LUND UNIVERSITY

Optimal Coordination of Homogeneous Agents Subject to Delayed Information Exchange

Madjidian, Daria; Mirkin, Leonid; Rantzer, Anders

2014

Total number of authors:
3

Optimal Coordination of Homogeneous Agents Subject to Delayed Information Exchange

Daria Madjidian, Leonid Mirkin, and Anders Rantzer

Abstract— We consider a class of large scale linear-quadratic coordination problems where the information exchange is subject to a time delay. We show that several previously known properties of the optimal solution to the delay-free problem extend to this case. In particular, the optimal control law comprises a diagonal (decentralized) term complemented by a rank-one coordination term, which can be implemented by a simple averaging operation. Moreover, the computational effort required to obtain the controller is independent of the number of agents.

I. INTRODUCTION

Control of large scale systems has been an established area of research for more than half a century and has received renewed attention over the last two decades. These systems are characterized by a very large number of interconnected subsystems, each with their own sensors and actuators. In such situations, fully centralized, structureless information processing becomes infeasible.

A common approach to limit the amount of information processing in the controller is to impose a sparsity pattern on the controller structure, which shapes permitted information exchange between subsystems [1]–[3]. However, the design of these type of structured controllers is generally a notoriously difficult problem [4], [5]. Considerable effort has been made in the control community to understand the nature of these difficulties and devise tools to address them, see [3], [6], [7] and the references therein.

Sparsity is not the only way to attain computational and implementational scalability. One example of a non-sparse, yet scalable, controller is the diagonal-plus-low rank configuration. This type of controller comprises a block-diagonal term, which is completely decentralized, complemented by a low-rank component that can be implemented via a few averaging operations. To the best of our knowledge, it first explicitly appeared in [8] as the structure of optimal controllers for symmetrically interconnected systems and then in [9] as a constraint imposed on some large-scale robust stability problems.

Recently, we showed that this type of control structure also appears naturally (i.e., without being imposed) as the optimal

D. Madjidian and A. Rantzer are with the Department of Automatic Control, Lund University, Box 118, SE–221 00 Lund, Sweden. E-mails: {daria, anders.rantzer}@control.lth.se.

L. Mirkin is with the Faculty of Mechanical Engineering, Technion—IIT, Haifa 32000, Israel. E-mail: mirkin@technion.ac.il.

solution to a class of large-scale linear-quadratic *coordination problems* [10]. Specifically, we studied a homogeneous group of autonomous agents that are coupled through a constraint on their average behavior. It was shown that the optimal *centralized* solution has a diagonal-plus-rank-one form. Although the rank-one term depends on information from all systems, the only centralized computation required to implement it is a single averaging operation, which scales well as the number of agents grows large. We also showed that the solution has an additional attractive property in terms of large scale applications: the computational effort required to obtain the solution is independent of the number of agents.

A potential limitation of the solution derived in [10], as well as of the earlier examples, is that it assumes immediate information exchange between the agents. This might not be feasible in some applications due to communication constraints. In this paper, we show that the properties of the optimal control law discussed above for the delay-free case, extend to the case with delayed information exchange. Moreover, we derive analytic expressions that quantify the performance deterioration due to the delay. An important implication of our result is that it adds insight into the class of diagonal-plus-low-rank controllers by enlarging the class of known problems for which they are well-suited.

Notation: The transpose of a matrix M is denoted by M'. By e_i we refer to the ith standard basis of a Euclidean space and by I_n to the $n \times n$ identity matrix (we drop the dimension subscript when the context is clear). The Frobenius norm of a matrix is $\|M\|_F := \sqrt{\operatorname{tr}(M'M)}$. The notation $\otimes$ is used for the Kronecker product of matrices:

$$A \otimes B := \begin{bmatrix} a_{11}B & \cdots & a_{1m}B \\ \vdots & \ddots & \vdots \\ a_{p1}B & \cdots & a_{pm}B \end{bmatrix},$$

where a_{ij} stands for the (i, j) entry of A. The $L^2(j\mathbb{R})$ norm [11, Ch. 4] of a system G is denoted $\|G\|_2$. For a set $\mathbb{A}$, the indicator function $\mathbb{1}_{\mathbb{A}}(t)$ is 1 if $t \in \mathbb{A}$ and 0 elsewhere.

