Optimal Distributed LQG State Feedback With Varying Communication Delay

Nikolai Matni and John C. Doyle

***Abstract*— This paper presents an explicit solution to a two player distributed LQG problem in which communication between controllers occurs across a communication link with varying delay. We extend known dynamic programming methods to accommodate this time-varying delay, and show that the resulting optimal controller has a piecewise linear structure, with control actions dictated by the current and next delay regime. By treating the next delay regime as a disturbance in the dynamic programming argument, we then derive the optimal causal controller.**

I. INTRODUCTION

Decentralized control problems arise when several decision makers, or controllers, need to determine their actions based only on a subset of the total information available about the system. These types of problems arise in areas as diverse as physiology, economics, and the power grid.

In the past decade, this field 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 excellent tutorial paper [1] for a thorough and 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 over the past few decades 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 [2], [3], [4]. For more complex delay patterns, the separation principle fails, making extensions beyond the state feedback case [5], [6] difficult, although semi-definite programming (SDP) [7], [8], vectorization [9], and spectral factorization [10] 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 [11]. Furthermore, recent work [12] provides two 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. For example, [13], [14], [15], [16], [17] identify minimum bit-rate conditions sufficient to stabilize a system in the almost sure sense through feedback control over a data-rate constrained channel. In [18], anytime capacity was identified as the correct metric for additionally ensuring moment stability. Performance and robustness results also exist, such as elegant extensions of the well known Bode integral formula to quantify the effect of channels in the loop [19] on performance limits.

Closer in spirit to the results in this paper, [20], [21], [16], [22], [23] address the design of stabilizing feedback controllers subject to either varying communication delay or packet drop outs. Closer still is the work by Gupta et. al. [24], which addresses optimal LQG control of a single plant over a packet dropping channel. We note that the previous references are far from exhaustive, and we refer the interested reader to the excellent survey paper [25] for a more complete listing.

What is lacking, however, is a combination of NCS theory with decentralized optimal control. We believe that there are two major classes of problems in this emerging field: (i) the design of communication networks well suited for decentralized optimal control, and (ii) explicitly accounting for realistic communication channels between plants. In [26], we make a first step at addressing the former issue, whereas this paper focusses on the latter.

We conjecture that due to the relatively high bandwidth of channels, and the aforementioned well established NCS theory, data rates and quantization will not be limiting factors in control performance, but rather that noise, packet

N. Matni and J. C. Doyle are with the Department of Control and Dynamical Systems, California Institute of Technology, Pasadena, CA. nmatni@caltech.edu.

This research was in part supported by NSF NetSE, AFOSR, the Institute for Collaborative Biotechnologies through grant W911NF-09-0001 from the U.S. Army Research Office, and from MURIs "Scalable, Data-Driven, and Provably-Correct Analysis of Networks" (ONR) and "Tools for the Analysis and Design of Complex Multi-Scale Networks" (ARO). The content does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

drops and the ensuing varying *delays* will be of paramount importance. This notion is the driving motivation behind the problem addressed in this paper, where we seek to extend the distributed state-feedback results in [5], [6] to accommodate varying delays. In addition to allowing for communication channels to be more explicitly accounted for in the control design procedure, the ability to accommodate varying delays provides flexibility in the coding design aspect of this problem – we are currently exploring the application of deadline based coding schemes [27], initially designed for real-time video streaming, to optimal decentralized control.

In this paper, we focus on a two plant system in which communication between controllers occurs across a communication link with time-varying delay. We extend the dynamic programming methods in [6] to accommodate this varying delay, and show that under suitable assumptions, the resulting optimal controller has a piecewise linear structure. In particular, within each linear mode, the control policy is determined by how information is shared at the current and *next* time step. We then modify the dynamic programming argument to treat the next delay regime as a disturbance, and derive the optimal causal controller.

This paper is structured as follows: in Section II we fix notation, and outline both the general problem we are considering, as well as introduce the specific one to be solved in the paper. Section III surveys relevant results from [6] and provides novel extensions to accommodate the varying delay. Section IV presents the optimal controller, comments on its structure, and demonstrates its efficacy through simulations. Section V presents the dynamic programming argument used to derive the controller, and Section VI provides conclusions and directions for future work.

II. PROBLEM FORMULATION

We begin by describing the general problem of interest, and then specialize the formulation to the particular case to be addressed in this paper.

Notation: For a matrix partitioned into blocks

$$M = \begin{bmatrix} M_{11} & M_{12} & M_{13} \\ M_{21} & M_{22} & M_{23} \\ M_{31} & M_{32} & M_{33} \end{bmatrix}$$

and $s, v \subset \{1, 2, 3\}$, we let $M^{s,v} = (M_{ij})_{i \in s, j \in v}$. For example

$$M^{\{1,2,3\},\{1,2\}} = \begin{bmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \\ M_{31} & M_{32} \end{bmatrix}.$$

We also denote the sequence $x(t_0), ..., x(t_0 + t)$ by $x(t_0 : t_0 + t)$. For a random variable x, we denote by $\mathbb{E}[x]$ its expectation, and for an event $\mathcal{A}$, we denote by $\mathbb{P}(\mathcal{A})$ its probability, and by $\mathbb{1}_{\{\mathcal{A}\}}$ its indicator function, such that $\mathbb{E}[\mathbb{1}_{\{\mathcal{A}\}}] = \mathbb{P}(\mathcal{A})$.

