

Strategies for Distributed Sensor Selection Using Convex Optimization

Fabian Altenbach, Steven Corroy, Georg Böcherer*, and Rudolf Mathar

Institute for Theoretical Information Technology, RWTH Aachen University, Germany

*Institute for Communications Engineering, Technische Universität München, Germany

Email: {altenbach, corroy, mathar}@ti.rwth-aachen.de, georg.boecherer@tum.de

Abstract—Consider the estimation of an unknown parameter vector in a linear measurement model. Centralized sensor selection consists in selecting a set of k_s sensor measurements, from a total number of m potential measurements. The performance of the corresponding selection is measured by the volume of an estimation error covariance matrix. In this work, we consider the problem of selecting these sensors in a distributed or decentralized fashion. In particular, we study the case of two leader nodes that perform naive decentralized selections. We demonstrate that this can degrade the performance severely. Therefore, two heuristics based on convex optimization methods are introduced, where we first allow one leader to make a selection, and then to share a modest amount of information about his selection with the remaining node. We will show that both heuristics clearly outperform the naive decentralized selection, and achieve a performance close to the centralized selection.

I. INTRODUCTION

Consider a linear model where a centralized collector estimates an n -dimensional parameter vector via an arrangement of m sensors. The sensor readings are affected by measurement noise. The noise samples are assumed to be realizations of independent identically distributed Gaussian random variables. Now suppose the collector is allowed to use k_s active sensors only, where $n \leq k_s < m$. We call such a situation a centralized sensor selection problem. The performance of a particular selection can be assessed by the volume of the estimation error covariance matrix [1, Sec. II.A]. Therefore, the objective of the sensor selection problem is to select k_s sensors such that this volume is minimized. For this purpose, the centralized collector must know the complete measurement matrix, which is needed for calculating the error covariance.

In contrast, consider the sensor arrangement as depicted in Fig. 1. We have a partition of all sensors into two groups. Each sensor group is associated with a specific leader node. The decentralized sensor selection problem consists in selecting a subset of sensors by the corresponding leader nodes individually. After that, the individual selections are transmitted to the centralized collector. The main advantage of such an approach is that we do not need to know the complete measurement matrix at one point, i.e., at the centralized collector. This can be motivated, for example, by limitations of the available

transmission bandwidth in a sensor network. However, there is no guarantee that individual selections minimize the volume of the error covariance matrix. The reason is that the decentralized leader nodes may choose jointly correlated measurements, without even knowing it. In this work, we propose two simple heuristic methods for decentralized sensor selection. Both heuristics try to avoid jointly correlated measurements by transmitting a modest amount of data between leader nodes. We will show by numerous numerical experiments that the performance can be very close to the centralized solution.

The mathematical form of the sensor selection used in this paper was introduced in [1], [2]. In particular, the authors in [1] study the sensor selection problem embedded in the framework of convex optimization. Throughout this paper, we will make extensive use of this approach. In [3], a multi-step sensor selection strategies based on the Kalman filter error covariance matrix is investigated. Other authors propose single sensor scheduling algorithms, e.g., [4], [5]. A different, but conceptually related approach is the selection of reliable sensors in the context of robust sensing [6].

The remainder of the paper is organized as follows. Section II introduces the centralized and decentralized sensor selection problem. To the best of our knowledge, the latter was not considered by others in this form. In Section III, two heuristics for decentralized sensor selection are motivated and developed. The solution of the (nonconvex) sensor selection problem is outlined in Section IV. In Section V, extensive numerical simulation illustrates the performance gains of the proposed decentralized heuristics.




Fig. 1. Arrangement of different sensors and leader nodes. Sensors (●) are associated with leader node 1, sensors (★) with leader node 2. Both leader nodes may share a (very) limited amount of information.

This work has been supported by the UMIC Research Center, RWTH Aachen University.

