

Optimal Two Player LQR State Feedback With Varying Delay^{1,2}

Nikolai Matni* Andrew Lamperski** John C. Doyle*

* *Department of Control and Dynamical Systems, California Institute of Technology, Pasadena, CA, 91125 USA (e-mail: {nmatni,doyle}@caltech.edu)*

** *Department of Engineering, University of Cambridge, Cambridge CB2 1PZ, UK (e-mail: a.lamperski@eng.cam.ac.uk)*

Abstract: This paper presents an explicit solution to a two player distributed LQR problem in which communication between controllers occurs across a communication link with varying delay. We extend known dynamic programming methods to accommodate this varying delay, and show that under suitable assumptions, the optimal control actions are linear in their information, and that the resulting controller has piecewise linear dynamics dictated by the current effective delay regime.

1. INTRODUCTION

In the past decade, optimal decentralized controller synthesis has seen an explosion of advances at the theoretical, algorithmic and practical levels. We provide a brief survey of the more directly relevant results to our paper in the following, and refer the reader to the tutorial paper by Mahajan et al. [2012] for a timely presentation of the current state of the art in optimal decentralized control subject to information constraints.

A particular class of decentralized control problems that has received a significant amount of attention is that of optimal \mathcal{H}_2 (or LQG) control subject to delay constraints. In this case, the information constraints can be interpreted as arising from a communication graph, in which edge weights between nodes correspond to the delay required to transmit information between them. For the special case of the one-step delay information sharing pattern, the \mathcal{H}_2 problem was solved in the 1970s using dynamic programming (Sandell and Athans [1974], Kurtaran and Sivan [1974], Yoshikawa [1975]). For more complex delay patterns, sufficient statistics are not easily identified, making extensions beyond the state feedback case (Lamperski and Doyle [2011, 2012]) difficult, although semi-definite programming (SDP) (Rantzer [2006], Gattami [2006]), vectorization (Rotkowitz and Lall [2006]), and spectral factorization (Lamperski and Doyle [2013]) based solutions do exist. It is worth noting that for specific systems, sufficient statistics and a generalized separation principle have been identified and successfully applied, as in the work by Feyzmahdavian et al. [2012]. Furthermore, recent work by

Nayyar et al. [2011, 2013] provides dynamic programming decompositions for the general delayed sharing model.

An underlying assumption in all of the above is that information, albeit delayed, can be transmitted *perfectly* across a communication network with a *fixed* delay. A realistic communication network, however, is subject to data rate limits, quantization, noise and packet drops – all of these issues result in possibly varying delays (due to variable decoding times) and imperfect transmission (due to data rate limits/quantization). The assumption that these delays are fixed necessarily introduces a significant level of conservatism in the control design procedure. In particular, to ensure that the delays under which controllers exchange information do not vary, worst case delay times must be used for control design, sacrificing performance and robustness in the process.

These issues have been addressed by the networked control systems (NCS) community, leading to a plethora of results for channel-in-the loop type problems: see the recent survey by Hespanha et al. [2007], and the references therein. Some of the more relevant results from this field include the work by Gupta et al. [2005,] and Garone et al. [2010], which address optimal LQG control of a single plant over a packet dropping channel. Very few results exist, however, that seek to combine NCS and decentralized optimal control. A notable exception is the work by Chang and Lall [2011], in which an explicit state space solution to a sparsity constrained two-player decentralized LQG state-feedback problem over a TCP erasure channel is solved.

We take a different view from these results, and suppress the underlying details of the communication network, and instead assume that packet drops, noise, and congestion manifest themselves to the controllers as varying delays. In particular, we seek to extend the distributed state-feedback results of Lamperski and Doyle [2011, 2012] and Lamperski and Lessard [2012] to accommodate varying delays. In addition to allowing for communication channels to be more explicitly accounted for in the control design

¹ This research was in part supported by NSF Career award CCF-1350590, and by the NSF, AFOSR, ARPA-E, and the Institute for Collaborative Biotechnologies through grant W911NF-09-0001 from the U.S. Army Research Office. The content does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

² A. Lamperski was supported by a grant from the Whitaker International Scholars Program.

procedure, the ability to accommodate varying delays provides flexibility in the coding design aspect of this problem.

In this paper, we focus on a two plant system in which communication between controllers occurs across a communication link with varying delay. In Matni and Doyle [2013], we solved a special case of this problem by extending the methods used in Lamperski and Doyle [2011] and Lamperski and Doyle [2012]. Here, we use a variant of the dynamic programming methods in Lamperski and Lessard [2012] to accommodate this varying delay, and show that under suitable assumptions, the optimal control actions are linear in their information, and that the resulting controller has piecewise linear dynamics dictated by the current effective delay regime.

This paper is structured as follows: in Section 2 we fix notation, and present the problem to be solved in the paper. Section 3 introduces the concepts of effective delay, partial nestedness (c.f. Ho and Chu [1972]) and a system's information graph (c.f. Lamperski and Lessard [2012]) before presenting our main result. Section 4 derives the optimal control actions and controller, and Section 5 ends with conclusions and directions for future work. Proofs of all intermediary results can be found in the extended report by Matni et al. [2014].

