1}{1 + \alpha_1 - \frac{1}{\sum_{i=0}^{n-1} \alpha_2^i}} \quad (19)$$

The proofs for Lemma 1 and Lemma 2 are given in Appendix-A and Appendix-B, respectively. We now prove each statement of the Theorem separately.

1) **Equation (11):** Here we invoke Lemma 1. Let $\pi \in \Pi^{k,m}$ and $\rho \in \Pi_{\text{ng}}^{k,m}$. Using ρ , sensors make choices according to (7), therefore for any I , let $\alpha = 1$. Then both expressions in (17) are equivalent: $f(x^\pi(I))/f(x^\rho(I)) \leq 2 + \min(m/k, n-1)$.

2) **Equation (12):** Fix $\pi \in \Pi^{k,m}$ and assume first that $(n-1)m \leq k$. Suppose that f and S_1, \dots, S_n are as represented in Figure 4a. Here the format of example is the same as in Figure 2. 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 sensor $i \in \{1, \dots, n-1\}$, all measurements in S_i are identical according to f , since there is no horizontal overlap among them, and none of these sensors is aware of the measurements in S_n .

Assume that π_i selects the blue measurements for x_i^π and the orange measurements for z_i^π , $i = 1, \dots, n-1$. Note that $|S_n| = m(n-1) \leq k$, so $x_n^\pi = S_n$ is feasible, and an optimal choice regardless of the previous sensors' decisions. This implies that $f(x^\pi(I)) = k(n-1) + m(n-1)$.

On the other hand, consider the decision set of some $\rho \in \Pi_{\text{ng}}^{k,m}$: the rectangles shaded in green for $i < n$, and the blue rectangles for $i = n$. Thus $f(x^\rho(I)) = 2k(n-1) + m(n-1)$, and for this problem instance I ,

$$\begin{aligned} \frac{f(x^\pi(I))}{f(x^\rho(I))} &= \frac{k(n-1) + m(n-1)}{2k(n-1) + m(n-1)} \\ &= \frac{1}{2 - \frac{(n-1)m/k}{n-1+(n-1)m/k}}. \end{aligned} \quad (20)$$

In the case where $m(n-1) \geq k$, consider the example in Figure 4b. Here S_i are the same as in Figure 4a, for $i = 1, \dots, n-1$, implying again without that a possible choice for $x_1^\pi, \dots, x_{n-1}^\pi$ are the respective blue measurements. In this case, however, $|S_n| = k$, but note that $f(x_1^\pi \cup \dots \cup x_{n-1}^\pi \cup x_n) = kn$ for any $x_n \in (S_n)^k$, thus $f(x^\pi(I)) = kn$. The green measurements are again a possible greedy policy selection, where different from the blue, showing that $f(x^\rho(I)) = 2k(n-1) + k$. In this case we see that

$$\frac{f(x^\pi(I))}{f(x^\rho(I))} = \frac{kn}{2k(n-1) + k} = \frac{1}{2 - 1/n}. \quad (21)$$

Equations (20) and (21) establish (12) for all cases.

3) Equation (13): We appeal to an example of the same style as in Figure 2, which is illustrated in Figure 4c. Here S_1 is the union of a set of k green measurements, a set of k blue measurements, and a set of m orange measurements. The sets S_2, \dots, S_{n-1} are all empty, and $S_n = \{s^{\text{big}}\}$. All measurements have value 1, except s^{big} , which has value k . The measurement s^{big} ‘‘covers’’ the set of green measurements, i.e., if A is the set of green measurements, then $f(B \cup \{s^{\text{big}}\}) = k$ for any $B \subseteq A$.

Assume that using a nominal greedy policy ρ , sensor 1 selects the k green measurements. Then $f(x^\rho(I)) = k$. However, there exists an augmented greedy policy π such that sensor 1 selects the k blue measurements as x_1^π and the m orange measurements as z_1^π . Since the remaining sensors have no other alternatives, $\min\{m, k(n-1) - 1\}$ orange measurements are chosen for x_2^π, \dots, x_n^π . This implies that $f(x^\pi(I)) = 2k + \min\{m, k(n-1) - 1\}$, and that for this problem instance I

$$\frac{f(x^\pi(I))}{f(x^\rho(I))} = 2 + \min\{m/k, n-1 - 1/k\}. \quad (22)$$

4) Equation (14): We invoke Lemma 2 by finding acceptable values of α_1, α_2 which hold for any I . For any $\pi \in \Pi_{\text{ag}}^{k,m}$, (10a) implies that $\alpha_1 = 1$ is a valid parameter choice. When

$m \geq k$, $\alpha_2 = 1$, since $\max_{z \in (z_i^\pi)^k} \Delta(\tilde{z}|x_{1:i}^\pi) = z_i^k$ from (10c). When $m < k$, the following holds:

$$\max_{z \in (z_i^\pi)^k} \Delta(\tilde{z}|x_{1:i}^\pi) = \max_{z \in (S_i)^m} \Delta(\tilde{z}|x_{1:i}^\pi) \quad (23a)$$

$$\geq (m/k) \cdot \max_{\tilde{x} \in (S_i)^k} \Delta(\tilde{x}|x_{1:i}^\pi). \quad (23b)$$

Therefore, $\alpha_2 = 1/\min(m/k, 1)$ is an acceptable value. Since $\rho \in \Pi_{\text{ng}}^{k,m}$ satisfies that $x_i^\rho \subseteq S_i$, one can use these values for α_1, α_2 (combined with some algebraic manipulation), so that Lemma 2 implies (14). \square

C. Benchmark Comparison to Greedy

In this section we compare augmented greedy policies to nominal greedy policies by comparing them each to a the optimal solution to (1), to see how measurement sharing increases performance guarantees. When π is a nominal greedy policy, as has been stated, the resulting solution is within a factor of 1/2 of this benchmark. Therefore, we are interested in how the ability to share measurements can increase the 1/2 bound:

Corollary 1. *Consider the measurement selection problem with n sensors. Then for any $k \geq 1$, $m \geq 0$, the worst-case performance guarantee associated with the optimal measurement passing policy within $\Pi^{k,m}$ satisfies*

$$\max_{\pi \in \Pi^{k,m}} \min_{I \in \mathcal{I}} \frac{f(x^\pi(I))}{\text{OPT}(I, k)} \leq \frac{1}{2 - \frac{\min((n-1)m/k, 1)}{n-1 + \min((n-1)m/k, 1)}}, \quad (24)$$

where $\text{OPT}(I, k)$ is the value of the solution to (1). When restricting attention to augmented greedy policies, the worst-case performance guarantee associated with any π satisfies

$$\min_{I \in \mathcal{I}} \frac{f(x^\pi(I))}{\text{OPT}(I, k)} \geq \frac{1}{2 - \frac{(\min(m/k, 1))^{n-1}}{\sum_{i=0}^{n-1} (\min(m/k, 1))^i}}. \quad (25)$$

where the bound in (25) becomes an equality of the form (24) when $m \geq k$.

This result follows from Theorem 1: (24) since the defining example in Figures 4a and 4b can be altered so that the green measurements are the solution to (1), and (25) since a policy $\rho \in \Pi^{k,m}$ which finds $\text{OPT}(I, k)$ satisfies the requirement to apply Lemma 2. Corollary 1 builds on Theorem 1 to give insight into another way to understand the benefits of measurement sharing. Assuming $m > 0$, (25) shows that any augmented greedy policy will yield stricter guarantees than a nominal greedy policy. For instance, when $n = 2$ and $m \geq k$, $f(x^\pi(I)) \geq 2/3\text{OPT}(I, k)$ for any I and $\pi \in \Pi_{\text{ag}}^{k,m}$, as compared to 1/2 for any nominal greedy policy. As n gets large this boost in performance guarantees is less pronounced: as an example, when $n = 100$ and $m = k$, the augmented greedy policy guarantee is $f(x^\pi(I)) \geq (1/1.99)\text{OPT}(I, k)$.

Equation (24) also illustrates, similar to Theorem 1, that any augmented greedy policy is near-optimal in the sense that it provides close to the strongest possible guarantees against OPT for any policy in $\Pi^{k,m}$. In fact, such a policy is optimal in this sense when the two bounds are the same, i.e., when

$m \geq k$ or when $n = 2$. Since these bounds are the same as those in Theorem 1, the remaining discussion on specific values of m, n, k follows directly.

IV. SUBOPTIMAL SELECTIONS

To implement an augmented greedy policy, each agent is required to solve the optimizations in (10a) and (10d), both of which are NP-Hard, in terms of k , assuming large $|S_i|$. This means that, when k and $|S_i|$ are large, executing an augmented greedy policy may become computationally infeasible—an observation which is well-known for the nominal greedy algorithm [41]. Such scenarios are typical for the motivating example in Section II, in which the flying vehicles may need to select a large number of images from a much larger set of total images taken.

Real-world implementations of the augmented greedy algorithm must thus approximate (10a)–(10d). We devote this section to understanding how approximating the solution to such optimizations affects the measurement sharing. The key observation from the results in that follow is that while this approximation increases the range of possible values for $f(x^\pi(I))/f(x^\rho(I))$ (as one might expect), the lower bound decreases (roughly) linearly as a factor of the approximation error. The idea of using approximate maximization for the nominal greedy algorithm has been used previously in the literature (see, for instance [41], [42]) for analyzing approximations to the nominal greedy algorithm, which model and results we extend here for augmented greedy policies.

Definition 2 (β -Greedy Policy). *A policy $\pi \in \Pi^{k,m}$ is a β -greedy policy for some $\beta \geq 1$ if each sensor $i \in \{1, \dots, n\}$ is associated with a selection rule π_i of the form*

$$\Delta(x_i^\pi | x_{1:i-1}^\pi) \geq \beta \cdot \max_{\tilde{x} \in (S_i)^k} \Delta(\tilde{x} | x_{1:i-1}^\pi). \quad (26)$$

Note that when $\beta = 1$, the nominal greedy policy defined by (7) is recovered.

An analogous approximation can be defined for the augmented greedy algorithm:

Definition 3 ((β_k, β_m) -Augmented Greedy Policy). *A policy $\pi \in \Pi^{k,m}$ is a (β_k, β_m) -greedy policy for some $\beta_k, \beta_m \geq 1$ if each sensor $i \in \{1, \dots, n\}$ is associated with a selection rule π_i of the form*

$$\Delta(x_i^\pi | x_{1:i-1}^\pi) \geq \beta_k \cdot \max_{\tilde{x} \in (S_i \cup z_1^\pi \cup \dots \cup z_{i-1}^\pi)^k} \Delta(\tilde{x} | x_{1:i-1}^\pi), \quad (27a)$$

$$z_i^\pi = z_i^k \cup z_i^{m-k}, \text{ where} \quad (27b)$$

$$\Delta(z_i^k | x_{1:i}^\pi) \geq \beta_m \cdot \max_{\tilde{z} \in (S_i)^{\min(m,k)}} \Delta(\tilde{z} | x_{1:i}^\pi), \quad (27c)$$

$$\Delta(z_i^{m-k} | x_{1:i}^\pi, z_i^k) \geq \beta_m \cdot \max_{\tilde{z} \in (S_i)^{\max(m-k,0)}} \Delta(\tilde{z} | x_{1:i}^\pi, z_i^k) \quad (27d)$$

In essence, a (β_k, β_m) -augmented greedy policy is a policy which approximates an augmented greedy policy by finding a solution to (10a) and (10c) within a factor of β_k and β_m , respectively, of the optimal. When $\beta_k = \beta_m = 1$, the original augmented greedy policy is recovered.