II. SENSOR SELECTION

A. System Model

A linear measurement model can be written as [1, Sec. II.A]

$$y_i = a_i^T x + v_i, \quad i = 1, \dots, m \quad (1)$$

where $x \in \mathbb{R}^n$ is an unknown parameter vector that we want to estimate, $y \in \mathbb{R}^m$ is the measurement vector, and $m > n$. Throughout this paper, we will use the terms sensor and measurement synonymously. The measurements are corrupted by noise v_1, \dots, v_m that is independent and identically distributed (iid) with $\mathcal{N}(0, \sigma^2)$. The measurement matrix

$$A = \begin{bmatrix} a_1^T \\ \vdots \\ a_m^T \end{bmatrix} \quad (2)$$

is assumed to have full column rank, i.e., $\text{rank}(A) = n$. The maximum-likelihood estimate of x is then given by

$$x_{\text{ML}} = \left(\sum_{i=1}^m a_i a_i^T \right)^{-1} \sum_{i=1}^m y_i a_i. \quad (3)$$

The covariance matrix of the estimation error $x - x_{\text{ML}}$ has the form

$$\Sigma = \sigma^2 (A^T A)^{-1} = \sigma^2 \left(\sum_{i=1}^m a_i a_i^T \right)^{-1}. \quad (4)$$

We measure the quality of the estimation by the volume of this matrix. It can be shown [1, Sec. II.A] that this measure is related to the log-volume of a confidence ellipsoid given by

$$\log \text{vol}(\mathcal{E}) = \text{const.} - \frac{1}{2} \log \det \left(\sum_{i=1}^m a_i a_i^T \right). \quad (5)$$

This volume is a scalar measure for how informative the measurements are or how uncertain we have to be about our estimate x_{ML} . In particular, a small volume corresponds to a small uncertainty, and vice versa.

B. Centralized Sensor Selection Problem

Now suppose we have a total number of m measurements. A central collector attempts to find a subset of $k_s < m$ measurements that minimizes the uncertainty about x_{ML} . This leads to the *centralized sensor selection problem* that is stated in [1, Sec. II.B] as

$$\begin{aligned} & \underset{z}{\text{maximize}} && f_{\text{cen}}(z) = \log \det \left(\sum_{i=1}^m z_i a_i a_i^T \right) \\ & \text{subject to} && \mathbf{1}^T z = k_s \\ & && z_i \in \{0, 1\}, \quad i = 1, \dots, m \end{aligned} \quad (6)$$

where $\mathbf{1}$ is a vector of appropriate dimension with all entries equal to one. Each Boolean variable z_i corresponds to a particular choice of a measurement. Whenever $z_i = 1$, the i th measurement is to be used. The linear constraint $\mathbf{1}^T z = k_s$ is

a budget constraint on the total number of active sensors. On occasion, we will rewrite the objective as

$$f_{\text{cen}}(z) = \log \det (A^T \text{diag}(z) A), \quad (7)$$

where the matrix $A^T \text{diag}(z) A$ is assumed to be positive definite [7, Ch. 7].

Due to the Boolean constraints in (6), the centralized sensor selection problem is a nonconvex optimization problem and is generally hard to solve. However, note that the objective is a concave function for $z_i \geq 0$ [8, Sec. 3.1.5]. Relaxing the Boolean constraints to $0 \leq z_i \leq 1$, the *relaxed centralized sensor selection problem* has the form

$$\begin{aligned} & \underset{z}{\text{maximize}} && f_{\text{cen}}(z) \\ & \text{subject to} && \mathbf{1}^T z = k_s \\ & && 0 \leq z_i \leq 1, \quad i = 1, \dots, m \end{aligned} \quad (8)$$

This problem is a convex optimization problem and hence can be solved efficiently, where the solution is denoted as z_{cen}^* . The relaxation gives us a global upper bound for the centralized sensor selection problem (6). The bound is given by

$$U_{\text{cen}} = f_{\text{cen}}(z_{\text{cen}}^*), \quad (9)$$

and is subsequently used as a global performance measure for decentralized strategies.

In order to obtain a feasible solution to (6) we apply a simple rounding scheme as suggested in [1]. In this scheme, the elements of z_{cen}^* are rearranged in descending order. After that, the k_s largest elements are set equal to 1, and the remaining elements to 0. This gives us a suboptimal solution \hat{z}_{cen} to (6), along with a lower bound $L_{\text{cen}} = f_{\text{cen}}(\hat{z}_{\text{cen}})$. This lower bound and the corresponding duality gap $\Delta_{\text{cen}} = U_{\text{cen}} - L_{\text{cen}}$ will be used later in order to compare centralized and decentralized methods. In particular, when this gap becomes sufficiently small, \hat{z}_{cen} is nearly optimal for problem (6). Note that it is possible to apply more sophisticated rounding schemes [1, Sec. III.E]. Since we focus on the comparison of centralized and decentralized strategies, we will use only the prescribed simple rounding.

