Approximate Supermodularity of Kalman Filter Sensor Selection

Luiz F. O. Chamon, George J. Pappas, and Alejandro Ribeiro

Abstract—This work considers the problem of selecting sensors in a large scale system to minimize the error in estimating its states. More specifically, the state estimation mean-square error (MSE) and worst-case error for Kalman filtering and smoothing. Such selection problems are in general NP-hard, i.e., their solution can only be approximated in practice even for moderately large problems. Due to its low complexity and iterative nature, greedy algorithms are often used to obtain these approximations by selecting one sensor at a time choosing at each step the one that minimizes the estimation performance metric. When this metric is supermodular, this solution is guaranteed to be (1 - 1/e)-optimal. This is however not the case for the MSE or the worst-case error. This issue is often circumvented by using supermodular surrogates, such as the $\log \det$, despite the fact that minimizing the log det is not equivalent to minimizing the MSE. Here, this issue is addressed by leveraging the concept of approximate supermodularity to derive near-optimality certificates for greedily minimizing the estimation mean-square and worst-case error. In typical application scenarios, these certificates approach the (1-1/e) guarantee obtained for supermodular functions, thus demonstrating that no change to the original problem is needed to obtain guaranteed good performance.

I. INTRODUCTION

Often, the outputs of a dynamical system can only be partially observed due to a limited sensing budget and/or because the system scale makes instrumenting all outputs impractical. This issue is critical for distributed systems, where power and communication constraints further limit the number of sensors available [2], [3], and is found in applications such as target tracking, field monitoring, power allocation, and biological systems analysis [4]–[9]. In these cases, we must choose which outputs to observe in order to best estimate the system internal states.

In general, solving this constrained output selection problem, often known as *sensor selection*, is NP-hard [10]–[16]. Still, there exist instances for which its solution can be approximated. For instance, when the goal is to find the smallest set of outputs that makes the system observable, greedy sensor selection is near-optimal [17]. When only structural descriptions are available, the problem is also known to be NP-hard, although a variant in which the output sensing can be designed freely admits a polynomial time solution [18]–[20]. Nevertheless, observability alone may not guarantee good state estimates because it does not take into account how hard each system mode is to observe. To do so, we turn to measures of the state estimation error.

In this setting, natural choices of objectives, especially in Kalman filtering applications, are the state estimation mean-square error (MSE) and the worst-case estimation error. However, approximation guarantees for these problems are scarce. In particular, typical guarantees for greedy search do not apply since these functions are not supermodular, a diminishing returns property used to obtain these results see [13]–[16] for counter-examples in the context of control theory. Moreover, though convex relaxations of these problems have been proposed, they too lack near-optimality certificates [15], [21]–[26]. In fact, system-independent guarantees cannot be provided for optimizing the state estimation MSE unless P = NP [16].

A typical way to address this issue is to use surrogate supermodular objectives, such as the log-determinant of the error covariance matrix or some information-theoretic measure relating the selected output set and its complement [10], [12], [17], [23], [27], [28]. Though effective, there is no direct relation between the estimation MSE or the worstcase estimation error and these cost functions. As a matter of fact, despite its usefulness in minimizing the volume of the confidence ellipsoid, the log-determinant is a poor proxy for the MSE—see Remark 1.

In this work, we take a page out of Tukey's book and try to solve the "right problem" [29], i.e., selecting sensors so as to directly minimize the estimation MSE or the worst-case estimation error. To make this selection scalable and obtain performance guarantees, we prove that greedy sensor selection is near-optimal for minimizing the estimation MSE and the worst-case estimation error in both filtering and smoothing settings. Our results show that when the measurement noise is large compared to the process noise, a scenario of practical relevance in Kalman filtering, greedy solutions to these problems approach the classical (1 - 1/e)-optimal certificate for supermodular objectives [30].

To derive these guarantees, we leverage two notions of *approximate supermodularity* in which limited violations of the diminishing return property are allowed (Section III). In the case of α -supermodular functions, violations are allowed up to a multiplicative factor α , whereas ϵ -supermodular functions violate supermodularity up to an additive constant ϵ . We prove that these functions can be minimized near-optimally using greedy search and that α and ϵ not only quantify how much supermodularity is violated, but also the loss in performance guarantee due to these violations (Theorems 1 and 2). We proceed to show that the estimation MSE is α -supermodular and

Department of Electrical and Systems Engineering, Uniof Pennsylvania. e-mail: luizf@seas.upenn.edu, versity {luizf, pappasg, aribeiro}@seas.upenn.edu. This work was supported by ARL DCIST CRA W911NF-17-2-0181. Part of the results in this paper appeared in [1].

that the worst-case estimation error—i.e., the spectral norm of the error covariance matrix—is ϵ -supermodular (Section IV). Combining these results yields performance guarantees for greedy sensor selection using the MSE and the worst-case error in filtering (Theorem 5) and smoothing (Theorem 6) state estimation problems (Section V). Note that the α supermodularity and corresponding guarantees for the MSE in filtering and smoothing problems were previously established in [1]. A field monitoring example is provided to illustrate these results (Section VI).

Notation: Lowercase boldface letters represent vectors (\boldsymbol{x}) , uppercase boldface letters are matrices (\boldsymbol{X}) , and calligraphic letters denote sets (\mathcal{X}) . We write $|\mathcal{X}|$ for the cardinality of \mathcal{X} and let \emptyset denote the empty set. To say \boldsymbol{X} is a positive semidefinite (PSD) matrix we write $\boldsymbol{X} \succeq 0$, so that for $\boldsymbol{X}, \boldsymbol{Y} \in \mathbb{R}^{n \times n}$, $\boldsymbol{X} \preceq \boldsymbol{Y} \Leftrightarrow \boldsymbol{b}^T \boldsymbol{X} \boldsymbol{b} \leq \boldsymbol{b}^T \boldsymbol{Y} \boldsymbol{b}$ for all $\boldsymbol{b} \in \mathbb{R}^n$. Similarly, we write $\boldsymbol{X} \succ 0$ when \boldsymbol{X} is positive definite. For a square matrix \boldsymbol{X} , we denote its spectral norm (ℓ_2 induced norm) by $\|\boldsymbol{X}\|$ and its maximum and minimum eigenvalues by $\lambda_{\max}(\boldsymbol{X})$ and $\lambda_{\min}(\boldsymbol{X})$ respectively.

II. PROBLEM FORMULATION

Consider an autonomous dynamical system whose internal states $\boldsymbol{x}_k \in \mathbb{R}^n$ evolve as

$$\boldsymbol{x}_{k+1} = \boldsymbol{F}\boldsymbol{x}_k + \boldsymbol{w}_k \tag{1}$$

and a set of sensors \mathcal{O} such that the output $y_{u,k} \in \mathbb{R}^{p_u}$ of sensor $u \in \mathcal{O}$ is given by

$$\boldsymbol{y}_{u,k} = \boldsymbol{H}_u \boldsymbol{x}_k + \boldsymbol{v}_{u,k}, \qquad (2)$$

where $\boldsymbol{F} \in \mathbb{R}^{n \times n}$ is the state transition matrix and $\boldsymbol{H}_u \in \mathbb{R}^{p_u \times n}$ are the sensor output matrices. The process noise \boldsymbol{w}_k and measurement noises $\boldsymbol{v}_{u,k}$ are zero-mean Gaussian random vectors with arbitrary covariances $\mathbb{E} \boldsymbol{w}_k \boldsymbol{w}_k^T = \boldsymbol{R}_w \succ$ 0 and $\mathbb{E} \boldsymbol{v}_{u,k} \boldsymbol{v}_{u,k}^T = \boldsymbol{R}_{v,u} \succ 0$ for all k. We assume $\{\boldsymbol{v}_{u_1,i}, \boldsymbol{v}_{u_2,j}, \boldsymbol{w}_i, \boldsymbol{w}_j\}$ are independent for all $i \neq j$ and $u_1, u_2 \in \mathcal{O}$. The initial states $\boldsymbol{x}_0 \sim \mathcal{N}(\bar{\boldsymbol{x}}_0, \boldsymbol{\Pi}_0)$ are assumed to be a Gaussian random vector with mean $\bar{\boldsymbol{x}}_0$ and covariance $\boldsymbol{\Pi}_0 \succ 0$.

We seek to estimate the states x_k from the observations $\{y_{u,i}\}_{u \in \mathcal{X}}$ for a sensing set $\mathcal{X} \subseteq \mathcal{O}$ and $i \leq k$. Depending on the observation history used to estimate the states at time k, we define the following estimation problems:

- (i) **Filtering**, where the current states are estimated using past observations, i.e., we seek $\hat{x}_k(\mathcal{X}) = \mathbb{E} [x_k \mid \{y_{u,i}\}_{u \in \mathcal{X}, i \leq k}].$
- (ii) **Smoothing**, where all states up to time k are estimated based on all observations. Formally, we seek $\tilde{\boldsymbol{x}}_k(\mathcal{X}) = \mathbb{E}[\bar{\boldsymbol{x}}_k \mid \{\boldsymbol{y}_{u,i}\}_{u \in \mathcal{X}, i \leq k}]$, where $\bar{\boldsymbol{x}}_k = \begin{bmatrix} \boldsymbol{x}_0^T & \boldsymbol{x}_1^T & \cdots & \boldsymbol{x}_k^T \end{bmatrix}^T$. As opposed to (i), states are estimated based on both past and future observations.

Note that we do not consider the problem of *prediction* in which the states at time k is estimated based on observations up to j < k, i.e., in which we evaluate $\check{x}_k(\mathcal{X}) = \mathbb{E}[\mathbf{x}_k | \{\mathbf{y}_{u,i}\}_{u \in \mathcal{X}, i \leq j < k}]$. Although filtering and prediction are different estimation problems, they do not differ from the viewpoint of sensor selection. Indeed, prediction only extrapolate the *j*-th filtered estimate using $\check{\boldsymbol{x}}_k = \boldsymbol{F}^{k-j} \hat{\boldsymbol{x}}_j$. Since it does not use additional measurements after the *j*-th iteration, the sensor selection problem for prediction is in fact the same as the one for filtering.

The performance of the estimation problems (i) and (ii) for a given sensing set $\mathcal{X} \subseteq \mathcal{O}$ is fully characterized by the error covariance matrix of the estimators \hat{x} and \tilde{x} . Explicitly,

$$\boldsymbol{P}_{k}(\boldsymbol{\mathcal{X}}) = \mathbb{E}\left[\boldsymbol{\hat{x}}_{k}(\boldsymbol{\mathcal{X}}) - \boldsymbol{x}_{k}\right]\left[\boldsymbol{\hat{x}}_{k}(\boldsymbol{\mathcal{X}}) - \boldsymbol{x}_{k}\right]^{T}$$
(3a)

$$\boldsymbol{Q}_{k}(\mathcal{X}) = \mathbb{E}\left[\boldsymbol{\tilde{x}}_{k}(\mathcal{X}) - \boldsymbol{\bar{x}}_{k}\right] \left[\boldsymbol{\tilde{x}}_{k}(\mathcal{X}) - \boldsymbol{\bar{x}}_{k}\right]^{T}$$
 (3b)

Observe that since the estimators depend on the sensing set \mathcal{X} , so do their error covariance matrices. In the propositions below, we show that the matrices in (3) can be written in a unified manner as a set function $Y_k : 2^{\mathcal{O}} \to \mathbb{S}_+$ of the form

$$\boldsymbol{Y}_{k}(\boldsymbol{\mathcal{X}}) = \left(\boldsymbol{M}_{\boldsymbol{\emptyset},k} + \sum_{u \in \boldsymbol{\mathcal{X}}} \boldsymbol{M}_{u,k}\right)^{-1}.$$
 (4)

where $M_{\emptyset,k} \succ 0$ is the *a priori* error covariance matrix that represents the error in the absence of sensor measurements and $M_{u,k} \succeq 0$ carries the information obtained by using sensor $u \in \mathcal{O}$ at time k.

Proposition 1. The error covariance matrix $P_k(\mathcal{X})$ of the filtered estimator \hat{x} from (i) can be written recursively using (4) with

$$\boldsymbol{M}_{\emptyset,k} = \boldsymbol{P}_{k|k-1}^{-1} \text{ and } \boldsymbol{M}_{u,k} = \boldsymbol{V}_{u}, \tag{5}$$

where $P_{k|k-1} = FP_{k-1}F^T + R_w$ is the prediction error covariance matrix, P_{k-1} is the estimation error covariance matrix at time k-1, $P_{0|-1} = \Pi_0$, and $V_u = H_u^T R_{v,u}^{-1} H_u$.

Proof. This result is obtained directly from the Kalman filter (KF) recursion (see, e.g., [31]). For ease of reference, we provide a derivation in the appendix.