Alternatively stated, β_m is related to how well one can approximate the maximum on S_i and β_k is related to how well one can approximate the maximum on $S_i \cup z_1 \cup \dots \cup z_{i-1}$.

Theorem 2. *Consider the measurement selection problem with n sensors. Then for any (β_k, β_m) -augmented greedy policy π , any β_k -greedy policy ρ , and any problem instance I ,*

$$\frac{f(x^\pi(I))}{f(x^\rho(I))} \leq \begin{cases} \frac{2}{\beta_k} + \frac{m}{k}, & \text{if } m \leq k \\ 1 + \frac{1}{\beta_k} (1 + \min(\frac{m}{k}, n-1)), & \text{if } m > k, \end{cases} \quad (28)$$

$$\frac{f(x^\pi(I))}{f(x^\rho(I))} \geq \frac{1}{1 + \frac{1}{\beta_k} - \frac{1}{\sum_{i=0}^{n-1} (\beta_m \min(m/k, 1))^{n-1}}}. \quad (29)$$

Observe that when $\beta_k = \beta_m = 1$, the results are equivalent to (11) and (14) in Theorem 1. Here we forgo analogous results to (12) and (13), since the emphasis of this theorem is that while the approximations to the local optimization problems increase the range of possible values for $f(x^\pi(I))/f(x^\rho(I))$, this range is still desirable. For instance, if $\beta_k = \beta_m = 1/2$, regardless of the number of sensors, Theorem 2 shows that $f(x^\pi(I))/f(x^\rho(I)) \geq 1/3$, as compared to the $1/2$ bound from Theorem 1, i.e., the potential loss in performance decreases only a moderate amount. On the other hand, $f(x^\pi(I))/f(x^\rho(I)) \leq 2n+1$, an increase over the $n+1$ bound from Theorem 1. In words, measurement passing can offer a potentially larger benefit without a much higher risk.

An example of a (β_k, β_m) -augmented greedy policy π is where agent i chooses $x_i^\pi = \{s_1, \dots, s_k\}$ by sequentially selecting element s_l with the following method:

$$s_l \in \arg \max_{\tilde{s} \in S_i \cup z_1^\pi \cup \dots \cup z_{l-1}^\pi} \Delta(\{\tilde{s}\}, s_{1:l-1} | x_{1:i-1}^\pi). \quad (30)$$

This is yet another variation of greedy algorithm, this time for choosing the k elements of x_i^π . The guarantees for this algorithm are such that $\beta_k = 1 - (1 - 1/k)^k$ [26]. Using a similar method to choose z_i^π yields $\beta_m = 1 - (1 - 1/m)^m$, so that both β_k and β_m are greater than $1 - 1/e \approx 0.63$. The $(1 - 1/e, 1 - 1/e)$ -augmented greedy algorithm can now be implemented using $O(mn^2 + \sum_i |S_i|)$ calls to f , and the $(1 - 1/e)$ -greedy algorithm can be implemented using $O(\sum_i |S_i|)$ calls to f .

Figure 5 illustrates the upper bound shown in (28) and the lower bound shown in (29) for various values of n and m when $k = 2$. Here we assume that (30) is used to implement both algorithms, i.e., $\beta_k = 1 - (1 - 1/k)^k$ and $\beta_m = 1 - (1 - 1/m)^m$. Since β_m decreases as m increases, the lower bound decreases when $m > k$. The upper bound, however, continues to increase in a similar manner to that of Theorem 1.

We now give the proof for Theorem 2:

Proof. We first show (29) by invoking Lemma 2: it suffices to show valid values for α_1, α_2 in order to show the lower bound in (29). An immediate consequence of (27a) is that $\alpha_1 = 1/\beta_k$ holds for all I . Likewise, one can use a similar argument to (23a)–(23b) to show that $\alpha_3 = 1/(\beta_m \min(m/k, 1))$ holds for all I . Then by Lemma 2 (since again ρ is such that $x_i^\rho \subseteq S_i$),

$$\frac{f(x^\pi(I))}{f(x^\rho(I))} \geq \frac{1}{1 + \frac{1}{\beta_k} - \frac{1}{\sum_{i=0}^{n-1} (\beta_m \min(m/k, 1))^{n-1}}}$$

m	$n = 2$		$n = 3$		$n = 4$		$n = 5$		$n = 6$	
	upper	lower	upper	lower	upper	lower	upper	lower	upper	lower
0	2.667	0.4286	2.667	0.4286	2.667	0.4286	2.667	0.4286	2.667	0.4286
2	3.667	0.5676	3.667	0.4978	3.667	0.4700	3.667	0.4556	3.667	0.4470
4	3.667	0.5580	5	0.4893	5	0.4628	5	0.4494	5	0.4419
6	3.667	0.5553	5	0.4870	6.333	0.4608	6.333	0.4479	6.333	0.4406
8	3.667	0.5540	5	0.4859	6.333	0.4600	7.667	0.4471	7.667	0.4400

Fig. 5: Some examples that showcase the results of Theorem 2 for $k = 2$. It is assumed that each algorithm is implemented with the sequential greedy rule in (30), i.e., $\beta_k = 1 - (1 - 1/k)^k$ and $\beta_m = 1 - (1 - 1/m)^m$.

$$= \frac{1}{1 + \frac{1}{\beta_k} - \frac{(\beta_m \min(m/k, 1))^{n-1}}{\sum_{i=0}^{n-1} (\beta_m \min(m/k, 1))^i}}.$$

Equation (28) can be shown similarly using Lemma 1, where π is the (β_k, β_m) -augmented greedy algorithm, ρ is the β_k -greedy algorithm, and, by (26), $\alpha = 1/\beta_k$ for all I . \square

V. NUMERICAL EXAMPLE

In this section, we present results for instances of the flying vehicles problem in Section II, where $n = 2$ flying vehicles move on a curved path, each carrying a side-looking camera with a 90° field of view, a 50 pixel focal length, and measurement noise in the image plane with standard deviation $\sigma = 1$ pixel. There are two stationary ground targets whose 2-D location we would like to estimate using the images collected by the flying vehicles. A large number of instances were created with the two targets uniformly randomly placed in the square $[-100, 100] \times [-100, 100]$. The start position, direction, and turn rate of the each flying vehicle's path were also chosen uniformly randomly. Each flying vehicle moves at a constant forward speed and collects 100 independent measurements uniformly along its path. Details of how to construct the corresponding matrices Q_0 and Q_s in (3) and (4) that quantify the information gain of camera measurements are found in [12]. See Figure 6 for an example.