C. Decentralized Sensor Selection Problem

Selecting sensors in a centralized fashion, the full measurement matrix A must be known at one point, e.g., the centralized collector. Now suppose we have two leader nodes that have access to half of the measurements $m/2 \geq n$ via the partition

$$A = \begin{bmatrix} A_1 \\ A_2 \end{bmatrix}, \quad \text{leader node 1: } A_1, \quad \text{leader node 2: } A_2. \quad (10)$$

where $A_1, A_2 \in \mathbb{R}^{m/2 \times n}$, and $\text{rank}(A_1) = \text{rank}(A_2) = n$. In the decentralized sensor selection problem considered here, both leader nodes are only allowed to select $k_s/2$ sensors each. For each leader node l with $l \in \{1, 2\}$, we first solve the

relaxed optimization problems

$$\begin{aligned} & \underset{z_l}{\text{maximize}} && f_l(z_l) = \log \det \left(\sum_{i=1}^{m/2} z_{li} a_{li} a_{li}^T \right) \\ & \text{subject to} && \mathbf{1}^T z_l = k_s/2 \\ & && 0 \leq z_{li} \leq 1, \quad i = 1, \dots, m/2 \end{aligned} \quad (11)$$

where z_l^* is the optimal solution of (11). We call this approach *naive decentralized sensor selection* or simply *decentralized sensor selection*. After computing z_l^* , both leader nodes apply the simple rounding scheme separately in order to obtain the selections \hat{z}_l . Finally, both nodes transmit their selections to the centralized collector.

D. Suboptimality and Performance

The (global) performance of any decentralized method has to be judged at the centralized collector based on the full problem, i.e., the centralized objective f_{cen} . The global upper bound is given by the expression in (9). For calculating the lower bound, we first stack the solution vectors

$$\hat{z}_{\text{dec}} = \begin{bmatrix} \hat{z}_1 \\ \hat{z}_2 \end{bmatrix}. \quad (12)$$

and insert them into the centralized objective

$$\begin{aligned} L_{\text{dec}} = f_{\text{cen}}(\hat{z}_{\text{dec}}) &= \log \det \left(\begin{bmatrix} A_1 \\ A_2 \end{bmatrix}^T \begin{bmatrix} \hat{z}_1 \\ \hat{z}_2 \end{bmatrix} \begin{bmatrix} A_1 \\ A_2 \end{bmatrix} \right) \\ &= \log \det (A^T \mathbf{diag}(\hat{z}_{\text{dec}}) A). \end{aligned} \quad (13)$$

Our measure for any decentralized strategy is then given by the suboptimality gap

$$\Delta_{\text{dec}} = U_{\text{cen}} - L_{\text{dec}}. \quad (14)$$

Furthermore, the set of feasible solutions of (11) is a subset of (8). Therefore, we can also conclude that $f_{\text{cen}}(z_{\text{cen}}^*) \geq f_{\text{cen}}(z_{\text{dec}}^*)$.

Interestingly, it is not possible to make general statements about the lower bounds L_{cen} and L_{dec} , respectively. The reason is that the prescribed simple rounding scheme produces *one* suboptimal solution for the centralized sensor selection problem (6). In principle, it is possible that the rounding from the decentralized leader nodes results in a different suboptimal solution, which in turn achieves a higher lower bound. However, as numerical evaluation suggests (see Sec. V) this effect does not occur very often.

III. APPROACH

In this section, we will introduce two methods for solving the centralized sensor selection problem (6) in a partially decentralized manner. Partially decentralized means that we are willing to transmit a negligible amount of data from one leader node to the other. In particular, we allow the transmission of $N \ll m$ vectors of dimension \mathbb{R}^n . Without loss of generality, we assume that leader node 1 shares some vectors with leader node 2. These methods can be seen as simple heuristics that attempt to improve the lower bound L_{dec} and, accordingly, the suboptimality gap (14).