Proposition 2. The error covariance matrix $Q_k(\mathcal{X})$ of the smoothed estimator \tilde{x} from (ii) can be written as in (4) with

$$M_{\emptyset,k} = \text{blkdiag}(\boldsymbol{\Pi}_0, \boldsymbol{I} \otimes \boldsymbol{R}_w)^{-1}$$

$$M_{u,k} = \boldsymbol{\Phi}_k^T \left(\boldsymbol{I} \otimes \boldsymbol{V}_u \right) \boldsymbol{\Phi}_k$$
(6)

where $V_u = H_u^T R_{v,u}^{-1} H_u$, blkdiag(X, Y) is the block diagonal matrix whose diagonal blocks are X and Y, and

$$\boldsymbol{\Phi}_{k} = \begin{bmatrix} \boldsymbol{I} & & & \\ \boldsymbol{F} & \boldsymbol{I} & & \\ \vdots & \vdots & \ddots & \\ \boldsymbol{F}^{k} & \boldsymbol{F}^{k-1} & \cdots & \boldsymbol{I} \end{bmatrix}.$$
 (7)

Proof. The proof follows classical stochastic least-squares results (see, e.g., [31]). For ease of reference, we provide a derivation in the appendix.

Propositions 1 and 2 show how the abstract set function (4) applies to both filtering and smoothing estimation problems. In the sequel, we formulate the sensor selection problem and derive theoretical guarantees directly in terms of the generic (4). Particular results for the estimation problems (i) and (ii) are provided in Section V.

A. Kalman filter sensor selection

Our goal is to study sensor selection problems in which we seek a sensing set $\mathcal{X} \subseteq \mathcal{O}$ constrained by a budget $|\mathcal{X}| \leq s$ so that the estimation error covariance matrix $Y_k(\mathcal{X})$ in (4) is minimized in some sense over a given time horizon N. Formally:

Problem (Kalman filter sensor selection). Take $N \in \mathbb{N}$ and let $\theta_k \geq 0$ for $k \in [0, N-1]$ be a set of nonnegative scalar weights and $h : \mathbb{S}_+ \to \mathbb{R}$ be a real-valued spectral function over the PSD cone. Kalman filter sensor selection can be written as the discrete optimization problem

$$\underset{\mathcal{X}\subseteq\mathcal{O},\,|\mathcal{X}|\leq s}{\text{minimize}} \quad \sum_{k=0}^{N-1} \theta_k h\Big[\mathbf{Y}_{m+k}(\mathcal{X}) \Big] - C_{\emptyset}, \qquad (\mathbf{P}^{\star})$$

where $C_{\emptyset} \triangleq \sum_{k=0}^{N-1} \theta_k h[\boldsymbol{M}_{\emptyset,m+k}^{-1}]$ is a constant.

It is worth noting that the constant C_{\emptyset} is included only so that the objective of (\mathbf{P}^*) vanishes when no sensor is selected—since $\mathbf{Y}_k(\emptyset) = \mathbf{M}_{\emptyset,k}^{-1}$, taking $\mathcal{X} = \emptyset$ in (\mathbf{P}^*) yields $\sum_{k=0}^{N-1} \theta_k h[\mathbf{Y}_{m+k}(\emptyset)] - C_{\emptyset} = 0$. This normalization has no effect on the optimization problem, but simplifies the statement of the approximation certificates presented in the following sections. Also observe that Propositions 1 and 2 do not take control inputs into account. Hence, though a sensing set obtained from (\mathbf{P}^*) can be used to perform state estimation in the presence of control inputs, it need not be optimal for the joint sensing set-controller design problem. This formulation is left for future works.

Different choices of N, θ_k , and h in (P^{*}) yield different sensor selection problems. For instance, (P^{*}) becomes a myopic sensor selection problem when N = 1 [27]. When N > 1, we can choose to optimize the *final estimation error*—corresponding to $\theta_k = 0$ for all k < N - 1and $\theta_{N-1} = 1$ —, the average error—by making $\theta_k = 1$ for all k—, or some weighted average error—e.g., a geometric discount using $\theta_k = \rho^{N-1-k}$ for $\rho < 1$. Observe that since the smoothing problem (ii) estimates all states up to k, these multi-step versions of (P^{*}) may not be useful in practice. We still consider them for the sake of symmetry.

The scalarization h is necessary because matrix minimization in the PSD cone is not well-posed [32] and its choice affects both the complexity and the practical usefulness of (P^{*}). For instance, (P^{*}) with $h(\mathbf{X}) = -\operatorname{Tr}(\mathbf{X}^{-1})$ is combinatorial but modular and can therefore be solved in linear time [33], [34]. This objective, however, is of little relevance in practice inasmuch as it has no direct relation to the estimation MSE. In general, three scalarizations are widely agreed to be of practical significance:

i. The trace defined as

$$\mathcal{T}^{\star} \in \operatorname*{argmin}_{\mathcal{T} \subseteq \mathcal{O}, \, |\mathcal{T}| \le s} f_T(\mathcal{T}) \triangleq \sum_{k=0}^{N-1} \theta_k \operatorname{Tr} \left[\mathbf{Y}_{m+k}(\mathcal{T}) \right] - C_{\emptyset},$$
(PI)

in which case the objective represents the estimation MSE.

ii. The spectral norm or the maximum eigenvalue

$$\mathcal{S}^{\star} \in \operatorname*{argmin}_{\mathcal{S} \subseteq \mathcal{O}, \, |\mathcal{S}| \le s} f_{\mathcal{S}}(\mathcal{S}) \triangleq \sum_{k=0}^{N-1} \theta_{k} \, \| \mathbf{Y}_{m+k}(\mathcal{S}) \| - C_{\emptyset},$$
(PII)

in which case the objective is a robust version of (PI), since it minimizes the worst-case estimation error.

iii. The log-determinant

$$\mathcal{D}^{\star} \in \operatorname*{argmin}_{\mathcal{D} \subseteq \mathcal{O}, \, |\mathcal{D}| \le s} \sum_{k=0}^{N-1} \theta_k \log \det \left[\mathbf{Y}_{m+k}(\mathcal{D}) \right] - C_{\emptyset},$$
(PIII)

in which case the objective corresponds to the volume of the confidence ellipsoid [22], [28].

These formulations are natural because they minimize different estimation error metrics related to the MSE, which is a customary objective in traditional state estimation, i.e., problems that do not consider sensor selection ($\mathcal{X} = \mathcal{O}$). Still, problems (PI)–(PIII) are challenging to solve. They are known to be NP-hard [10], [12]–[15], [17], [33], so their solutions can only be approximated in practice, typically by means of greedy algorithms [15], [17].

In the case of problem (PIII), the log-determinant is a monotonically decreasing and supermodular set function [17], [23], [27], [28], i.e., it displays a "diminishing returns" property that yields near-optimal performance guarantees for its greedy minimization [30]. The MSE and the worst-case estimation error, on the other hand, are not supermodular in general [12]– [15], [33]. In fact, stringent conditions on the system matrices are needed to make the estimation MSE supermodular [35], [36]. These facts notwithstanding, greedy sensor selection has been observed to perform well when solving (PI) and (PII) in both control [13], [15], [17] and other contexts [37], [38].

Our goal in the remainder of this paper is to show that in certain noise regimes of interest these objectives are approximately submodular, so that the greedy solutions of (PI) and (PII) come with performance guarantees analogous to those of (PIII), thereby reconciling empirical observations with our theoretical understanding of these problems. To obtain these near-optimal certificates, the following section develops a theory of *approximately supermodular* functions, studying operations that preserve approximate supermodularity and showing that these functions can be near-optimally minimized. We then show that the trace and the spectral norm are approximately supermodular scalarizations of Y_k so as to finally provide explicit suboptimality bounds for greedy sensor selection for filtering/smoothing.

Remark 1. Besides its intrinsic value as a minimizer of the volume of the confidence ellipsoid [22], [28], the log det is often put forward as a supermodular surrogate for the MSE. It is important to point out that this figure of merit is not directly related to the MSE or the worst-case error and that, in general, it is not a good surrogate for either of them. In fact, the log determinant is a good substitute for the MSE only when the problem has some inherent structure that constraints the confidence ellipsoid to be somewhat symmetric. Otherwise, since the volume of an ellipsoid can be reduced by decreasing

the length of a single principal axis, using the log det can lead to designs that perform well—in the MSE sense—along a few directions of the parameter space and poorly along all others. Formally, this can be seen by comparing the variation of the log-determinant and trace functions with respect to the eigenvalues of the PSD matrix Y,

$$rac{\partial \log \det(oldsymbol{Y})}{\partial \lambda_j(oldsymbol{Y})} = rac{1}{\lambda_j(oldsymbol{Y})} \qquad ext{and} \qquad rac{\partial \operatorname{Tr}(oldsymbol{Y})}{\partial \lambda_j(oldsymbol{Y})} = 1.$$

The gradient of the log-determinant is largest in the direction of the smallest eigenvalue of the error covariance matrix. In contrast, the MSE gives equal weight to all directions of the space. The latter yields balanced designs that are similar to the former only if those are forced to be balanced by the problem structure. Even for random dynamical systems with 10 states and 10 outputs (see Section VI for details), the MSE error achieved by a sensing set (s = 3) selected by minimizing the log det was up to 13% larger than a sensing set selected by minimizing the MSE itself (100 realizations).

III. APPROXIMATELY SUPERMODULAR FUNCTION MINIMIZATION

Supermodularity (submodularity) encodes a "diminishing returns" property of certain set functions that implies nearoptimality bounds on their greedy minimization (maximization). Well-known representatives of this class include the rank or log det of a sum of PSD matrices, the Shannon entropy, and the mutual information [33], [34]. Still, supermodularity is a stringent condition. In particular, it does not hold for the objectives of (PI) or (PII) [12]–[14], [33].

The purpose of *approximate supermodularity* (*submodularity*) is to relax the original "diminishing returns" property while controlling the magnitude of the violations. The rationale is that if a function is "almost" supermodular, then it should behave similar to a supermodular function. In what follows, we formalize and quantify these statements.

Consider a set function $f : 2^{\mathcal{O}} \to \mathbb{R}$ whose value for an arbitrary set $\mathcal{X} \subseteq \mathcal{O}$ is denoted by $f(\mathcal{X})$. We say f is *normalized* if $f(\emptyset) = 0$ and f is *monotone decreasing* if for all sets $\mathcal{A} \subseteq \mathcal{B} \subseteq \mathcal{O}$ it holds that $f(\mathcal{A}) \ge f(\mathcal{B})$. Observe that if a function is normalized and monotone decreasing it must be that $f(\mathcal{X}) \le 0$ for all $\mathcal{X} \subseteq \mathcal{O}$. Define

$$\Delta_u f(\mathcal{X}) = f(\mathcal{X}) - f(\mathcal{X} \cup \{u\})$$
(8)

to be the variation in the value of f incurred by adding the element $u \in \mathcal{O} \setminus \mathcal{X}$ to the set \mathcal{X} . Then, a set function f is *supermodular* if for all sets $\mathcal{A} \subseteq \mathcal{B} \subseteq \mathcal{O}$ and elements $u \in \mathcal{O} \setminus \mathcal{B}$ it holds that

$$\Delta_u f(\mathcal{A}) \ge \Delta_u f(\mathcal{B}). \tag{9}$$

A function f is submodular if -f is supermodular.

The relevance of supermodular functions in this work is due to the celebrated bound on the suboptimality of their greedy minimization [30]. Specifically, consider the generic cardinality constrained optimization problem

$$\mathcal{X}^{\star} \in \operatorname*{argmin}_{|\mathcal{X}| \le s} f(\mathcal{X}), \tag{PIV}$$

and construct its greedy solution by starting with $\mathcal{G}_0 = \emptyset$ and incorporating the elements from \mathcal{O} one at a time so as to maximize the gain at each step. Explicitly, at step j we do

$$\mathcal{G}_{j+1} = \mathcal{G}_j \cup \{u\},$$

with $u = \operatorname*{argmin}_{w \in \mathcal{O} \setminus \mathcal{G}_j} f\left(\mathcal{G}_j \cup \{w\}\right).$ (10)

The recursion in (10) is repeated for s steps to obtain a greedy solution with s elements. If f is monotone decreasing and supermodular [30], then

$$f(\mathcal{G}_s) \le (1 - e^{-1}) f(\mathcal{X}^\star). \tag{11}$$

The guarantee in (11), however, no longer applies when f is not supermodular. To provide guarantees in these cases, we leverage two measures of approximate supermodularity and derive near-optimality bounds for each of them. It is worth noting that though intuitive, such results are not trivial. In fact, [39] showed that for another measure of proximity, functions δ -close to supermodular cannot be optimized in polynomial time unless δ is small.

We start with a multiplicative relaxation of the supermodular property (9).

Definition 1 (α -supermodularity). A set function $f : 2^{\mathcal{O}} \to \mathbb{R}$ is α -supermodular, for $\alpha \in \mathbb{R}$, if for all sets $\mathcal{A} \subseteq \mathcal{B} \subseteq \mathcal{O}$ and all $u \in \mathcal{O} \setminus \mathcal{B}$ it holds that

$$\Delta_u f(\mathcal{A}) \ge \alpha \, \Delta_u f(\mathcal{B}). \tag{12}$$

For $\alpha \ge 1$, (12) reduces the original definition of supermodularity (9), in which case we refer to the function simply as supermodular [33], [34]. On the other hand, when $\alpha < 1$, fis said to be *approximately supermodular*. Notice that if f is decreasing, then (12) always holds for $\alpha = 0$. We are therefore interested in the largest α for which (12) holds, i.e.,

$$\alpha = \min_{\substack{\mathcal{A} \subset \mathcal{B} \subseteq \mathcal{O} \\ u \in \mathcal{O} \setminus \mathcal{B}}} \frac{\Delta_u f(\mathcal{A})}{\Delta_u f(\mathcal{B})}$$
(13)

This concept first appeared in the context of auction design [40], although it has been rediscovered in the context of discrete optimization, estimation, and control [1], [41]–[43]. It is worth noting that α in Definition 1 is also related to the *submodularity ratio* γ introduced in [11]. However, the proposed bounds on γ depended on the sparse eigenvalues of a matrix, that are NP-hard to compute. The first explicit (Pcomputable) bounds on α were obtained in [1], [42], though the same bounds were more recently derived for γ for a special case of (PI) (m = 1, $M_{\emptyset} = \sigma^2 I$, and rank-one M_u) [43]. These results were particularized for the problem of controlling linear systems in [44].

