Proceedings of the
46th IEEE Conference on Decision and Control
New Orleans, LA, USA, Dec. 12-14, 2007

On the Period of Communication Policies for Networked Control Systems, and the Question of Zero-Order Holding

Dimitrios Hristu-Varsakelis†

Abstract— **We discuss stochastic, linear networked control systems (NCSs) in which only a limited number of the plant's sensors and actuators may communicate with the controller at any one time. We explore the problem of designing an LQG controller and an accompanying periodic communication policy, using recent results that forgo optimal communication for the sake of lowering the complexity of the joint (control-communication) problem. We show that the period of the policies in question can be shorter than previously established, and that policies designed under a simpler NCS model, in which sensors and actuators are "ignored" by the controller when they are not actively communicating, can also be effective in the more complex setting which includes zero-order holding (ZOH). Interestingly, the inclusion of a ZOH — although sometimes practical — does not always lead to better performance.**

I. INTRODUCTION

This work is concerned with Networked Control Systems (NCSs) in which only a limited number of the plant's sensors and actuators may access the shared medium simultaneously, in order to communicate with the controller. This type of medium access constraint arises "naturally" in simple laboratory-scale networks [1], [2] and more sophisticated Fieldbus and CAN-based networks [3], among others. In that context, it is only meaningful to specify a controller in conjunction with a communication policy [4], [5] which prescribes the times at which the plant's sensors and actuators are to be granted medium access. For example, input data sent through the network must be "bound" to the specific actuator(s) they are meant for. Thus, the choice of communication influences the performance of the controller, and vise versa, leading to high complexity if one insists on seeking jointly optimal solutions [6], [7], [8]. One way to manage that complexity is to make strong assumptions regarding the underlying plant, such as block-diagonal dynamics [9], or "one-sided" access constraints [10], [11]. This paper is a continuation of a complementary approach to NCS controller design [12], [13], that relaxes the requirement for optimal communication for the sake of being able to compose straightforward solutions.

We pursue a recently-developed controller design method for linear NCS with medium access constraints and delays, whereby the control and communication subproblems are "decoupled". This is accomplished as proposed in [12], [14], by: i) restricting communication to periodic sequences, and ii) using existing techniques to design a controller for the resulting periodic NCS. This last step requires that the communication policy be such that it preserves the structural properties (e.g., controllability and observability) of the underlying plant in the presence of communication constraints [13], [15]. We investigate the existence of such policies under an NCS model that forgoes the use of ZOH; instead, the controller ignores sensors and the plant turns off actuators when they are not actively communicating. Compared to existing methods, [13], [15], [16], our approach simplifies the construction of effective communication sequences, and allows us to establish a less conservative upper bound on their period, that bound being the dimension of the state vector. Furthermore, we show that if the plant's parameters satisfy an orthogonality condition then sequences designed for our non-ZOH model are also effective in NCSs which *include* ZOH elements. Being able to use the same communication pattern in both settings allows for direct comparisons between the two. In the context of LQG control, the inclusion of ZOH does not always lead to better performance.

For the NCSs we have in mind here, the basic constraint — lack of simultaneous access to all sensors and actuators — is handled by time-multiplexing. Related work in the same setting includes [5], [9], [13] on stabilization, and [7], [8] on optimal control. Recently, [17] addressed the question of how much time should be devoted to measuring versus controlling the underlying plant. There is also a significant body of literature on the effects of quantization and limited channel throughput [18], [19], [20], [21], probabilistic data losses [22], [23], as well as the effects of transmission delays in the feedback loop [24], [25], [26]. See [27] for a fuller review.

In the next Section we discuss an NCS model without the ZOH elements typically attached to the plant's inputs and outputs; the NCS is later transformed to an equivalent periodic system by imposing periodic communication. Section III explores the problem of choosing communication policies that preserve the structural properties of the NCS, and discusses their period and their effectiveness when a ZOH is included. Section IV applies our technique to the problem of designing an LQG controller for an NCS with access constraints and known delays, and discusses a numerical example.