A. General problem

We consider a cyber-physical system comprised of n linear time invariant (LTI) sub-systems, which interact with each other according to two overlaid, but distinct, topologies: (1) a physical interaction graph and (2) a communication graph. To encode these interactions, we define the possibly time varying cyber-physical interaction graph $\mathcal{G} = (\mathcal{X}, E^p, E^c)$. We denote by $i \in \mathcal{X}$ the i^{th} node in the graph, and by x_i the state of the corresponding sub-system. We assume that each plant $i \in \mathcal{X}$ has its own control input u_i and centered Gaussian process noise w_i (satisfying $\mathbb{E}[w_i w_i^T] = W_i$ and $\mathbb{E}[w_i w_j^T] = 0 \ \forall i \neq j$), and that plants physically interact with each other according to E^p. In particular, an edge $e^p_{ij} \in E^p$ encodes the delay with which the state x_j directly affects x_i, and is non-zero if and only if sub-systems i and j share a direct coupling through their dynamics. This allows for the dynamics of each sub-plant to be described by the following difference equation

$$\begin{aligned} x_i(t+1) = A_{ii} x_i(t) + \textstyle\sum_{e^p_{ij} \neq 0} A_{ij} x_j(t - (e^p_{ij} - 1)) \\ + B_i u_i(t) + w_i(t) \end{aligned}$$

with initial conditions $x_i(0) = 0$.

Thus the distributed nature of the dynamics of the plant are captured by E^p. Finally, we define P_{ij} to be the minimum weight path in E^p from plant j to plant i. This quantity measures how long it takes for the state at plant j to affect the state at plant i, even if i and j are not direct neighbors; i.e. the propagation delay from j to i.

Similarly, E^c defines the delay with which information can be shared between neighboring controllers, and therefore imposes the information constraints on the control inputs. In particular, an edge $e^c_{ij} \in E^c$ encodes the delay with which control input u_i can use state x_j, and is non-zero if and only if sub-systems i and j share a direct link between controllers. We assume that E^c is a strongly connected graph, and letting d_{ij} be the minimum weight path from node j to i (note d_{ii} =0), we have that u_i must be of the form

$$u_i(t) = \gamma_i(x_1(0 : t - d_{i1}), \ldots, x_n(0 : t - d_{in}))$$

for some Borel measurable function γ_i.

In our setting, we assume $\mathcal{X}$ and E^p are fixed, but allow E^c to be time-varying – i.e. $E^c = E^c(t)$. We impose only one additional assumption on $E^c(t)$, namely that for all $t \geq 0$, we have that $d_{ij}(t) \leq P_{ij}$ for all $i, j \in \mathcal{X}$; i.e. that information is transmitted at least as fast as dynamics propagate through the plant, regardless of the delay regime. This condition is trivially satisfied if for all i, j such that $e^p_{ij} \neq 0$, we have that $e^c_{ij}(t) \in \{1, \ldots, e^p_{ij}\}$ for all t. These conditions ensure that the information constraints are *partially nested*, which implies that the optimal control inputs are linear in the associated information [28]. This is a restatement of the now familiar notion that in order to preserve the linearity of the optimal control input, there must not be any incentive to signal through the plant, and is strongly linked to notions of quadratic invariance [9] and poset causality [29].

Defining $x = (x_i)_{i \in \mathcal{X}}$ and $u = (u_i)_{i \in \mathcal{X}}$, the control

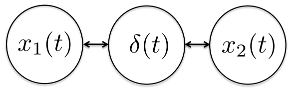

Fig. 1: The distributed plant considered in (3). A dummy node $\delta(t) = x(t-1)$ is introduced to maintain compatibility with the results in [6], making explicit the propagation delay of $P_{12} = P_{21} = 2$ between plants.

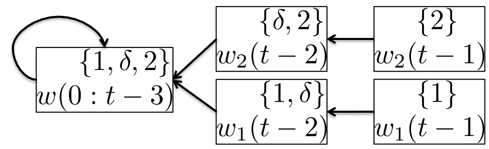

Fig. 2: Information hierarchy graph associated with system (3). The labeling of the nodes indicates the information available to to all players in the node, but none of the other players, under the worst case delay regime $d(t) = (2, 2)$. We associate with each node $v \in \mathcal{V}$ a state and control input component, ζ_v and φ_v, respectively – by construction they are pairwise independent.

problem then becomes to minimize the average stage cost

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

subject to the system's inputs respecting the communication constraints dictated by $E_c(t)$. The weight matrices are assumed to be partitioned into blocks of appropriate dimension (i.e. $Q = (Q_{ij})_{i,j\in\mathcal{X}}$ and $R = (R_{ij})_{i,j\in\mathcal{X}}$), conforming to the partitions of x and u. We assume Q to be positive semi-definite and R to be positive definite.

B. The two-player problem

Although a general problem was defined, this paper focuses on a two plant system with physical propagation delay of $P_{ij} = 2$ between plants, and stochastic time-varying communication delays $d_{ij}(t) \in \{1, 2\}$, where each $d_{ij}(t)$ is assumed to be independent and identically distributed according to some probability mass function (pmf) $\{p, 1-p\}$, $p \in [0, 1]$. To ease notation, we let $d(t) := (d_{12}(t), d_{21}(t))$. The assumption of identical distributions across time is introduced for convenience, and can easily be removed.