II. PRELIMINARIES

In this section we review the main results of [10], which studies a coordination problem among ν uncoupled homogeneous systems (agents)

$$\Sigma_i : \dot{x}_i(t) = Ax_i(t) + B_u u_i(t), \quad x_i(0) = x_{i0}$$

where the state vectors x_i can be measured. Associated with each system is the cost function

$$\mathcal{J}_i = \int_0^\infty \big(x_i'(t)Q_\alpha x_i(t) + u_i'(t)u_i(t)\big)\mathrm{d}t.$$

Fig. 1. Overall optimal controller corresponding to (3)

Fig. 2. Aggregate standard state-feedback H^2 problem

The coordination among the agents is imposed by constraining the behavior of their *center of mass*, defined as the system

$$\bar{\Sigma} : \dot{\bar{x}}(t) = A\bar{x}(t) + B_u\bar{u}(t), \quad \bar{x}(0) = \bar{x}_0,$$

which connects the signals

$$\bar{u}(t) := \sum_{i=1}^{\nu} \mu_i u_i(t) \quad \text{and} \quad \bar{x}(t) := \sum_{i=1}^{\nu} \mu_i x_i(t), \quad (1)$$

where $\mu_i \neq 0$ may be thought of as the mass of the ith system. We then require that $\bar{\Sigma}$ evolves according to

$$\dot{\bar{x}}(t) = (A + B_u\bar{F})\bar{x}(t), \quad \bar{x}(0) = \bar{x}_0$$

for a given *stabilizing* $\bar{F}$ (in other words, we require that $\bar{u} = \bar{F}\bar{x}$). The local objectives and coordination requirement are then combined in the following constrained optimization problem:

$$\text{minimize} \quad \mathcal{J} := \sum_{i=1}^{\nu} \mathcal{J}_i \quad (2a)$$

$$\text{subject to} \quad \Sigma_i, \quad i = 1, \dots, \nu \quad (2b)$$

$$\bar{u} - \bar{F}\bar{x} = 0 \quad (2c)$$

Problem (2) is a constrained LQR problem, which can, in principle, be reduced to a standard, unconstrained, LQR problem via resolving (2c) in one of u_i's. Yet this approach is not readily scalable if the number of agents ν grows. Instead, [10] applied a coordinate transformation[1], which splits problem (2) into a set of ν uncoupled problems. Using this approach, it was shown that the optimal control law has the following form:

$$u_i(t) = F_\alpha x_i(t) + \mu_i(\bar{F} - F_\alpha)\bar{x}(t), \quad i = 1, \dots, \nu, \quad (3)$$

where F_α is the optimal LQR gain corresponding to the local, unconstrained, problem with the plant Σ_i and the cost function $\mathcal{J}_i$.

The control law (3) has two attractive scalability properties. First, the computational effort required to calculate its parameters does not depend on the number of agents ν, only one local unconstrained LQR problem needs to be solved. Second, the information exchange between the subsystems Σ_i requires only the knowledge of the state of the center of mass, $\bar{x}$. Computing this quantity only requires a single averaging operation, which is far less demanding, from both computation and communication viewpoints, than full centralized information processing. The overall controller then has the block diagonal-plus-rank-one structure, sketched out in Fig. 1.

[1]This transformation happens to be similar to that proposed in [12] and, earlier, in the literature on symmetric systems, see the references in [8].
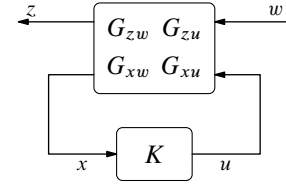
A potential limitation of the control law (3) is that each system needs to have immediate access to $\bar{x}$. This might not be feasible in applications, where communication resources are limited. Motivated by this, below we revisit the coordination problem (2) under the additional constraint that the information exchange between the agents is subject to a time-delay and show that the scalability properties discussed above extend to this case.

III. PROBLEM FORMULATION

The problem setup considered in this paper is slightly altered with respect to that studied in [10] (see Section II).

First, instead of the LQR formulation (nonzero initial conditions, no disturbances), we choose an H^2 formulation (zero initial conditions, exogenous disturbances). Accordingly, the local dynamics of each agent read

$$\dot{x}_i(t) = Ax_i(t) + B_{w_i}w_i(t) + B_u u_i(t) \quad (4)$$

for $i = 1, \dots, \nu$, where $x_i(t) \in \mathbb{R}^n$ are (measured) state vectors, $u_i(t) \in \mathbb{R}^m$ are control inputs, and $w_i(t) \in \mathbb{R}^{p_i}$ are exogenous disturbances. Local requirements are quantified in terms of the H^2 problems associated with the regulated outputs

$$z_i(t) := C_z x_i(t) + D_{zu}u_i(t). \quad (5)$$

Aggregating (4)–(5) for all $i = 1, \dots, \nu$, the local problems can be cast as the standard state-feedback H^2 problem [11, §14.8.1] depicted in Fig. 2. Here w, z, u, and x are the aggregate disturbance, regulated output, control input, and measured state vector, respectively (e.g., $w := \sum_{i=1}^{\nu} e_i \otimes w_i$), and the generalized plant

$$G(s) = \begin{bmatrix} G_{zw}(s) & G_{zu}(s) \\ G_{xw}(s) & G_{xu}(s) \end{bmatrix} = \left[\begin{array}{c|cc} I_\nu \otimes A & B_w & I_\nu \otimes B_u \\ \hline I_\nu \otimes C_z & 0 & I_\nu \otimes D_{zu} \\ I_{\nu n} & 0 & 0 \end{array} \right]$$

with $B_w := \sum_{i=1}^{\nu} (e_i e_i') \otimes B_{w_i}$. Each sub-block of this G is block diagonal. Hence, if no other constraints were imposed, the optimal solution K would be block diagonal as well.