Interestingly, α not only measures how much f violates supermodularity, but it also quantifies the loss in performance guarantee incurred from these violations.

Theorem 1. Let f be a normalized, monotone decreasing, and α -supermodular set function (i.e., $f(\mathcal{X}) \leq 0$ for all $\mathcal{X} \subseteq \mathcal{O}$). Then, the solution obtained after r steps of the greedy algorithm in (10) obeys

$$f(\mathcal{G}_r) \le (1 - e^{-\alpha r/s}) f(\mathcal{X}^\star). \tag{14}$$

$$\begin{array}{c|c} & \text{Theorem 1} \\ \hline f(\mathcal{X}^{\star}) & f(\mathcal{G}_r) & (1 - e^{-\alpha r/s})f(\mathcal{X}^{\star}) & 0 \end{array}$$

Figure 1. Illustration of the near-optimal guarantee from Theorem 1.

Proof. See appendix.

Theorem 1 bounds the suboptimality of the greedy solution from (10) when its objective is α -supermodular. Indeed, since f is a non-positive function, it guarantees that $f(\mathcal{G}_r)$ cannot be too far from the optimal value (Figure 1). At the same time, it quantifies the effect of relaxing the supermodularity hypothesis typically used to provide performance guarantees in these settings. In fact, if f is supermodular ($\alpha = 1$) and r = s, we recover the guarantee (11) from [30]. On the other hand, for an approximately supermodular function ($\alpha < 1$), the result in (14) shows that the same 63% guarantee is recovered by greedily selecting a set of size s/α . Hence, α not only measures how much f violates supermodularity, but also gives a factor by which a solution set must increase to maintain supermodular-like near-optimality. It is worth noting that, as with the original bound in [30], (14) is not tight and that better results are typically obtained in practice (see Section VI).

Although α -supermodularity yields a multiplicative approximation factor, finding meaningful bounds on α can be challenging for certain set functions, such as the objective of (PII). It is therefore useful to look at approximate supermodularity from a different perspective, as proposed in [45].

Definition 2 (ϵ -supermodularity). A set function $f : 2^{\mathcal{O}} \to \mathbb{R}$ is ϵ -supermodular, for $\epsilon \in \mathbb{R}$, if for all multisets $\mathcal{A} \subseteq \mathcal{B} \subseteq \mathcal{O}$ and all $u \in \mathcal{O} \setminus \mathcal{B}$ it holds that

$$\Delta_u f(\mathcal{A}) \ge \Delta_u f(\mathcal{B}) - \epsilon. \tag{15}$$

Again, we say f is supermodular if $\epsilon \leq 0$ and approximately supermodular otherwise. As with α , we want the best ϵ that satisfies (15), which is given by

$$\epsilon = \max_{\substack{\mathcal{A} \subseteq \mathcal{B} \subseteq \mathcal{O} \\ u \in \mathcal{O} \setminus \mathcal{B}}} \Delta_u f(\mathcal{B}) - \Delta_u f(\mathcal{A}).$$
(16)

In contrast to α -supermodularity, we obtain an additive approximation guarantee for the greedy minimization of ϵ -supermodular functions.

Theorem 2. Let f be a normalized, monotone decreasing, and ϵ -supermodular set function (i.e., $f(\mathcal{X}) \leq 0$ for all $\mathcal{X} \subseteq \mathcal{O}$). Then, the solution obtained after r steps of the greedy algorithm in (10) obeys

$$f(\mathcal{G}_r) \le (1 - e^{-r/s}) \left[f(\mathcal{X}^*) + s \cdot \epsilon \right]. \tag{17}$$

Proof. See appendix.

As before, ϵ quantifies the loss in performance guarantee due to relaxing supermodularity. Indeed, (17) reveals that ϵ supermodular functions have the same guarantees as a supermodular function up to an additive factor of $\Theta(s\epsilon)$. In fact, if $\epsilon \leq (es)^{-1}|f(\mathcal{X}^*)|$ (recall that $f(\mathcal{X}^*) \leq 0$ due to normalization), then taking r = 3s recovers the supermodular 63% approximation factor. This same factor is obtained for $(\alpha \ge 1/3)$ -supermodular functions.

Although Theorems 1 and 2 characterize the loss in suboptimality incurred from violating supermodularity, their performance certificates depend on the specific values of α and ϵ . However, (13) and (16) reveal that finding α and ϵ for a general function is a combinatorial problem. To give actual nearoptimal guarantees for greedy solution of the sensor selection problems (PI) and (PII), the next section bounds the values of α and ϵ for different scalarizations of (4). Notice that we do not need to tackle the objectives of (PI) and (PII) directly thanks to the following lemma.

Lemma 1. Consider the set functions $f_i : 2^{\mathcal{O}} \to \mathbb{R}, i \in \mathbb{N}$. Then, for $\theta_i \ge 0$ and $b \in \mathbb{R}$,

- (i) if the f_i are α_i -supermodular, then $g = \sum_i \theta_i f_i + b$ is $\min(\alpha_i)$ -supermodular;
- (ii) if the f_i are ϵ_i -supermodular, then $g = \sum_i \theta_i f_i + b$ is $(\sum_i \theta_i \epsilon_i)$ -supermodular.

Proof. See appendix.

IV. APPROXIMATELY SUPERMODULAR SCALARIZATIONS

Theorems 1 and 2 apply to set functions that are (i) normalized, (ii) monotone decreasing, and (iii) approximately supermodular. By construction, the objective of (P^*) is normalized [(i)]. In this section, we obtain properties (ii) and (iii) by leveraging Lemma 1 and focusing only on the scalarizations underlying the objectives of (PI) and (PII), namely the trace

$$t(\mathcal{X}) = \operatorname{Tr}[\mathbf{Y}_k(\mathcal{X})] \tag{18}$$

and the spectral norm

$$e(\mathcal{X}) = \|\mathbf{Y}_k(\mathcal{X})\|, \qquad (19)$$

for Y_k defined as in (4).

To establish that set functions t and e are monotone decreasing [(ii)], we prove that $Y_k(\mathcal{X})$ is a decreasing set function in the PSD cone (Lemma 2). The definition of Loewner order and the monotonicity of the trace [46] imply the desired result. The monotone decreasing property of the objectives of (PI) and (PII) then follows immediately from $\theta_i \geq 0$.

Lemma 2. The matrix-valued set function $Y_k(\mathcal{X})$ in (4) is monotonically decreasing with respect to the PSD cone, i.e., $\mathcal{A} \subseteq \mathcal{B} \subseteq \mathcal{O} \Rightarrow Y_k(\mathcal{A}) \succeq Y_k(\mathcal{B}).$

Proof. The monotone decreasing nature of Y_k stems from the fact that matrix inversion is an operator antitone function, i.e., that for $X, Y \succeq 0$, it holds that $X \succeq Y \Leftrightarrow X^{-1} \preceq Y^{-1}$ [47]. To see this, write (4) as $Y_k(\mathcal{X}) = \mathbf{R}(\mathcal{X})^{-1}$, where $\mathbf{R}(\mathcal{X}) = \mathbf{M}_{\emptyset,k} + \sum_{u \in \mathcal{X}} \mathbf{M}_{u,k}$. Then, notice that since \mathbf{R} is a sum of PSD matrices, it holds that $\mathbf{R}(\mathcal{A}) \succeq 0$ for all $\mathcal{A} \subseteq \mathcal{O}$. Moreover, \mathbf{R} is a modular (additive) function, i.e., $\mathbf{R}(\mathcal{A} \cup \mathcal{B}) = \mathbf{R}(\mathcal{A}) + \mathbf{R}(\mathcal{B})$. Hence, for $\mathcal{A} \subseteq \mathcal{B} \subseteq \mathcal{O}$, we obtain

$$\boldsymbol{R}(\boldsymbol{\mathcal{B}}) = \boldsymbol{R}(\boldsymbol{\mathcal{A}}) + \boldsymbol{R}(\boldsymbol{\mathcal{B}} \setminus \boldsymbol{\mathcal{A}}) \succeq \boldsymbol{R}(\boldsymbol{\mathcal{A}}).$$
(20)

It is straightforward to obtain from (20) that for $\mathcal{A} \subseteq \mathcal{B} \subseteq \mathcal{O}$

$$oldsymbol{R}(\mathcal{A}) \preceq oldsymbol{R}(\mathcal{B}) \Leftrightarrow oldsymbol{R}(\mathcal{A})^{-1} \succeq oldsymbol{R}(\mathcal{B})^{-1} \Leftrightarrow oldsymbol{Y}_k(\mathcal{A}) \succeq oldsymbol{Y}_k(\mathcal{B}).$$

The remainder of this section is dedicated to showing that the trace and the spectral norm are approximately supermodular scalarizations of Y_k , starting with the set trace function (18). In what follows, we omit the dependence of Y, M_{\emptyset} , and M_u on k for clarity.

Theorem 3. Let t be the set function in (18). Then, t is α -supermodular with

$$\alpha \ge \frac{\mu_{\min}}{\mu_{\max}} > 0, \tag{21}$$

where

$$0 < \mu_{\min} \leq \lambda_{\min} \left[\boldsymbol{M}_{\boldsymbol{\emptyset}} \right] \leq \lambda_{\max} \left[\boldsymbol{M}_{\boldsymbol{\emptyset}} + \sum_{u \in \mathcal{O}} \boldsymbol{M}_{u} \right] \leq \mu_{\max}.$$

Remark 2. Although there exist examples for which $t(\mathcal{X})$ is not supermodular, the general statement of Theorem 3 does not allow us to claim that $\alpha < 1$. A simple counter-example involves the case in which $M_{\emptyset} = \mu_0 I$ and $M_u = \mu_u I$, $\mu_0, \mu_u \ge 0$, so that t becomes effectively a scalar function of μ_0, μ_u . Since scalar convex functions of positive modular functions are supermodular [48], we have $\alpha \ge 1$ in this case.

The proof of Theorem 3 relies on the following bounds on the variation $\Delta_u t(\mathcal{X})$:

Lemma 3. For all $\mathcal{X} \subset \mathcal{O}$ and $u \in \mathcal{O} \setminus \mathcal{X}$, it holds that

 $\lambda_{\min} \left[\boldsymbol{Y}(\mathcal{X}) \right] \operatorname{Tr} \left[\boldsymbol{M}_{u} \boldsymbol{Y} \left(\mathcal{X} \cup \{u\} \right) \right] \leq \Delta_{u} t(\mathcal{X}) \leq \\\lambda_{\max} \left[\boldsymbol{Y}(\mathcal{X}) \right] \operatorname{Tr} \left[\boldsymbol{M}_{u} \boldsymbol{Y} \left(\mathcal{X} \cup \{u\} \right) \right].$ (22)

Proof. See appendix.

Theorem 3 then follows readily from this technical lemma.

Proof of Theorem 3. Notice from (8) and (13) that α can be written as

$$\alpha = \min_{\substack{\mathcal{A} \subset \mathcal{B} \subseteq \mathcal{O} \\ u \in \mathcal{O} \setminus \mathcal{B}}} \frac{\Delta_u t(\mathcal{A})}{\Delta_u t(\mathcal{B})}.$$

From Lemma 3, we then obtain

$$\alpha \geq \frac{\lambda_{\min} \left[\mathbf{Y}(\mathcal{A}) \right]}{\lambda_{\max} \left[\mathbf{Y}(\mathcal{B}) \right]} \times \frac{\operatorname{Tr} \left[\mathbf{M}_u \mathbf{Y} \left(\mathcal{A} \cup \{u\} \right) \right]}{\operatorname{Tr} \left[\mathbf{M}_u \mathbf{Y} \left(\mathcal{B} \cup \{u\} \right) \right]}.$$
 (23)

To simplify (23), recall from Lemma 2 that Y is monotone decreasing in the PSD cone. Since $\mathcal{A} \subset \mathcal{B} \subseteq \mathcal{O}$, it holds that $Y(\mathcal{A} \cup \{u\}) \succeq Y(\mathcal{B} \cup \{u\})$. The ordering of the PSD cone (Loewner order) gives us that the second term in (23) is always greater than one, which yields

$$\alpha \geq \frac{\lambda_{\min} \left[\boldsymbol{Y}(\mathcal{A}) \right]}{\lambda_{\max} \left[\boldsymbol{Y}(\mathcal{B}) \right]}.$$

The lower bound in (21) is readily obtained by observing that the decreasing nature of Y implies that for any set $\mathcal{X} \subseteq \mathcal{O}$:

$$\lambda_{\min} \left[\boldsymbol{Y}(\mathcal{O}) \right] \leq \lambda_{\min} \left[\boldsymbol{Y}(\mathcal{X}) \right] \leq \lambda_{\max} \left[\boldsymbol{Y}(\mathcal{X}) \right] \leq \lambda_{\max} \left[\boldsymbol{Y}(\emptyset) \right].$$

Theorem 3 gives a deceptively simple bound on the α supermodularity of the set function t in (18) depending on the spectrum of the M_{\emptyset}, M_i . This bound can be interpreted geometrically in terms of the "range" of the error covariance matrix Y. To see this, define the *numerical range* of the set function Y as

$$W_{\mathcal{O}}(\boldsymbol{Y}) = W \left[\bigoplus_{\mathcal{X} \subseteq \mathcal{O}} \boldsymbol{Y}(\mathcal{X}) \right], \qquad (24)$$

where $A \oplus B$ = blkdiag(A, B) is the direct sum of matrices A and B and $W(M) = \{x^T M x \mid ||x||_2 = 1\}$ is the classical numerical range [46]. Since the numerical range is a convex set, we can define its relative diameter as

$$\Delta = \max_{\mu,\eta \in W_{\mathcal{O}}(\mathbf{Y})} \left| \frac{\mu - \eta}{\mu} \right|.$$
(25)

Then, the following holds:

Proposition 3. The set functions t in (18) is α -supermodular with

$$\alpha \ge 1 - \Delta_{\alpha}$$

where Δ is the relative diameter of the numerical range of Y in (25).