A. Main Idea

The main idea behind both heuristics can be described as follows. Assume we have measurements that are approximately collinear, i.e., rows from the matrix A are weakly correlated. In the case of two rows, it follows from the error covariance matrix (4) that

$$a_j a_j^T + a_k a_k^T \approx (1 + \gamma) a_j a_j^T, \quad \gamma \in \mathbb{R}. \quad (15)$$

In the above expression, we have the sum of two rank-1 matrices that can be approximately be rewritten as a scaled version of one rank-1 matrix. This means, we have (effectively) lost one rank. Now recall that each binary variable \hat{z}_i corresponds to a specific selection out of all rows from A . Since $\sum_{i=1}^m \hat{z}_i a_i a_i^T$ is symmetric, it is orthogonally diagonalizable and we can rewrite the objective (6) as

$$\log \det \left(\sum_{i=1}^m \hat{z}_i a_i a_i^T \right) = \log \det (U \Lambda U^T) = \sum_{i=1}^n \log \lambda_i,$$

where λ_i are the (positive) eigenvalues of $A^T \mathbf{diag}(\hat{z}) A$. Selecting many weakly correlated measurements via \hat{z} will reduce the effective rank of the above matrix. This increases the number of small eigenvalues, and therefore increases the total volume of the corresponding confidence ellipsoid.

Unless we are allowed to use $k = m$ measurements, the centralized sensor selection will avoid correlated measurements since they do not add significantly to the objective in (6). This situation is different for the decentralized sensor selection problem. Both leader nodes maximize their own ellipsoid, using only the local data A_1 and A_2 , respectively. However, it is possible and likely that their individual selections are jointly correlated, which in turn leads to a smaller global ellipsoid. Now suppose leader node 1 shares some information about the largest contribution to its own ellipsoid. Leader node 2 would avoid picking the same contribution. This is the main idea behind both heuristics.

To be more specific, suppose only leader node 1 has solved the relaxed decentralized sensor selection problem (11). The main contribution to the volume of the local ellipsoid is given by the largest eigenvalues of the matrix $A_1^T \mathbf{diag}(\hat{z}_1) A_1$. Denote $\lambda_1, \dots, \lambda_N$ the N largest eigenvalues and u_1, \dots, u_N the associated eigenvectors with that matrix. We will now introduce two heuristics that capitalize the influence of these eigenvalues and -vectors.

B. Focused Diversity Method

Leader node 2 modifies its data matrix A_2 and selection vector z_2 such that

$$A_{\text{fdm}} = \begin{bmatrix} A_2 \\ \lambda_1 u_1^T \\ \vdots \\ \lambda_N u_N^T \end{bmatrix}, \quad \tilde{z} = \begin{bmatrix} z_2 \\ \tilde{z}_{m/2+1} \\ \vdots \\ \tilde{z}_{m/2+N} \end{bmatrix}. \quad (16)$$

We can now rewrite objective of the decentralized sensor selection problem for leader node 2 as

$$\log \det \left(A_{\text{fdm}}^T \mathbf{diag}(\tilde{z}) A_{\text{fdm}} \right) \quad (17)$$

$$= \log \det \left(\sum_{i=1}^{m/2} \tilde{z}_i a_{2i} a_{2i}^T + \sum_{i=1}^N \tilde{z}_{i+m/2} \lambda_i^2 u_i u_i^T \right). \quad (18)$$

Leader node 2 must avoid the same directions or contributions that were already made by leader node 1. One way to achieve this is to set $\tilde{z}_{m/2+1} = \dots = \tilde{z}_{m/2+N} = 1$. Basically, leader node 1 has already made a decision for the second node. The objective for leader node 2 has then the form

$$f_{2,\text{fdm}}(z_2) = \log \det \left(\sum_{i=1}^{m/2} z_{2i} a_{2i} a_{2i}^T + \sum_{i=1}^N \lambda_i^2 u_i u_i^T \right). \quad (19)$$