Second, the coordination constraint are also adjusted, to keep it in line with the system-based formulation. To this end, introduce the vector

$$\mu := \begin{bmatrix} \mu_1 & \cdots & \mu_\nu \end{bmatrix}'$$

and rewrite (2c) as

$$\begin{aligned} 0 = \bar{u} - \bar{F}\bar{x} &= (\mu' \otimes I_m)u - (\mu' \otimes \bar{F})x \\ &= ((\mu' \otimes I_m)K(s) - \mu' \otimes \bar{F})x \\ &= (\mu' \otimes I_m)(K(s) - I_\nu \otimes \bar{F})x. \end{aligned}$$

We require this equality to hold for every x (restrictive compared to (2c) only if part of the state space is not excited by w). This implies $(\mu' \otimes I_m)(K(s) - I_\nu \otimes \bar{F}) = 0$, which is signal independent and obviously guarantees (2c).

The last, and only nontrivial, deviation from [10] is that we now impose an additional constraint on the controller structure. We require that

$$K(s) = \begin{bmatrix} K_{11}(s) & e^{-sh}K_{12}(s) & \cdots & e^{-sh}K_{1\nu}(s) \\ e^{-sh}K_{21}(s) & K_{22}(s) & \cdots & e^{-sh}K_{2\nu}(s) \\ \vdots & \vdots & \ddots & \vdots \\ e^{-sh}K_{\nu1}(s) & e^{-sh}K_{\nu2}(s) & \cdots & K_{\nu\nu}(s) \end{bmatrix} \quad (6)$$

for some *proper* $K_{ij}(s)$ and $h > 0$. This structure reflects our assumption that any information exchange between different agents is delayed by at least h time units.

The formal statement of the problem considered in this paper is then as follows:

$$\text{minimize} \quad \|T_{zw}\|_2 \quad (7a)$$

$$\text{subject to} \quad (\mu' \otimes I_m)(K(s) - I_\nu \otimes \bar{F}) = 0 \quad (7b)$$

$$K \text{ is of the form (6)} \quad (7c)$$

We address (7) under the following assumptions:

$\mathcal{A}_1$: (A, B_u) is stabilizable,

$\mathcal{A}_2$: $\begin{bmatrix} A - j\omega I & B_u \\ C_z & D_{zu} \end{bmatrix}$ has full column rank $\forall \omega \in \mathbb{R}$,

$\mathcal{A}_3$: $D'_{zu}D_{zu} = I$,

$\mathcal{A}_4$: $\mu'\mu = 1$ and all entries in μ are non-zero,

$\mathcal{A}_5$: the matrix $\bar{A} := A + B_u\bar{F}$ is Hurwitz.

Assumptions $\mathcal{A}_{1,2}$ are necessary for the well-posedness of the unconstrained local problems and $\mathcal{A}_5$ is necessary for the stabilizability of the overall system. The normalization assumptions in $\mathcal{A}_{3,4}$ are introduced to simplify the exposition and can be relaxed. Finally, if $\mu_i = 0$, then the ith system is not a part of the coordination problem and can therefore be excluded from the analysis.

IV. PROBLEM SOLUTION

To formulate the solution, we need the following algebraic Riccati equation (ARE):

$$A'X_\alpha + X_\alpha A + C'_z C_z \\ - (X_\alpha B_u + C'_z D_{zu})(B'_u X_\alpha + D'_{zu} C_z) = 0, \quad (8)$$

This is the state-feedback ARE associated with the uncoordinated version of the problem. It is known [11, Cor. 13.10] that $\mathcal{A}_{1-3}$ guarantee that its stabilizing solution $X_\alpha \geq 0$ exists. We also need the Lyapunov equation

$$\bar{A}'Y + Y\bar{A} + (\bar{F} - F_\alpha)'(\bar{F} - F_\alpha) = 0, \quad (9)$$

where $F_\alpha := -B'_u X_\alpha - D'_{zu} C_z$, and the matrix function

$$Y_h := Y - e^{\bar{A}'h}Y e^{\bar{A}h} = \int_0^h e^{\bar{A}'\theta}(\bar{F} - F_\alpha)'(\bar{F} - F_\alpha)e^{\bar{A}\theta}d\theta.$$

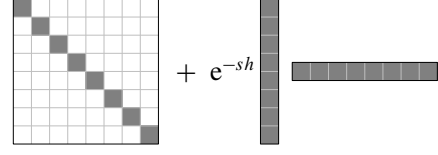$(\lim_{h\to 0} Y_h = 0$ and, as $\bar{A}$ is Hurwitz, $\lim_{h\to\infty} Y_h = Y)$.

Fig. 3. Overall optimal controller corresponding to (11)

The main result of this paper, whose proof is postponed to §IV-B, is formulated below.

