LUND UNIVERSITY

Gradient methods for iterative distributed control synthesis

Mårtensson, Karl; Rantzer, Anders

2009

[Link to publication](#)

Citation for published version (APA):

Mårtensson, K., & Rantzer, A. (2009). *Gradient methods for iterative distributed control synthesis*. Paper presented at 48th IEEE Conference on Decision and Control, Shanghai, China.

Total number of authors:

2

General rights

Unless other specific re-use rights are stated the following general rights apply:

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

Gradient methods for iterative distributed control synthesis

Karl Mårtensson and Anders Rantzer

Abstract—In this paper we present a gradient method to iteratively update local controllers of a distributed linear system driven by stochastic disturbances. The control objective is to minimize the sum of the variances of states and inputs in all nodes.

We show that the gradients of this objective can be estimated distributively using data from a forward simulation of the system model and a backward simulation of the adjoint equations. Iterative updates of local controllers using the gradient estimates gives convergence towards a locally optimal distributed controller.

I. INTRODUCTION

Decision making when the decision makers have access to different information concerning underlying uncertainties has been studied since the late 1950s [7], [8]. The subject is sometimes called team theory, sometimes decentralized or distributed control. The theory was originally static, but work on dynamic aspects was initiated by Witsenhausen [13], who also pointed out a fundamental difficulty in such problems. Some special types of team problems were solved in the 1970's [12], [5], but the problem area has recently gain renewed interest. Spatial invariance was exploited in [1], [2], conditions for closed loop convexity were derived in [11], [10] and methods using linear matrix inequalities were given in [6], [9], [3].

In this paper we develop a distributed gradient method to update local linear feedback controllers in a distributed system. The objective is to minimize a global quadratic cost of the system. For a centralized control problem, the method would be a special case of iterative feedback tuning [4]. However, in our case local feedback laws are updated for each agent using information only about local dynamics and measurements from neighboring agents. This makes the complexity of the scheme linear in the number of agents.

Section II describes the distributed systems structure and notations are defined. The method to update the feedback laws is given in Section III. Here we give an expression of the gradient of the cost function and we see how it can be estimated using only local information. We give some examples in Section IV.

The authors are with Automatic Control LTH, Lund University, Box 118, SE-221 00 Lund, Sweden, {karl, rantzer}@control.lth.se

II. PROBLEM FORMULATION

Consider the system

$$x(t+1) = Ax(t) + Bu(t) + w(t) \quad (1)$$

where w is white noise with variance W , and $w(t)$ is independent of $x(s)$ for $s \leq t$. When we consider distributed systems, there is usually an associated graph structure. Let the graph consist of n agents i , $1 \leq i \leq n$, and the edge set E , such that $(i, j) \in E$ if there is an edge between agent i and j (by convention we let $(i, i) \in E$ for all i). We call i and j neighboring agents if there is an edge connecting them. Let E_i contain the indices of the neighboring agents of i , i.e.

$$E_i = \{j \mid \text{if } (i, j) \in E\}$$

Now, the dynamics matrix has a sparsity structure which resembles the graph structure of the distributed system, i.e.

$$A_{ij} = 0 \quad \text{if } (i, j) \notin E$$

where the notation A_{ij} means the block associated with how agent i affects agent j (throughout the paper, subscripts i, j will refer to blocks associated with agents i and j). In this paper we assume that each agent has one set of distinct control signals, i.e. each control signal affects only one agent directly. This means that B is a block-diagonal matrix. The case that an agent does not have an input signal could be modeled as letting the corresponding block in B be zero. This is not necessary, and the columns corresponding to such zero entries in B will be removed. One example of the complete setup is found in Figure 1.

The system is closed using state feedback

$$u(t) = -Lx(t) \quad (2)$$

for some L . When we consider a distributed setup, we limit the feedback matrix to have a structure that matches the system. The calculation of the input $u_i(t)$ in agent i should only require measurements of the neighboring agents. Hence, if exactly p agents have inputs, letting

$$L = \begin{bmatrix} L_{11} & L_{12} & \cdots & L_{1n} \\ L_{21} & L_{22} & \cdots & L_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ L_{p1} & L_{p2} & \cdots & L_{pn} \end{bmatrix}$$

Fig. 1. Graphical representation of a distributed system. The arrows shows how each agent affects the others. The set $E = \{(1,1), (2,2), (3,3), (4,4), (1,2), (2,1), (1,3), (3,1), (2,3), (3,2), (3,4), (4,3)\}$

where u_k is the control signal in agent i_k . Then we have that

$$L_{kj} = 0 \quad \text{if } (i_k, j) \notin E$$

With this setup the dynamics matrix $A - BL$ of the closed loop system will have the same structure as A .