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

$$\begin{aligned} x_1(t+1) &= A_{11}x_1(t) + A_{12}x_2(t-1) \\ &\quad + B_1u_1(t) + w_1(t) \\ x_2(t+1) &= A_{21}x_1(t-1) + A_{22}x_2(t) \\ &\quad + B_2u_2(t) + w_2(t) \end{aligned} \tag{1}$$

with initial conditions $x_1(0) = x_2(0) = 0$, and input constraints given by

$$\begin{aligned} u_1(t) &= \gamma_1(x_1(0:t), x_2(0:t-d_{12}(t))) \\ u_2(t) &= \gamma_2(x_1(0:t-d_{21}(t)), x_2(0:t)). \end{aligned} \tag{2}$$

In order to build on the results in [6], we model the two plant system as a *three node graph*, with a "dummy delay" node introduced in the middle to enforce the propagation delay between plants. Specifically, letting $\delta(t) = [x_1(t-1)^T, x_2(t-1)^T]^T$, where δ is the state of the dummy node, we obtain the following state space representation for the system

$$\begin{aligned} \begin{bmatrix} x_1(t+1) \\ \delta(t+1) \\ x_2(t+1) \end{bmatrix} &= \begin{bmatrix} A_{11} & A_{1\delta} & 0 \\ A_{\delta 1} & 0 & A_{\delta 2} \\ 0 & A_{2\delta} & A_{22} \end{bmatrix} \begin{bmatrix} x_1(t) \\ \delta(t) \\ x_2(t) \end{bmatrix} + \\ &\quad \begin{bmatrix} B_1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & B_2 \end{bmatrix} \begin{bmatrix} u_1(t) \\ 0 \\ u_2(t) \end{bmatrix} + \begin{bmatrix} w_1(t) \\ 0 \\ w_2(t) \end{bmatrix} \end{aligned} \tag{3}$$

with $A_{1\delta} = [\ 0 \quad A_{12}\ ]$, $A_{2\delta} = [\ A_{21} \quad 0\ ]$, $A_{\delta 1} = [\ I \quad 0]^T$, and $A_{\delta 2} = [\ 0 \quad I]^T$. The physical topology of the plant is illustrated in Figure 1.

To condense notation, we let $\bar{x}^T = [x_1^T, \delta^T, x_2^T]^T$, $\bar{u}^T = [u_1^T, 0, u_2^T]^T$ and $\bar{w}^T = [w_1^T, 0, w_2^T]^T$, allowing us to write (3) as

$$\bar{x}(t+1) = A\bar{x}(t) + B\bar{u}(t) + \bar{w}(t) \tag{4}$$

for appropriately defined A and B matrices. 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.

Example 1: The need to take delay into account: Consider the system (3) with two scalar plants (i.e. $x_1, x_2 \in \mathbb{R}$), $A_{11} = A_{22} = 2$, $A_{12} = 3$, $A_{21} = 4$ and $B_1 = B_2 = 1$. Let $Q = I_4$, $R = I_3$ and $W_1 = W_2 = 1$. Then the costs (computed as in Section IV, Theorem 1) for $d(t) \equiv (1, 1)$ and $d(t) \equiv (2, 2)$ are 35.20 and 301.86, respectively. A *significant* difference in performance, which can be made arbitrarily large by varying A, exists between the two regimes.

Suppose now that the actual delay processes $d_{ij}(t)$ are governed by the pmf $\{.9, .1\}$. A conservative approach that simply implements the $d_{ij}(t) \equiv 2$ controller is unacceptable, as presumably a performance much closer to that of the $d_{ij}(t) \equiv 1$ regime should be possible.

III. State Decompositions and Globally Available Information

A. Information Hierarchy Graph

In [6], the concept of an information hierarchy graph is introduced in order to decompose the state and control actions into pairwise independent components, yielding solutions to the optimal control problem that obey the information constraints of the system by construction.

The *information hierarchy graph* associated with system (3), denoted by $\mathcal{I} = \{\mathcal{V}, \mathcal{E}\}$, is given as in Figure 2. The labeling of the nodes indicates the information (i.e. the components of the noise vector $w(t)$) available to to all players in the node, but none of the other players, under the worst case delay regime $d(t) = (2, 2)$. We can then associate with each node $v \in \mathcal{V}$ a state component, ζ_v, and a control component φ_v that depends solely on ζ_v. These components are a linear function of their respective noise labels, and thus by construction are pairwise independent. Additionally, they satisfy

$$x(t)=\zeta_{1,\delta,2}(t)+\begin{bmatrix}\zeta_{1,\delta}(t)\\0\end{bmatrix}+\begin{bmatrix}0\\\zeta_{\delta,2}(t)\end{bmatrix}+\begin{bmatrix}\zeta_1(t)\\0\\\zeta_2(t)\end{bmatrix} \quad (5)$$

$$u(t)=\varphi_{1,\delta,2}(t)+\begin{bmatrix}\varphi_{1,\delta}(t)\\0\end{bmatrix}+\begin{bmatrix}0\\\varphi_{\delta,2}(t)\end{bmatrix}+\begin{bmatrix}\varphi_1(t)\\0\\\varphi_2(t)\end{bmatrix} \quad (6)$$

It can be shown [6] that the update equations for the state decomposition $\zeta_v(t)$ components are given by

$$\begin{aligned}\zeta_s(t+1) &= \sum_{r:(r,s)\in\mathcal{E}} (A^{s,r}\zeta_r(t)+B^{s,r}\varphi_r(t)) \\ &\quad \text{for } s\in\mathcal{V} \text{ with } |s|>1 \\ \zeta_i(t+1) &= w_i(t) \text{ for } i=1,2\end{aligned} \quad (7)$$

with initial conditions $\zeta_s(0)=0$ for all $s\in\mathcal{V}$.