Hence the remaining relaxed optimization problem to be solved is given by

$$\begin{aligned} & \underset{z_2}{\text{maximize}} && f_{2,\text{fdm}}(z_2) \\ & \text{subject to} && \mathbf{1}^T z_2 = k_s/2 \\ & && 0 \leq z_{2i} \leq 1, \quad i = 1, \dots, m/2 \end{aligned} \quad (20)$$

where the solution is denoted as $z_{2,\text{fdm}}^* := z_2^*$. We call this heuristic *focused diversity method*. Note that the underlying maximization problem is still concave. After solving, leader node 2 performs the simple rounding scheme in order to obtain $\hat{z}_{2,\text{fdm}}$. The upper and lower bounds for the focused diversity method are then calculated based on the vectors

$$z_{\text{fdm}}^* = \begin{bmatrix} z_1^* \\ z_{2,\text{fdm}}^* \end{bmatrix}, \quad \hat{z}_{\text{fdm}} = \begin{bmatrix} \hat{z}_1 \\ \hat{z}_{2,\text{fdm}} \end{bmatrix}. \quad (21)$$

C. Linear Penalty Method

Another way to force leader node 2 to avoid the main directions from leader node 1 is to introduce a linear penalty for choosing similar measurements. This can be accomplished by adding an additional term to the objective given in (11). Consider the quantity

$$\underbrace{\left| \frac{a_{2i}^T u_j}{\|a_{2i}\|_2} \right|}_{\text{similarity}} \cdot \underbrace{\frac{\lambda_j}{\|a_{2i}\|_2}}_{\text{relevance}} = |a_{2i}^T \lambda_j u_j| \cdot \frac{1}{\|a_{2i}\|_2^2}, \quad (22)$$

where $i = 1, \dots, m/2$, and $j = 1, \dots, N$. Consider the LHS first. Whenever a measurement a_{2i} has a direction that is similar to a main direction from leader node 1, the absolute value of the normalized inner product between a_{2i} and u_j will be 'large'. In order to account for a possibly higher contribution from A_2 , we put an additional weight on this similarity, which is called relevance. Note that we only need to calculate the product $\lambda_j u_j$. Therefore, the RHS of (22) does not violate the restriction that only N vectors from leader node 1 can be shared.

Using the quantity (22) we can now construct a penalty term. Denote the cost vector

$$c_i = \sum_{j=1}^N |a_{2i}^T \lambda_j u_j| \cdot \frac{1}{\|a_{2i}\|_2^2} \geq 0, \quad i = 1, \dots, m/2. \quad (23)$$

We rewrite the objective for leader node 2 as

$$f_{2,\text{lpm}}(z_2) = \log \det \left(\sum_{i=1}^{m/2} z_{2i} a_{2i} a_{2i}^T \right) - \sum_{i=1}^{m/2} c_i z_{2i} \quad (24)$$

$$= f_2(z_2) - c^T z_2, \quad (25)$$

where $c^T = [c_1 \dots c_{m/2}]$. The resulting heuristic is called *linear penalty method*, associated with the (relaxed) concave optimization problem

$$\begin{aligned} & \underset{z_2}{\text{maximize}} && f_{2,\text{lpm}}(z_2) \\ & \text{subject to} && \mathbf{1}^T z_2 = k_s/2 \\ & && 0 \leq z_{2i} \leq 1, \quad i = 1, \dots, m/2 \end{aligned} \quad (26)$$

The solution is called $z_{2,\text{lpm}}^* := z_2^*$, and the resulting vectors for calculating the upper and lower bound are given by

$$z_{\text{lpm}}^* = \begin{bmatrix} z_1^* \\ z_{2,\text{lpm}}^* \end{bmatrix}, \quad \hat{z}_{\text{lpm}} = \begin{bmatrix} \hat{z}_1 \\ \hat{z}_{2,\text{lpm}} \end{bmatrix}. \quad (27)$$

IV. SOLVING A SENSOR SELECTION PROBLEM

In order to solve a concave maximization problem similar to (8), several methods are at hand. One could resort to optimization software like CVX [9]. In our case, we have implemented a logarithmic barrier method, see [8, Sec. 11.3]. A full implementation will be made available on our website. The reason for choosing a barrier method is that it can be implemented without great effort. However, other interior-point methods, which are used in practice more often (for example, primal-dual methods), may solve the above problems within fewer iterations and higher accuracy.