Figure 7 summarizes the results in terms of the ratio between the performance of a $(1 - (1 - 1/k)^k, 1 - (1 - 1/m)^m)$ -augmented greedy policy and a $1 - (1 - 1/k)^k$ -greedy policy, where $k = m = n = 2$ and $|S_1| = |S_2| = 100$. Because the number of measurements is very large, the optimizations in (10a), (10c), and (7) are all approximated by the sequential algorithm in (30), where all ties are broken by the index of the measurement.

Additional trials were run for different combinations of the number of flying vehicles and targets. Each combination was repeated 10^5 times with random position and paths. The results are summarized in Figure 8. The setup is the same as above, with the exception that, for the purposes of computation the flying vehicles gather 20 measurements over the course of the simulation rather than 100. One can see that when there are more sensors, the mean benefit of measurement passing decreases, since the measurement passing is most useful when a large percentage of flying vehicles cannot observe any targets. This is more likely to happen with fewer flying vehicles. However, one can also observe from the column labeled “% > 1” that measurement passing is more likely to give some benefit (although smaller) when there are

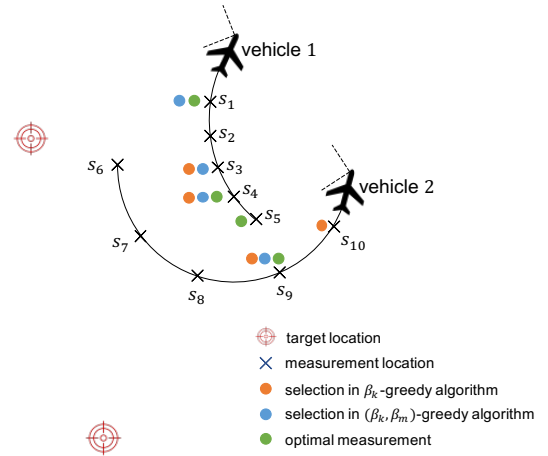


Fig. 6: An instance of the flying vehicles problem where $n = 2$, $k = 2$, $m = 1$, and $|S_1| = |S_2| = 5$. The black solid lines indicate each vehicle's path, and the black dashed lines indicate each vehicles field of view. The black 'x's are where measurements of the red targets are taken; of course, both targets are not always in view of both vehicles. The β_k -greedy (orange dots) and (β_k, β_m) -augmented (blue dots) greedy policies are used in conjunction with the sequential selection rule in (30). The green dots indicate measurements that are optimal in the sense that they are a solution to (9).

more vehicles, simply because more measurements are being passed, increasing the likelihood that some vehicle will use another's measurement. Likewise, it can be observed from the last column that systems with many sensors are less likely to see a decrease in performance due to measurement passing.

While measurement selection is based on the submodular function (6), it is worth noting that the actual improvement in target estimation, such as that described by the commonly used D-optimality estimation criterion of $\det(Q_0 + \sum_{s \in S} Q_s)$ (without the log function), can be much greater than that shown in Figure 7.

VI. CONCLUSION

In this paper we have shown how measurement passing affects the performance guarantees of a group of sensors sequentially selecting measurements. We also showed that the augmented greedy policy gives optimal performance guarantees for some problem instances and near-optimal guarantees for others. Using such a policy, this paper explored how much performance could increase or decrease for any problem instance, and showed how both of these results are affected by inability of each sensor to solve its own local optimization problem.

Future work will continue to explore measurement passing, first by asking which is more important: relaying a sensor's decision or measurement passing? Some preliminary exploration has shown that the sensor's decision is more impactful, but there is not a formal proof of this. Another direction is to apply these results to situations where sensors can only observe the measurements of a subset of previous sensors, and again ask questions about what measurements should be shared and selected.




Fig. 7: Histogram of the relative performance for 10^6 random simulations of the 2 vehicle measurement selection problem. A large number of samples fall at 1, which is interpreted as the two policies performing the same. Instances where the value is greater than one indicate the (β^k, β^m) -augmented greedy performed better than the β^k -greedy in the simulated scenario. The height of the bin corresponding to the ratio of 1 has been cropped. The solid blue bars represent the range of values for $f(x^\pi(I))/f(x^\rho(I))$ as determined by Theorem 2 and in Figure 5.

Trials Measuring $f(x^\pi(I))/f(x^\rho(I))$

# Vehicles	# Targets	Mean	% > 1	% ≥ 1
3	10	1.05	0.625	0.97
3	3	1.045	0.667	0.985
4	4	1.044	0.835	0.972
10	4	1.024	0.996	0.997
2	10	1.041	0.33	1
5	10	1.042	0.895	0.945

Fig. 8: Results from many different types of trials comparing nominal greedy to augmented greedy. Each row represents a different set of numbers of agents and targets, and there were 10^5 trials run for each row. One can observe that measurement passing helps on average more when there are fewer agents, however, many agents ensure that measurement passing helps more frequently.