Proof. Since $\mathbf{Y} \succ 0$, the numerical range in (24) is the bounded convex hull of the eigenvalues of $\mathbf{Y}(\mathcal{X})$ for all $\mathcal{X} \subseteq \mathcal{O}$ [46]. We can therefore simplify (25) using the fact that it is monotonically increasing in μ and decreasing in η . Explicitly,

$$\Delta = \max_{\mathcal{X}, \mathcal{Y} \subseteq \mathcal{O}} \left| \frac{\lambda_{\max} \left[\mathbf{Y}(\mathcal{Y}) \right] - \lambda_{\min} \left[\mathbf{Y}(\mathcal{X}) \right]}{\lambda_{\max} \left[\mathbf{Y}(\mathcal{Y}) \right]} \right|$$

Using the fact that Y is monotonically decreasing (Lemma 2), this maximum is achieved for

$$\Delta = rac{\lambda_{\max}\left[oldsymbol{Y}(\emptyset)
ight] - \lambda_{\min}\left[oldsymbol{Y}(\mathcal{O})
ight]}{\lambda_{\max}\left[oldsymbol{Y}(\emptyset)
ight]}.$$

The bound in (21) thus becomes

$$\alpha \geq \frac{\lambda_{\min} \left[\mathbf{Y}(\mathcal{O}) \right]}{\lambda_{\max} \left[\mathbf{Y}(\emptyset) \right]} = 1 - \Delta.$$

Hence, (21) bounds how much t deviates from a supermodular function (as quantified by α) in terms of the numerical range of its underlying function Y. The shorter the range of Y, the more supermodular-like (18) will be.

Notice from Lemma 1, that evaluating the α -supermodularity of the objective of (PI) is a straightforward corollary of Theorem 3.

Corollary 1. The objective of (PI) is α -supermodular for

$$\alpha \ge \min_{0 \le k \le N-1} \frac{\lambda_{\min} \left[\boldsymbol{M}_{\boldsymbol{\emptyset}, m+k} \right]}{\lambda_{\max} \left[\boldsymbol{M}_{\boldsymbol{\emptyset}, m+k} + \sum_{u \in \mathcal{O}} \boldsymbol{M}_{u, m+k} \right]}.$$
 (26)

Proceeding, we now bound the ϵ -supermodularity of (19), once again temporarily omitting the dependences on k.

Theorem 4. Let e be defined as in (19). Then, e is ϵ -supermodular with

$$\epsilon \leq \frac{\lambda_{\max}(\sum_{u \in \mathcal{O}} M_u)}{\lambda_{\min} \left[M_{\emptyset} \right]^2}.$$
(27)

Proof. The proof follows a homotopy argument, i.e., we define a continuous map between $\Delta_u e(\mathcal{A})$ and $\Delta_u e(\mathcal{B})$ and bound its derivative using spectral bounds on the sum of Hermitian matrices. The inequality in (27) then follows from the fundamental theorem of calculus.

Start by defining the homotopy, with $t \in [0, 1]$,

$$h_{\mathcal{AB}}(t) = \left\| \mathbf{Z}(t)^{-1} \right\| - \left\| \left(\mathbf{Z}(t) + \mathbf{M}_u \right)^{-1} \right\|$$
 (28)

for $\mathcal{A} \subset \mathcal{B} \subseteq \mathcal{O}$, $\mathbf{Z}(t) = \mathbf{R}(\mathcal{A}) + t[\mathbf{R}(\mathcal{B}) - \mathbf{R}(\mathcal{A})]$, and $\mathbf{R}(\mathcal{X}) = \mathbf{M}_{\emptyset} + \sum_{u \in \mathcal{X}} \mathbf{M}_{u}$. Note that since $\mathbf{Y}(\mathcal{X}) = \mathbf{R}(\mathcal{X})^{-1}$, we have $h_{\mathcal{A}\mathcal{B}}(0) = \Delta_{u}e(\mathcal{A})$ and $h_{\mathcal{A}\mathcal{B}}(1) = \Delta_{u}e(\mathcal{B})$. If $\dot{h}_{\mathcal{A}\mathcal{B}}(t)$ is the derivative of $h_{\mathcal{A}\mathcal{B}}$ with respect to t, it therefore holds that

$$\Delta_u e(\mathcal{B}) = \Delta_u e(\mathcal{A}) + \int_0^1 \dot{h}_{\mathcal{A}\mathcal{B}}(t) dt.$$
(29)

Comparing (29) to the definition of ϵ in (16), we obtain

$$\epsilon = \max_{\substack{\mathcal{A} \subset \mathcal{B} \subseteq \mathcal{O} \\ u \in \mathcal{O} \setminus \mathcal{B}}} \int_{0}^{1} \dot{h}_{\mathcal{A}\mathcal{B}}(t) dt.$$
(30)

We now proceed by evaluating \dot{h} . We omit the dependence on \mathcal{A} and \mathcal{B} for conciseness. First, recall from matrix analysis that for any $\mathbf{X} \succ 0$ we have

$$\frac{d}{dt} \left\| \boldsymbol{X}(t)^{-1} \right\| = \boldsymbol{q}(t)^T \left[\frac{d}{dt} \boldsymbol{X}(t)^{-1} \right] \boldsymbol{q}(t)$$

$$= -\boldsymbol{q}(t)^T \boldsymbol{X}(t)^{-1} \dot{\boldsymbol{X}}(t) \boldsymbol{X}(t)^{-1} \boldsymbol{q}(t),$$
(31)

where q(t) is the eigenvector relative to the maximum eigenvalue of X(t). To obtain (31), we used the fact that $||X|| = \lambda_{\max}(X)$ for $X \succeq 0$ and $\frac{d}{dt}\lambda_{\max}[X(t)] = q(t)^T \dot{X}(t)q(t)$. We then used $\frac{d}{dt}X(t)^{-1} = -X(t)^{-1}\dot{X}(t)X(t)^{-1}$ [47]. In view of (28) and (31), we obtain

$$\dot{h}(t) = \boldsymbol{w}(t)^{T} \left[\boldsymbol{Z}(t) + \boldsymbol{M}_{u} \right]^{-1} \left[\boldsymbol{R}(\boldsymbol{\mathcal{B}}) - \boldsymbol{R}(\boldsymbol{\mathcal{A}}) \right] \\ \times \left[\boldsymbol{Z}(t) + \boldsymbol{M}_{u} \right]^{-1} \boldsymbol{w}(t) \qquad (32) \\ - \boldsymbol{u}(t)^{T} \boldsymbol{Z}(t)^{-1} \left[\boldsymbol{R}(\boldsymbol{\mathcal{B}}) - \boldsymbol{R}(\boldsymbol{\mathcal{A}}) \right] \boldsymbol{Z}(t)^{-1} \boldsymbol{u}(t),$$

where $\boldsymbol{u}(t)$ and $\boldsymbol{w}(t)$ are the eigenvectors relative to the maximum eigenvalues of $\boldsymbol{Z}(t)^{-1}$ and $[\boldsymbol{Z}(t) + \boldsymbol{M}_u]^{-1}$ respectively.

We now proceed by finding a bound for \dot{h} in (32) that is independent of t. To do so, observe that $\mathbf{R}(\mathcal{B}) \succeq \mathbf{R}(\mathcal{A})$, since $\mathcal{A} \subseteq \mathcal{B}$ and \mathbf{R} is monotone increasing (Lemma 2). By the Loewner ordering, the second term in (32) is therefore non-positive. Thus, using Rayleigh's spectral inequality and the fact $||\mathbf{w}(t)|| = 1$ for all t yields

$$\dot{h}(t) \leq \left\| \left(\boldsymbol{Z}(t) + \boldsymbol{M}_{u} \right)^{-1} \left[\boldsymbol{R}(\boldsymbol{B}) - \boldsymbol{R}(\boldsymbol{A}) \right] \left(\boldsymbol{Z}(t) + \boldsymbol{M}_{u} \right)^{-1} \right\|.$$

Next, we once again use the fact that $\mathbf{R}(\mathcal{B}) \succeq \mathbf{R}(\mathcal{A})$ to obtain $\mathbf{Z}(t) \succeq \mathbf{Z}(0) = \mathbf{R}(\mathcal{A})$, effectively removing the dependence on t. Using Cauchy-Schwartz then yields

$$\dot{h}(t) \leq \lambda_{\max} \left[\left(\boldsymbol{R}(\boldsymbol{A}) + \boldsymbol{M}_u \right)^{-2} \right] \lambda_{\max} \left[\boldsymbol{R}(\boldsymbol{B}) - \boldsymbol{R}(\boldsymbol{A}) \right],$$

which can be used in (30) to get

$$\epsilon \leq \max_{\substack{\mathcal{A} \subseteq \mathcal{B} \subseteq \mathcal{O} \\ u \in \mathcal{O} \setminus \mathcal{B}}} \frac{\lambda_{\max} \left[\boldsymbol{R}(\mathcal{B}) - \boldsymbol{R}(\mathcal{A}) \right]}{\lambda_{\min} \left[\boldsymbol{R}(\mathcal{A}) + \boldsymbol{M}_{u} \right]^{2}},$$
(33)

Algorithm 1 Cardinality r greedy set selection for (P^*)

Initialize
$$\mathcal{G}_0 \leftarrow \emptyset$$

for $j = 1, \dots, r$
 $u \leftarrow \underset{w \in \mathcal{O} \setminus \mathcal{G}_j}{\operatorname{argmin}} \sum_{k=0}^{N-1} \theta_k h \Big[\mathbf{Y}_{m+k}(\mathcal{G}_j \cup \{w\}) \Big]$
 $\mathcal{G}_{j+1} \leftarrow \mathcal{G}_j \cup \{u\}$
end

The inequality in (27) is then obtained using $\|\mathbf{R}(\mathcal{B}) - \mathbf{R}(\mathcal{A})\| \leq \lambda_{\max}(\sum_{u \in \mathcal{O}} \mathbf{M}_u)$ and $\lambda_{\min} [\mathbf{R}(\mathcal{A}) + \mathbf{M}_u] \geq \lambda_{\min} [\mathbf{M}_{\emptyset}].$

From Theorem 4 and Lemma 1, we obtain:

Corollary 2. The objective of (PII) is ϵ -supermodular for

$$\epsilon \leq \sum_{k=0}^{N-1} \theta_k \frac{\lambda_{\max}(\sum_{u \in \mathcal{O}} M_{u,m+k})}{\lambda_{\min} \left[M_{\emptyset,m+k} \right]^2},$$
(34)

In the sequel, we particularize these results for the estimation problems (i) and (ii) in Section II, showing that we we obtain supermodular-like guarantees when a certain measure of signal-to-noise ratio (SNR) is small. We also argue that this is in fact the typical setting in which KFs are useful in practice.

V. NEAR-OPTIMAL SENSOR SELECTION

In this section, we obtain approximation guarantees for the greedy optimization of the MSE and worst-case error in filtering and smoothing problems. More specifically, we provide near-optimal certificates for the solution \mathcal{G}_r obtained using Algorithm 1 with Y taken from Proposition 1 (filtering) or Proposition 2 (smoothing) and with the trace or spectral norm scalarizations h (as in Section II-A).

A. Greedy sensor selection for filtering is near-optimal

The main result of this section stems from applying Proposition 1 to Corollaries 1 and 2.