II. NCS MODEL AND PROBLEM FORMULATION

Our model for the NCS under consideration follows that in [13], [15] (see Fig. 1). We will take $x \in \mathbb{R}^n$, $u \in \mathbb{R}^m$, and $y \in \mathbb{R}^p$ to be the state, input, and output vectors of the underlying plant, respectively. The communication medium imposes an upper bound on the number of sensors, $w_\sigma < p$,

† Department of Applied Informatics, University of Macedonia, 156 Egnatia St., Thessaloniki 54006, Greece, dcv@uom.gr.

1-4244-1498-9/07/$25.00 ©2007 IEEE.

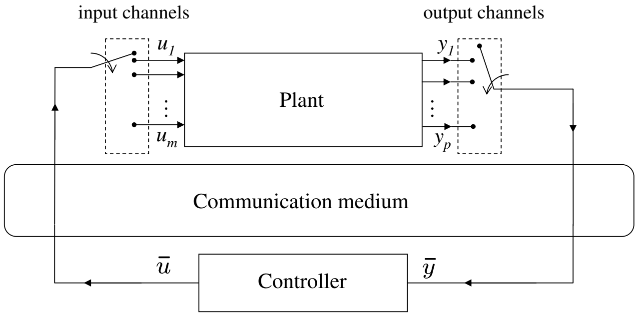

Fig. 1. A Networked Control System. The communication medium cannot facilitate simultaneous communication between all sensors/actuators and the controller; it may also impose transmission delays. The "open" or "closed" status of the switches indicates the medium access status of the corresponding sensors or actuators.

and actuators, $w_\rho < m$, which may communicate simultaneously with the controller, so that medium access must be time-multiplexed. Moreover, controller-plant communication may be subject to transmission delays, which we will take to be known. We will not consider quantization or bit-rate constraints.

We will assume linear stochastic plant

$$\begin{aligned} x(k+1) &= Ax(k) + Bu(k) + v(k) \\ y(k) &= Cx(k) + w(k), \qquad k = 0, 1, 2, \dots, \end{aligned} \tag{1}$$

where $v(\cdot)$, $w(\cdot)$, are both Gaussian, i.i.d., with $v(\cdot) \sim N(0, G)$, $G = G^T > 0$, and $w(\cdot) \sim N(0, I)$, and with $x(0) \sim N(x_0, \Sigma)$, $\Sigma = \Sigma^T > 0$. The problem under consideration is to find a medium access policy for the plant's sensors and actuators and a control policy that minimizes an LQG-type cost, subject to the access constraints (w_ρ, w_σ) and delays. Here, the LQG problem is used mainly for the sake of concreteness; in the sequel, it will become apparent that our approach applies to other control design problems as well. We will begin by discussing the deterministic, delay-free counterpart of (1), and will return to the delay-inclusive, stochastic NCS later.

We will use the notion of a *communication sequence* [4], [5] to describe the medium access status of the plant's inputs and outputs across time.

Definition 1: Let $M, N \in \mathbb{N}$ with $N \leq M$. An M-to-N *communication sequence* is a map, $\sigma(k) : \mathbb{Z}_+ \mapsto \{0,1\}^M$, satisfying $\|\sigma(k)\|^2 = N$, $\forall k$.

We will take $\sigma(k)$, $k = 0, 1, 2, \dots$, to be a p-to-w_σ communication sequence, with $\sigma_i(k) = 1$ if the i-th sensor is to access the controller at time k, and $\sigma_i(k) = 0$ otherwise. Similarly, $\rho(k)$ will be a w_σ-to-m communication sequence prescribing the times at which the actuators are to communicate with the controller. For a p-to-w_σ communication sequence, σ, we will write $\mu_\sigma(k)$ to denote the matrix that results by deleting the $p - w_\sigma$ all-zero rows from the $p \times p$ matrix $M_\sigma(k) \triangleq diag(\sigma(k))$.