B. Globally Available Information

We will denote by $\mathcal{G}(d(t))$ the *globally available information* at time t, given the current delay regime. In particular based on the discussion in the previous section, we have that $\mathcal{G}(2,2):=\{1,\delta,2\}$ – this follows from the fact that $\{1\}$ and $\{2\}$ are both in $\{1,\delta,2\}$, but not simultaneously in any other node. We also have that the corresponding control action $\varphi_{1,\delta,2}$ is the action taken by the controller based on this globally available information.

In other delay regimes, the globally available information set increases to absorb other components of the state decomposition. In particular, the labeling of the nodes allows us to enumerate the globally available information for the remaining delay regimes by inspection as

$$\begin{aligned}\mathcal{G}(1,1) &:= \{\{1,\delta,2\},\{1,\delta\},\{\delta,2\}\} \\ \mathcal{G}(2,1) &:= \{\{1,\delta,2\},\{1,\delta\}\} \\ \mathcal{G}(1,2) &:= \{\{1,\delta,2\},\{\delta,2\}\}\end{aligned} \quad (8)$$

Denoting by $\zeta_{\mathcal{G}(d(t))}$ the component of $x(t)$ that is globally available under the delay regime $d(t)$, and by $\varphi_{\mathcal{G}(d(t))}$ the global action taken based on this globally available component, we may rewrite (5) and (6) as

$$x(t) = \zeta_{\mathcal{G}(d(t))}(t)+\sum_{v\in\mathcal{V}\setminus\mathcal{G}(d(t))} I_x^{\{1,\delta,2\},v}\zeta_v(t) \quad (9)$$

$$u(t) = \varphi_{\mathcal{G}(d(t))}(t)+\sum_{v\in\mathcal{V}\setminus\mathcal{G}(d(t))} I_u^{\{1,\delta,2\},v}\varphi_v(t), \quad (10)$$

where

$$\zeta_{\mathcal{G}(d(t))} := \sum_{v\in\mathcal{G}(d(t))} I_x^{\{1,\delta,2\},v}\zeta_v(t), \quad (11)$$

and I_x and I_u are identity matrices partitioned into blocks conforming to the partitions of $x(t)$ and $u(t)$, respectively.

In this way, we make explicit the increase in globally available information due to more favorable delay regimes. By adopting the convention that $\varphi_{1,\delta,2}=\varphi_{\mathcal{G}(d(t))}$, and that $\varphi_{i,\delta}=0$ if $\{i,\delta\}\in\mathcal{G}(d(t))$, for $i=1,2$, the update dynamics (7) remain consistent.

IV. Optimal Controller Structure and Examples

This section provides the two main results of the paper, namely the optimal solution for the case of known and unknown $d(t+1)$. Additionally, we revisit the motivating Example 1 from Section II-B.

Note that in the following we assume that the current global delay regime $d(t)=(d_{12}(t),d_{21}(t))$ is known – in practice, this requires a negative acknowledgment (NACK) mechanism. Future work will explore extensions of our results to when NACKs are not available.

A. Known $d(t+1)$

Let $X_{1,\delta,2}$ be the stabilizing solution to the discrete-time algebraic Riccati equation

$$S=Q+A^TSA-A^TSB(R+B^TSB)^{-1}B^TSA \quad (12)$$

and define the global gain $K_{1,\delta,2}$ to be the standard LQR gain

$$K_{1,\delta,2}=(R+B^TSB)^{-1}B^TSA. \quad (13)$$

For $r\neq\{1,\delta,2\}$, let s be the unique node such that $(r,s)\in\mathcal{E}$. Assume that X_s has already been defined, and define $X_r(t)$ by

$$X_r(t)=\begin{cases}X_s'^{r,r} \text{ if } r\in\mathcal{G}(d(t)) \\ Q^{r,r}+A^{s,rT}X_s'A^{s,r}-A^{s,rT}X_s'B^{s,r}\times\dots \\ (R^{r,r}+B^{s,rT}X_s'B^{s,r})^{-1}B^{s,rT}X_s'A^{s,r} \text{ otherwise}\end{cases} \quad (14)$$

where we write $X_s':=X_s(t+1)$ to save space.