2. PROBLEM FORMULATION

Notation: For a matrix partitioned into blocks

$$M = \begin{bmatrix} M_{11} & \cdots & M_{1N} \\ \vdots & \ddots & \vdots \\ M_{N1} & \cdots & M_{NN} \end{bmatrix}$$

and $s, v \subset \{1, \dots, N\}$, we let $M^{s,v} = (M_{ij})_{i \in s, j \in v}$. For example, $M^{\{3\}\{1,2\}} = [M_{31} \ M_{32}]$. We denote the sequence $x_{t_0}, \dots, x_{t_0+t}$ by $x_{t_0:t_0+t}$, and given the history of a random process $r_{0:t}$, we denote the conditional probability of an event \mathcal{A} occurring given this history by $\mathbb{P}_{r_{0:t}}(\mathcal{A})$. If $\mathcal{Y} = \{y^1, \dots, y^M\}$ is a set of random vectors (possibly of different sizes), we say that $z \in \text{lin}(\mathcal{Y})$ if there exist appropriately sized real matrices C^1, \dots, C^M such that $z = \sum_{i=1}^M C^i y^i$.

The two-player problem: This paper focuses on a two plant system with physical propagation delay of D between plants, and stochastically varying communication delays $d_t^i \in \{0, \dots, D\}$ – to ease notation, we let $d_t := (d_t^1, d_t^2)$. We impose some additional assumptions on the stochastic process d_t in Section 3 such that the infinite horizon solution is well defined.

The dynamics of the sub-system i are then captured by the following difference equation:

$$x_{t+1}^i = A_{ii}x_t^i + A_{ij}x_{t-(D-1)}^j + B_i u_t^i + w_t^i \quad (1)$$

with mutually independent Gaussian initial conditions and noise vectors

$$x_0^i \sim \mathcal{N}(\mu_0^i, \Sigma_0^i), \quad w_t^i \sim \mathcal{N}(0, W_t^i) \quad (2)$$

We may describe the information available to controller i at time t , denoted by \mathcal{I}_t^i , via the following recursion:

$$\begin{aligned} \mathcal{I}_0^i &= \{x_0^i\} \\ \mathcal{I}_{t+1}^i &= \mathcal{I}_t^i \cup \{x_{t+1}^i\} \cup \{x_k^j : 1 \leq k \leq t+1 - d_{t+1}^j\} \end{aligned} \quad (3)$$



Fig. 1. The distributed plant considered in (6), shown here for $D = 4$.

The inputs are then constrained to be of the form

$$u_t^i = \gamma_t^i(\mathcal{I}_t^i) \quad (4)$$

for Borel measurable γ_t^i .

In order to build on the results in Lamperski and Lessard [2012], we model the two plant system as a $D + 1$ node graph, with “dummy delay” nodes introduced to explicitly enforce the propagation delay between plants. Specifically, letting

$$\delta_t^i = \begin{bmatrix} x_{t-i}^1 \\ x_{t-(D-i)}^2 \end{bmatrix}, \quad i = 1, \dots, D-1 \quad (5)$$

where δ^i is the state of the i^{th} dummy node, we obtain the following state space representation for the system

$$x_{t+1} = Ax_t + Bu_t + w_t \quad (6)$$

where, to condense notation, we let

$$x = \begin{bmatrix} x^1 \\ \delta^1 \\ \vdots \\ \delta^{D-1} \\ x^2 \end{bmatrix}, \quad u = \begin{bmatrix} u^1 \\ 0 \\ \vdots \\ 0 \\ u^2 \end{bmatrix}, \quad w = \begin{bmatrix} w^1 \\ 0 \\ \vdots \\ 0 \\ w^2 \end{bmatrix}, \quad (7)$$

and A and B are such that (6) is consistent with (1) and (5). The physical topology of the plant is illustrated in Figure 1.

Problem 2.1. Given the linear time invariant (LTI) system described by (1), (5) and (6), with disturbance statistics (2), minimize the infinite horizon expected cost

$$\lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E} \left[\sum_{t=1}^N x(t)^T Q x(t) + u(t)^T R u(t) \right] \quad (8)$$

subject to the input constraints (4).

The weight matrices are assumed to be partitioned into blocks of appropriate dimension, i.e. $Q = (Q_{ij})$ and $R = (R_{ij})$, conforming to the partitions of x and u . We assume Q to be positive semi-definite, and R to be positive definite, and in order to guarantee existence of the stabilizing solution to the corresponding Riccati equation, we assume (A, B) to be stabilizable and $(Q^{\frac{1}{2}}, A)$ to be detectable.