Theorem 4.1: Let assumptions $\mathcal{A}_{1-5}$ be satisfied. Then the optimal performance in (7) is

$$\|T_{zw}\|_2^2 = \sum_{i=1}^\nu \text{tr}(B'_{w_i} X_\alpha B_{w_i}) + \sum_{i=1}^\nu \mu_i^2 \, \text{tr}(B'_{w_i} Y B_{w_i}) \\ + \sum_{i=1}^\nu (1 - \mu_i^2) \, \text{tr}(B'_{w_i} Y_h B_{w_i}) \quad (10)$$

and it is attained by

$$u_i(t) = \bar{F}x_i(t) + (F_\alpha - \bar{F})\hat{x}_i(t) - \mu_i(F_\alpha - \bar{F})e^{\bar{A}h}\bar{x}(t-h), \quad (11)$$

where

$$\hat{x}_i(t) := e^{\bar{A}h}x_i(t-h) \\ + \int_{t-h}^t e^{\bar{A}(t-\theta)} B_u\big(u_i(\theta) - \bar{F}x_i(\theta)\big)d\theta, \quad (12)$$

and $\bar{x}$ is the state of the center of mass defined in (1). ▽

A. Discussion

In this subsection, we discuss some properties of the optimal control law in Theorem 4.1.

1) Computational scalability and structure: An important consequence of Theorem 4.1 is that the two scalability properties of the solution to (2) discussed in Section II extend to the case with delayed information exchange. First, as in the case of $h = 0$, we only need to solve ARE (8) to form the optimal control law in Theorem 4.1. The computational effort is thus independent of the number of agents ν. Second, the optimal control law (11) comprises three terms, where the the first two are completely decentralized. As for the last term, the only global computation needed to form it is a single (scaled) averaging operation as in (1).

The structure of the overall controller is shown in Fig. 3. This structure is identical to that presented in Fig. 1 modulo the delay e^{-sh} in the rank-one coordination term.

2) Interpretation: It may be useful to view (11) in terms of the signals

$$v_i := u_i - \bar{F}x_i, \quad i = 1, \ldots, \nu.$$

With $v_i = 0$, $\forall i = 1, \ldots, \nu$, constraint (7b) would be satisfied without any need to exchange information between the agents. The signals v_i may thus be viewed as additional degrees of freedom in the control, which are brought about by the possibility to exchange information with other agents. Substituting $u_i = \bar{F}x_i + v_i$ into (4), the dynamics of the ith agent can be rewritten as

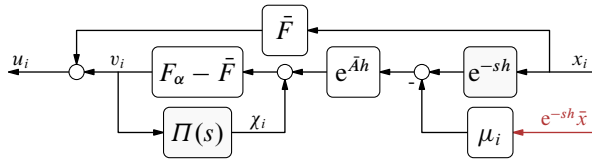$$\dot{x}_i(t) = \bar{A}x_i(t) + B_{w_i} w_i(t) + B_u v_i(t). \quad (4')$$

Fig. 4. A realization of the optimal controller for the ith system

Fig. 5. Equivalent reformulation with a uniform loop delay and biased control input via $u = (I_m \otimes \bar{F})x + v$

This form helps to explain the signal $\hat{x}$ defined by Eqn. (12). Indeed, in terms of v_i this equation reads

$$\hat{x}_i(t) = e^{\bar{A}h}x_i(t-h) + \int_{t-h}^{t} e^{\bar{A}(t-\theta)}B_u v_i(\theta)d\theta, \quad (12')$$

which is the *mean-squared prediction* of $x_i(t)$ based on $x_i(\tau)$, $\tau \le t - h$, cf. [13, Lesson 16].

Now, rewrite the control law (11) in terms of v_i:

$$v_i(t) = (F_\alpha - \bar{F})\hat{x}_i(t) - \mu_i(F_\alpha - \bar{F})e^{\bar{A}h}\bar{x}(t-h) \quad (11')$$

$$= (F_\alpha - \bar{F})\hat{x}_i(t) - \mu_i \sum_{j=1}^{v} \mu_j(F_\alpha - \bar{F})\hat{x}_j(t)$$

(the last expression can be verified by totting up the right-hand side of (12') and using the fact that $\bar{v} := \sum_j \mu_j v_j = 0$). This structure bears resemblance with the control law in the delay-free ($h = 0$) case, (3), which can be presented as

$$v_i(t) = \underbrace{(F_\alpha - \bar{F})x_i(t)}_{v_{\text{loc},i}(t)} - \mu_i \sum_{j=1}^{v} \mu_j \underbrace{(F_\alpha - \bar{F})x_j(t)}_{v_{\text{loc},j}(t)}, \quad (3')$$

where $v_{\text{loc},i}$ is the optimal control law for the ith subsystem in the absence of the coordination constraint (2c). The second term can then be interpreted as the least harmful correction to locally optimal strategy to enforce (2c) and it is a weighted average of the locally optimal control signals of all agents. If $h > 0$, (3') cannot be applied because $v_{\text{loc},j}(t)$ is available to agents $i \ne j$ only after h time units. The control law (11') then substitutes all $v_{\text{loc},j}(t)$ with their best (in the mean-squared sense) predictions, which are available to all agents.