It will be convenient to initially assume that at each time, k, the controller "ignores" sensors and actuators which are not actively communicating, and computes inputs using only the w_σ elements of $y(k)$ for which $\sigma_i(k) = 1$, i.e.,

$$\bar{y}(k) = \mu_\sigma(k) y(k). \tag{2}$$

The controller produces an input $\bar{u}(k) \in \mathbb{R}^{w_\rho}$, containing those elements of $u(k)$ for which the corresponding actuators are granted medium access at time k. This leaves us to specify what happens to the actuators which are *not* updated at k. As in [13], we will choose to have them *turn off* until communication with the controller is re-established, i.e., $u_i(k) = 0$ while $\sigma_i(k) = 0$. Thus, the plant input is given by

$$u(k) = \mu_\rho(k)^T \bar{u}(k). \tag{3}$$

Perhaps a more practical assumption would be to apply ZOH to the plant's inputs, so that an actuator maintains its input level until it receives new data. We will discuss such a model in the sequel. For now, we will take advantage of the simplicity resulting by forgoing the ZOH in order to more easily construct "effective" periodic communication sequences. Our choice will also have implications for their period and for their suitability when a ZOH *is* in place.

By combining (1)-(3) we obtain a linear time-varying system with w_ρ inputs and w_σ outputs:

$$\begin{aligned} x(k+1) &= Ax(k) + B\mu_\rho(k)^T \bar{u}(k) \\ \bar{y}(k) &= \mu_\sigma(k) C x(k). \end{aligned} \tag{4}$$

These equations, termed *the extended plant* [13], describe the NCS "from the controller's point of view" and incorporate the dynamics of the plant together with the access status of the communication medium.

III. Choosing effective communication sequences

The communication policies ρ and σ determine the time-varying dynamics of the extended plant (4). Thus, in any control design problem involving the NCS under consideration, we can only say that a controller is optimal for a given communication policy. Optimizing with respect to both control and communication is generally difficult [6], [7], [8] and often involves combinatorial complexity. Instead, we will relax the requirement for joint optimality and show how to select communication sequences that are "good enough", in the sense that they guarantee the existence of an accompanying optimal controller, and are easy to generate. In particular, we will look for sequences which preserve important structural properties of the underlying LTI plant in (4), including controllability and observability[1]. Such sequences are not unique. However, effective examples can be easily constructed based on the plant's controllability/observability indices, as we will show.

[1] Here we have assumed a discrete-time (or sampled-data) plant. It is worth noting that if one adopts a continuous-time model instead (e.g., plant inputs can change continuously, but not all simultaneously), then the communication policy selection problem becomes much simpler. In that case, "round robin" policies preserve the plant's structural properties, as does any periodic policy that devotes some time to every actuator (sensor), in any order [28]. As this work and its predecessors illustrate, the situation is not as simple in the discrete-time setting.

Definition 2: The NCS (4) is *controllable on* $[k_0, k_f]$ if, $\forall\ x_0$, $\exists\ \bar{u}(\cdot)$ that steers (4) from $x(k_0) = x_0$ to the origin at time k_f. We say that (4) is *l-step controllable*, or simply *controllable* if there exists an integer $l > 0$ such that (4) is controllable on $[k, k+l]\ \forall\ k$.

Definition 3: The NCS (4) is *observable on* $[k_0, k_f]$ if any initial condition at k_0 can be uniquely determined from $\bar{y}(k)$, $k \in [k_0, k_f]$. We say that (4) is *l-step observable*, or simply *observable* if there exists an integer $l > 0$ such that (4) is observable on $[k, k+l]\ \forall\ k$.
Reconstructability and detectability are defined in a similar manner.

For convenience, and because of space limitations, we will assume that the matrix A in (4) is invertible, so that controllability and reachability of (4) are equivalent, as are observability and reconstructability. For the case where A is singular see [14]. The following theorem, from [13], concerns the existence of controllability- (observability-) preserving communication sequences.