Theorem 5. Let $Y_k(\mathcal{A}) = P_k(\mathcal{A})$, the filtering error covariance from Proposition 1, in (PI)–(PII) and denote their solutions as $\widehat{\mathcal{T}}^*$ and $\widehat{\mathcal{S}}^*$, respectively. Let $\widehat{\mathcal{T}}_g$ and $\widehat{\mathcal{S}}_g$ be the r-elements greedy solutions obtained by using Algorithm 1 on (PI) and (PII). Then,

$$f_T(\widehat{\mathcal{T}}_g) \le (1 - e^{-\widehat{\alpha}r/s}) f_T(\widehat{\mathcal{T}}^*)$$

$$f_S(\widehat{\mathcal{S}}_g) \le (1 - e^{-r/s}) \left[f_S(\widehat{\mathcal{S}}^*) + s \cdot \widehat{\epsilon} \right]$$

for f_T and f_S the objectives of (PI) and (PII) respectively, $\hat{\alpha} \ge \min_{k \le N-1} \hat{\alpha}_k$, and $\hat{\epsilon} \le \sum_{k=0}^{N-1} \theta_k \hat{\epsilon}_k$ with

$$\hat{\alpha}_{k} \triangleq \frac{\lambda_{\min} \left[\boldsymbol{P}_{m+k|m+k-1}^{-1} \right]}{\lambda_{\max} \left[\boldsymbol{P}_{m+k|m+k-1}^{-1} + \sum_{u \in \mathcal{O}} \boldsymbol{V}_{u} \right]}$$
(35)

$$\hat{e}_{k} \triangleq \lambda_{\max} \left(\sum_{u \in \mathcal{O}} \boldsymbol{V}_{u} \right) \lambda_{\max} \left(\boldsymbol{P}_{m+k|m+k-1} \right)^{2} \qquad (36)$$



Figure 2. α -supermodularity bounds (Theorem 5) for different process noise variances and $||F||_2$.

where $V_u = H_u^T R_{v,u}^{-1} H_u$.

Theorem 5 shows that greedy sensor selection for state filtering is near-optimal. In fact, it gives explicit bounds on the suboptimality of greedy sensing sets as a function of parameters of the dynamical system. When we do not allow violations of the cardinality constraint, i.e., for r = s, and (35) is close to one or (36) is close to zero, the guarantees from Theorem 5 approach the (1 - 1/e)-optimality of greedy supermodular minimization.

To see when this occurs specifically, notice that (35) can be rewritten using the Riccati equation (48) as

$$\hat{\alpha}_{k} = \frac{\lambda_{\min}\left[\boldsymbol{P}_{m+k|m+k-1}^{-1}\right]}{\lambda_{\max}\left[\boldsymbol{P}_{m+k}^{-1}(\mathcal{O})\right]} = \frac{\lambda_{\min}\left[\boldsymbol{P}_{m+k}(\mathcal{O})\right]}{\lambda_{\max}\left[\boldsymbol{P}_{m+k|m+k-1}\right]}$$

i.e., as the ratio between the *a posteriori* error using all sensors and the *a priori* error. Hence, α_j in (35) is large when the *a priori* and *a posteriori* error covariance matrices are similar. To get additional insight, consider the particular case in which $\Pi_0 = \mathbf{R}_w = \sigma_w^2 \mathbf{I}$ and $\mathbf{R}_v = \sigma_v^2 \mathbf{I}$. Also, suppose that each state is measured directly, i.e., that $\mathcal{O} = \{1, \dots, p\}$ and that $\mathbf{H}_u = \mathbf{e}_u^T$, $u \in \mathcal{O}$, where \mathbf{e}_u is the *u*-th column of a $p \times p$ identity matrix. Then, from (47) and (48), the matrices of interest take the form

$$\boldsymbol{P}_{k|k-1} = \boldsymbol{F} \boldsymbol{P}_{k-1} \boldsymbol{F}^T + \sigma_w^2 \boldsymbol{I}$$
(37a)

$$\boldsymbol{P}_{k}(\mathcal{O}) = \left(\boldsymbol{P}_{k|k-1}^{-1} + \sigma_{v}^{-2}\boldsymbol{I}\right)^{-1}$$
(37b)

Notice that for (35) to approach one, it must be that $P_k(\mathcal{O}) \approx P_{k|k-1}$. Since $P_{k|k-1}$ depends directly on σ_w^2 , it is clear from (37b) that this occurs when $\sigma_v^2 \gg \sigma_w^2$. Additionally, we need $P_{k|k-1}$ to be well-conditioned. When σ_w^2 is large, this stems directly from (37a). On the other hand, if $\sigma_w^2 \ll ||F||$, then $P_{k|k-1} \approx FP_{k-1}F^T$ and its condition number depends directly on F, since $\kappa(P_{k|k-1}) \leq \kappa(P_{k-1})\kappa(F)^2$, where κ is the condition number with respect



Figure 3. ϵ -supermodularity bounds (Theorem 5) for different process noise variances and $\|F\|_2$.

to the spectral norm [46]. A similar argument applies to (36), which in this simplified setting reduces to

$$\hat{\epsilon}_k = \sigma_v^{-2} \, \lambda_{ ext{max}} \left(oldsymbol{F} oldsymbol{P}_{m+k-1} oldsymbol{F}^T + \sigma_w^2 oldsymbol{I}
ight)^2$$
 .

Hence, similar conditions as above also guarantee that $\hat{\epsilon}$ is small. In particular, this occurs when σ_v^2 is much larger than σ_w^2 and $\|F\|$.

Figures 2 and 3 illustrate these observations by showing the values of $\hat{\alpha}$ and $\hat{\epsilon}$, averaged over 100 dynamical systems realizations, for different values of σ_w^2/σ_v^2 and $\|m{F}\|$ (see Section VI for details). Indeed, the guarantees from Theorem 5 are stronger when $\sigma_v^2 \gg \sigma_w^2, \sigma_v^2 \gg \|F\|$, and when the system has modes with similar rates ($\kappa(F) \approx 1$). It is important to highlight that this is the scenario in which KFs are most useful: if the process noise dominates the estimation error, the system trajectory is mostly random and the choice of sensing subset has little impact on the estimation performance [31], [49]. Also notice from Figure 2 that for $\sigma_v^2 > 10\sigma_w^2$ and $\|\mathbf{F}\| = 0.3$, we can obtain supermodular-like guarantees for greedy sensor selection by violating the cardinality constraint by 35%. Since guarantees for problem (PII) are additive, meaningful performance bounds are obtained on a narrower range of parameters (Figure 3).

Remark 3. In [16], it was shown that it is impossible to obtain a non-trivial, universal (system-independent) bound on $\hat{\alpha}$ unless P = NP. To do so, they provide a reduction from the problem of *exact cover by 3-sets* (X3C) that relies on a dynamical system with noiseless outputs. Notice that this is in line with the results from Theorem 5 given that $\hat{\alpha} \rightarrow 0$ as $\sigma_v^2 \rightarrow 0$, i.e., the guarantees based on approximate supermodularity become vacuous for noise-free outputs. When this is not the case, however, Theorem 5 provides non-trivial near-optimal certificates for greedy sensor selection as a function of the dynamical system parameters.



Figure 4. α -supermodularity bounds (Theorem 6) for different process noise variances and $||F||_2$.

B. Greedy sensor selection for smoothing is near-optimal

We obtain similar results for smoothing by taking the values of $M_{\emptyset,k}$ and $M_{u,k}$ from Proposition 2.

Theorem 6. Let $Y_k(\mathcal{A}) = Q_k(\mathcal{A})$, the smoothing error covariance matrix from Proposition 2, in (PI)–(PII) and take $\tilde{\mathcal{T}}^*$ and $\tilde{\mathcal{S}}^*$ to be their respective solutions. Let $\tilde{\mathcal{T}}_g$ and $\tilde{\mathcal{S}}_g$ be the *r*-elements greedy solutions obtained by using Algorithm 1 on (PI) and (PII). Then,

$$f_T(\widetilde{\mathcal{T}}_g) \le (1 - e^{-\tilde{\alpha}r/s}) f_T(\widetilde{\mathcal{T}}^\star)$$
(38)

$$f_S(\widetilde{\mathcal{S}}_g) \le (1 - e^{-r/s}) \left[f_S(\widetilde{\mathcal{S}}^\star) + s \cdot \widetilde{\epsilon} \right]$$
(39)

for f_T and f_S the objectives of (PI) and (PII) respectively, $\tilde{\alpha} \ge \min_{k \le N-1} \tilde{\alpha}_k$ and $\tilde{\epsilon} \le \sum_{k=0}^{N-1} \theta_k \tilde{\epsilon}_k$ with

$$\tilde{\alpha}_{k} \triangleq \frac{1/\ell_{\max}}{\lambda_{\max} \left[\boldsymbol{C}^{-1} + \sum_{u \in \mathcal{O}} \boldsymbol{\Phi}_{m+k}^{T} \left(\boldsymbol{I} \otimes \boldsymbol{V}_{u} \right) \boldsymbol{\Phi}_{m+k} \right]} \quad (40)$$

$$\tilde{\epsilon}_{k} \triangleq \lambda_{\max} \left[\sum_{u \in \mathcal{O}} \boldsymbol{\Phi}_{m+k}^{T} \left(\boldsymbol{I} \otimes \boldsymbol{V}_{u} \right) \boldsymbol{\Phi}_{m+k} \right] \ell_{\max}^{2}$$
(41)

where $V_u = H_u^T R_{v,u}^{-1} H_u$, $C = \text{blkdiag}(\Pi_0, I \otimes R_w)$, and $\ell_{\max} = \max [\lambda_{\max}(\Pi_0), \lambda_{\max}(R_w)]$.

Theorem 6 shows that greedy sensor selection is also nearoptimal for state smoothing. The suboptimality guarantees, however, vary in strength depending on the parameters of the dynamical system. Nevertheless, as with the filtering case from Section V-A, the certificates improve in typical practical settings. For simplicity, let us again analyze the case where $\Pi_0 = \mathbf{R}_w = \sigma_w^2 \mathbf{I}$, $\mathbf{R}_v = \sigma_v^2 \mathbf{I}$, and $\mathbf{H}_j = \mathbf{e}_j^T$, for $j \in \mathcal{O} = \{1, \dots, p\}$. Immediately, $\ell_{\text{max}} = \sigma_w^2$, $\mathbf{C} = \sigma_w^2 \mathbf{I}$, and (40) simplifies to

$$\tilde{\alpha}_k = \frac{1}{1 + \sigma_w^2 / \sigma_v^2 \,\lambda_{\max}\left(\boldsymbol{\Phi}_{m+k}^T \boldsymbol{\Phi}_{m+k}\right)}.\tag{42}$$

Notice that a noise ratio appears explicitly in the denominator of (42). Hence, as for filtering, it is ready that $\tilde{\alpha} \approx 1$



Figure 5. ϵ -supermodularity bounds (Theorem 6) for different process noise variances and $||F||_2$.

when $\sigma_v^2 \gg \sigma_w^2$. Moreover, $\tilde{\alpha}_j$ depends on the system dynamics through the matrix $\boldsymbol{\Phi}$. In fact, this dependence can be made explicit by bounding the norm of $\boldsymbol{\Phi}$ in terms of its blocks. Explicitly, $\lambda_{\max} \left(\boldsymbol{\Phi}_{m+k}^T \boldsymbol{\Phi}_{m+k} \right) = \| \boldsymbol{\Phi}_{m+k} \| \leq \sqrt{\sum_{j=0}^{m+k} \| \boldsymbol{F} \|^{2j}}$. Hence, stronger guarantees are obtained when $\sigma_v^2 \gg \| \boldsymbol{F} \|$. Once again, similar arguments hold for (41), which reduces to

$$\tilde{\epsilon}_k = \sigma_w^4 / \sigma_v^2 \,\lambda_{\max}\left[\Phi_{m+k}^T \Phi_{m+k} \right]. \tag{43}$$

We illustrate the values of $\tilde{\alpha}$ and $\tilde{\epsilon}$ for different values of σ_w^2/σ_v^2 and $\|\mathbf{F}\|$ in Figures 4 and 5 (see Section VI for details). As (42) and (43) suggest, Theorem 6 provides better guarantees when $\|\mathbf{F}\| \sigma_w^2/\sigma_v^2 \ll 1$. In other words, when the measurement noise dominates over both the process noise and the decay rate of the system modes, a scenario of practical value in KF applications [31], [49]. Similar to the filtering case, Figure 4 shows that if we violate the cardinality constraints by 35%, we can obtain (1 - 1/e)-optimality for $\|\mathbf{F}\| = 0.3$ and $\sigma_v^2 > 75\sigma_w^2$. As in the filtering case, good performance guarantees for problem (PII) are obtained for a stricter range of parameters (Figure 5).

VI. NUMERICAL EXAMPLES

We begin by giving details on the simulations in Figures 2– 5. We considered dynamical systems with n = p = 100in which the elements of F were drawn randomly from a standard Gaussian distribution, $H_u = e_u^T$, where $e_u \in \mathbb{R}^p$ is a vector of zeros except for the *u*-th element, which is one, $\Pi_0 = 10^{-2}I$, $R_w = \sigma_w^2 I$, and $R_v = I$. The state transition matrix F was normalized so that its spectral norm is in [0.1, 0.9] and the process noise σ_w^2 varied in [0.01, 100]. For Figures 2 and 3 (filtering), we minimize the 10-steps average error [N = 10 and $\theta_k = 1$ in (P^{*})]. For Figures 4 and 5 (smoothing), we minimize the error at the end of the window [$N = 10, \theta_9 = 1$, and $\theta_k = 0$ for k < 9 in (P^{*})].