3) Implementation: Denote by Π the operator $v_i \mapsto \hat{x}_i$ in (12'). This is an LTI distributed-delay system, whose impulse response is $e^{\bar{A}t}B_u \mathbb{1}_{[0,h]}(t)$. As such, its transfer function is

$$\Pi(s) = \int_0^h e^{-(sI - \bar{A})\theta}d\theta B_u. \quad (13)$$

A possible implementation of the optimal controller of Theorem 4.1 can then be as depicted in Fig. 4. This implementation contains two infinite-dimensional dynamical elements: the pure delay e^{-sh}, which is easy to implement, and the distributed-delay system Π. The latter is an intrinsic part of many optimal control strategies, see, e.g., [14]–[18], which study problems with a single delay.

Although distributed-delay systems can be safely implemented [19]–[21], their implementation might be numerically involved. The implementation in our case, however, is simplified because the matrix $\bar{A}$ is Hurwitz (by $\mathcal{A}_5$). Indeed, it is readily seen that (13) can be equivalently written as

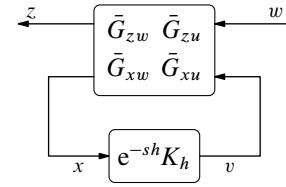$$\Pi(s) = (sI - \bar{A})^{-1}(B_u - e^{\bar{A}h}B_u e^{-sh}), \quad (13')$$

whose singularities at the eigenvalues of $\bar{A}$ are removable. This transfer function can be implemented as

$$\dot{\chi}_i(t) = \bar{A}\chi_i(t) + B_u v_i(t) - e^{\bar{A}h}B_u v_i(t-h), \quad (14)$$

which is a combination of a (stable) finite-dimensional system and a pure delay element, whose implementations are standard. Although this implementation involves pole-zero cancellations of all eigenvalues of $\bar{A}$, the cancellations are *stable*. Hence, the implementation via (14) is internally stable and thus valid.

Finally, the only matrix exponential involved in implementing the blocks in Fig. 4 is $e^{\bar{A}h}$. Because $\bar{A}$ is Hurwitz, the computation of the optimal controller parameters is numerically stable even for large h.

B. Proof of Theorem 4.1

The first step in solving (7) is to reduce it to a constrained H^2 problem, in which the unorthodox delay structure in (6) is replaced by a single delay element that applies to all measurement channels uniformly. This is done in the following lemma:

Lemma 4.2: All K satisfying (7b) and (7c) are of the form

$$K(s) = I_v \otimes \bar{F} + e^{-sh}K_h(s),$$

where $K_h(s)$ is any proper transfer function verifying

$$(\mu' \otimes I_m)K_h = 0. \quad (15)$$

Proof: If K is of the form (6), the jth block-column of (7b) reads ($\mu_j \ne 0$ by $\mathcal{A}_4$)

$$\mu_j K_{jj}(s) + e^{-sh}\sum_{i \ne j}\mu_i K_{ij}(s) - \mu_j \bar{F} = 0$$

It is readily seen that the latter condition holds iff

$$K_{jj}(s) = \bar{F} + e^{-sh}\tilde{K}_{jj}(s)$$

for some $\tilde{K}_{jj}$ satisfying $\mu_j \tilde{K}_{jj} + \sum_{i \ne j}\mu_i K_{ij} = 0$. The result follows by repeating these arguments for every j. ∎

We now substitute K with the right-hand side of the expression in Lemma 4.2, which is equivalent to shifting $u_i = \bar{F}x_i + v_i$, where v_i is a new control signal. This step replaces constraint (7c) with the unified loop delay e^{-sh}, transforming (7) into the H^2 problem depicted in Fig. 5, where K_h is now subject to (15), and the generalized plant is given by

$$\begin{bmatrix} \bar{G}_{zw} & \bar{G}_{zu} \\ \bar{G}_{xw} & \bar{G}_{xu} \end{bmatrix} = \left[\begin{array}{c|cc} I_v \otimes \bar{A} & B_w & I_v \otimes B_u \\ \hline I_v \otimes \bar{C}_z & 0 & I_v \otimes D_{zu} \\ I_{vn} & 0 & 0 \end{array} \right], \quad (16)$$

where $\bar{C}_z := C_z + D_{zu}\bar{F}$.

With the delay applied uniformly to all components of the controller, the problem in Fig. 5 subject to (15) can be solved using the approach in [10]. Due to space limitations, we only provide an outline of the steps involved in deriving the control law (11). For details we refer to [10, §III.B].