REFERENCES

- [1] D. Grimsman, M. R. Kirchner, J. P. Hespanha, and J. R. Marden, "The impact of message passing in agent-based submodular maximization," in *IEEE Conference on Decision and Control*, 2020.
- [2] D. M. Stipanović, G. Inalhan, R. Teo, and C. J. Tomlin, "Decentralized overlapping control of a formation of unmanned aerial vehicles," *Automatica*, vol. 40, no. 8, pp. 1285–1296, 2004.
- [3] M. DeBord, W. Hönig, and N. Ayanian, "Trajectory planning for heterogeneous robot teams," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 7924–7931.
- [4] M. R. Kirchner, R. Mar, G. Hewer, J. Darbon, S. Osher, and Y. Chow, "Time-optimal collaborative guidance using the generalized Hopf formula," *IEEE Control Systems Letters*, vol. 2, no. 2, pp. 201–206, 2018.
- [5] G. Arslan, J. R. Marden, and J. S. Shamma, "Autonomous vehicle-target assignment: A game-theoretical formulation," *Journal of Dynamic Systems, Measurement, and Control*, vol. 129, no. 5, pp. 584–596, 2007.
- [6] H.-L. Choi, L. Brunet, and J. P. How, "Consensus-based decentralized auctions for robust task allocation," *IEEE Transactions on Robotics*, vol. 25, no. 4, pp. 912–926, 2009.
- [7] M. R. Kirchner, M. J. DeBord, and J. P. Hespanha, "A Hamilton–Jacobi formulation for optimal coordination of heterogeneous multiple vehicle systems," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 11 623–11 630.
- [8] J. R. Marden and A. Wierman, "Distributed welfare games with applications to sensor coverage," in *Decision and Control, 2008. CDC 2008. 47th IEEE Conference on*. IEEE, 2008, pp. 1708–1713.
- [9] M. R. Kirchner, "Towards general performance bounds of the distributed sensor coverage problem," *Electrical, Computer & Energy Engineering Graduate Theses & Dissertations*, no. 71, 2013.
- [10] H. Ferraz and J. P. Hespanha, "Iterative algorithms for distributed leader-follower model predictive control," in *2019 IEEE 58th Conference on Decision and Control (CDC)*. IEEE, 2019, pp. 3533–3539.
- [11] J. Xie, L. R. G. Carrillo, L. Jin, and J. P. Hespanha, "A robust state estimator for multi-agent systems under impulsive noise and missing measurements," in *2019 IEEE 58th Conference on Decision and Control (CDC)*. IEEE, 2019, pp. 6019–6024.
- [12] M. R. Kirchner, J. P. Hespanha, and D. Garagić, "Heterogeneous measurement selection for vehicle tracking using submodular optimization," in *IEEE Aerospace Conference*, 2020, pp. 1–10.
- [13] R. Iyer, N. Khargonkar, J. Bilmes, and H. Asnani, "Submodular Combinatorial Information Measures with Applications in Machine Learning," in *Proceedings of Machine Learning Research*, vol. 132, 2021, pp. 1–33.
- [14] 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.
- [15] A. Badanidiyuru, B. Mirzasoleiman, A. Karbasi, and A. Krause, "Streaming submodular maximization: Massive data summarization on the fly," in *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2014, pp. 671–680.
- [16] 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*, 2007, pp. 420–429.
- [17] D. Kempe, J. Kleinberg, and É. Tardos, "Maximizing the spread of influence through a social network," in *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2003, pp. 137–146.
- [18] M. Gomez-Rodriguez, J. Leskovec, and A. Krause, "Inferring networks of diffusion and influence," *ACM Transactions on Knowledge Discovery from Data*, vol. 5, no. 4, 2012.
- [19] H. Lin and J. Bilmes, "A class of submodular functions for document summarization," in *Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, vol. 1, 2011, pp. 510–520.
- [20] 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.
- [21] G. Qu, D. Brown, and N. Li, "Distributed greedy algorithm for multi-agent task assignment problem with submodular utility functions," *Automatica*, vol. 105, no. February, pp. 206–215, 2019.
- [22] 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*, 2007, pp. 2204–2211.
- [23] A. Clark and R. Poovendran, "A submodular optimization framework for leader selection in linear multi-agent systems," in *IEEE Conference on Decision and Control*. IEEE, 2011, pp. 3614–3621.
- [24] J. R. Marden, "The role of information in distributed resource allocation," *IEEE Transactions on Control of Network Systems*, vol. 4, no. 3, pp. 654–664, 2017.
- [25] L. Lovász, "Submodular functions and convexity," in *Mathematical Programming The State of the Art*, 1983, pp. 235–257.
- [26] 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.
- [27] M. L. Fisher, G. L. Nemhauser, and L. A. Wolsey, "An analysis of approximations for maximizing submodular set functions-II," *Polyhedral Combinatorics*, pp. 73–87, 1978.
- [28] 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.
- [29] M. Gairing, "Covering games: Approximation through non-cooperation," in *Lecture Notes in Computer Science*, vol. 5929, 2009, pp. 184–195.
- [30] U. Feige, "A threshold of $\ln n$ for approximating set cover," *Journal of the ACM*, vol. 45, no. 4, pp. 634–652, 1998.
- [31] A. Clark, B. Alomair, L. Bushnell, and R. Poovendran, *Submodularity in dynamics and control of networked systems*, A. Isidori, J. H. van Schuppen, E. D. Sontag, and M. Krstic, Eds. Springer, 2016.
- [32] A. Robey, A. Adibi, B. Schlotfeldt, G. J. Pappas, and H. Hassani, "Optimal Algorithms for Submodular Maximization with Distributed Constraints," no. Cdc, 2019.
- [33] W. Luo, S. S. Khatib, S. Nagavalli, N. Chakraborty, and K. Sycara, "Distributed knowledge leader selection for multi-robot environmental sampling under bandwidth constraints," *IEEE International Conference*