From [1, Sec. III D], the approximate objective of the relaxed sensor selection is given by

$$\psi_{\text{cen}}(z) = \log \det \left(\sum_{i=1}^m z_i a_i a_i^T \right) + \kappa \sum_{i=1}^m (\log(z_i) + \log(1 - z_i)) \quad (28)$$

$$= \log \det \left(\sum_{i=1}^m z_i a_i a_i^T \right) + \phi_\kappa(z). \quad (29)$$

The authors also give explicit expressions for the gradient $\nabla \psi_{\text{cen}}(z)$ and Hessian $\nabla^2 \psi_{\text{cen}}(z)$. This is needed in order to calculate the Newton step in the inner iteration of the barrier method [8]. In our case, leader node 2 has a modified approximate objective, depending on the method (e.g., focused diversity or linear penalty method). For completeness, the explicit expressions are given below.

For the focused diversity method (20), the approximate objective of leader node 2 has the form

$$\psi_{\text{fdm}}(z_2) = \log \det \left(\sum_{i=1}^{m/2} z_{2i} a_{2i} a_{2i}^T + \sum_{i=1}^N \lambda_i^2 u_i u_i^T \right) + \phi_\kappa(z_2), \quad (30)$$

where $\phi_\kappa : \mathbb{R}^{m/2} \rightarrow \mathbb{R}$. Therefore, the gradient is given by

$$\nabla \psi_{\text{fdm}}(z_2) = \mathbf{diag} \left(A_2^T W_{\text{fdm}} A_2 \right) + \nabla \phi_\kappa(z_2), \quad (31)$$

where

$$W_{\text{fdm}} = \left(\sum_{i=1}^{m/2} z_{2i} a_{2i} a_{2i}^T + \sum_{i=1}^N \lambda_i^2 u_i u_i^T \right)^{-1}. \quad (32)$$

The Hessian is given by

$$\nabla^2 \psi_{\text{fdm}}(z_2) = - (A_2^T W_{\text{fdm}} A_2) \odot (A_2^T W_{\text{fdm}} A_2) + \nabla^2 \phi_{\kappa}(z_2), \quad (33)$$

where \odot is the Hadamard product.

In the case of the linear penalty method (26), the approximate objective reads as

$$\psi_{\text{lpm}}(z_2) = \log \det \left(\sum_{i=1}^{m/2} z_{2i} a_{2i} a_{2i}^T \right) - c^T z_2 + \phi_{\kappa}(z_2). \quad (34)$$

Hence the gradient has the form

$$\nabla \psi_{\text{lpm}}(z_2) = \mathbf{diag} (A_2^T W_{\text{lpm}} A_2) - c + \nabla \phi_{\kappa}(z_2), \quad (35)$$

where

$$W_{\text{lpm}} = \left(\sum_{i=1}^{m/2} z_{2i} a_{2i} a_{2i}^T \right)^{-1}. \quad (36)$$

Since we have a linear penalty term, the Hessian remains unchanged when compared to [1, Sec. III D], except the different dimension $m/2 \times m/2$.

V. NUMERICAL EXAMPLE

We will now compare the naive decentralized sensor selection (Sec. II-C) with the proposed heuristics, e.g., focused diversity (Sec. III-B) and linear penalty method (Sec. III-C). Our benchmark for all comparisons is the solution from the centralized sensor selection problem (Sec. II-B).

Throughout this section we investigate the following modeling setup. We use $m = 100$ measurements, $n = 40$ unknown parameters, and a total number of $k_s = 40, \dots, 60$ sensors. The measurement matrix is partitioned as given in (10). The entries of the submatrices A_1 and A_2 are iid with $\mathcal{N}(0, 1)$. In order to create weakly correlated measurements, we pick randomly two different rows from A_1 and A_2 , say row i from A_1 and row j from A_2 . Both rows are then modified as follows

$$a_{1i}^T = \sqrt{1 - \sigma^2} \cdot b^T + \sigma w_i^T, \quad (37)$$

$$a_{2j}^T = \sqrt{1 - \sigma^2} \cdot b^T + \sigma w_j^T, \quad (38)$$

where $b_1, \dots, b_n, w_{i1}, \dots, w_{in}$, and w_{j1}, \dots, w_{jn} are iid with $\mathcal{N}(0, 1)$. In this formulation, σ represents of the strength of correlation. In our simulation we modified a total number of 30 rows and used $\sigma = 0.1$.