Figures 2–5, along with Theorems 5 and 6, illustrate the wide range of parameters over which good performance certificates can be provided for greedy sensor selection. It is worth



Figure 6. Relative suboptimality of greedy sensor selection using the MSE [(PI)] for filtering (1000 system realizations).



Figure 7. Relative suboptimality of greedy sensor selection using the MSE [(PI)] for smoothing (1000 system realizations).

noting, however, that these are worst-case guarantees and that better results are common in practice. To illustrate this point, we evaluate the relative suboptimality of greedily selected sensing sets for both filtering and smoothing over 1000 system realizations. Explicitly, we evaluate

$$\nu^{\star}(\mathcal{G}) = \frac{f(\mathcal{X}^{\star}) - f(\mathcal{G})}{f(\mathcal{X}^{\star})},$$

where \mathcal{G} and \mathcal{X}^* are the greedy and optimal solutions of (P^{*}), respectively. In all experiments, we take r = s, i.e., we do not consider violations of the cardinality constraint. Since ν^* depends on the optimal sensing set, we restrict ourselves to small dynamical systems [n = 10 states, p = 10 outputs,and s = 4 in P^{*}]. However, we now randomly draw the elements of both F and H from standard Gaussian distributions and normalize F so that its norm is 0.9. We keep $\sigma_w^2 = 10^{-2}$ fixed, but draw the measurement noise variances at each output uniformly at random from $[10^{-2}, 1]$.



Figure 8. Relative suboptimality of greedy sensor selection using the spectral norm [(PII)] for filtering (1000 system realizations).



Figure 9. Relative suboptimality of greedy sensor selection using the spectral norm [(PII)] for smoothing (1000 system realizations).

In Figures 6 and 7, we show the value of ν^* for the result of greedily solving problem (PI) for filtering and smoothing. Notice that the values obtained are considerably lower than the guarantee given in Theorems 5 and 6. In fact, they are lower than the $1 - e^{-1} \approx 0.63$ bound for supermodular functions. For both problem, greedy sensor selection found the optimal sampling set in over 96% of the realizations.

In Figures 8 and 9, we show results for problem (PII). To make sure the system is observable over the time window selected, we use a smaller system $[n = 5 \text{ states}, p = 10 \text{ outputs}, \text{ and } s = 5 \text{ in } (P^*)]$. Moreover, since the guarantees for the greedy solution of (PII) are additive and therefore not as strong as those for (PI), we take $\sigma_w^2 = 10^{-3}$. For the filtering problem, the results are similar to those observed in Figure 6 for minimizing the MSE. In over 98% of the realizations, greedy sensor selection actually obtained the optimal sensing set. On the other hand, the smoothing problem is more challenging for the worst-case error and the results obtained are poorer. This is especially due to the fact that $f(\tilde{S}^*) \ll s \cdot \tilde{\epsilon}$



Figure 10. Amazon basin, sensor pool (light circles), and greedy sensing set (dark circles).

in (39) for this example, so that the guarantees we obtain are less strict. Still, greedy sensor selection found the optimal sensing set in 77% of the realizations.

A. Application: sensing for the Amazon basin

We conclude these numerical examples by illustrating the use of greedy sensor selection for water monitoring in the Amazon basin. The Amazon drainage basin, showed in green in Figure 10, covers 7.5 million km^2 and is composed of over 7000 tributaries. In our experiments, we use a simplified network of the basin (blue curve in Figure 10) obtained by smoothing the original map [50].

Using this river network, we construct a weighted directed tree $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. The set of vertices \mathcal{V} is composed of the possible sensor positions shown in Figure 10 and an additional node midway between each sensor ($|\mathcal{V}| = 127$). The set of edges \mathcal{E} is composed of ordered pairs from $\mathcal{V} \times \mathcal{V}$ such that $(i, j) \in \mathcal{E}$ if water flows from node i to node j. We assume all river flow toward the ocean. The adjacency matrix of \mathcal{G} is then given by

$$[\mathbf{A}]_{ij} = \begin{cases} \exp\left(\left\|\mathbf{z}_i - \mathbf{z}_j\right\|^2 / \sigma^2\right), & \text{if } (i, j) \in \mathcal{E} \\ 0, & \text{otherwise} \end{cases}$$
(44)

where $z_i \in \mathbb{R}^2$ is the position of node *i* on the map and $\sigma^2 = 10$ is a smoothness parameter. We also define its Laplacian as L = D - A, where D is a diagonal matrix whose elements are the sum of the columns of A, and its symmetrized Laplacian as $L' = D' - (A + A^T)$, where D' is a diagonal matrix whose elements are the sum of the columns of $A + A^T$.

Using A from (44) and the Laplacians defined above, we construct a dynamical system whose states x_k evolve according to (1) with

$$\boldsymbol{F} = 0.9 \exp(-\boldsymbol{L}\Delta t) + 0.099 \exp(-\boldsymbol{L}'\Delta t), \quad (45)$$

where $\Delta t = 0.1$ is the sampling period. The dynamics in (45) are a combination of two processes: the first term corresponds to the advection process by which water flows to the ocean



Figure 11. Estimated unobserved states along the Amazon river. States positions are indicated in Figure 10.

and the second term corresponds to a diffusion process. The combination coefficients are chosen so that the system is marginally stable (||F|| < 1.02). The process noise w_k is drawn from a zero-mean Gaussian random variable with covariances $R_w = \sigma_w^2 I$, for $\sigma_w^2 = 10^{-4}$. The true initial state of the system is a spike at the beginning of the Amazon river (the red curve in Figure 10), i.e., x_0 is a vector of zeros except at position 51, where it is 10. For the KF, however, we assume the initial state is a zero-mean Gaussian random variable with covariance $\Pi_0 = I$.

We greedily select s = 10 sensors (dark circles in Figure 10) from a pool of p = 64 sensors whose positions are shown in Figure 10. The sensors observe a noisy version of their respective states, i.e., $y_{u,k} = e_u^T x_k + v_{u,k}$ for $u \in \mathcal{O}$, where e_u is a vector of zeros except for the *u*-th element that is one and $v_{u,k}$ is a zero-mean Gaussian random variable independent of $v_{w,\ell}$ for all $w \neq u$ or $\ell \neq k$ with variance $\sigma_v^2 = 10^{-1}$.

The state trajectories estimated using Kalman filtering are shown in Figure 11. We show only the values at 7 unobserved portions along the Amazon river (indicated in Figure 10). We compare the performance of a greedy sensing set selected to optimize the average 200-steps MSE and a sensing set containing one randomly selected sensor for each of the first ten levels of the Amazon basin tree. Notice from Figure 11 that by portion C, the results using the greedy sensing set track closely the true states of the system, whereas the random sensing only detects the change in the water on the second half of the river (position E). In terms of error, the average 200steps MSE for a full sensing solution ($\mathcal{X} = \mathcal{O}$) is 185.67. The greedy sensing set, which contains only 15% of the possible sensors, achieves a cost of 1005.03, an almost 50% reduction over the 1978.22 obtained by the random set. Since the modes of this system have similar decays (the condition number of Fis 1.3), the sensing set obtain by greedily solving (PIII) differs in a single sensor from that obtained using (PI). Thus, they achieve similar MSE performances. As discussed in Remark 1, this is not necessarily the case in general.

VII. CONCLUSION

This work studied the filtering/smoothing sensor selection problems for both the MSE and the worst-case error and provided near-optimal guarantees to their greedy solutions. To do so, it relied on two concepts of approximate supermodularity, α - and ϵ -supermodularity, and derived approximation bounds for the greedy minimization of these functions. By bounding α and ϵ for the MSE and the spectral norm of the error covariance matrix, it showed that filtering/smoothing sensor selection problem can be solved near-optimally. In fact, in typical application scenarios, the guarantees approach the 1-1/e certificate of supermodular functions. These results justify the use of greedy sensor selection without the need for surrogate cost functions, such as the log det. We expect that other cost functions have similar approximate supermodular behavior and that this theory can be used to give approximation certificates for other control problems. Of particular interest is a formulation that takes the controller inputs into account when designing the sensing set. The underlying theory developed here could also potentially be applied to problems involving joint state/input estimation and/or input and state constraints.

REFERENCES

- L.F.O. Chamon, G.J. Pappas, and A. Ribeiro, "The mean square error in Kalman filtering sensor selection is approximately supermodular," in *Conf. on Decision and Contr.*, 2017, pp. 343–350.
- [2] Y. Mo, R. Ambrosino, and B. Sinopoli, "Sensor selection strategies for state estimation in energy constrained wireless sensor networks," *Automatica*, vol. 47[7], pp. 1330–1338, 2011.
- [3] H. Rowaihy, S. Eswaran, M. Johnson, D. Verma, A. Bar-Noy, T. Brown, and T. La Porta, "A survey of sensor selection schemes in wireless sensor networks," *Proc. of the Int. Soc. for Optics and Photonics*, vol. 6562, 2007.
- [4] F. Zhao, J. Shin, and J. Reich, "Information-driven dynamic sensor collaboration," *IEEE Signal Process. Mag.*, vol. 19[2], pp. 61–72, 2002.
- [5] H. Zhang, J.M.F. Moura, and B. Krogh, "Dynamic field estimation using wireless sensor networks: Tradeoffs between estimation error and communication cost," *IEEE Trans. Signal Process.*, vol. 57[6], pp. 2383– 2395, 2009.
- [6] A.O. Hero and D. Cochran, "Sensor management: Past, present, and future," *IEEE Sensors J.*, vol. 11[12], pp. 3064–3075, 2011.
- [7] L. Chen, N. Li, L. Jiang, and S.H. Low, *Optimal demand response: Problem formulation and deterministic case*, pp. 63–85, Springer New York, 2012.
- [8] Y.-Y. Liu, J.-J. Slotine, and A.-L. Barabási, "Observability of complex systems," *Proc. of the National Academy of Sciences*, vol. 110[7], pp. 2460–2465, 2012.
- [9] G. Thatte and U. Mitra, "Spatio-temporal sensor management for environmental field estimation," *Signal Process.*, vol. 128, pp. 369–381, 2016.
- [10] A. Krause, A. Singh, and C. Guestrin, "Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies," J. Mach. Learning Research, vol. 9, pp. 235–284, 2008.
- [11] A. Das and D. Kempe, "Submodular meets spectral: Greedy algorithms for subset selection, sparse approximation and dictionary selection," in *Int. Conf. on Mach. Learning*, 2011, pp. 1057–1064.
- [12] J. Ranieri, A. Chebira, and M. Vetterli, "Near-optimal sensor placement for linear inverse problems," *IEEE Trans. Signal Process.*, vol. 62[5], pp. 1135–1146, 2014.
- [13] V. Tzoumas, M.A. Rahimian, G.J. Pappas, and A. Jadbabaie, "Minimal actuator placement with bounds on control effort," *IEEE Trans. Contr. Netw. Syst.*, vol. 3[1], pp. 67–78, 2016.
- [14] A. Olshevsky, "On (non)supermodularity of average control energy," *IEEE Trans. Contr. Netw. Syst.*, vol. 5[3], pp. 1177–1181, 2018.
- [15] H. Zhang, R. Ayoub, and S. Sundaram, "Sensor selection for Kalman filtering of linear dynamical systems: Complexity, limitations and greedy algorithms," *Automatica*, vol. 78, pp. 202–210, 2017.