- on *Intelligent Robots and Systems*, vol. 2016–November, pp. 5751–5757, 11 2016.
- [34] B. Ghahsifard and S. L. Smith, “Distributed submodular maximization with limited information,” *IEEE Transactions on Control of Network Systems*, vol. 5, no. 4, pp. 1635–1645, 2017.
- [35] D. Grimsman, M. S. Ali, J. P. Hespanha, and J. R. Marden, “The impact of information in distributed submodular maximization,” *IEEE Transactions on Control of Network Systems*, vol. 6, no. 4, pp. 1334–1343, 2019.
- [36] D. Grimsman, J. P. Hespanha, and J. R. Marden, “Strategic Information Sharing in Greedy Submodular Maximization,” *IEEE Conference on Decision and Control*, pp. 2722–2727, 2018.
- [37] M. Corah and N. Michael, “Distributed matroid-constrained submodular maximization for multi-robot exploration: theory and practice,” *Autonomous Robots*, vol. 43, no. 2, pp. 485–501, 2019.
- [38] —, “Distributed submodular maximization on partition matroids for planning on large sensor networks,” in *IEEE Conference on Decision and Control*, 2018, pp. 6792–6799.
- [39] A. Pázman, *Foundations of optimum experimental design*. Springer, 1986, vol. 14.
- [40] T. H. Summers, F. L. Cortesi, and J. Lygeros, “On Submodularity and Controllability in Complex Dynamical Networks,” *IEEE Transactions on Control of Network Systems*, vol. 3, no. 1, pp. 91–101, 2016.
- [41] P. Goundan and A. Schulz, “Revisiting the greedy approach to submodular set function maximization,” *Optimization online*, no. 1984, pp. 1–25, 2007. [Online]. Available: http://www.optimization-online.org/DB_FILE/2007/08/1740.pdf
- [42] B. Lehmann, D. Lehmann, and N. Nisan, “Combinatorial auctions with decreasing marginal utilities,” *Games and Economic Behavior*, vol. 55, no. 2, pp. 270–296, 2006.

APPENDIX

A. Proof for Lemma 1

Begin with the following:

$$f(x^\pi) \leq f(x^\rho) + \Delta(x^\pi | x^\rho) \quad (31a)$$

$$\leq f(x^\rho) + \Delta(z_{1:n-1}^\pi \cap x^\pi | x^\rho) + \Delta(x^\pi \setminus z_{1:n-1}^\pi | x^\rho) \quad (31b)$$

$$= f(x^\rho) + \Delta(z_{1:n-1}^\pi \cap x^\pi | x^\rho) + \sum_i \Delta(x_i^\pi \setminus z_{1:n-1}^\pi | x^\rho, x_{1:i-1}^\pi \setminus z_{1:n-1}^\pi) \quad (31c)$$

$$\leq f(x^\rho) + \Delta(z_{1:n-1}^\pi \cap x^\pi | x^\rho) + \sum_i \Delta(x_i^\pi \setminus z_{1:n-1}^\pi | x_{1:i-1}^\rho) \quad (31d)$$

$$\leq f(x^\rho) + \Delta(z_{1:n-1}^\pi \cap x^\pi | x^\rho) + \sum_i \alpha \Delta(x_i^\rho | x_{1:i-1}^\rho) \quad (31e)$$

$$= (1 + \alpha) f(x^\rho) + \Delta(z_{1:n-1}^\pi \cap x^\pi | x^\rho) \quad (31f)$$

where (31a), (31b), (31d) are true by submodularity of f , and (31e) is true by (16).

We denote \tilde{z}_i^π to mean $z_i^\pi \cap x^\pi$, and suppose that there exists β such that $\Delta(\tilde{z}_i^\pi | x_{1:i}^\rho) \leq \beta \Delta(x_i^\rho | x_{1:i-1}^\rho)$. Then the second term in (31f) can be upper bounded by the following:

$$\Delta(\tilde{z}_{1:n-1}^\pi | x^\rho) \leq \sum_{i=1}^{n-1} \Delta(\tilde{z}_i^\pi | \tilde{z}_{1:i-1}^\pi, x^\rho) \quad (32a)$$

$$\leq \sum_{i=1}^{n-1} \Delta(\tilde{z}_i^\pi | x_{1:i}^\rho) \quad (32b)$$

$$\leq \sum_{i=1}^{n-1} \beta \Delta(x_i^\rho | x_{1:i}^\rho) \quad (32c)$$

$$= \beta f(x_{1:n-1}^\rho) \leq \beta f(x^\rho), \quad (32d)$$

where (32b) and (32d) are true by the submodularity of f . Substituting this upper bound back into (31f), we see that

$$f(x^\pi) \leq (1 + \alpha + \beta) f(x^\rho).$$

We now show that such a β exists and define it for two cases: when $m \leq k$ and when $m \geq k$. First suppose that $m \leq k$. Denote $x_i^{k-m} \in \arg \max_{\tilde{x} \in (x_i^\rho)^{k-m}} \Delta(\tilde{x} | x_{1:i-1}^\rho)$. Then

$$\Delta(x_i^\rho | x_{1:i-1}^\rho) \geq (1/\alpha) \Delta(\tilde{z}_i^\pi \cup x_i^{k-m} | x_{1:i-1}^\rho) \quad (33a)$$

$$= (1/\alpha) \Delta(\tilde{z}_i^\pi | x_{1:i-1}^\rho, x_i^{k-m}) + (1/\alpha) \Delta(x_i^{k-m} | x_{1:i-1}^\rho) \quad (33b)$$

$$\geq (1/\alpha) \Delta(\tilde{z}_i^\pi | x_{1:i}^\rho) + (1/\alpha) \Delta(x_i^{k-m} | x_{1:i-1}^\rho) \implies \quad (33c)$$

$$\Delta(\tilde{z}_i^\pi | x_{1:i}^\rho) \leq \alpha \Delta(x_i^\rho | x_{1:i-1}^\rho) - \Delta(x_i^{k-m} | x_{1:i-1}^\rho) \quad (33d)$$