1) Decoupling: The first step is to decouple the local problems and the coordination constraint. To this end, let U be a *unitary* matrix such that $U\mu = e_1$ (this is the singular value decomposition of μ, so U exists) and set

$$\tilde{x} := (U \otimes I_n)x, \quad \tilde{v} := (U \otimes I_m)v, \quad \tilde{z} := (U \otimes I_q)z.$$

This coordinate transformation decouples the problem into $\nu-1$ decoupled delayed H^2 problems and one problem with the predefined control law $\tilde{v}_1 = 0$. Although the transformation couples the disturbances in the H^2 formulation, this does not affect the solution of the state-feedback problem.

2) Single-delay H^2 problems in transformed coordinates: Following the solution steps in [17] (which addresses the output-feedback version of the problem), we end up with

$$\tilde{v}_i(t) = F\left(e^{\bar{A}h}\tilde{x}_i(t-h) + \int_{t-h}^{t} e^{\bar{A}(t-\theta)}B_u\tilde{v}_i(\theta)d\theta\right)$$

for $i = 2, \ldots, \nu$, where F is the optimal feedback gain of the delay-free H^2 problem associated with the generalized plant (16). It can be shown that $F = F_\alpha - \bar{F}$. The overall control law in the transformed coordinates is then

$$\tilde{v}(t) = \left((I_\nu - e_1 e_1') \otimes (F_\alpha - \bar{F})\right)\left((I_\nu \otimes e^{\bar{A}h})\tilde{x}(t-h)\right.$$
$$\left. + \int_{t-h}^{t}\left(I_\nu \otimes (e^{\bar{A}(t-\theta)}B_u)\right)\tilde{v}(\theta)d\theta\right)$$
$$= \left(I_\nu \otimes (F_\alpha - \bar{F})\right)\left((I_\nu \otimes e^{\bar{A}h})\tilde{x}(t-h)\right.$$
$$\left. + \int_{t-h}^{t}\left(I_\nu \otimes (e^{\bar{A}(t-\theta)}B_u)\right)\tilde{v}(\theta)d\theta\right)$$
$$- \left((e_1 e_1') \otimes ((F_\alpha - \bar{F})e^{\bar{A}h})\right)\tilde{x}(t-h),$$

where the fact that $(e_1' \otimes I_m)\tilde{v} = 0$ was used to obtain the last equality.

3) Return to the original coordinates: This last stage just follows the steps of the proof of [10, Thm. 3.1].

To derive (10) note that since U is unitary,

$$\|T_{zw}\|_2^2 = \|T_{\tilde{z}w}\|_2^2 = \sum_i \|T_{\tilde{z}_i w}\|_2^2.$$

Set $B_i := ((e_i'U) \otimes I)B_w$ and note that the observability Gramian associated with $(\bar{A}, \bar{C}_z)$ is $X_\alpha + Y$. Then

$$\|T_{\tilde{z}_i w}\|_2^2 = \begin{cases} \text{tr}(B_i(X_\alpha + Y)B_i) & \text{if } i = 1 \\ \text{tr}(B_i(X_\alpha + Y_h)B_i) & \text{otherwise} \end{cases}$$

(cf. [17, Lemma 9]). Hence,

$$\|T_{zw}\|_2^2 = \text{tr}\big(B_w'(U' \otimes I_n)\big((e_1 e_1') \otimes (X_\alpha + Y) + (I_\nu - e_1 e_1')(X_\alpha + Y_h)\big)(U' \otimes I_n)B_w\big)$$
$$= \text{tr}\big(B_w'\big((\mu\mu') \otimes (Y - Y_h) + I_\nu \otimes (X_\alpha + Y)\big)B_w\big),$$

where the last equality follows by the facts that U is unitary and $Ue_1 = \mu$. Straightforward algebra gives then (10). ∎

V. COST OF COORDINATION PER AGENT

In this section we study the performance of each system in (4) under the optimal control law (11). Namely, we quantify the H^2 norm of the closed-loop system $T_{z_i w}$ from the aggregate disturbance w to the local regulated signal z_i defined in (5).

Proposition 5.1: Let $\mathcal{J}_i(h) := \|T_{z_i w}\|_2^2$. Then

$$\mathcal{J}_i(h) = \text{tr}(B_{w_i}'X_\alpha B_{w_i}) + \text{tr}(B_{w_i}'Y_h B_{w_i})$$
$$+ \mu_i^2 \sum_{j=1}^{\nu} \mu_j^2 \text{tr}\big(B_{w_j}'e^{\bar{A}'h}Y e^{\bar{A}h}B_{w_j}\big) \quad (17)$$

and terms containing B_{w_j} reflect the effect of w_j on z_i.