3. MAIN RESULT

Effective delay: The information constraint sets (3) are defined in such a way that controllers do not forget information that they have already received. This leads to the x^j component of the information set \mathcal{I}_t^i being a function of the *effective* delay seen by the controller, as opposed to the current delay value of the communication channel d_t^j .

Definition 3.1. Let

$$\begin{aligned} e_t^j &:= \min\{d_t^j, d_{t-1}^j + 1, d_{t-2}^j + 2, \dots, \\ &\quad d_{t-(D-2)}^j + (D-2), d_{t-(D-1)}^j + (D-1)\} \end{aligned} \quad (9)$$

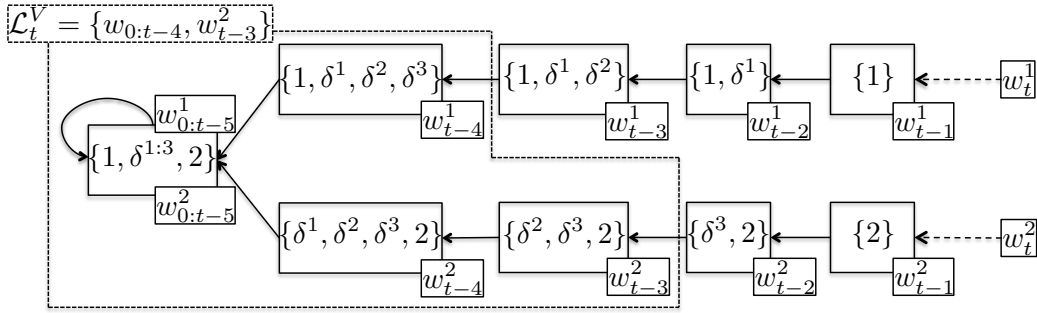


Fig. 2. The information graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, and label sets $\{\mathcal{L}_t^s\}_{s \in \mathcal{V}}$, for system (6), shown here for $D = 4$, and $e_t = (3, 2)$.

be the *effective delay* in transmitting information from controller j to controller i .

Lemma 3.1. The information set available to controller i at time t may be written as

$$\mathcal{I}_t^i = \mathcal{I}_{t-1}^i \cup \{x_t^i\} \cup \mathcal{I}_{t-e_t^j}^j \quad (10)$$

In order to ensure that the infinite horizon solution is well defined, we assume that the stochastic delay process d_t induces an effective delay process such that

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{P}_{d_{0:t}}(e_{t+1}^i \leq d) \quad (11)$$

exists for any integer d .

Partial Nestedness: Here we show that the information constraints (4) and system (6) are *partially nested* (c.f. Ho and Chu [1972]), and hence that the optimal control policies γ_t^i are linear in their information set.

Definition 3.2. A system (6) and information structure (4) is partially nested if, for every admissible policy γ , whenever u_τ^i affects \mathcal{I}_τ^j , then $\mathcal{I}_\tau^i \subset \mathcal{I}_\tau^j$.

Lemma 3.2. (see Ho and Chu [1972]) Given a partially nested information structure, the optimal control law that minimizes a quadratic cost of the form (8) exists, is unique, and is linear.

Using partial nestedness, the following lemma shows that the optimal state and input lie in the linear span of \mathcal{I}_t^i and \mathcal{H}_t , where \mathcal{H}_t is the *noise history* of the system given by

$$\mathcal{H}_t = \{x_0, w_{0:t-1}\} \quad (12)$$

Lemma 3.3. The system (6) and information structure (4) is partially nested, and for any linear controller, we have that

$$x_t^i, u_t^i \in \text{lin}(\mathcal{I}_t^i), \quad x_t, u_t \in \text{lin}(\mathcal{H}_t) \quad (13)$$

Information Graph and Controller Coordinates: Lemma 3.3 indicates that each \mathcal{I}_t^i is a subspace of \mathcal{H}_t : in this section, we exploit this observation to define pairwise independent controller coordinates. An explicit characterization of these subspaces is given in Section 4.

We begin by defining the *information graph*, as in Lamperski and Lessard [2012], associated with system (6) by $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, with

$$\begin{aligned} \mathcal{V} &:= \{\{1\}, \{1, \delta^1\}, \dots, \{1, \delta^1, \dots, \delta^{D-1}\}\} \cup \\ &\quad \{\{2\}, \{\delta^{D-1}, 2\}, \dots, \{\delta^1, \dots, \delta^{D-1}, 2\}\} \cup V \\ \mathcal{E} &:= \{(r, s) \in \mathcal{V} \times \mathcal{V} : |s| = |r| + 1\} \cup \{(V, V)\} \end{aligned} \quad (14)$$

where $V := \{1, \delta^1, \dots, \delta^{D-1}, 2\}$. For the case of $D = 4$, the graph \mathcal{G} is illustrated in Figure 2.