Theorem 1 ([13]): Let the pair (A, B) be controllable, where B is $n \times m$, and A is invertible. For any integer $1 \le w_\rho < m$, there exist integers $l, N > 0$ and an N-periodic[2] m-to-w_ρ communication sequence $\rho(\cdot)$, with $N \le \left\lceil \frac{n}{w_\rho} \right\rceil \cdot n$, such that the extended plant (4) is controllable on $[k, k+l]$ for all k, and thus controllable.
The proof of Th. 1 provides an explicit sequence construction algorithm which selects columns from $[B, AB, \cdots, A^{N-1}B]$, w_ρ at a time, and requires at most $\left\lceil \frac{n}{w_\rho} \right\rceil \cdot n$ steps in order to build a full-rank collection of columns, in the worst case. The same upper bound is reported in [16].

If our objective is LQG control, then it suffices to guarantee the weaker properties of stabilizability and detectability in the NCS. The proof of the next result can be found in [28], [14].

Corollary 1: Let (A, B) be stabilizable, where B is $n \times m$. For any integer $1 \le w_\rho < m$, there exists an integer $N \le n$ and an N-periodic $m - to - w_\rho$ communication sequence $\rho(\cdot)$ such that the extended plant is stabilizable.
Corollary 1 and its dual suggest that in order to obtain a communication sequence that maintains stabilizability (detectability) in the presence of communication constraints, it is sufficient to identify the sequence that does the job for the non-singular part of the controllable (reconstructible) subsystem [14].

A. Shorter-period communication sequences

The algorithms given in previous work [12], [13], are known to produce communication sequences whose period is usually far shorter than the upper bound given in Th. 1. The following result shows that in fact a period of $N = n$ steps is sufficient.

Theorem 2: Let (A, B) be controllable, where B is $n \times m$, and A is invertible. For any integer $1 \le w_\rho < m$, there exists an n-periodic m-to-w_ρ communication sequence $\rho(\cdot)$ such that the extended plant (4) is controllable. In particular, if $n_1, n_2, ..., n_m$ are the controllability indices of (A, B), then for $w_\rho = 1$, the sequence whose first period is $\{\underbrace{e_m, e_m ..., e_m}_{n_m\ times}, \underbrace{e_{m-1}, e_{m-1} ..., e_{m-1}}_{n_{m-1}\ times}, ... \underbrace{e_1, e_1 ... e_1}_{n_1\ times}\}$, where e_i is the i-th standard basis vector, makes (4) controllable.
Sketch of proof:
It is sufficient to prove the theorem for the worst-case scenario, $w_\rho = 1$. Let

$$R \triangleq [B\mu_\rho^T(n-1), AB\mu_\rho^T(n-2), \cdots, A^{n-1}B\mu_\rho^T(0)]. \quad (5)$$

The NCS (4) is controllable on $[0, n]$ (and thus controllable [13]), iff $\text{rank}(R) = n$. As before, $\mu_\rho^T(k)$ has the effect of "selecting" one of m columns from $A^{n-k-1}B$. It is well-known that if (A, B) is controllable then the following collection of n columns from R,

$$\begin{aligned} \mathcal{R}_1 \triangleq\ & [b_1, Ab_1, ..., A^{n_1-1}b_1,\ b_2, Ab_2, ..., A^{n_2-1}b_2, ..., \\ & b_m, Ab_m, ..., A^{n_m-1}b_m] \end{aligned}$$

is full-rank, where b_i is the i-th column of B, and $n_1, ..., n_m$ are the controllability indices of (A, B), satisfying $\sum_1^m n_i = n$. Observe that we can use the last n_1 terms of the communication sequence (equivalently, μ_ρ) to select the first n_1 columns of $\mathcal{R}_1$ in R, but that after that we are unable to select b_2, Ab_2 and so on because we are restricted to choosing terms with increasing powers of A.