Define the gain $K_r(t)$ by

$$K_r(t)=(R^{r,r}+B^{s,rT}X_s'B^{s,r})^{-1}B^{s,rT}X_s'A^{s,r} \quad (15)$$

and the control input components as

$$\begin{aligned}\varphi_{1,\delta,2}(t) &= -K_{1,\delta,2}\zeta_{\mathcal{G}(d(t))}(t) \\ \varphi_r(t) &= \begin{cases}-K_r(t)\zeta_r(t) & \text{for } r\notin\mathcal{G}(d(t)) \\ 0 & \text{otherwise}\end{cases}\end{aligned} \quad (16)$$

where the update dynamics for $\zeta_r(t)$ are given by (7), and $\zeta_{\mathcal{G}(d(t))}$ is given by (11).

Theorem 1: The optimal controller to the two player problem defined in Section II-B for known $d(t+1)$ is given by

$$u(t)=\varphi_{1,\delta,2}(t)+\sum_{r\in\mathcal{V}\setminus\mathcal{G}(d(t))} I_u^{\{1,\delta,2\},r}\varphi_r(t), \quad (17)$$

where the $\varphi_r(t)$ are given as in (16). The steady state cost is given by

$$\sum_{d\in\{1,2\}^2} p_d \mathrm{Tr}(W_1X_1^d+W_2X_2^d) \quad (18)$$

where X_i^d corresponds to the bottom level recursion matrix for a constant delay pattern $d(t)\equiv d$, and $p_d:=\mathbb{P}(d(t)=d)$ is the probability of that delay regime occurring.

Remark 1: The need to know $d(t+1)$ at time t arises from the matrix recursions (14) – $X_i(t)$ is a function of $X_{i,\delta}(t+1)$, which in turn is a function of $d(t+1)$. This corresponds to local actions being different depending on the amount of globally available information at the next time step. Equation (18) shows that in this setting the infinite horizon cost is a convex combination of the individual delay regimes' costs, weighted by the probability of each regime.

Remark 2: In the case of more favorable delay regimes, what can be seen as redundant layers are added to the matrix recursions (14) to maintain compatibility with the delay-invariant information hierarchy graph. In particular, if $d_{ij}(t+1) = 1$, then $X_{j,\delta}(t+1)$ is simply a sub-block of $X_{1,\delta,2}$, and thus $X_j(t)$ is directly a function of $X_{1,\delta,2}$ as well.

B. Unknown $d(t+1)$

As will be shown in the next section, Theorem 1 can be adapted to the case where $d(t+1)$ is not known by treating it as a disturbance in the dynamic programming argument.

Theorem 2: The optimal controller to the two player problem defined in Section II-B for unknown $d(t+1)$ is given by

$$u(t) = \varphi_{1,\delta,2}(t) + \sum_{r\in\mathcal{V}\setminus\mathcal{G}(d(t))} I_u^{\{1,\delta,2\},r}\varphi_r(t), \qquad (19)$$

where the $\varphi_r(t)$ are given as in (16), but where now $X'_{i,\delta} = \mathbb{E}_{d(t+1)}[X_{i,\delta}(t+1)]$ in (14). The steady state cost is then given by

$$\mathrm{Tr}(W_1X_1 + W_2X_2) \qquad (20)$$

with X_1 and X_2 as in (29).

Remark 3: Not knowing the next delay regime leads to a "hedging" type controller, where local actions are taken based on the *expected* amount of globally available information at the next time step.

Remark 4: Due to the structure of the controller, Theorem 2 indicates that if the distribution of $d(t)$ is not known, it may be estimated in real time online, and so long as the estimator is consistent, the same steady state cost will be achieved.

C. Examples

We revisit the system introduced in Example 1, and consider the following controllers:

1) $d(t+1)$ is known at time t – the steady state cost predicted by Theorem 1 is 61.87.
2) $d(t+1)$ is not know at time t, but $d(t)$'s distribution is known – the steady state cost predicted by Theorem 2 is 66.77.
3) $d(t+1)$ is not know at time t, and $d(t)$'s distribution is approximated by its local empirical value in real time (i.e. $\hat{p}_i(t) = (t)^{-1}\sum_{i=1}^t \mathbb{1}_{\{d_{ij}(t)=1\}}$) – as this estimator is consistent by the i.i.d. assumption on $d(t)$ and the Strong Law of Large Numbers, the steady state cost predicted is once again 66.77.

Illustrated in Figures 3(a)-(c) are the empirical costs, which can be seen to converge to the predicted ones. The important thing to note here is that all three methods perform *significantly* better than the conservative approach of ignoring early arriving information and implementing the $d(t) \equiv (2,2)$ controller.

V. Controller Derivation

A. Known $d(t+1)$

We follow [6], and begin by considering the finite horizon problem of minimizing the expected cost of

$$\sum_{t=1}^{N-1}(x(t)^TQx(t) + u(t)^TRu(t)) + x(N)^TQ_Nx(N). \qquad (21)$$

Denoting the optimal expected cost to go function by $\mathbb{E}[J(\zeta,t)]$, where the expectation is over the noise process $w(t)$, and recalling the state decomposition (9), we may write

$$\begin{aligned}\mathbb{E}[J(\zeta,N)] &= \mathbb{E}[x^TQ_Nx] \\ &= \mathbb{E}[\zeta_{\mathcal{G}(d(t))}^TQ_N\zeta_{\mathcal{G}(d(t))}] + \dots \\ &\qquad \sum_{s\in\mathcal{V}\setminus\mathcal{G}(d(t))}\mathbb{E}[\zeta_s^TQ_N^{s,s}\zeta_s]\end{aligned} \qquad (22)$$

where the last equality follows from the pairwise independence of the ζ_s.