We first fix the number of shared vectors $N = 5$, and vary only the number of sensors to be used. In Fig. 2(a), the lower bounds L are shown for this case. We can clearly see that the focused diversity and the linear penalty method outperform the naive decentralized selection, notably for a low number of





Fig. 2. Lower bounds (a) and relative suboptimality gap (b) for the four different sensor selection strategies, where $N = 5$. Note that the centralized sensor selection serves as a benchmark for all decentralized strategies.

sensors. Since this gives only a lower bound we introduce the relative suboptimality gap

$$\Delta_{\text{rel},i} = 100 \cdot \frac{|U_{\text{cen}} - L_i|}{|U_{\text{cen}}|}, \quad i \in \{\text{cen}, \text{dec}, \text{fdm}, \text{lpm}\}. \quad (39)$$

This gap measures how far we are away from the optimum of the (nonconvex) centralized sensor selection problem (6). Note that the relative suboptimality gap is a worst-case measure. This means in practice we are often closer to this optimum than suggested by this gap. The corresponding results are depicted in Fig. 2(b). When compared to the naive decentralized selection, a considerable performance gain for both heuristics can be observed.

As mentioned in Sec. II-D, it is in principle not clear how the simple rounding scheme affects the different lower bounds, and hence the relative suboptimality gap. In order to get a meaningful interpretation, we run a simulation with 10^4 different random realizations of the aforementioned modeling

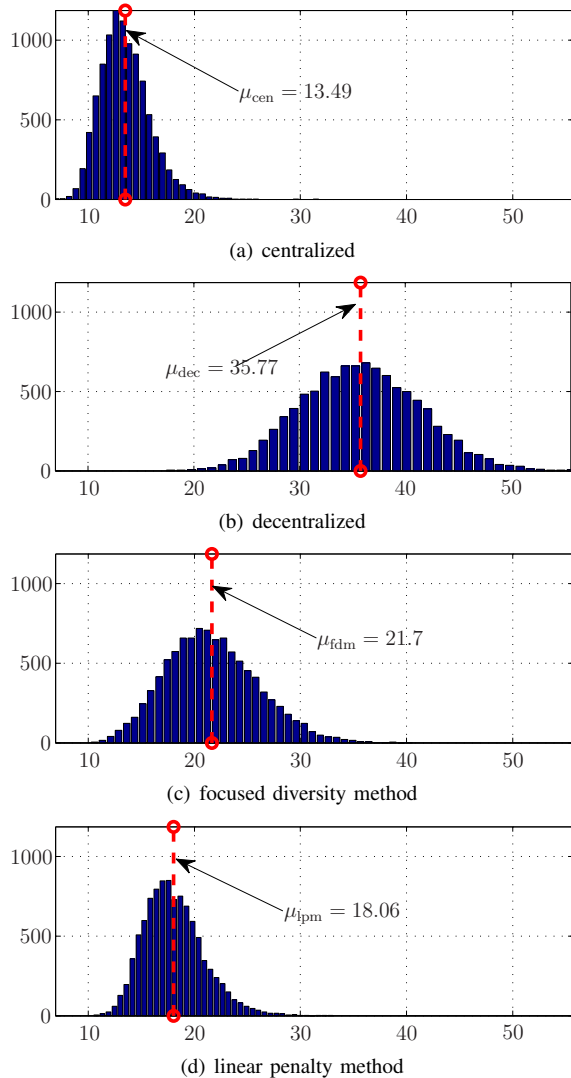


Fig. 3. Relative suboptimality gap for the four different sensor selection strategies, where $k_s = 40$ and $N = 5$.

setup. The results for $k_s = 40$ and $N = 5$ can be seen in Fig. 3. As shown by the histograms, the linear penalty method performs on average slightly better than the focused diversity method. It is also evident that the decentralized sensor selection is far away from being optimal. We have also plotted the empirical means μ_i , of all relative suboptimality gaps.