Now, consider the matrix $\mathcal{R}_2$, obtained from (5) by using the sequence suggested in the theorem's statement.

$$\begin{aligned} \mathcal{R}_2 = \ & [b_1, Ab_1, ..., A^{n_1-1}b_1, \\ & A^{n_1}b_2, A^{n_1+1}b_2, ..., A^{n_1+n_2-1}b_2, ..., \quad (6) \\ & A^{n_1+...+n_{m-1}}b_m, A^{n_1+...+n_{m-1}+1}b_m, ..., A^{n-1}b_m] \end{aligned}$$

It is sufficient to show that $rank(\mathcal{R}_2) = n$. For a system with only two inputs ($m = 2$), $\mathcal{R}_1$ and $\mathcal{R}_2$ specialize to

$$\begin{aligned} R_1 &= [b_1, Ab_1, ..., A^{n_1-1}b_1,\ b_2, Ab_2, ..., A^{n_2-1}b_2], \quad (7) \\ R_2 &= [b_1, Ab_1, ..., A^{n_1-1}b_1,\ A^{n_1}b_2, A^{n_1+1}b_2, ..., A^{n-1}b_2], \end{aligned}$$

with $rank(R_1) = n$ by the controllability of (A, B). Because A is assumed to be invertible, $rank(A^{n_1}R_1) = n$ as well. One can then show that $rank(R_2) = n$ by comparing the columns of $A^{n_1}R_1$ and R_2. In particular, the first n_1 columns of $A^{n_1}R_1$ are linear combinations of the first n_1 columns of R_2 (n_1 is the first controllability index of (A, B)). Thus, $n = rank(A^{n_1}R_1) \le rank(R_2)$, implying that $rank(R_2) = n$ for a system with two inputs. The theorem is established via successive applications of the same argument for $m = 3$ (to show that $n = rank(A^{n_1+n_2}R_1) \le rank(R_2)$) and higher, extending R_1 and R_2 one controllability index at a time. □

The duality of controllability and observability yields an analogous result regarding sequences that preserve the observability of the extended plant.

[2] A communication sequence $\sigma(\cdot)$ will be called N-periodic if $\sigma(k) = \sigma(k+N)$ for all $k \ge 0$.

B. NCS with an input ZOH

Under the NCS model adopted thus far, actuators "turn off" when they are not communicating. This choice has the effect of avoiding any "enlargement" of the state vector [4], [5] but may result in inputs with large step changes. A more practical approach might be to amend the extended plant model to include zero-order holding of inputs [15], [16], by enlarging the state to include the contents of the ZOH, i.e.,

$$\begin{aligned} x(k+1) &= Ax(k) + B(I - M_\rho(k))u_{ZOH}(k) \\ &\quad + B\mu_\rho^T(k)\bar{u}(k), \qquad (8) \\ u_{ZOH}(k+1) &= (I - M_\rho(k))u_{ZOH}(k) + \mu_\rho^T(k)\bar{u}(k), \\ \bar{y}(k) &= \mu_\sigma(k)Cx(k). \end{aligned}$$

where $M_\rho(k) \triangleq \mathrm{diag}(\rho(k))$. In the following, we will refer to (8) as the *ZOH-inclusive* model; we will also use $\mathcal{X} \subset \mathbb{R}^{n+m}$ to denote the span of the first n standard basis vectors in $\mathbb{R}^{n+m}$ (i.e., the subspace corresponding to the x-states in (8)).

Recent work [16] has proved Th. 1 for the ZOH-inclusive model (8) by retracing the steps of [12]. The next result addresses the question of whether the controllability-preserving sequences obtained using the plant's controllability indices (Th. 2) remain effective if the NCS includes a ZOH.

Corollary 2: Let the pair (A, B) be controllable, where B is $n \times m$, and A is invertible. Let $n_1, ..., n_m$ be the controllability indices of (A, B). For $i = 2, ..., m-1$, denote with $\mathcal{C}_i$ the matrix that contains the first $\sum_{t=1}^{i} n_t$ columns of $\mathcal{R}_2$ in (6). Let $j(i) = \sum_{t=1}^{i-1} n_t$, and $V_i \in \mathbb{R}^n$ be a unit vector orthogonal to the range of $\mathcal{C}_i$ with column $A^{j(i)}b_i$ deleted. If

$$V_i^T \sum_{k=0}^{j(i)} A^k b_i \neq 0 \quad \forall i = 2, ..., m-1, \qquad (9)$$

then there exists an N-periodic controllability-preserving communication sequence ρ, with $N \leq n$, such that the controllable subspace of the ZOH-inclusive NCS (8) contains $\mathcal{X}$ (the span of the first n states in (8))