III. CONTROL SYNTHESIS

For every stabilizing feedback matrix L , define

$$J(L) = \mathbf{E} (|x|_Q^2 + |u|_R^2) \quad (3)$$

At every iteration, we wish to change L in a descent direction of J . Obviously, if the gradient $\nabla_L J$ of the cost is non-zero, we can change the feedback matrix to get a lower cost by introducing

$$L_{\text{new}} = L - \gamma \nabla_L J$$

where γ is sufficiently small. With this new feedback matrix we recompute the gradient and keep iterating.

Next, we will derive an expression for the gradient of $J(L)$.

Proposition 1: Given matrices A and B , consider L such that $A - BL$ has all eigenvalues inside the unit circle. Consider stationary stochastic process satisfying (1) and (2) where w is white noise with covariance W . Then $J(L)$, defined by (3), has the gradient

$$\nabla_L J = 2 [RL - B^T P(A - BL)] X \quad (4)$$

where X and P satisfy the Lyapunov equations

$$X = (A - BL)X(A - BL)^T + W \quad (5)$$

$$P = (A - BL)^T P(A - BL) + Q + L^T RL \quad (6)$$

Proof. It is known that $J = \text{tr}(XQ) + \text{tr}(LX L^T R) = \text{tr}(PW)$. To calculate the differential of P , let

$$\begin{aligned} A_L &= A - BL \\ M &= [L^T R - A_L^T P B] dL \end{aligned}$$

Differentiating (6) shows that dP satisfies the Lyapunov equation

$$dP = A_L^T dP A_L + M + M^T$$

Hence

$$\begin{aligned} dP &= \sum_{k=0}^{\infty} (A_L^T)^k (M + M^T) A_L^k \\ \text{tr}(dPW) &= \text{tr} \left(2M^T \sum_{k=0}^{\infty} A_L^k W (A_L^T)^k \right) \\ &= \text{tr} (2dL^T (RL - B^T P A_L) X) \end{aligned}$$

This concludes the proof. $\square$

However, we do not only want an expression for $\nabla_L J$, we want one suitable for distributed computations. It turns out that this can be achieved by introducing the adjoint system:

Proposition 2: Under the conditions of Proposition 1, consider the stationary stochastic process λ defined by the backwards iteration

$$\lambda(t-1) = (A - BL)^T \lambda(t) - (Q + L^T RL)x(t) \quad (7)$$

where $x(t)$ are the states of the original system. Then

$$\nabla_L J = 2 (RLExx^T + B^T E \lambda x^T)$$

Proof. For simplicity, let $Q_L = Q + L^T RL$. For any j ,

$$\begin{aligned} \lambda(j) &= - \sum_{k=j+1}^{\infty} (A_L^T)^{k-j-1} Q_L x(k) \\ &= - \sum_{k=0}^{\infty} (A_L^T)^k Q_L A_L^{k+1} x(j) \\ &\quad + \Psi\{w(j), w(j+1), \dots\} \end{aligned}$$

where Ψ is the appropriate linear operator on the sequence $(w(j), w(j+1), \dots)$. Hence

$$\begin{aligned} \mathbf{E} \lambda(j)x(j)^T &= -\mathbf{E} \left[\sum_{k=0}^{\infty} (A_L^T)^k Q_L A_L^{k+1} x(j)x(j)^T \right] \\ &= -PA_L X \end{aligned}$$

Fitting this into (4) gives the desired result. $\square$

The proposition gives a way of estimating an update direction for the feedback matrix L . But as it is posed, it cannot be used in a distributed way. In the calculation of $\nabla_L J$ according to the equation in Proposition 2, the covariance between all states and between all states and all adjoint states needs to be determined. But with the appropriate projection of the gradient, we will find a distributed scheme to update the feedback matrix. To do this we restrict Q and R to be block-diagonal, with blocks fitting the size of the state space and number of inputs of each agent, respectively.