Before proceeding, we define the following sets, which will help us state the main result. Let

$$\begin{aligned} v_t^{i,+} &:= \{s \in \mathcal{V} \setminus V \mid i \in s, |s| \geq e_t^i\} \\ v_t^{i,++} &:= \{s \in \mathcal{V} \setminus V \mid i \in s, |s| > e_t^i\} \end{aligned} \quad (15)$$

and similarly define $v_t^{i,-}$ and v_t^{i--} as in (15), but with the (strict) inequality reversed.

Theorem 3.1. Consider Problem 2.1, and let $\mathcal{G}(\mathcal{V}, \mathcal{E})$ be the associated information graph. Let

$$\begin{aligned} X^V &= Q + A^\top X^V A + A^\top X^V B K^V \\ K^V &:= -(R + B^\top X^V B)^{-1} B^\top A, \end{aligned} \quad (16)$$

be the stabilizing solution to the discrete algebraic Riccati equation, and the centralized LQR gain, respectively. Now, assume that X^s is given, and let $r \neq s \in \mathcal{V}$ be the unique node such that $(r, s) \in \mathcal{E}$. Define the matrices

$$\begin{aligned} \Lambda^r &= Q^{rr} + p^r (A^{Vr})^\top X^V A^{Vr} + q^r (A^{sr})^\top X^s A^{sr} \\ \Psi^r &= R^{rr} + p^r (B^{Vr})^\top X^V B^{Vr} + q^r (B^{sr})^\top X^s B^{sr} \\ \Omega^r &= p^r (A^{Vr})^\top X^V B^{Vr} + q^r (A^{sr})^\top X^s B^{sr} \\ X^r &= \Lambda^r + \Omega^r K^r \\ K^r &= -(\Psi^r)^{-1} (\Omega^r)^\top \end{aligned} \quad (17)$$

where p^r is given by

$$p^r := \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{P}_{d_{0:t}}(r \in v_{t+1}^{i,++}) \quad (18)$$

and $q^r = 1 - p^r$.

The optimal control decisions then satisfy

$$\begin{aligned} \zeta_{t+1}^V &= A \zeta_t^V + B \varphi_t^V + \\ &\quad \sum_{i=1}^2 \sum_{r \in v_{t+1}^{i,++}} (A^{Vr} \zeta_t^r + B^{Vr} \varphi_t^r) \\ \zeta_{t+1}^s &= \begin{cases} A^{sr} \zeta_t^r + B^{sr} \varphi_t^r & \text{if } s \in \cup_i v_{t+1}^{i,-}, (r, s) \in \mathcal{E} \\ 0 & \text{otherwise} \end{cases} \\ \zeta_{t+1}^i &= w_t^i, \quad \zeta_0^i = x_0^i \\ u_t^i &= \varphi_t^i + \sum_{s \in v_{t+1}^{i,-}} I^{V,s} \varphi_t^s, \quad \varphi_t^r = K^r \zeta_t^r \end{aligned} \quad (19)$$

and the corresponding infinite horizon expected cost is

$$\sum_{i=1}^2 \text{Trace}(X^{(i)} W^i) \quad (20)$$

Proof. See Section 4.

Remark 3.1. Notice that the global action taken based on ζ^V must be taken simultaneously by *both* players. In other words, it is assumed that an acknowledgment mechanism is in place such that e_t is known to both players; relaxing this assumption will be the subject of future work.

Remark 3.2. The probabilities p^r and q^r can be computed directly if we assume the $\{d_t\}$ to be independently and identically distributed. In this case, e_t^j evolves according to an irreducible and aperiodic Markov chain with transition probability matrix computable directly from the definition of effective delay and the pmf of d_t . As such, p^r and q^r can be computed from the chain's stationary distribution, which is guaranteed to exist. Future work will explore what additional distributions on d_t will lead to closed form expressions for p^r and q^r . Failing the existence of closed form expressions for these asymptotic distributions, computing estimates via simulation should be a feasible option for many interesting delay processes.

4. CONTROLLER DERIVATION

Controller States and Decoupled Dynamics: As mentioned previously, each \mathcal{I}_t^i is a subspace of \mathcal{H}_t : in this section, we aim to explicitly characterize these subspaces by assigning label sets $\{\mathcal{L}_{0:t}^s\}_{s \in \mathcal{V}}$ to the graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ as defined by (14). In particular, they are defined recursively as:

$$\begin{aligned} \mathcal{L}_0^s &= \emptyset, \text{ for } |s| > 1, & \mathcal{L}_0^i &= \{x_0^i\}, & \mathcal{L}_{t+1}^i &= \{w_t^i\} \\ \mathcal{L}_{t+1}^s &= \mathcal{L}_t^r, \text{ for } (r, s) \in \mathcal{E}, & 1 < |s| < D+1 \\ \mathcal{L}_{t+1}^V &= \mathcal{L}_t^V \cup_i \cup_{s \in v_t^{i,+}} \mathcal{L}_t^s \end{aligned} \tag{21}$$

where we have let \cup_i denote $\cup_{i=1}^2$ to lighten notational burden. An example of these label sets for the case of $D = 4$ is illustrated in Figure 2.