- [16] L. Ye, S. Roy, and S. Sundaram, "On the complexity and approximability of optimal sensor selection for Kalman filtering," in *American Contr. Conf.*, 2018, pp. 5049–5054.
- [17] T.H. Summers, F.L. Cortesi, and J. Lygeros, "On submodularity and controllability in complex dynamical networks," *IEEE Trans. Contr. Netw. Syst.*, vol. 3[1], pp. 91–101, 2016.
- [18] C. Commault and J.-M. Dion, "Input addition and leader selection for the controllability of graph-based systems," *Automatica*, vol. 49[11], pp. 3322–3328, 2013.
- [19] S. Pequito, S. Kar, and A. P. Aguiar, "On the complexity of the constrained input selection problem for structural linear systems," *Automatica*, vol. 62, pp. 193–199, 2015.
- [20] S. Pequito, S. Kar, and A.P. Aguiar, "A framework for structural input/output and control configuration selection in large-scale systems," *IEEE Trans. Autom. Control*, vol. 61[2], pp. 303–318, 2016.
- [21] J.E. Weimer, B. Sinopoli, and B.H. Krogh, "A relaxation approach to dynamic sensor selection in large-scale wireless networks," in *Int. Conf.* on Distributed Computing Syst., 2008, pp. 501–506.
- [22] S. Joshi and S. Boyd, "Sensor selection via convex optimization," *IEEE Trans. Signal Process.*, vol. 57[2], pp. 451–462, 2009.
- [23] N. Moshtagh, L. Chen, and R. Mehra, "Optimal measurement selection for any-time Kalman filtering with processing constraints," in *Conf. on Decision and Contr.*, 2009, pp. 5074–5079.
- [24] Y. Wang, M. Sznaier, and F. Dabbene, "A convex optimization approach to worst-case optimal sensor selection," in *Conf. on Decision and Contr.*, 2013, pp. 6353–6358.
- [25] N.K. Dhingra, M.R. Jovanović, and Z.Q. Luo, "An ADMM algorithm for optimal sensor and actuator selection," in *Conf. on Decision and Contr.*, 2014, pp. 4039–4044.
- [26] S. Liu, S.P. Chepuri, M. Fardad, E. Maşazade, G. Leus, and P.K. Varshney, "Sensor selection for estimation with correlated measurement noise," *IEEE Trans. Signal Process.*, vol. 64[13], pp. 3509–3522, 2016.
- [27] M. Shamaiah, S. Banerjee, and H. Vikalo, "Greedy sensor selection: Leveraging submodularity," in *Conf. on Decision and Contr.*, 2010, pp. 2572–2577.
- [28] V. Tzoumas, A. Jadbabaie, and G.J. Pappas, "Sensor placement for optimal Kalman filtering: Fundamental limits, submodularity, and algorithms," in *American Contr. Conf.*, 2016, pp. 191–196.
- [29] G.E.P. Box, J.S. Hunter, and W.G. Hunter, *Statistics for Experimenters: Design, Innovation, and Discovery*, Springer, 2005.
- [30] G.L. Nemhauser, L.A. Wolsey, and M.L. Fisher, "An analysis of approximations for maximizing submodular set functions—I," *Mathematical Programming*, vol. 14[1], pp. 265–294, 1978.
- [31] T. Kailath, A.H. Sayed, and B. Hassibi, *Linear estimation*, Prentice-Hall, 2000.
- [32] S. Boyd and L. Vandenberghe, *Convex optimization*, Cambridge University Press, 2004.
- [33] A. Krause and D. Golovin, "Submodular function maximization," in *Tractability: Practical Approaches to Hard Problems*. Cambridge University Press, 2014.
- [34] F. Bach, "Learning with submodular functions: A convex optimization perspective," *Foundations and Trends in Machine Learning*, vol. 6[2-3], pp. 145–373, 2013.
- [35] A. Das and D. Kempe, "Algorithms for subset selection in linear regression," in ACM Symp. on Theory of Comput., 2008, pp. 45–54.
- [36] P. Singh, M. Chen, L. Carlone, S. Karaman, E. Frazzoli, and D. Hsu, "Supermodular mean squared error minimization for sensor scheduling in optimal Kalman filtering," in *American Contr. Conf.*, 2017, pp. 5787– 5794.
- [37] H. Avron and C. Boutsidis, "Faster subset selection for matrices and applications," *SIAM Journal on Matrix Analysis and Applications*, vol. 34[4], pp. 1464–1499, 2013.
- [38] L.F.O. Chamon and A. Ribeiro, "Approximate supermodularity bounds for experimental design," in *Advances in Neural Information Processing Systems*, 2017, pp. 5403–5412.
- [39] T. Horel and Y. Singer, "Maximization of approximately submodular functions," in Advances in Neural Information Processing Systems, 2016, pp. 3045–3053.
- [40] B. Lehmann, D. Lehmann, and N. Nisan, "Combinatorial auctions with decreasing marginal utilities," *Games and Economic Behavior*, vol. 55[2], pp. 270–296, 2006.
- [41] M. Sviridenko, J. Vondrák, and J. Ward, "Optimal approximation for submodular and supermodular optimization with bounded curvature," in *SIAM Symposium on Discrete Algorithms*, 2014, pp. 1134–1148.
- [42] L.F.O. Chamon and A. Ribeiro, "Near-optimality of greedy set selection in the sampling of graph signals," in *Global Conf. on Signal and Inform. Process.*, 2016, pp. 1265–1269.

- [43] A. Bian, J.M. Buhmann, A. Krause, and S. Tschiatschek, "Guarantees for greedy maximization of non-submodular functions with applications," in *Int. Conf. on Mach. Learning*, 2017, pp. 498–507.
- [44] T. Summers and M. Kamgarpour, "Performance guarantees for greedy maximization of non-submodular set functions in systems and control," in *European Contr. Conf.*, 2019, pp. 2796–2801.
- [45] A. Krause and V. Cevher, "Submodular dictionary selection for sparse representation," in *Int. Conf. on Mach. Learning*, 2010, pp. 567–574.
- [46] R.A. Horn and C.R. Johnson, *Matrix analysis*, Cambridge University Press, 2013.
- [47] R. Bhatia, Matrix analysis, Springer, 1997.
- [48] L. Lovász, "Submodular functions and convexity," in *Mathematical Programming: The State of the Art.* Springer-Verlag, 1982.
- [49] M.S. Grewal and A.P. Andrews, Kalman filtering: Theory and practice using MATLAB, Wiley, 2014.
- [50] F. Muller, F. Seyler, and J.-L. Guyot, "Utilisation d'imagerie radar (ROS) JERS-1 pour l'obtention de réseaux de drainage. Exemple du Rio Negro (Amazonie)," in *Hydrological and Geochemical Processes* in Large Scale River Basins, 1999, http://www.ore-hybam.org/index. php/eng/Data/Cartography/Amazon-basin-hydrography.
- [51] H.V. Henderson and S.R. Searle, "On deriving the inverse of a sum of matrices," *SIAM Review*, vol. 23[1], pp. 53–60, 1981.
- [52] B. Wang and F. Zhang, "Some inequalities for the eigenvalues of the product of positive semidefinite Hermitian matrices," *Linear Algebra and its Applications*, vol. 160, pp. 113–118, 1992.

APPENDIX

PROOF OF PROPOSITION 1

Proof. The filtered estimate $\hat{x}_{k}(\mathcal{X})$ can be obtained from a previous estimate \hat{x}_{k-1} using the KF iteration

$$\hat{\boldsymbol{x}}_{k}(\boldsymbol{\mathcal{X}}) = \boldsymbol{F}\hat{\boldsymbol{x}}_{k-1} + \sum_{u \in \boldsymbol{\mathcal{X}}} \boldsymbol{K}_{u,k} \left[\boldsymbol{y}_{u,k} - \boldsymbol{H}_{u} \boldsymbol{F} \hat{\boldsymbol{x}}_{k-1} \right], \quad (46)$$

where $K_{u,k} = P_{k|k-1}H_u^T E_{u,k}^{-1}$ is the Kalman gain of measurement $u \in \mathcal{O}$, with $E_{u,k} = H_u P_{k|k-1}H_u^T + R_{v,u}$ denoting the innovation covariance matrix and

$$\boldsymbol{P}_{k|k-1} = \boldsymbol{F}\boldsymbol{P}_{k-1}\boldsymbol{F}^T + \boldsymbol{R}_w \tag{47}$$

denoting the *a priori* error covariance matrix. Then, $P_k(\mathcal{X})$ is obtained directly from the KF recursion as

$$\boldsymbol{P}_{k}(\boldsymbol{\mathcal{X}}) = \left[\boldsymbol{P}_{k|k-1}^{-1} + \sum_{u \in \boldsymbol{\mathcal{X}}} \boldsymbol{H}_{u}^{T} \boldsymbol{R}_{v,u}^{-1} \boldsymbol{H}_{u}\right]^{-1}, \quad (48)$$

where $P_{k-1} = \mathbb{E} \left[\hat{x}_{k-1} - x_{k-1} \right] \left[\hat{x}_{k-1} - x_{k-1} \right]^T$ is the error covariance matrix of the state estimate and $P_{0|-1} = \Pi_0$ [31].

PROOF OF PROPOSITION 2

Proof. Instead of relying on a Kalman smoother, we obtain the error covariance matrix by analyzing this problem in the batch setting. Note that both methods solve the same problem and therefore lead to the same solution. Proceeding as in [28], note from (1) that estimating $\bar{\boldsymbol{x}}_k$, i.e., all states \boldsymbol{x}_j for $j \leq k$, is equivalent to estimating the first state \boldsymbol{x}_0 and the process noises $\{\boldsymbol{w}_j\}_{j\leq k-1}$. Indeed, we can recover the smoothed states using $\tilde{\boldsymbol{x}}_k = \boldsymbol{\Phi}_k \bar{\boldsymbol{z}}_k$, for $\bar{\boldsymbol{z}}_k = \begin{bmatrix} \boldsymbol{x}_0^T & \boldsymbol{w}_0^T & \cdots & \boldsymbol{w}_{k-1}^T \end{bmatrix}^T$ and $\boldsymbol{\Phi}_k$ as in (7).

The estimation of \bar{z}_k can be cast as a stochastic estimation problem by lifting. Formally, using the stacked $\bar{y}_{n,k}$ = $\begin{bmatrix} \boldsymbol{y}_{u,0}^T & \cdots & \boldsymbol{y}_{u,k}^T \end{bmatrix}^T$ and $\bar{\boldsymbol{v}}_{u,k} = \begin{bmatrix} \boldsymbol{v}_{u,0}^T & \cdots & \boldsymbol{v}_{u,k}^T \end{bmatrix}^T$, it is straightforward that

$$\begin{bmatrix} \bar{\boldsymbol{y}}_{u_1,k} \\ \vdots \\ \bar{\boldsymbol{y}}_{u_s,k} \end{bmatrix} = \underbrace{\begin{bmatrix} \boldsymbol{I} \otimes \boldsymbol{H}_{u_1} \\ \vdots \\ \boldsymbol{I} \otimes \boldsymbol{H}_{u_s} \end{bmatrix}}_{\boldsymbol{O}_k(\boldsymbol{\mathcal{X}})} \boldsymbol{\Phi}_k \, \bar{\boldsymbol{z}}_k + \begin{bmatrix} \bar{\boldsymbol{v}}_{u_1,k} \\ \vdots \\ \bar{\boldsymbol{v}}_{u_s,k} \end{bmatrix}, \quad (49)$$

where $\mathcal{X} = \{u_1, \ldots, u_s\}$ is the observed output subset and \otimes denotes the Kronecker product. The minimum MSE incurred from estimating \bar{z} from \bar{y} is then given by the trace of

$$\boldsymbol{K}(\mathcal{X}) = \left[\boldsymbol{C}^{-1} + \boldsymbol{O}_{k}(\mathcal{X})^{T} \text{ blkdiag} \left(\boldsymbol{I} \otimes \boldsymbol{R}_{v,u}\right)^{-1} \boldsymbol{O}_{k}(\mathcal{X})\right]^{-1}$$
(50)

for $u \in \mathcal{X}$, O_k in (49), $C = \text{blkdiag}(\Pi_0, I \otimes R_w)$, and blkdiag(X, Y) denoting the block diagonal matrix whose diagonal blocks are X and Y [31].

To obtain the form in (4), use the fact that $blkdiag(\mathbf{A}_i)^{-1} = blkdiag(\mathbf{A}_i^{-1})$ and the mixed product property of the Kronecker product [46] to get

$$\boldsymbol{O}_{k}(\mathcal{X})^{T} \text{ blkdiag} \left(\boldsymbol{I} \otimes \boldsymbol{R}_{v,u}\right)^{-1} \boldsymbol{O}_{k}(\mathcal{X}) = \sum_{u \in \mathcal{X}} \boldsymbol{\Phi}_{k}^{T} \left(\boldsymbol{I} \otimes \boldsymbol{H}_{u}^{T} \boldsymbol{R}_{v,u}^{-1} \boldsymbol{H}_{u}\right) \boldsymbol{\Phi}_{k}.$$
(51)

Substituting (51) in (50) yields (6).

PROOF OF THEOREM 1

Proof. Since f is monotone decreasing, it holds for every set \mathcal{G}_j that

$$f(\mathcal{X}^{\star}) \ge f(\mathcal{X}^{\star} \cup \mathcal{G}_j)$$

= $f(\mathcal{G}_j) - \sum_{i=1}^{s} \left[f(\mathcal{T}_{i-1}) - f(\mathcal{T}_{i-1} \cup v_i^{\star}) \right],$ (52)

where $\mathcal{T}_0 = \mathcal{G}_j$, $\mathcal{T}_i = \mathcal{G}_j \cup \{v_1^*, \ldots, v_i^*\}$, $i = 1, \ldots, s$, and v_i^* is the *i*-th element of \mathcal{X}^* . Notice that this holds regardless of the order in which the v_i^* are taken. Since f is α -supermodular and $\mathcal{G}_j \subseteq \mathcal{T}_i$ for all *i*, the incremental gains in (52) can be bounded using (9) to get

$$f(\mathcal{X}^{\star}) \ge f(\mathcal{G}_j) - \alpha^{-1} \sum_{i=1}^{s} \left[f(\mathcal{G}_j) - f(\mathcal{G}_j \cup v_i^{\star}) \right].$$
(53)

To proceed, we use the fact that $\mathcal{G}_{j+1} = \mathcal{G}_j \cup \{u\}$ is constructed by the greedy procedure in (10) so as to maximize the incremental gains in (53). It therefore holds that

$$f(\mathcal{X}^{\star}) \ge f(\mathcal{G}_j) - \alpha^{-1} s \left[f(\mathcal{G}_j) - f(\mathcal{G}_{j+1}) \right].$$
(54)

By adding and subtracting $f(\mathcal{X}^*)$ in the brackets of (54), (54) yields a recursion for the distance to the optimal value $\delta_j = f(\mathcal{X}^*) - f(\mathcal{G}_j)$ given by

$$\delta_j \ge \alpha^{-1} s \left[\delta_j - \delta_{j+1} \right] \Rightarrow \delta_{j+1} \ge \left(1 - \frac{1}{\alpha^{-1} s} \right) \delta_j.$$

Note that since f is non-positive, $\delta_j \leq 0$ for all j by the optimality of \mathcal{X}^* .