Set $X_s(N) = Q_N^{s,s}$ for all $s\in\mathcal{V}$, and define $J(\zeta,N) = \sum_{s\in\mathcal{V}}\zeta_s^TX_s(N)\zeta_s$. Inductively assume that for some $t+1 \leq N$, $J(\zeta,t+1)$ is given by

$$J(\zeta,t+1) = \sum_{s\in\mathcal{V}}\zeta_s^TX_s(t+1)\zeta_s + \sum_{k=t+2}^{N}\sum_{i=1}^{2}\mathrm{Tr}(W_iX_i(k)) \qquad (23)$$

We compute the optimal expected cost-to-go function at time t by solving the Bellman equation:

$$\mathbb{E}[J(\zeta,t)] = \min_{\varphi}\mathbb{E}[x^TQx + u^TRu + J(\zeta',t+1)], \qquad (24)$$

where the ζ'_s are updates of ζ_s given by equations (7).

Recalling the pairwise independent decompositions of $x(t)$ and $u(t)$ given by (9) and (10), and the update dynamics (7), given $d(t)$, we can rewrite (24) as

$$\begin{aligned}&\min_{\varphi_{\mathcal{G}(d(t))}}\mathbb{E}[\zeta_{\mathcal{G}(d(t))}^TQ\zeta_{\mathcal{G}(d(t))} + \varphi_{\mathcal{G}(d(t))}^TR\varphi_{\mathcal{G}(d(t))} + \dots \\ &(A\zeta_{\mathcal{G}(d(t))} + B\varphi_{\mathcal{G}(d(t))})^TX'_{1,\delta,2}(A\zeta_{\mathcal{G}(d(t))} + B\varphi_{\mathcal{G}(d(t))})] \\ &+ \sum_{r\in\mathcal{V}\setminus\mathcal{G}(d(t))}\min_{\varphi_r}\mathbb{E}[\zeta_r^TQ^{r,r}\zeta_r + \varphi_r^TR^{r,r}\varphi_r + \dots \\ &\qquad (A^{s,r}\zeta_r + B^{s,r}\varphi_r)^TX'_s(A^{s,r}\zeta_r + B^{s,r}\varphi_r)] + \dots \\ &\qquad\quad \sum_{k=t+1}^{N}\sum_{i=1}^{2}\mathrm{Tr}(W_iX_i(k))\end{aligned} \qquad (25)$$

where we have once again used X'_s to denote $X_s(t+1)$ to save space.

Standard quadratic minimization yields (15) and (16), where we adopt the convention that $\varphi_{i,\delta}(t) = 0$ if $\{i,\delta\} \in \mathcal{G}(d(t))$. Plugging in the computed inputs, and recalling that $\zeta_{\mathcal{G}(d(t))}$ can be expanded in terms of ζ_v, $v \in \mathcal{G}(d(t))$, as in (11), we see that $J(\zeta,t)$ has the form

$$J(\zeta,t) = \sum_{r\in\mathcal{V}}\zeta_r^TX_r(t)\zeta_r + \sum_{k=t+1}^{N}\sum_{i=1}^{2}\mathrm{Tr}(W_iX_i(k)) \qquad (26)$$

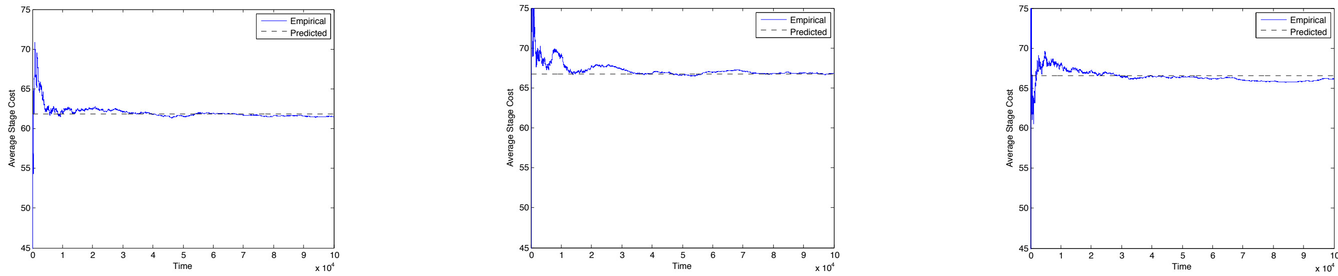

(a) Empirical and predicted costs for $d(t+1)$ known.

(b) Empirical and predicted costs for $d(t+1)$ unknown, but known distribution on $d(t)$.

(c) Empirical and predicted costs for $d(t+1)$ unknown and unknown distribution on $d(t)$.

Fig. 3: Empirical and predicted costs for the three classes of delay patterns considered in Section IV.

where the matrices $X_r(t)$ are computed as in (14), with the exception of $X_{1,\delta,2}$, which is given by the standard discrete time Riccati recursion

$$\begin{aligned} X_{1,\delta,2}(t) = Q + A^T X_{1,\delta,2}(t+1)A - A^T X_{1,\delta,2}(t+1)B\times \\ (R + B^T X_{1,\delta,2}(t+1)B)^{-1} B^T X_{1,\delta,2}(t+1)A. \end{aligned} \tag{27}$$

Noting that this top level recursion is invariant with respect to the delay regime, we may consider the infinite horizon solution by letting $N \to \infty$. In particular, under the assumptions made, we have that $X_{1,\delta,2}(t)$ converges to the stabilizing solution of (12), which we denote by $X_{1,\delta,2}$ – similarly, we denote the the resulting LQR gain by $K_{1,\delta,2}$.