Sketch of Proof:

Consider, $\rho(\cdot)$, an n-periodic controllability-preserving communication sequence constructed via Th. 2. The state evolution after n steps, with the ZOH in place, is of the form:

$$x(n) = \Lambda \cdot [\bar{u}^T(n-1), \bar{u}^T(n-2), ..., \bar{u}^T(0)]^T, \qquad (10)$$

after assuming without loss of generality that $u_{ZOH}(0) = 0$ in (8). To prove the theorem, one must expand Λ in terms of A, B and μ_ρ, and notice that μ_ρ selects the same columns from R as in Th. 2, but there will now be added terms selected at the same time in each column of Λ. If ρ is constructed as per Th. 2 then whenever $\rho(n-k) = \rho(n-k-1)$, $k = 1, ..., n$, the additional terms in the $(k+1)$st column of Λ vanish. The only columns of Λ for which the additional terms do not vanish are those which correspond to a change in $\rho(k)$ over the previous step, $k-1$, i.e., those in columns n_1+1, n_1+n_2+1, ..., $n_1+n_2+...+n_{m-1}+1$.

By constructing Λ for $m = 2, 3, ...$, and so on, we can show that the additional terms in the columns of Λ are linear combinations of columns already in R_2 (7), and that if condition (9) holds then the addition of those columns cannot result in loss of rank, thus $rank(\Lambda) = rank(R) = n$. □

The dual statement to Cor. 2 also holds. If (9) is not satisfied, one can still obtain a controllability-preserving communication sequence, using [16]. In that case, the upper bound on the sequence length is $\lceil n/w_\rho \rceil n$, rather than n. At this time, there does not seem to be an obvious way of modifying the proof in [16] to reduce the upper bound on the communication sequence length without requiring (9).

IV. Example: LQG control

We proceed to combine our communication sequence selection approach with existing tools for periodic systems in order to design an LQG controller for the NCS. Consider the stochastic version of the extended plant,

$$\begin{aligned} x(k+1) &= Ax(k) + \bar{B}(k)\bar{u}(k) + v(k) \qquad (11) \\ \bar{y}(k) &= \bar{C}(k)x(k) + \bar{w}(k), \end{aligned}$$

where $\bar{B}(k) \triangleq B\mu_\rho(k)^T$, $\bar{C}(k) \triangleq \mu_\sigma(k)C$, $\bar{w}(k) \triangleq \mu_\sigma(k)w(k)$, and $v(\cdot)$, $w(\cdot)$ and $x(0)$ are as in (1). Having lifted the requirement for optimal communication, we would like to solve the following problem.

Problem 1: Given a pair of communication sequences $(\rho(\cdot), \sigma(\cdot))$ which preserve the stabilizability and detectability of (1) in the NCS (11), and $Q = Q^T > 0$, design a controller that minimizes

$$J = \mathcal{E}\left\{ \sum_{k=0}^{N_f} x^T(k)Qx(k) + \bar{u}^T(k)\bar{u}(k) \right\}, \qquad (12)$$

If $\rho(\cdot)$ and $\sigma(\cdot)$) are constructed as per Sec. III-A and taken to be periodic, then Problem 1 becomes a standard LQG problem for a periodic stochastic plant. Its solution consists of i) a Kalman filter that estimates $x(k)$ from the output $\bar{y}(0), \ldots \bar{y}(k)$, and ii) an LQ optimal feedback controller, obtained by solving a deterministic LQ problem with perfect state information [29]. Computing the Kalman filter and controller feedback gains involves solving a pair of discrete-time periodic Riccati equations (DPREs). From [30], if the stochastic extended plant is stabilizable and detectable (by choice of the communication sequences ρ, σ), then each of the DPREs has a unique symmetric periodic positive semidefinite solution to which the DPRE converges, one being the optimal Kalman gain which stabilizes the estimation error dynamics and the other being the optimal LQ gain which stabilizes the stochastic extended plant. See [31] for a more detailed summary of these results in the context of the NCS studied here.