Proof: Omitted because of space limitations. ∎

Two quantities that appear in the right-hand side of (17) are important for understanding properties of $\mathcal{J}_i(h)$. The term $\alpha_i := \text{tr}(B_{w_i}'X_\alpha B_{w_i})$ is the performance of the ith agent in the absence of the coordination constraint (7b). The term

$$\delta_i(h) := \text{tr}(B_{w_i}'Y_h B_{w_i}) = \int_0^h \|(\bar{F} - F_\alpha)e^{\bar{A}\theta}B_{w_i}\|_F^2 d\theta$$

would be the deterioration of the local uncoordinated H^2 performance due to the delay h in the control law $u_i(t) = \bar{F}x_i(t) + v_i(t-h)$ with free v_i. In other words, $\delta_i(h)$ reflects the "H^2 sensitivity" of the generalized plant in Fig. 5 to a loop delay. It is an increasing function of h. Moreover, if $(\bar{F} - F_\alpha)(sI - \bar{A})^{-1}B_{w_i} \neq 0$, $\delta_i(h)$ is strictly increasing for almost all h.

A. Discussion

In the remaining of the section we investigate the effect of the delay time h and the number of agents ν on $\mathcal{J}_i(h)$.

1) Gains and losses due to delay: The effect of the interaction delay on the performance of the ith agent is

$$\mathcal{J}_i(h) - \mathcal{J}_i(0) = \delta_i(h) - \mu_i^2 \sum_{j=1}^{\nu} \mu_j^2 \delta_j(h) \quad (18a)$$
$$= (1 - \mu_i^4)\delta_i(h) - \mu_i^2 \sum_{j \neq i} \mu_j^2 \delta_j(h). \quad (18b)$$

Because $\mu_i < 1$, the sensitivity of z_i to the local w_i, quantified by the first term in (18b), generically increases as h grows. At the same time, the effect of disturbances applied to the other subsystems weakens.

Although the total performance of the agents cannot be improved due to a delay (cf. (10)), it is not necessarily true for individual agents. It may happen that $\mathcal{J}_i(h) < \mathcal{J}_i(0)$, i.e., the presence of a delay in information exchange may actually benefit some agents. Intuitively, these are the agents that experience relatively low level of exogenous disturbances, so they might not even benefit from exchanging information with the other agents. Furthermore, the second term in the right-hand side of (18a) can be thought of as the average delay sensitivity across the agents, scaled by μ_i^2. Thus, $\mathcal{J}_i(h) < \mathcal{J}_i(0)$ if the delay sensitivity of the ith agent is smaller than the scaled average one.

2) Cost of coordination in large groups: The quantity

$$\sigma_i(h) := \mathcal{J}_i(h) - \alpha_i$$

may be interpreted as the *cost of coordination* for the ith agent. It quantifies the deterioration of the local performance due to the need to satisfy the coordination constraint (7b). To see a generic behavior of $\sigma_i(h)$ as a function of ν in the situation when all $\mu_i \to 0$ as $\nu \to \infty$, consider the case of

$$\mu_i = \mu_j = 1/\sqrt{\nu} \quad \text{and} \quad B_{w_i} = B_{w_j}$$

$\forall j = 1, \ldots, \nu$ and assume that $(\bar{F} - F_\alpha)(sI - \bar{A})^{-1} B_{w_i} \neq 0$. The cost of coordination is then

$$\sigma_i(h) = \delta_i(h) + \frac{1}{\nu} \operatorname{tr}\!\big(B'_{w_i} \mathrm{e}^{\bar{A}'h} Y \mathrm{e}^{\bar{A}h} B_{w_i}\big),$$
$$= \delta_i(h) + \frac{1}{\nu} \int_h^\infty \|(\bar{F} - F_\alpha)\mathrm{e}^{\bar{A}\theta} B_{w_i}\|_{\mathrm{F}}^2 \, \mathrm{d}\theta$$

which is a decreasing function of ν. In the delay-free case $\sigma_i(0)$ vanishes as $\nu \to \infty$ (see also [10, §III-C.3]). This is no longer true if $h > 0$, in which case $\sigma_i(h)$ is lowerbounded by $\delta_i(h) > 0$. In fact, the cost of coordination with $h > 0$ is higher than that in the delay-free case for all $\nu \geq 2$. Indeed, the increment of the cost of coordination due to h is

$$\sigma_i(h) - \sigma_i(0) = \mathcal{J}_i(h) - \mathcal{J}_i(0) = \big(1 - \tfrac{1}{\nu}\big)\delta_i(h) > 0$$

and it is a strictly increasing function of ν.

Finally, as $h \to \infty$, the cost of coordination approaches $\delta_i(\infty) = \operatorname{tr}(B'_{w_i} Y B_{w_i})$, which is (expectedly) the cost of satisfying (7b) in the case when no coordination between agents is allowed (so that each agent has to use the feedback gain $\bar{F}$ to satisfy (7b)).

VI. Conclusions

We have studied a large-scale coordination problem, in which a homogeneous group of autonomous agents are coupled through a constraint on their average state. In order to satisfy the coordination requirement, the agents must coordinate their actions over a delayed communication channel.