The infinite horizon cost is calculated by noting that, by the assumption of i.i.d. $d(t)$,

$$\lim_{N\to\infty} \frac{1}{N} \sum_{t=1}^{N} \sum_{i=1}^{2} \mathrm{Tr}(W_i X_i(t)) = \sum_{i=1}^{2} \mathrm{Tr}(W_i \mathbb{E}_{d(t+1)}[X_i(t)]). \tag{28}$$

Explicitly computing the expectation given the recursions (14) gives (18).

B. Unknown $d(t+1)$

Based on the previous argument, the dynamic programming argument can be modified to solve for $\varphi(t)$ that are causal with respect to the delay pattern by treating $d(t+1)$ as a disturbance. In particular, we now take expectations over both the noise and delay processes, $w(t)$ and $d(t+1)$, in (22) and (24).

Under our assumptions, we have that for all $r \in \mathcal{V}$,

1) $\zeta_r(t)$ is independent of $d(t+1)$ (this follows from the $d(t)$ being independent).
2) $\varphi_r(t)$ is now assumed to be causal, and must therefore be completely determined by $\zeta_r(0:t)$ and $d(0:t)$. This implies that $\varphi_r(t)$ must be independent of $d(t+1)$.
3) $\varphi_r(t)$ is, much as in the centralized LQG state feedback case, independent of the noise process. This implies that the control policy, and consequently the X_s''s, are independent of the noise process .

This allows us to take expectations with respect to $d(t+1)$ *prior* to solving for the optimal inputs $\varphi_r(t)$, effectively replacing the X_s''s in (25), and consequently in (14), with $\mathbb{E}_{d(t+1)}[X_s(t+1)]$, from which Theorem 2 follows immediately.

In particular, since $X_{1,\delta,2}$ is invariant with respect to delay, we see that the only change occurs in the recursions for $X_i(t)$. Letting $r = \{i\}$, $s = \{i,\delta\}$, and $V = \{1,\delta,2\}$, we have that

$$\begin{aligned} X_i \quad \equiv \quad & Q^{r,r} + A^{s,rT} X_s' A^{s,r} - A^{s,rT} X_s' B^{s,r} \times \dots \\ & (R^{r,r} + B^{s,rT} X_s' B^{s,r})^{-1} B^{s,rT} X_s' A^{s,r} \end{aligned} \tag{29}$$

with

$$\begin{aligned} X_{i,\delta}' &= \mathbb{E}_{d(t)}[X_{i,\delta}(t)] = pX_{i,\delta}^{d_{ji}=1} + (1-p)X_{i,\delta}^{d_{ji}=2} \\ X_{i,\delta}^{d_{ji}=1} &= X_{1,\delta,2}^{\{i,\delta\},\{i,\delta\}} \\ X_{i,\delta}^{d_{ji}=2} &= Q^{s,s} + A^{V,sT} X_V A^{V,s} - A^{V,sT} X_V B^{V,s} \times \dots \\ & \quad (R^{s,s} + B^{V,sT} X_V B^{V,s})^{-1} B^{V,sT} X_V A^{V,s} \end{aligned} \tag{30}$$

VI. 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 be time-varying, but impose that they preserve partial nestedness. It was seen that the time-varying delay pattern induces a piecewise linear structure in the resulting optimal controller, with mode switches being dictated by both the current and next delay regime. By treating the next delay regime as a disturbance in the dynamic programming argument, we derived a hedging-type controller to deal with this non-causality. Finally, simulations were performed demonstrating the effectiveness of the proposed method.

Future work will be to extend the results to more general delay patterns, and removing the assumption of strong connectedness, much as was done in [30] for the case of constant delays. Additionally, we will seek to relax the assumptions of independence of $d(t)$, and to deal with the setting in which NACKs are not present. We are also currently exploring a principled integration of these results with recent deadline based coding techniques [27].

VII. Acknowledgements

We would like to thank Andrew Lamperski for helpful discussions regarding the technical results of this paper.

References

[1] A. Mahajan, N. Martins, M. Rotkowitz, and S. Yuksel, "Information structures in optimal decentralized control," in *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*, 2012, pp. 1291–1306.

[2] J. Sandell, N. and M. Athans, "Solution of some nonclassical LQG stochastic decision problems," *Automatic Control, IEEE Transactions on*, vol. 19, no. 2, pp. 108 – 116, apr 1974.

[3] B.-Z. Kurtaran and R. Sivan, "Linear-quadratic-gaussian control with one-step-delay sharing pattern," *Automatic Control, IEEE Transactions on*, vol. 19, no. 5, pp. 571 – 574, oct 1974.

[4] T. Yoshikawa, "Dynamic programming approach to decentralized stochastic control problems," *Automatic Control, IEEE Transactions on*, vol. 20, no. 6, pp. 796 – 797, dec 1975.