$$\leq \alpha \Delta(x_i^\rho | x_{1:i-1}^\rho) - \frac{k-m}{k} \Delta(x_i^\rho | x_{1:i-1}^\rho) \quad (33e)$$

$$= (\alpha - 1 + m/k) \Delta(x_i^\rho | x_{1:i-1}^\rho), \quad (33f)$$

where (33a) is true by (16), (33c) is true by submodularity of f . We conclude that when $m \leq k$, $\beta = (\alpha - 1 + m/k)$, implying that for this case

$$f(x^\pi) \leq (2\alpha + m/k) f(x^\rho).$$

Next suppose that $m \geq k$. Observe that $|\tilde{z}_i^\pi| \leq k(n-1)$, since using the approximated augmented greedy policy, no more than $k(n-1)$ elements of z_i^π can be chosen by other agents. This implies the following:

$$\Delta(\tilde{z}_i^\pi | x_{1:i}^\rho) \leq (1/\min(k/|\tilde{z}_i^\pi|, 1)) \cdot \max_{z \in (\tilde{z}_i^\pi)^k} \Delta(z | x_{1:i}^\rho) \quad (34a)$$

$$\leq \max(|\tilde{z}_i^\pi|/k, 1) \cdot \max_{z \in (\tilde{z}_i^\pi)^k} \Delta(z | x_{1:i}^\rho) \quad (34b)$$

$$\leq \max(\min(k(n-1), m)/k, 1) \cdot \max_{z \in (\tilde{z}_i^\pi)^k} \Delta(z | x_{1:i}^\rho) \quad (34c)$$

$$= \min(n-1, m/k) \cdot \max_{z \in (\tilde{z}_i^\pi)^k} \Delta(z | x_{1:i}^\rho) \quad (34d)$$

$$\leq \alpha \cdot \min(n-1, m/k) \Delta(x_i^\rho | x_{1:i-1}^\rho) \quad (34e)$$

We conclude that when $m \geq k$, $\beta = \alpha \cdot \min(n-1, m/k)$, implying that for this case:

$$f(x^\pi) \leq (1 + \alpha(1 + \min(n-1, m/k))) f(x^\rho).$$

B. Proof for Lemma 2

We begin with

$$f(x^\rho) \leq f(x_{1:n-1}^\pi) + \Delta(x^\rho | x_{1:n-1}^\pi) \quad (35a)$$

$$= f(x_{1:n-1}^\pi) + \Delta(x_n^\rho | x_{1:n-1}^\pi, x_{1:n-1}^\pi) + \sum_{i=1}^{n-1} \Delta(x_i^\rho | x_{1:i-1}^\rho, x_{1:n-1}^\pi) \quad (35b)$$

$$\leq f(x_{1:n-1}^\pi) + \Delta(x_n^\rho | x_{1:n-1}^\pi) + \sum_{i=1}^{n-1} \Delta(x_i^\rho | x_{1:i}^\pi) \quad (35c)$$

$$\leq f(x_{1:n-1}^\pi) + \alpha_1 \Delta(x_n^\pi | x_{1:n-1}^\pi) + \sum_{i=1}^{n-1} \Delta(x_i^\rho | x_{1:i}^\pi) \quad (35d)$$

where (35a) and (35c) follow from submodularity of f , (35b) follows from the definition of $\Delta(\cdot)$, and (35d) follows

from (18a). Focusing on the sum in (35d), for any $0 \leq \varepsilon_1, \dots, \varepsilon_{n-1} \leq 1$ (and defining $\varepsilon_0 = 0$), we see that

$$\sum_{i=1}^{n-1} \Delta(x_i^\rho | x_{1:i}^\pi) = \sum_{i=1}^{n-1} (1 - \varepsilon_i) \Delta(x_i^\rho | x_{1:i}^\pi) + \sum_{i=1}^{n-1} \varepsilon_i \Delta(x_i^\rho | x_{1:i}^\pi) \quad (36a)$$

$$\leq \sum_{i=1}^{n-1} (1 - \varepsilon_i) \Delta(x_i^\rho | x_{1:i-1}^\pi) + \sum_{i=1}^{n-1} \alpha_2 \varepsilon_i \Delta^k(z_i^\pi | x_{1:i}^\pi) \quad (36b)$$

$$\leq \sum_{i=1}^{n-1} \alpha_1 (1 - \varepsilon_i) \Delta(x_i^\pi | x_{1:i-1}^\pi) + \sum_{i=1}^{n-1} \alpha_1 \alpha_2 \varepsilon_i \Delta(x_{i+1}^\pi | x_{1:i}^\pi) \quad (36c)$$

$$= \alpha_1 \alpha_2 \varepsilon_{n-1} \Delta(x_n^\pi | x_{1:n-1}^\pi) + \sum_{i=1}^{n-1} \alpha_1 (1 - \varepsilon_i + \alpha_2 \varepsilon_{i-1}) \Delta(x_i^\pi | x_{1:i-1}^\pi), \quad (36d)$$

where (36b) is true by submodularity of f (1st term) and (18b) (2nd term), (36c) is true by (18a), and (36d) is just a rearrangement of the terms. Applying this to (35d) yields

$$f(x^\rho) \leq f(x_{1:n-1}^\pi) + (\alpha_1 + \alpha_1 \alpha_2 \varepsilon_{n-1}) \Delta(x_n^\pi | x_{1:n-1}^\pi) + \sum_{i=1}^{n-1} (\alpha_1 - \alpha_1 \varepsilon_i + \alpha_1 \alpha_2 \varepsilon_{i-1}) \Delta(x_i^\pi | x_{1:i-1}^\pi) \quad (37a)$$

$$= (\alpha_1 + \alpha_1 \alpha_2 \varepsilon_{n-1}) \Delta(x_n^\pi | x_{1:n-1}^\pi) + \sum_{i=1}^{n-1} (1 + \alpha_1 - \alpha_1 \varepsilon_i + \alpha_1 \alpha_2 \varepsilon_{i-1}) \Delta(x_i^\pi | x_{1:i-1}^\pi). \quad (37b)$$