The expression in (14) yields directly from this recursion by noticing that since f is normalized, we have that $\delta_0 =$ $f(\mathcal{X}^{\star}) - f(\emptyset) = f(\mathcal{X}^{\star})$. Immediately, since Algorithm 1 is used for r iterations, we obtain

$$f(\mathcal{X}^{\star}) - f(\mathcal{G}_r) \ge \left(1 - \frac{\alpha}{s}\right)^r f(\mathcal{X}^{\star}).$$
(55)

Using the fact that $1 - x \le e^{-x}$, the expression in (55) can be rearranged into (14).

PROOF OF THEOREM 2

Proof. Using the assumption that f is monotone decreasing, we write

$$f(\mathcal{X}^{\star}) \ge f(\mathcal{X}^{\star} \cup \mathcal{G}_j)$$

= $f(\mathcal{G}_j) - \sum_{i=1}^{s} \left[f(\mathcal{T}_{i-1}) - f(\mathcal{T}_{i-1} \cup \{v_i^{\star}\}) \right],$ (56)

where again $\mathcal{T}_0 = \mathcal{G}_j$, $\mathcal{T}_i = \mathcal{G}_j \cup \{v_1^{\star}, \dots, v_i^{\star}\}$, for $i = 1, \dots, s$, and v_i^{\star} is the *i*-th element of \mathcal{X}^{\star} . Observe that the order in which the v_i^{\star} occur in the telescopic sum is irrelevant. Since fis ϵ -supermodular and $\mathcal{G}_j \subseteq \mathcal{T}_i$ for all i, (15) can be used to bound the incremental gains in (56). Explicitly,

$$f(\mathcal{X}^{\star}) \ge f(\mathcal{G}_j) - \sum_{i=1}^{s} \left[f(\mathcal{G}_j) - f(\mathcal{G}_j \cup \{v_i^{\star}\}) + \epsilon \right].$$

Since $\mathcal{G}_{j+1} = \mathcal{G}_j \cup \{u\}$ is constructed in (10) so as to minimize $f(\mathcal{G}_{j+1})$, it holds that

$$f(\mathcal{X}^{\star}) \ge f(\mathcal{G}_j) - s\left[f(\mathcal{G}_j) - f(\mathcal{G}_{j+1}) + \epsilon\right].$$
(57)

Finally, a recursion for the optimality gap $\delta'_j = f(\mathcal{G}_j) - f(\mathcal{X}^*)$ is obtained from (57) by adding and subtracting $f(\mathcal{X}^*)$ in the brakets. Explicitly,

$$\delta'_{j} \leq s \left(\delta'_{j} - \delta'_{j+1} + \epsilon\right) \Rightarrow \delta'_{j+1} \leq \left(1 - \frac{1}{s}\right)\delta'_{j} + \epsilon.$$
 (58)

Notice that since f is non-positive, $\delta'_j \ge 0$ for all j due to the optimality of \mathcal{X}^* . The solution of (58) after r steps is

$$\delta'_r \le \left(1 - \frac{1}{s}\right)^r \delta'_0 + \epsilon \sum_{j=0}^{r-1} \left(1 - \frac{1}{s}\right)^j$$
 (59)

Evaluating the geometric series in (59) yields

$$\delta'_r \le \left(1 - \frac{1}{s}\right)^r \delta'_0 + s \left[1 - \left(1 - \frac{1}{s}\right)^r\right] \epsilon$$

and under the assumption that f is normalized, i.e., for $\delta'_0 = f(\emptyset) - f(\mathcal{X}^*) = -f(\mathcal{X}^*) \ge 0$, we obtain

$$f(\mathcal{G}_r) \leq \left[1 - \left(1 - \frac{1}{s}\right)^r\right] \left[f(\mathcal{X}^\star) + s \cdot \epsilon\right].$$

Using once again the fact that $1 - x \le e^{-x}$ yields (17).

PROOF OF LEMMA 1

Proof. We proceed with the proof case-by-case:

 (i) From the definition of α-supermodularity in (12), for *A* ⊆ *B* ⊆ *O* and *u* ∈ *O* \ *B* it holds that

$$g(\mathcal{A}) - g(\mathcal{A} \cup \{u\}) = \sum_{i} \theta_{i} \left[f_{i}(\mathcal{A}) - f_{i}(\mathcal{A} \cup \{u\})\right]$$
$$\geq \sum_{i} \alpha_{i} \theta_{i} \left[f_{i}(\mathcal{B}) - f_{i}(\mathcal{B} \cup \{u\})\right],$$

where the constant factors of the affine transformations cancel out. Since $\alpha_i \ge \min(\alpha_i)$, we obtain

$$g(\mathcal{A}) - g(\mathcal{A} \cup \{u\}) \ge \min(\alpha_i) \sum_i \theta_i \left[f_i(\mathcal{B}) - f_i(\mathcal{B} \cup \{u\}) \right] \ge \min(\alpha_i) \left[g(\mathcal{B}) - g(\mathcal{B} \cup \{u\}) \right].$$

(ii) Similarly, using the definition of ϵ -supermodularity in (15), for $\mathcal{A} \subseteq \mathcal{B} \subseteq \mathcal{V}$ and $u \in \mathcal{O} \setminus \mathcal{B}$ we obtain

$$g\left(\mathcal{A}\right) - g\left(\mathcal{A} \cup \{u\}\right) \geq \sum_{i} \theta_{i} \left[f_{i}\left(\mathcal{B}\right) - f_{i}\left(\mathcal{B} \cup \{u\}\right)\right] - \sum_{i} \theta_{i} \epsilon_{i} = g\left(\mathcal{B}\right) - g\left(\mathcal{B} \cup \{u\}\right) - \sum_{i} \theta_{i} \epsilon_{i},$$

where again the constant factors b canceled out.

Proof of Lemma 3

Proof. Start by simplifying $\Delta_u t(\mathcal{X}) = t(\mathcal{X}) - t(\mathcal{X} \cup \{u\})$ using the matrix inversion lemma. To do so, write

$$\boldsymbol{Y}\left(\mathcal{X}\cup\{u\}\right) = \left[\boldsymbol{R}(\mathcal{X}) + \boldsymbol{M}_{u}\right]^{-1}, \quad (60)$$

letting once again $R(\mathcal{X}) = M_{\emptyset} + \sum_{u \in \mathcal{X}} M_u$. Since $R(\mathcal{X}) \succ 0$, but the matrices M_u need not be invertible, we use an alternative form of the matrix inversion lemma [51] to get

$$\boldsymbol{Y}\left(\mathcal{X}\cup\{u\}\right) = \boldsymbol{Y}(\mathcal{X}) - \boldsymbol{Y}(\mathcal{X})\boldsymbol{M}_{u}\boldsymbol{Y}\left(\mathcal{X}\cup\{u\}\right).$$
(61)

where we used (60) to write $Y(\mathcal{X}) = R(\mathcal{X})^{-1}$. Using (61) and the linearity of the trace [46], we can write $\Delta_u t(\mathcal{X})$ as

$$\Delta_u t(\mathcal{X}) = \operatorname{Tr} \left[\mathbf{Y}(\mathcal{X}) \mathbf{M}_u \mathbf{Y} \left(\mathcal{X} \cup \{u\} \right) \right].$$
 (62)

To proceed, let $\tilde{M}_u = M_u + \epsilon I \succ 0$ for $\epsilon > 0$ and define the perturbed version of (62) as

$$\widetilde{\Delta}_{u}t(\mathcal{X}) = \operatorname{Tr}\left[\boldsymbol{Y}(\mathcal{X})\tilde{\boldsymbol{M}}_{u}\left(\boldsymbol{R}(\mathcal{X}) + \tilde{\boldsymbol{M}}_{u}\right)^{-1}\right].$$
 (63)

Notice that $\widetilde{\Delta}_u t \to \Delta_u t$ as $\epsilon \to 0$. Using the invertibility of \widetilde{M}_u and $R(\mathcal{X})$, we obtain

$$\widetilde{\Delta}_{u}t(\mathcal{X}) = \operatorname{Tr}\left[\boldsymbol{Y}(\mathcal{X})\left(\boldsymbol{Y}(\mathcal{X}) + \widetilde{\boldsymbol{M}}_{u}^{-1}\right)^{-1}\boldsymbol{Y}(\mathcal{X})\right].$$

Since $Y(\mathcal{X}) \succ 0$, its square-root $Y(\mathcal{X})^{1/2}$ is well-defined and unique [46]. We can therefore use the circular commutation property of the trace to get

$$\Delta_u t(\mathcal{X}) = \operatorname{Tr} \left[\boldsymbol{Y}(\mathcal{X}) \boldsymbol{Z}(\mathcal{X}, u) \right], \tag{64}$$

with $Z(\mathcal{X}, u) = Y(\mathcal{X})^{1/2} \left[Y(\mathcal{X}) + \tilde{M}_u^{-1} \right]^{-1} Y(\mathcal{X})^{1/2}$. Since both matrices in (64) are positive definite, we can use the bound from [52] to get

$$\lambda_{\min}\left[\boldsymbol{Y}(\mathcal{X})\right] \operatorname{Tr}\left(\boldsymbol{Z}\right) \leq \widetilde{\Delta}_{u} t(\mathcal{X}) \leq \lambda_{\max}\left[\boldsymbol{Y}(\mathcal{X})\right] \operatorname{Tr}\left(\boldsymbol{Z}\right).$$

Reversing the manipulations used to obtain (64) yields

$$\lambda_{\min} \left[\boldsymbol{Y}(\mathcal{X}) \right] \operatorname{Tr} \left[\tilde{\boldsymbol{M}}_{u} \left(\boldsymbol{Y}(\mathcal{X}) + \tilde{\boldsymbol{M}}_{u} \right)^{-1} \right] \leq \tilde{\Delta}_{u} t(\mathcal{X}) \leq \lambda_{\max} \left[\boldsymbol{Y}(\mathcal{X})^{-1} \right] \operatorname{Tr} \left[\tilde{\boldsymbol{M}}_{u} \left(\boldsymbol{Y}(\mathcal{X}) + \tilde{\boldsymbol{M}}_{u} \right)^{-1} \right].$$

The result in (22) is obtained by continuity as $\epsilon \to 0$.



Luiz F. O. Chamon (S'12) received the B.Sc. and M.Sc. degree in electrical engineering from the University of São Paulo, São Paulo, Brazil, in 2011 and 2015. In 2009, he was an undergraduate exchange student at the Masters in Acoustics of the École Centrale de Lyon, Lyon, France. He is currently working toward the Ph.D. degree in electrical and systems engineering at the University of Pennsylvania (Penn), Philadelphia. In 2009, he was an Assistant Instructor and Consultant on nondestructive testing at INSACAST Formation Continue. From

2010 to 2014, he worked as a Signal Processing and Statistical Consultant on a project with EMBRAER. His research interest include signal processing, optimization, statistics, and control.



George J. Pappas (F'09) received the Ph.D. degree in electrical engineering and computer sciences from the University of California, Berkeley, CA, USA, in 1998. He is currently the Joseph Moore Professor and Chair of the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA, USA. He also holds a secondary appointment with the Department of Computer and Information Sciences and the Department of Mechanical Engineering and Applied Mechanics. He is a Member of the GRASP Lab and the PRECISE

Center. He had previously served as the Deputy Dean for Research with the School of Engineering and Applied Science. His research interests include control theory and, in particular, hybrid systems, embedded systems, cyberphysical systems, and hierarchical and distributed control systems, with applications to unmanned aerial vehicles, distributed robotics, green buildings, and biomolecular networks. Dr. Pappas has received various awards, such as the Antonio Ruberti Young Researcher Prize, the George S. Axelby Award, the Hugo Schuck Best Paper Award, the George H. Heilmeier Award, the National Science Foundation PECASE award and numerous best student papers awards at ACC, CDC, and ICCPS.



Alejandro Ribeiro received the B.Sc. degree in electrical engineering from the Universidad de la Republica Oriental del Uruguay, Montevideo, in 1998 and the M.Sc. and Ph.D. degree in electrical engineering from the Department of Electrical and Computer Engineering, the University of Minnesota, Minneapolis in 2005 and 2007.

From 1998 to 2003, he was a member of the technical staff at Bellsouth Montevideo. After his M.Sc. and Ph.D studies, in 2008 he joined the University of Pennsylvania (Penn), Philadelphia, where he is

currently the Rosenbluth Associate Professor at the Department of Electrical and Systems Engineering. His research interests are in the applications of statistical signal processing to the study of networks and networked phenomena. His focus is on structured representations of networked data structures, graph signal processing, network optimization, robot teams, and networked control.

Dr. Ribeiro received the 2014 O. Hugo Schuck best paper award, and paper awards at CDC 2017, 2016 SSP Workshop, 2016 SAM Workshop, 2015 Asilomar SSC Conference, ACC 2013, ICASSP 2006, and ICASSP 2005. His teaching has been recognized with the 2017 Lindback award for distinguished teaching and the 2012 S. Reid Warren, Jr. Award presented by Penn's undergraduate student body for outstanding teaching. Dr. Ribeiro is a Fulbright scholar class of 2003 and a Penn Fellow class of 2015.