First we have to make sure that the adjoint system can be simulated locally, i.e. agent i should be able to

simulate its corresponding λ_i using only local information. Lets examine the state equation

$$\lambda_i(t-1) = [A_L^T \lambda(t)]_i - [(Q + L^T R L)x(t)]_i \quad (8)$$

By previous assumptions, it is known that the structure of A_L satisfies

$$(A_L)_{ij} = 0 \quad \text{if } (i, j) \notin E$$

Also, if $(i, j) \in E$ then $(A_L)_{ij}$ is considered to be known to both agent i and j . The first term of (8) will be

$$[A_L^T \lambda(t)]_i = \sum_{j \in E_i} (A_L)_{ji}^T \lambda_j(t)$$

The second term of (8) is simplified to

$$[(Q + L^T R L)x(t)]_i = Q_i x_i(t) - \sum_{j \in E_i} L_{ji}^T R_j u_j(t)$$

This shows that the adjoint state equation (8) can be simulated in each agent using only local information.

Since the structure of L satisfies that

$$L_{ij} = 0 \quad \text{if } (i, j) \notin E$$

the actual direction we update the feedback matrix must also satisfy this structure. Hence, we project the gradient $\Delta_L J$ to the subspace equivalent to the structure. Letting G be the update direction, we have that

$$G_{ij} = (\nabla_L J)_{ij} \quad \text{if } (i, j) \in E$$

$$G_{ij} = 0 \quad \text{otherwise}$$

Assuming that the projected gradient G is non-zero, $-G$ is a descent direction of $J(L)$. Now, this means that to update the feedback matrix, an agent i needs only to determine the gradient in the blocks corresponding to the neighboring agents. This requires that both $(R L E x x^T)_{ij}$ and $(B^T E \lambda x^T)_{ij}$ can be estimated locally. The first term can be simplified to

$$(R L E x x^T)_{ij} = -R_i E u_i x_j^T$$

which obviously can be estimated locally. With the assumed structure on B , the second term can be written as

$$(B^T E \lambda x^T)_{ij} = B_i^T E \lambda_i x_j^T$$

which is also possible to estimate locally. With this analysis we understand that the update of a feedback matrix in a distributed system can be made locally. The method is summarised in the following update scheme.

Algorithm 1: At time t_k , let the state feedback law be $u(t) = -L^{(k)}x(t)$. To update the feedback matrix in agent i :

- 1) Simulate the states $x_i(t)$ of the system (1) for times $t = t_k, \dots, t_k + N$ by communicating states from and to neighboring agents.

$$x_i(t+1) = \sum_{j \in E_i} (A - B L)_{ij} x_j(t) + w_i(t)$$

- 2) Simulate the adjoint states $\lambda_i(t)$ of the system (7) for times $t = t_k, \dots, t_k + N$ in the backwards direction, by communicating states from and to neighboring agents.

$$\lambda_i(t-1) = \sum_{j \in E_i} (A - B L)_{ji}^T \lambda_j(t) - \left(Q_i x_i(t) - \sum_{j \in E_i} L_{ji}^T R_j u_j(t) \right)$$

- 3) For every neighboring agent j , calculate the estimates of $E u_i x_j^T$ and $E \lambda_i x_j^T$ by

$$(E u_i x_j^T)_{\text{est}} = \frac{1}{N+1} \sum_{t=t_k}^{t_k+N} u_i(t) x_j(t)^T$$

$$(E \lambda_i x_j^T)_{\text{est}} = \frac{1}{N+1} \sum_{t=t_k}^{t_k+N} \lambda_i(t) x_j(t)^T$$

- 4) The estimate of the i, j -block of the gradient becomes

$$G_{ij} = -2 \left[R_i (E u_i x_j^T)_{\text{est}} + B_i^T (E \lambda_i x_j^T)_{\text{est}} \right]$$

- 5) For each neighboring agent j , update $L_{ij}^{(k+1)} = L_{ij}^{(k)} - \gamma G_{ij}$ for some step length γ .
- 6) Let $t_{k+1} = t_k + N$, increase k by one and go to 1).

We denote N by the iteration time, i.e. the length of the time interval where the system is controlled using a constant feedback matrix.

An important property of the posed scheme is that the complexity regarding the number of agents is linear. Introducing more agents to the system does not in principal change the calculations that are made in the old agents. Hence it does not involve much effort to add more agents to existing system.