Before delving in to the technical justification for these label sets, we provide some intuition. The information graph \mathcal{G} characterizes how the effect of noise terms spread through the system, and labels are introduced as a means of explicitly tracking this spreading. As can be seen in Figure 2, for each $(r, s) \in \mathcal{E}$, with $|r| < D+1$, we have that $|s| = |r| + 1$, and additionally, that $|s|$ measures exactly how delayed the information in the label set is. We also see that the graph is naturally divided into two *disjoint* branches, with each branch corresponding to information about a specific plant. Finally, the label corresponding to the root node V can be interpreted as the information available to both controllers – this is reflected by its explicit dependence on the effective delay e_t^i .

Remark 4.1. Note that in contrast to Lamperski and Lessard [2012], the label sets as defined will in general *not* be disjoint. However, as will be made explicit in Lemma 4.2, an effective delay dependent subset of the label sets will indeed form a partition (i.e. a pairwise disjoint cover) of the noise history.

We may now characterize the subspaces of \mathcal{H}_t that are associated with each \mathcal{I}_t^i . This characterization will be shown to depend on the effective delay e_t^j seen at node i , and will lead to an intuitive partitioning of both the state and the control input.

We begin by pointing out the following useful facts that will be used repeatedly in the derivation to come

Lemma 4.1. Let $v_t^{i,*}$, $*$ $\in \{-, --\}$, be given as in (15). Then, for a fixed i , we have that

$$\cup_{s \in v_t^{i,-}} \mathcal{L}_{t+1}^s = \cup_{r \in v_t^{i,-}} \mathcal{L}_t^r \cup \mathcal{L}_{t+1}^i, \tag{22}$$

and for integers $a, b \in \{0, \dots, D-1\}$

$$\cup_{a < |s| \leq b+1} \mathcal{L}_{t+1}^s = \cup_{a \leq |r| \leq b} \mathcal{L}_t^r \tag{23}$$

Proof. Follows immediately by applying the recursion rules (21) and the fact that for each $(r, s) \in \mathcal{E}$, with $|r| < D+1$, we have that $|s| = |r| + 1$.

Lemma 4.2. Consider the information graph \mathcal{G} as defined in equation (14), and the label sets defined as in (21). We then have that

- (i) For all $t \geq 0$, a subset of the labels form a partition of the noise history. In particular, we have that

$$\mathcal{H}_t = \mathcal{L}_t^V \cup_i \cup_{s \in v_t^{i,-}} \mathcal{L}_t^s \tag{24}$$

where the union is disjoint, i.e. $\mathcal{L}_t^V \cap \mathcal{L}_t^s = \emptyset$ if $s \in v_t^{i,-}$, and $\mathcal{L}_t^s \cap \mathcal{L}_t^{s'} = \emptyset$ for any $s \neq s'$, $s, s' \in \cup_i v_t^{i,-}$.

- (ii) For $i = 1, 2$

$$\text{lin}(\mathcal{I}_t^i) = \text{lin}(\mathcal{L}_t^V \cup_{s \in v_t^{i,-}} \mathcal{L}_t^s). \tag{25}$$

Remark 4.2. Although the proof of this result is notationally cumbersome, it is mainly an exercise in bookkeeping. The idea is illustrated in Figure 2: labels for nodes $v \neq V$ track the propagation of a disturbance through the plant, whereas the label for V selects those labels corresponding to globally available information, as dictated by the effective delay.

With the previous lemmas at our disposal, we may now write

$$\begin{aligned} x_t &= \zeta_t^V + \sum_{i=1}^2 \sum_{s \in v_t^{i,-}} I^{V,s} \zeta_t^s \\ u_t &= \varphi_t^V + \sum_{i=1}^2 \sum_{s \in v_t^{i,-}} I^{V,s} \varphi_t^s \end{aligned} \tag{26}$$

where each $\zeta_t^s, \varphi_t^s \in \text{lin}(\mathcal{L}_t^s)$. We may accordingly derive update dynamics for these state and control components.

Lemma 4.3. If the control components are such that $\varphi_s^t \in \text{lin}(\mathcal{L}_t^s)$, then the state components $\{\zeta_t^s\}$ satisfy the following update dynamics

$$\begin{aligned} \zeta_{t+1}^V &= A \zeta_t^V + B \varphi_t^V + \\ &\quad \sum_{i=1}^2 \sum_{r \in v_t^{i,++}} (A^{Vr} \zeta_t^r + B^{Vr} \varphi_t^r) \\ \zeta_{t+1}^s &= \begin{cases} A^{sr} \zeta_t^r + B^{sr} \varphi_t^r & \text{if } s \in \cup_i v_t^{i,-}, (r, s) \in \mathcal{E} \\ 0 & \text{otherwise} \end{cases} \\ \zeta_{t+1}^i &= w_t^i, \quad \zeta_0^i = x_0^i \end{aligned} \tag{27}$$

In particular, notice that the dynamics (27) imply $\zeta_t^s = 0$ for all $s \in \cup_i v_t^{i,++}$, allowing us to rewrite the decomposition for x_t as

$$x_t = \sum_{s \in \mathcal{V}} I^{Vs} \zeta_t^s, \tag{28}$$

where we have simply added the zero valued state components to the expression in (26).