It has been demonstrated that several key properties of the solution to the delay-free version of the problem, which was derived in [10], extend to the case with delayed information exchange. Namely, the optimal control law is decomposed into a diagonal term complemented by a rank-one component, which is now delayed. While the first term is completely decentralized, the latter can be implemented by a single averaging operation. Moreover, to form the solution, we only need to solve an optimal control problem for a stand-alone agent. This means that the computational effort required to obtain the solution is independent of the number of agents.

Two other properties of the derived optimal solution are worth mentioning. First, adding the delay constraint on the information exchange does not introduce any additional dynamics to the rank-one coordination term of the optimal controller. The only complexity brought about by the delay is the introduction of a dead-time compensation element into the local control law for each agent. Second, we have derived an analytic expression for the performance of each agent under the optimal control law. Unlike the delay-free case, where the cost of satisfying the coordination constraint vanishes as the number of agents increases, the presence of delays renders the cost of coordination lowerbounded by a (generically) nonzero quantity, which grows with the delay.

References

[1] D. D. Šiljak, *Large-Scale Dynamic Systems: Stability and Structure*. NY: North-Holland, 1978.

[2] P. Shah and P. A. Parrilo, "$\mathcal{H}_2$-optimal decentralized control over posets: A state-space solution for state-feedback," *IEEE Trans. Automat. Control*, vol. 58, no. 12, pp. 3084–3096, 2013.

[3] A. Mahajan, N. C. Martins, M. C. Rotkowitz, and S. Yüksel, "Information structures in optimal decentralized control," in *Proc. 51st IEEE Conf. Decision and Control*, Maui, HW, 2012, pp. 1291–1306.

[4] H. S. Witsenhausen, "A counterexample in stochastic optimum control," *SIAM J. Control*, vol. 6, no. 1, pp. 131–147, 1968.

[5] J. N. Tsitsiklis and M. Athans, "On the complexity of decentralized decision making and detection problems," *IEEE Trans. Automat. Control*, vol. 30, no. 5, pp. 440–446, 1985.

[6] T. Tanaka and C. Langbort, "The bounded real lemma for internally positive systems and H-infinity structured static state feedback," *IEEE Trans. Automat. Control*, vol. 56, no. 9, pp. 2218–2223, 2011.

[7] S. Yüksel and T. Başar, *Stochastic Networked Control Systems: Stabilization and Optimization under Information Constraints*. Basel: Birkhäuser, 2013.

[8] M. Hovd and S. Skogestad, "Control of symmetrically interconnected plants," *Automatica*, vol. 30, no. 6, pp. 957–973, 1994.

[9] A. I. Zečević and D. D. Šiljak, "Global low-rank enhancement of decentralized control for large-scale systems," *IEEE Trans. Automat. Control*, vol. 50, no. 5, pp. 740–744, 2005.

[10] D. Madjidian and L. Mirkin, "Distributed control with low-rank coordination," *IEEE Trans. Control of Network Systems*, vol. 1, no. 1, pp. 53–63, 2014.

[11] K. Zhou, J. C. Doyle, and K. Glover, *Robust and Optimal Control*. Englewood Cliffs, NJ: Prentice-Hall, 1995.

[12] J. A. Fax and R. M. Murray, "Information flow and cooperative control of vehicle formations," *IEEE Trans. Automat. Control*, vol. 49, no. 9, pp. 1465–1476, 2004.

[13] J. M. Mendel, *Lessons in Estimation Theory for Signal Processing, Communication, and Control*. Upper Saddle River, NJ: Prentice-Hall, 1995.

[14] D. L. Kleinman, "Optimal control of linear systems with time-delay and observation noise," *IEEE Trans. Automat. Control*, vol. 14, no. 5, pp. 524–527, 1969.

[15] G. Tadmor, "The standard H_∞ problem in systems with a single input delay," *IEEE Trans. Automat. Control*, vol. 45, no. 3, pp. 382–397, 2000.

[16] G. Meinsma and H. Zwart, "On $\mathcal{H}_\infty$ control for dead-time systems," *IEEE Trans. Automat. Control*, vol. 45, no. 2, pp. 272–285, 2000.

[17] L. Mirkin and N. Raskin, "Every stabilizing dead-time controller has an observer-predictor-based structure," *Automatica*, vol. 39, no. 10, pp. 1747–1754, 2003.

[18] L. Mirkin, "On the dead-time compensation from L^1 perspectives," *IEEE Trans. Automat. Control*, vol. 51, no. 6, pp. 1069–1073, 2006.

[19] P. K. S. Tam and J. B. Moore, "Stable realization of fixed-lag smoothing equations for continuous-time signals," *IEEE Trans. Automat. Control*, vol. 19, no. 1, pp. 84–87, 1974.

[20] L. Mirkin, "On the approximation of distributed-delay control laws," *Syst. Control Lett.*, vol. 55, no. 5, pp. 331–342, 2004.

[21] J. R. Partington and P. M. Mäkilä, "Rational approximation of distributed-delay controllers," *Int. J. Control*, vol. 78, no. 16, pp. 1295–1301, 2005.