[5] A. Lamperski and J. 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*, dec. 2011, pp. 6901 –6906.

[6] ——, "Dynamic programming solutions for decentralized state-feedback LQG problems with comminication delays," in *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*, jun. 2012.

[7] A. Rantzer, "A separation principle for distributed control," in *Decision and Control, 2006 45th IEEE Conference on*, dec. 2006, pp. 3609 – 3613.

[8] A. Gattami, "Generalized linear quadratic control theory," in *Decision and Control, 2006 45th IEEE Conference on*, dec. 2006, pp. 1510 –1514.

[9] M. Rotkowitz and S. Lall, "A characterization of convex problems in decentralized control," *Automatic Control, IEEE Transactions on*, vol. 51, no. 2, pp. 274–286, 2006.

[10] A. Lamperski and J. C. Doyle, "Output feedback $\mathcal{H}_2$ model matching for decentralized systems with delays," in *2013 IEEE American Control Conference, To appear.*, sept. 2012.

[11] H. R. Feyzmahdavian, A. Gattami, and M. Johansson, "Distributed output-feedback LQG control with delayed information sharing," in *IFAC Worskhop on Distributed Estimation and Control of Networked Systems*, vol. 3, no. 1, 2012, pp. 192–197.

[12] A. Nayyar, A. Mahajan, and D. Teneketzis, "Optimal control strategies in delayed sharing information structures," *Automatic Control, IEEE Transactions on*, vol. 56, no. 7, pp. 1606 –1620, july 2011.

[13] N. Elia and S. Mitter, "Stabilization of linear systems with limited information," *Automatic Control, IEEE Transactions on*, vol. 46, no. 9, pp. 1384–1400, 2011.

[14] J. Hespanha, A. Ortega, and L. Vasudevan, "Towards the control of linear systems with minimum bit-rate," in *In Proc. of the Int. Symp. on the Mathematical Theory of Networks and Syst*, 2002.

[15] G. N. Nair and R. J. Evans, "Exponential stabilisability of finite-dimensional linear systems with limited data rates," *Automatica*, vol. 39, no. 4, pp. 585 – 593, 2003. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0005109802002856

[16] S. Tatikonda and S. Mitter, "Control under communication constraints," *Automatic Control, IEEE Transactions on*, vol. 49, no. 7, pp. 1056–1068, 2004.

[17] W. S. Wong and R. Brockett, "Systems with finite communication bandwidth constraints. ii. stabilization with limited information feedback," *Automatic Control, IEEE Transactions on*, vol. 44, no. 5, pp. 1049–1053, 1999.

[18] A. Sahai, "Anytime information theory," Ph.D. dissertation, Massachusetts Institute of Technology, 2001.

[19] N. Martins and M. Dahleh, "Feedback control in the presence of noisy channels: "bode-like" fundamental limitations of performance," *Automatic Control, IEEE Transactions on*, vol. 53, no. 7, pp. 1604–1615, 2008.

[20] H. Lin and P. Antsaklis, "Persistent disturbance attenuation properties for networked control systems," in *Decision and Control, 2004. CDC. 43rd IEEE Conference on*, vol. 1, 2004, pp. 953–958 Vol.1.

[21] P. Naghshtabrizi and J. Hespanha, "Designing an observer-based controller for a network control system," in *Decision and Control, 2005 and 2005 European Control Conference. CDC-ECC '05. 44th IEEE Conference on*, 2005, pp. 848–853.

[22] D. Yue, Q.-L. Han, and C. Peng, "State feedback controller design of networked control systems," *Circuits and Systems II: Express Briefs, IEEE Transactions on*, vol. 51, no. 11, pp. 640–644, 2004.

[23] W. Zhang, M. Branicky, and S. Phillips, "Stability of networked control systems," *Control Systems, IEEE*, vol. 21, no. 1, pp. 84–99, 2001.

[24] V. Gupta, D. Spanos, B. Hassibi, and R. Murray, "On LQG control across a stochastic packet-dropping link," in *American Control Conference, 2005. Proceedings of the 2005*, 2005,, pp. 360–365.

[25] J. Hespanha, P. Naghshtabrizi, and Y. Xu, "A survey of recent results in networked control systems," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 138–162, 2007.

[26] N. Matni, "Communication delay co-design in $\mathcal{H}_2$ decentralized control using atomic norm minimization," in *To appear in the IEEE Conf. on Decision and Control, 2013*, 2013.

[27] D. Leong and T. Ho, "Erasure coding for real-time streaming," in *Information Theory Proceedings (ISIT), 2012 IEEE International Symposium on*, 2012, pp. 289–293.

[28] Y.-C. Ho and K.-C. Chu, "Team decision theory and information structures in optimal control problems–part i," *Automatic Control, IEEE Transactions on*, vol. 17, no. 1, pp. 15–22, 1972.

[29] P. Shah and P. A. Parrilo, "$\mathcal{H}_2$-optimal decentralized control over posets: A state space solution for state-feedback," in *Decision and Control (CDC), 2010 49th IEEE Conference on*. IEEE, 2010, pp. 6722–6727.

[30] A. Lamperski and L. Lessard, "Optimal state-feedback control under sparsity and delay constraints," in *3rd IFAC Workshop on Distributed Estimation and Control in Networked Systems*, 2012, pp. 204–209.