Next we address the issue of finding a step length in which the updated feedback matrix actually gives a lower cost than the previous. In the equation for the gradient

$$\nabla_L J = 2 (R L E x x^T + B^T E \lambda x^T)$$

we notice that if the closed system with the initial feedback matrix has eigenvalues near the boundary of instability, the magnitude of the gradient tends to be large and the step length is required to be short to assure that the updated feedback matrix is stabilizing. With some assumptions we can find a step length that assures that the updated feedback matrix actually reduces the cost.

Proposition 3: Consider the system and the cost function in Proposition 1 with a stabilizing feedback matrix L_0 . For a descent direction Δ (of $J(L)$), define $L_h = L_0 + h\Delta$ and X_h through the equation

$$X_h = (A - B L_h) X_h (A - B L_h)^T + W$$

The cost is $\hat{J}(h) = J(L_h) = \text{tr}(X_h Q) + \text{tr}(L_h X_h L_h^T R)$. Let α be such that $X_h \leq \alpha W$ for all h that satisfy $\hat{J}(h) \leq \hat{J}(0)$. Furthermore, choose μ and ξ such that

$$\alpha B \Delta W \Delta^T B^T \leq (\mu - \alpha + 1)W$$

$$\Delta^T R \Delta \leq \xi Q$$

Let $\beta = 2(\mu^2 + \mu - \alpha + 1)$ and $\nu \geq 0$ solve the equation

$$\frac{4\mu^2}{\nu} + 2 = \frac{\nu}{\xi}$$

i.e. $\nu = \xi + \sqrt{\xi^2 + 4\xi\mu^2}$. For all $h \in [0, h_0]$, where

$$h_0 = \frac{-\text{tr}(\Delta^T \nabla_L J)}{(\beta + \nu)\hat{J}(0)}$$

then $\hat{J}'(h) \leq 0$. (Note that all matrix inequalities are with respect to the positive semidefinite cone).

Proof. In the proof we use the following relation on a number of occasions

$$M^T N + N^T M \leq a M^T M + a^{-1} N^T N \quad (9)$$

where M and N are square matrices and a a positive scalar.

Define $A_h = A - B L_h$. Assume that h is such that $\hat{J}(h) \leq \hat{J}(0)$. First we find a bound on $\frac{d}{dh} X_h$

$$X'_h = A_h X'_h A_h^T - B \Delta X_h A_h^T - A_h X_h \Delta^T B^T \quad (10)$$

Examining the last two terms of the right hand side, we get

$$-B \Delta X_h A_h^T - A_h X_h \Delta^T B^T \leq A_h X_h A_h^T + B \Delta X_h \Delta^T B^T$$

$$\leq (\alpha - 1)W + \alpha B \Delta W \Delta^T B^T \leq \mu W$$

and by the Lyapunov equation (10) we understand that $X'_h \leq \mu X_h$. To bound the second derivative of X_h , we proceed in a similar manner.

$$X''_h = A_h X''_h A_h^T - 2B \Delta X'_h A_h^T - 2A_h X'_h \Delta^T B^T$$

$$+ 2B \Delta X_h \Delta^T B^T$$

Using the bound on X'_h , the last three terms are bounded

$$-2B \Delta X'_h A_h^T - 2A_h X'_h \Delta^T B^T + 2B \Delta X_h \Delta^T B^T \leq$$

$$\leq 2\mu A_h X_h A_h^T + 2(\mu + 1)B \Delta X_h \Delta^T B^T \leq$$

$$\leq 2\mu(\alpha - 1)W + 2(\mu + 1)(\mu - \alpha + 1)W =$$

$$= 2(\mu^2 + \mu - \alpha + 1)W = \beta W$$

the bound on the second derivative will be $X''_h \leq \beta X_h$. We now use these bounds to find a bound on $\hat{J}''(h)$

$$\hat{J}'(h) = \text{tr}(X'_h Q) + \text{tr}(L_h X'_h L_h^T R) + 2\text{tr}(L_h X_h \Delta^T R)$$

$$\hat{J}''(h) = \text{tr}(X''_h Q) + \text{tr}(L_h X''_h L_h^T R) + 4\text{tr}(L_h X'_h \Delta^T R)$$

$$+ 2\text{tr}(\Delta X_h \Delta^T R) \leq$$

$$\leq \beta \hat{J}(h) + 4\mu \text{tr}(L_h X_h \Delta^T R) + 2\text{tr}(\Delta X_h \Delta^T R) \leq$$

$$\leq \beta \hat{J}(h) + \nu \text{tr}(L_h X_h L_h^T R) + \frac{\nu}{\xi} \text{tr}(\Delta X_h \Delta^T R) \leq$$

$$\leq (\beta + \nu) \hat{J}(h) \leq (\beta + \nu) \hat{J}(0)$$

Fig. 2. Graphical representation of the system in the example. The arrows shows how each agent affects the others.