The influence of the number of shared vectors on the relative suboptimality gap is illustrated in Fig. 4. We assume that all vectors are transmitted perfectly, i.e., no transmission errors occur. Again, we have used total number of 10^4 realizations and $k_s = 40$ sensors. Each point in Fig. 4 corresponds to a mean $\mu_i(N)$. As a reasonable deviation measure, we have included standard deviations depicted as error bars around the corresponding means. Note that we can not make the gap between the centralized selection and decentralized heuristics arbitrarily small. The reason is that leader node 1 already made a decision, which in turn can be suboptimal from a centralized point of view.

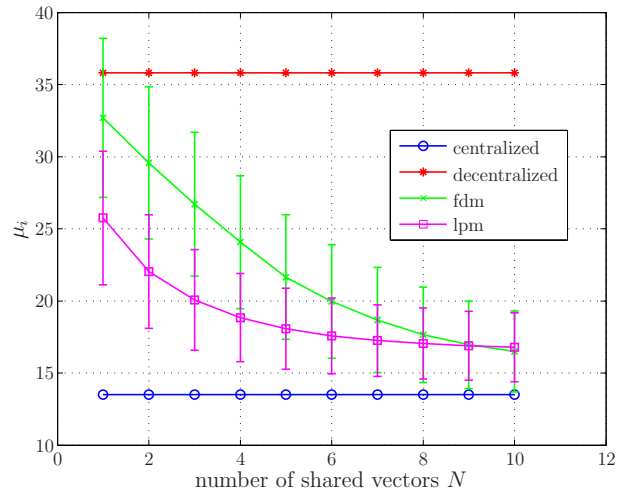


Fig. 4. Influence of the number of shared vectors N on the empirical mean of the relative suboptimality gap, where $k_s = 40$. The error bars around the mean values represent the corresponding standard deviations.

VI. CONCLUSION

We have introduced the decentralized sensor selection problem for a partition of m sensors between two leader nodes. A naive decentralized solution solves the sensor selection problem separately for both leader nodes. It was shown, that this solution is considerably worse when compared to a centralized approach. Therefore, we proposed two simple heuristics that allowed a very limited transmission of information from one leader to the other. These heuristics try to avoid similar contributions from both leader nodes, or equivalently, to exploit diversity in possible sensor selections. As suggested by extensive numerical simulations, the heuristics outperformed the naive decentralized selection by a substantial margin.

REFERENCES

- [1] S. Joshi and S. Boyd, "Sensor selection via convex optimization," *Trans. Sig. Proc.*, vol. 57, pp. 451–462, February 2009.
- [2] L. Yao, W. Sethares, and D. Kammer, "Sensor placement for on-orbit modal identification of large space structure via a genetic algorithm," in *Systems Engineering, 1992., IEEE International Conference on*, sep 1992, pp. 332–335.
- [3] Y. Mo, R. Ambrosino, and B. Sinopoli, "A convex optimization approach of multi-step sensor selection under correlated noise," in *Communication, Control, and Computing, 2009. Allerton 2009. 47th Annual Allerton Conference on*, 30 2009-oct. 2 2009, pp. 186–193.
- [4] V. Gupta, T. H. Chung, B. Hassibi, and R. M. Murray, "On a stochastic sensor selection algorithm with applications in sensor scheduling and sensor coverage," *Automatica*, vol. 42, no. 2, pp. 251–260, 2006.
- [5] M. Vitus, W. Zhang, A. Abate, J. Hu, and C. Tomlin, "On efficient sensor scheduling for linear dynamical systems," in *American Control Conference (ACC), 2010*, 30 2010-july 2 2010, pp. 4833–4838.
- [6] V. Kekatos and G. Giannakis, "Selecting reliable sensors via convex optimization," in *Signal Processing Advances in Wireless Communications (SPAWC), 2010 IEEE Eleventh International Workshop on*, june 2010, pp. 1–5.
- [7] R. A. Horn and C. R. Johnson, *Matrix Analysis*. Cambridge University Press, 1990.
- [8] S. Boyd and L. Vandenberghe, *Convex Optimization*. New York, NY, USA: Cambridge University Press, 2004.
- [9] M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming, version 1.22," <http://cvxr.com/cvx>, Feb. 2012.