We now have all of the elements required to solve for the optimal control law via dynamic programming.

Finite Horizon Dynamic Programming Solution: Let $\gamma_t = \{\gamma_t^s\}_{s \in \mathcal{V}}$ be the set of policies at time t . By Lemma 3.3, we may assume the γ_t^s to be linear. Define the cost-to-go

$$V_t(\gamma_{0:t-1}) = \min_{\gamma_{t:T-1}} \mathbb{E}^{\gamma \times d} \left(\sum_{k=t}^{T-1} x_k^\top Q x_k + u_k^\top R u_k + x_T^\top Q_T x_T \right) \quad (29)$$

where the expectation is taken with respect to the joint probability measure on $(x_{t:T}, u_{t:T-1}) \times (d_{t:T-1})$ induced by the choice of $\gamma = \gamma_{0:T-1}$ (note that the d_t component is assumed to be independent of the policy choice).

Remark 4.3. Following Nayyar et al. [2013], we adopt the common information formalism and define our cost-to-go function in terms of the control policy γ to be chosen by a ‘‘centralized coordinator.’’ Noting that these policies can in fact be computed off-line and in a centralized manner (it is only their *implementation* that requires measurement of the state components $\{\zeta^s\}$), this in effect reduces the dynamic programming argument to a standard full-information setting.

Via the dynamic programming principle, we may iterate the minimizations and write a recursive formulation for the cost-to-go:

$$V_t(\gamma_{0:t-1}) = \min_{\gamma_{t:T-1}} \mathbb{E}^{\gamma \times d} (x_t^\top Q x_t + u_t^\top R u_t + V_{t+1}(\gamma_{0:t-1}, \gamma_t)). \quad (30)$$

We begin with the terminal time-step, T , and use the decomposition (28) to obtain

$$V_T(\gamma_{0:T-1}) = \mathbb{E}^{\gamma \times d} (x_T^\top Q_T x_T) = \mathbb{E}^\gamma \sum_{s \in \mathcal{V}} (\zeta_T^s)^\top Q_T^{ss} (\zeta_T^s), \quad (31)$$

where in the last step we have used the pairwise independence of the coordinates ζ_T^s . By induction, we shall show that the value function, for some $t \geq 0$, always takes the form

$$V_{t+1}(\gamma_{0:t}) = \mathbb{E}^\gamma \sum_{s \in \mathcal{V}} ((\zeta_{t+1}^s)^\top X_{t+1}^s (\zeta_{t+1}^s) + c_{t+1}) \quad (32)$$

where $\{X_{t+1}^s\}_{s \in \mathcal{V}}$ is a set of matrices and c_{t+1} is a scalar. We now solve for $V_t(\gamma_{0:t-1})$ via the recursion (30). Given e_t , apply (28) and the independence result to write

$$V_t(\gamma_{0:t-1}) = \min_{\gamma_t} \mathbb{E}^{\gamma \times d} \left(\sum_{s \in \mathcal{V}} (\zeta_t^s)^\top Q^{ss} (\zeta_t^s) + (\varphi_t^s)^\top R^{ss} (\varphi_t^s) + \sum_{s \in \mathcal{V}} (\zeta_{t+1}^s)^\top X_{t+1}^s (\zeta_{t+1}^s) + c_{t+1} \right) \quad (33)$$

We now substitute the update equations (27), average over d_{t+1} and use independence to obtain

$$V_t(\gamma_{0:t-1}) = \min_{\gamma_t} \mathbb{E}^\gamma \left(\sum_{r \in \mathcal{V}} \begin{bmatrix} \zeta_t^r \\ \varphi_t^r \end{bmatrix}^\top \Gamma_t^r \begin{bmatrix} \zeta_t^r \\ \varphi_t^r \end{bmatrix} + c_t \right) \quad (34)$$

where $\Gamma_{0:T-1}^r$ and $c_{0:T-1}$ are given by:

$$\Gamma_t^r = \begin{bmatrix} Q^{rr} & 0 \\ 0 & R^{rr} \end{bmatrix} + \mathbb{P}_{d_{0:t}}(r \in v_{t+1}^{i,++}) [A^{Vr} \ B^{Vr}]^\top X_{t+1}^V [A^{Vr} \ B^{Vr}] + \mathbb{P}_{d_{0:t}}(r \in v_{t+1}^{i,-}) [A^{sr} \ B^{sr}]^\top X_{t+1}^s [A^{sr} \ B^{sr}] \quad (35)$$

$$c_t = c_{t+1} + \sum_{i=1}^2 \text{Trace} \left(X_{t+1}^{\{i\}} W^i \right). \quad (36)$$

The terminal conditions are $c_T = 0$ and $\Gamma^r = Q_T^{rr}$, and s is the unique node such that $(r, s) \in \mathcal{E}$.