where in the third last inequality, the relation (9) has been used with $M = L X_h^{1/2}$, $N = \Delta X_h^{1/2}$ and a and ν solves the equations

$$a \geq 0$$

$$2\mu a = \nu$$

$$2\mu a^{-1} + 2 = \frac{\nu}{\xi}$$

Now, for $h > 0$, we know that there exists $\theta \in [0, h]$ such that

$$\hat{J}'(h) = \hat{J}'(0) + \theta \hat{J}''(\theta)$$

Then for all $h \leq h_0$, with

$$h_0 = \frac{-\hat{J}'(0)}{(\beta + \nu)\hat{J}(0)}$$

we must have that $\hat{J}'(h) \leq 0$. To understand this, let the interval $[0, h_-]$ be the maximal interval such that $\hat{J}'(h) \leq 0$ for all h in the interval. Then $\hat{J}'(h_-) = 0$ and $\hat{J}(h_-) \leq \hat{J}(0)$. If $h_0 > h_-$, then

$$\hat{J}'(h_-) = \hat{J}'(0) + \theta \hat{J}''(\theta) <$$

$$< \hat{J}'(0) + \frac{-\hat{J}'(0)}{(\beta + \nu)\hat{J}(0)}(\beta + \nu)\hat{J}(0) = 0$$

Hence $h_0 \leq h_-$. The fact that $\hat{J}'(0) = \text{tr}(\Delta^T \nabla_L J)$ concludes the proof. $\square$

IV. EXAMPLE

The system

$$x(t+1) = Ax(t) + Bu(t) + w(t)$$

that is considered, consists of 10 agents, where the agents are connected in a linear fashion, see Figure 2. This leads to a tri-diagonal dynamics matrix, which, in this example, is

$$A = \begin{bmatrix} 0.5 & 0.5 & & & & & & & & \\ -0.5 & 0.1 & -0.3 & & & & & & & \\ & 0.4 & -0.2 & -0.5 & & & & & & \\ & & -0.4 & -0.5 & 0.2 & & & & & \\ & & & 0.2 & 0.3 & -0.1 & & & & \\ & & & & -0.3 & 0.1 & 0.3 & & & \\ & & & & & 0.2 & -0.4 & -0.4 & & \\ & & & & & & 0.2 & -0.2 & 0.3 & \\ & & & & & & & 0.5 & -0.5 & 0.3 \\ & & & & & & & & -0.1 & -0.1 \end{bmatrix}$$

and with the remaining entries equal to zero. We allow each agent to have an input and set $B = I$. The white noise w has unit covariance. We wish to minimize the cost

$$J(L) = \mathbf{E}(|x|_Q^2 + |u|_R^2)$$

where $u = -Lx$ and $Q = R = I$.

Fig. 3. The performance of the scheme when the case of full information is considered. The solid line shows the cost $J(L)$ when the value of gradient of J is approximately correct in each iteration. The dashed line shows the cost when N is small, resulting in that the estimation of the gradient is worse.

The magnitude of the maximal eigenvalue of A , $\rho(A) \approx 0.81$, hence we can initially let the system be uncontrolled, i.e. let $L = 0$.

To start with, we assume that each agent has full knowledge of the global system and gets measurements from every other node. If we let the iteration time N , i.e. the time interval when L is fixed, be large, the estimates of $\mathbf{E}ux^T$ and $\mathbf{E}\lambda x^T$ will be more accurate, meaning that the approximation of $\nabla_L J$ will be more accurate. Hence, the expectation is that the algorithm updates L in a way such that the cost $J(L)$ approaches the optimal cost $J(L_{\text{opt}})$. As the iteration time N is reduced, the approximation of $\nabla_L J$ becomes worse. We now expect the approach towards the optimal cost to be worse or even not seem to converge. Simulations illustrating this can be found in the plot in Figure 3.