Suppose that for a particular choice of ε_i , we let

$$\varepsilon_i = \frac{(1/\alpha_1) \sum_{j=0}^{i-1} \alpha_2^j}{\sum_{j=0}^{n-1} \alpha_2^j}. \quad (38)$$

Since $\alpha_1, \alpha_2 \geq 1$, this satisfies the requirement that $0 \leq \varepsilon_i \leq 1$ for $i \in \{1, \dots, n-1\}$. Then

$$-\alpha_1 \varepsilon_i + \alpha_1 \alpha_2 \varepsilon_{i-1} = -\frac{\sum_{j=0}^{i-1} \alpha_2^j}{\sum_{j=0}^{n-1} \alpha_2^j} + \frac{\sum_{j=1}^{i-1} \alpha_2^j}{\sum_{j=0}^{n-1} \alpha_2^j} \quad (39)$$

$$= -\frac{1}{\sum_{j=0}^{n-1} \alpha_2^j} \quad (40)$$

Likewise

$$\alpha_1 \alpha_2 \varepsilon_{n-1} = \frac{\alpha_2 \sum_{j=0}^{n-2} \alpha_2^j}{\sum_{j=0}^{n-1} \alpha_2^j} = \frac{\sum_{j=1}^{n-1} \alpha_2^j}{\sum_{j=0}^{n-1} \alpha_2^j} \quad (41)$$

$$= 1 - \frac{1}{\sum_{j=0}^{n-1} \alpha_2^j} \quad (42)$$

Applying (40) and (42) to (37b) yields

$$f(x^\rho) \leq \left(\alpha_1 + 1 - \frac{1}{\sum_{j=0}^{n-1} \alpha_2^j} \right) \Delta(x_n^\pi | x_{1:n-1}^\pi) + \sum_{i=1}^{n-1} \left(1 + \alpha_1 - \frac{1}{\sum_{j=0}^{n-1} \alpha_2^j} \right) \Delta(x_i^\pi | x_{1:i-1}^\pi) \quad (43)$$

$$= \left(1 + \alpha_1 - \frac{1}{\sum_{j=0}^{n-1} \alpha_2^j} \right) f(x^\pi) \quad (44)$$



David Grimsman is a Ph.D. candidate in the Electrical and Computer Engineering Department at UC Santa Barbara. He completed BS in Electrical and Computer Engineering at Brigham Young University in 2006 as a Heritage Scholar, and with a focus on signals and systems. After working for BrainStorm, Inc. for several years as a trainer and IT manager, he returned to Brigham Young University and earned an MS in Computer Science in 2016. His research interests include multi-agent systems, game theory, distributed optimization, network science, linear systems theory, and security of cyberphysical systems.



Matthew R. Kirchner received his B.S. in Mechanical Engineering from Washington State University in 2007 and his M.S. in Electrical Engineering from the University of Colorado at Boulder in 2013. In 2007 he joined the Naval Air Warfare Center Weapons Division in the Navigation and Weapons Concepts Develop Branch and in 2012 transferred into the Image and Signal Processing Branch in the Research and Intelligence Department, Code D5J1000. He is currently a Ph.D. student in the Electrical and Computer Engineering Department at the University of California, Santa Barbara. His research interests include level set methods for optimal control, differential games, and reachability; multi-vehicle robotics; nonparametric signal and image processing; and navigation and flight control. He was the recipient of a Naval Air Warfare Center Weapons Division Graduate Academic Fellowship from 2010 to 2012 and in 2011 was named a Paul Harris Fellow by Rotary International. Matthew is a student member of the IEEE.



João P. Hespanha was born in Coimbra, Portugal, in 1968. He received the Licenciatura in electrical and computer engineering from the Instituto Superior Técnico, Lisbon, Portugal in 1991 and the Ph.D. degree in electrical engineering and applied science from Yale University, New Haven, Connecticut in 1998. From 1999 to 2001, he was Assistant Professor at the University of Southern California, Los Angeles. He moved to the University of California, Santa Barbara in 2002, where he currently holds a

Professor position with the Department of Electrical and Computer Engineering.

Dr. Hespanha is the recipient of the Yale University's Henry Prentiss Becton Graduate Prize for exceptional achievement in research in Engineering and Applied Science, a National Science Foundation CAREER Award, the 2005 best paper award at the 2nd Int. Conf. on Intelligent Sensing and Information Processing, the 2005 Automatica Theory/Methodology best paper prize, the 2006 George S. Axelby Outstanding Paper Award, and the 2009 Ruberti Young Researcher Prize. Dr. Hespanha is a Fellow of the International Federation of Automatic Control (IFAC) and of the IEEE. He was an IEEE distinguished lecturer from 2007 to 2013.

His current research interests include hybrid and switched systems; multi-agent control systems; game theory; optimization; distributed control over communication networks (also known as networked control systems); the use of vision in feedback control; stochastic modeling in biology; and network security.



Jason R. Marden is an associate professor in the Department of Electrical and Computer Engineering at the University of California, Santa Barbara. He received the B.S. degree in 2001 and the Ph.D. degree in 2007 (under the supervision of Jeff S. Shamma), both in mechanical engineering from the University of California, Los Angeles, where he was awarded the Outstanding Graduating Ph.D. Student in Mechanical Engineering. After graduating, he was a junior fellow in the Social and Information Sciences Laboratory at the California Institute of Technology until 2010 and then an assistant professor at the University of Colorado until 2015. He is a recipient of an ONR Young Investigator Award (2015), an NSF Career Award (2014), the AFOSR Young Investigator Award (2012), the SIAM CST Best Sicon Paper Award (2015), and the American Automatic Control Council Donald P. Eckman Award (2012). His research interests focus on game-theoretic methods for the control of distributed multiagent systems.