Let $p_t^r := \mathbb{P}_{d_{0:t}}(r \in v_{t+1}^{i,++})$ and $q_t^r := \mathbb{P}_{d_{0:t}}(r \in v_{t+1}^{i,-})$, and introduce the following matrices:

$$\begin{aligned} \Lambda_{t+1}^r &= Q^{rr} + p_t^r (A^{Vr})^\top X_{t+1}^V A^{Vr} + q_t^r (A^{sr})^\top X_{t+1}^s A^{sr} \\ \Psi_{t+1}^r &= R^{rr} + p_t^r (B^{Vr})^\top X_{t+1}^V B^{Vr} + q_t^r (B^{sr})^\top X_{t+1}^s B^{sr} \\ \Omega_{t+1}^r &= p_t^r (A^{Vr})^\top X_{t+1}^V B^{Vr} + q_t^r (A^{sr})^\top X_{t+1}^s B^{sr} \end{aligned} \quad (37)$$

Then each expression of the sum in (34) can be written as $(\zeta_t^r)^\top \Lambda_{t+1}^r (\zeta_t^r) + (\varphi_t^r)^\top \Psi_{t+1}^r (\varphi_t^r) + 2(\zeta_t^r)^\top \Omega_{t+1}^r (\varphi_t^r)$. (38)

Due to the definitions of $\{\zeta^r\}$ and $\{\varphi^r\}$, it is clear that the terms (38) are pairwise independent and hence can be optimized independently. Removing the information constraints, and optimizing over φ_t^r , we see that the optimal action is given by

$$\varphi_t^r = -(\Psi_{t+1}^r)^{-1} (\Omega_{t+1}^r)^\top \zeta_t^r \quad (39)$$

which, by construction, satisfies the information constraints \mathcal{I}_t^i . Substituting this solution back in to (38), we see that the matrices X_t^r must satisfy

$$\begin{aligned} X_t^r &= \Lambda_{t+1}^r + \Omega_{t+1}^r K_t^r \\ K_t^r &:= -(\Psi_{t+1}^r)^{-1} (\Omega_{t+1}^r)^\top \end{aligned} \quad (40)$$

The finite horizon optimal cost is then given by

$$\begin{aligned} V_0 &= \mathbb{E} \sum_{i=1}^2 (x_0^i)^\top X^{\{i\}} (x_0^i) + c_0 \\ &= \mathbb{E} \sum_{i=1}^2 (\mu_0^i)^\top X_0^{\{i\}} (\mu_0^i) + \text{Trace} \left(X_0^{\{i\}} \Sigma_0^i \right) + c_0 \end{aligned} \quad (41)$$

where c_0 can be computed according to (36) beginning with terminal conditions $c_T = 0$.

Infinite Horizon Solution: In order to determine the infinite horizon solution, we first notice that for $r = V$, $p_t^V = 1$, $q_t^V = 0$ and that the recursions (40) for $r = V$ are then simply given by

$$\begin{aligned} X_t^V &= Q + A^\top X_{t+1}^V A + A^\top X_{t+1}^V B K_t^V \\ K_t^V &:= (R + B^\top X_{t+1}^V B)^{-1} B^\top A, \end{aligned} \quad (42)$$

that is to say the standard discrete algebraic Riccati recursion/gain. By assumption, we have that $(X_t^V, K_t^V) \rightarrow (X^V, K^V)$, where X^V and K^V are, respectively, the stabilizing solution the discrete algebraic riccati equation, and the centralized LQR gain.

Now assume that X_t^s is defined, and let $r \neq s \in \mathcal{V}$ be the unique node such that $(r, s) \in \mathcal{E}$. Much as in the finite horizon case, define the following matrices:

$$\begin{aligned} \Lambda^r &= Q^{rr} + p^r (A^{Vr})^\top X^V A^{Vr} + q^r (A^{sr})^\top X^s A^{sr} \\ \Psi^r &= R^{rr} + p^r (B^{Vr})^\top X^V B^{Vr} + q^r (B^{sr})^\top X^s B^{sr} \\ \Omega^r &= p^r (A^{Vr})^\top X^V B^{Vr} + q^r (A^{sr})^\top X^s B^{sr} \end{aligned} \quad (43)$$

where we have let

$$(p^r, q^r) = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} (p_t^r, q_t^r). \quad (44)$$

Note that these limits are well defined by the assumption (11).