Now consider the case when we impose the same structure on L as the dynamics matrix A , i.e. restricting L to be tri-diagonal. In accordance with the notation in Algorithm 1, the direction G in which L is updated is the projected version of $\nabla_L J$.

Results of this scheme is presented in Figure 4. As the iteration time N is decreased, the initial reduction rate in cost is increased. A drawback of having short iteration time is that noise affects both the cost and the feedback matrix L in a larger extent. The value where the cost settles is also increased. As the time N increases, the convergence becomes slower, but instead the influence of the noise is reduced and the final cost approaches the optimal value.

In Figure 5 a comparison of the performance of the algorithm, between the case of full information and the distributed system, can be found. Here we see that initially the distributed setup actually performs better. At first this can seem counter-intuitive. The

Fig. 4. The performance of the scheme when the case of tri-diagonal state feedback is considered. The dashed line shows the case when $N = 10$, the solid line shows the case when $N = 3$ and the dotted line shows the case when $N = 20$.

Fig. 5. A comparison between the case of full information and tri-diagonal state feedback. The solid line represents the case of full information and the dashed line shows the case tri-diagonal state feedback. In both cases, the iteration time $N = 10$.

explanation of this effect is that the estimates of the covariance of non-neighboring agents will be quite poor when using a short iteration time N . Hence, in the case of full information we are introducing errors in elements that are not close to the diagonal.

V. CONCLUDING REMARKS

By determining the gradient of the cost in each agent with only local information, we have obtained a method to iteratively change the feedback laws to improve the global performance of a distributed system. One important property of the method is that the complexity is linear. Also, including new agents to an already existing system, does not change the calculations in previously existing nodes except the neighboring ones.

VI. ACKNOWLEDGMENT

The research leading to these results has received funding from the European Community's Seventh Framework Programme under grant agreement number 224428, acronym CHAT.

This work was also supported by a Linnaeus Grant from the Swedish Research Council.

REFERENCES

- [1] Bassam Bamieh, Fernando Paganini, and Munther A. Dahleh. Distributed control of spatially invariant systems. *IEEE Transactions on Automatic Control*, 47(7), July 2002.
- [2] Bassam Bamieh and Petros G. Voulgaris. A convex characterization of distributed control problems in spatially invariant systems with communication constraints. *Systems & Control Letters*, 54(6):575–583, June 2005.
- [3] Ather Gattami. Generalized linear quadratic control theory. In *Proceedings of the 45th IEEE Conference on Decision and Control*, San Diego, CA, December 2006.
- [4] Håkan Hjalmarsson. Iterative feedback tuning –an overview. *International Journal on Adaptive Control and Signal Processing*, January 2002.
- [5] Yu-Chi Ho and Kai-Ching Chu. Team decision theory and information structures in optimal control problems–Part I. *IEEE Transactions on Automatic Control*, 17(1):15–22, Feb 1972.
- [6] Cédric Langbort, Ramu Sharat Chandra, and Raffaello D’Andrea. Distributed control design for systems interconnected over an arbitrary graph. *IEEE Transactions on Automatic Control*, 49(9):1502–1519, September 2004.
- [7] J. Marshak. Elements for a theory of teams. *Management Sci.*, 1:127–137, 1955.
- [8] R. Radner. Team decision problems. *Ann. Math. Statist.*, 33(3):857–881, 1962.
- [9] Anders Rantzer. A separation principle for distributed control. In *Proceedings of the 45th IEEE Conference on Decision and Control*, San Diego, CA, December 2006.
- [10] M. Rotkowitz and S. Lall. A characterization of convex problems in decentralized control. *IEEE Trans. on Automatic Control*, 51(2):274–286, Feb 2006.
- [11] Michael Rotkowitz and Sanjay Lall. Decentralized control information structures preserved under feedback. In *Proceedings of the IEEE Conference on Decision and Control*, December 2002.
- [12] Nils R. Sandell and Michael Athans. Solution of some nonclassical LQG stochastic decision problems. *IEEE Transactions on Automatic Control*, 19(2):108–116, April 1974.
- [13] H.S. Witsenhausen. A counterexample in stochastic optimum control. *SIAM Journal on Control*, 6(1):138–147, 1968.