We then have that

$$X^r = \Lambda^r + \Omega^r K^r, \quad K^r := -(\Psi^r)^{-1} (\Omega^r)^\top. \quad (45)$$

What remains to be computed is the infinite horizon average cost, which is given by (ignoring without loss the cost incurred by the uncertainty in the initial conditions)

$$\frac{1}{N} \sum_{t=1}^N \sum_{i=1}^2 \text{Trace} \left(X_t^{\{i\}} W^i \right) \xrightarrow{N \rightarrow \infty} \sum_{i=1}^2 \text{Trace} \left(X^{\{i\}} W^i \right) \quad (46)$$

5. CONCLUSION

This paper presented extensions of a Riccati-based solution to a distributed control problem with communication delays – in particular, we now allow the communication delays to vary, but impose that they preserve partial nestedness. It was seen that the varying delay pattern induces piecewise linear dynamics in the state of the resulting optimal controller, with changes in dynamics dictated by the current *effective* delay regime.

Future work will be to extend the results to systems with several players and more general delay patterns, and to remove the assumption of strong connectedness, much as was done in Lamperski and Lessard [2012] for the case of constant delays. We will also seek to identify conditions on the delay process d_t such that assumption (11) holds. Additionally, we will explore the setting in which the global delay regime is not known.

REFERENCES

- Chung-Ching Chang and Sanjay Lall. Synthesis for optimal two-player decentralized control over tcp erasure channels with state feedback. In *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*, pages 3824–3829. IEEE, 2011.
- Hamid Reza Feyzmahdavian, Ather Gattami, and Mikael Johansson. Distributed output-feedback LQG control with delayed information sharing. In *IFAC Workshop on Distributed Estimation and Control of Networked Systems*, volume 3, pages 192–197, 2012.
- Emanuele Garone, Bruno Sinopoli, and Alessandro Casavola. LQG control over lossy TCP-like networks with probabilistic packet acknowledgements. *International Journal of Systems, Control and Communications*, 2(1):55–81, 2010.
- A. Gattami. Generalized linear quadratic control theory. In *Decision and Control, 2006 45th IEEE Conference on*, pages 1510–1514, dec. 2006.
- V. Gupta, D. Spanos, B. Hassibi, and R.M. Murray. On LQG control across a stochastic packet-dropping link. In *American Control Conference, 2005. Proceedings of the 2005*, pages 360–365, 2005.
- J.P. Hespanha, P. Naghshtabrizi, and Yonggang Xu. A survey of recent results in networked control systems. *Proceedings of the IEEE*, 95(1):138–162, 2007.
- Yu-Chi Ho and K.-C. Chu. Team decision theory and information structures in optimal control problems—part i. *Automatic Control, IEEE Transactions on*, 17(1):15–22, 1972.
- B.-Z. Kurtaran and R. Sivan. Linear-quadratic-gaussian control with one-step-delay sharing pattern. *Automatic Control, IEEE Transactions on*, 19(5):571–574, oct 1974.
- Andrew Lamperski and John C. Doyle. On the structure of state-feedback LQG controllers for distributed systems with communication delays. In *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*, pages 6901–6906, dec. 2011.
- Andrew Lamperski and John C. Doyle. Dynamic programming solutions for decentralized state-feedback LQG problems with communication delays. In *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*, jun. 2012.
- Andrew Lamperski and John C. Doyle. Output feedback \mathcal{H}_2 model matching for decentralized systems with delays. In *2013 IEEE American Control Conference.*, sept. 2013.
- Andrew Lamperski and Laurent Lessard. Optimal state-feedback control under sparsity and delay constraints. In *3rd IFAC Workshop on Distributed Estimation and Control in Networked Systems*, pages 204–209, 2012.
- A. Mahajan, N.C. Martins, M.C. Rotkowitz, and S. Yucsel. Information structures in optimal decentralized control. In *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*, pages 1291–1306, 2012.
- Nikolai Matni and John C. Doyle. Optimal distributed LQG state feedback with varying communication delay. In *IEEE Conf. on Decision and Control, 2013*, 2013.
- Nikolai Matni, Andrew Lamperski, and John C. Doyle. Optimal two player LQR state feedback with varying delay. *CoRR*, arXiv:1403.7790, 2014. URL <http://arxiv.org/abs/1403.7790>.
- A. Nayyar, A. Mahajan, and D. Teneketzis. Optimal control strategies in delayed sharing information structures. *Automatic Control, IEEE Transactions on*, 56(7):1606–1620, july 2011.
- A. Nayyar, A. Mahajan, and D. Teneketzis. Decentralized stochastic control with partial history sharing: A common information approach. *Automatic Control, IEEE Transactions on*, 58(7):1644–1658, 2013.
- A. Rantzer. A separation principle for distributed control. In *Decision and Control, 2006 45th IEEE Conference on*, pages 3609–3613, dec. 2006.
- Michael Rotkowitz and Sanjay Lall. A characterization of convex problems in decentralized control. *Automatic Control, IEEE Transactions on*, 51(2):274–286, 2006.
- Jr. Sandell, N. and M. Athans. Solution of some non-classical LQG stochastic decision problems. *Automatic Control, IEEE Transactions on*, 19(2):108–116, apr 1974.
- T. Yoshikawa. Dynamic programming approach to decentralized stochastic control problems. *Automatic Control, IEEE Transactions on*, 20(6):796–797, dec 1975.