Remark: If there are known transmission delays between controller and plant, then a delay compensator [32] can be combined with the LQG controller and communication sequences developed previously. A detailed description of how to do this is given in [14]. Briefly, for a discrete-time plant, the extended plant with delays will be of the form

$$\begin{aligned} x(k+1) &= Ax(k) + \bar{B}(k-\Delta_2)\bar{u}(k-\Delta_2) + v(k) \qquad (13) \\ \bar{y}(k) &= \bar{C}(k-\Delta_1)x(k-\Delta_1) + \bar{w}(k-\Delta_1), \end{aligned}$$

where $\Delta_1,\ \Delta_2 > 0$ are the (integer) plant-to-controller and controller-to-plant delays, respectively. See also [15], [14] for the case of continuous-time NCSs. The delay compensator operates by estimating the delayed plant state (using a Kalman filter) and then propagating that estimate forward in time using the plant dynamics in order to guess the state at the time the controller's currently-generated output is to reach the plant. When the delay compensator is applied, the communication sequences as well as the Kalman and LQ optimal gains computed for the delay-free case can be re-used, by merely time-shift the Kalman filter gains by an appropriate number of steps.

A. Numerical Results

To illustrate our approach, we simulated the 2-input, 2-output, 4th order unstable LTI plant with parameters

$$A = \begin{bmatrix} -1.05 & 0 & 0 & 0 \\ 0 & 0.75 & 0 & 0 \\ 0 & 1.05 & 1.05 & 0 \\ 0 & 0 & -2 & 0.5 \end{bmatrix},$$

$$B = \begin{bmatrix} 0.5 & 0 \\ 0 & 0 \\ 0 & 0.5 \\ 0 & 0 \end{bmatrix}, C = \begin{bmatrix} 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \end{bmatrix}.$$

The disturbance terms in (1) were $v(\cdot) \sim N(0, 0.35 I_{5\times5})$ and $w(\cdot) \sim N(0, I_{2\times2})$. We assumed a plant-to-controller delay of $\Delta_1 = 2$, and a controller-to-plant delay of $\Delta_2 = 3$ steps. We formulated an infinite-horizon version of the LQG problem, with $Q = 4I_{4\times4}$ and initial conditions $x(0) = [-5, 5, 2, -10]^T$, initial state estimate $\hat{x}(0) = 0$, and $\Sigma(0) = 0.2 \cdot I_{4\times4}$. The plant was controlled through a shared communication medium with only one input and one output channels ($w_\rho = w_\sigma = 1$). The plant's Kalman decomposition reveals that the plant is indeed stabilizable and detectable, with a 1-dimensional stable uncontrollable/observable subsystem, and a 3-dimensional unstable controllable/observable subsystem.

Using Cor. 1 and Th. 2, we designed a pair of communication sequences that preserved the stabilizability and detectability of the NCS as follows. We computed the controllability and observability indices corresponding to the (three-dimensional) controllable and observable subsystem; they were $n_1 = 1, n_2 = 2$. Consequently, the input communication sequence was taken to be 3-periodic, with $\rho(k) = \{e_2, e_2, e_1,\ e_2, e_2, e_1, ...\}$. The observability indices for the first and second outputs were $n_1 = 2,\ n_2 = 1$, indicating that a communication sequence that maintains the reconstructibility of the NCS is $\sigma(k) = \{e_2, e_1, e_1,\ e_2, e_1, e_1,\ ...\}$. Having thus guaranteed that the NCS is stabilizable and detectable, we constructed a delay compensator and LQG controller [14]. This involved computing the 3-periodic SPPS solutions and of the DPREs associated with the Kalman filter and optimal LQ controller, respectively, and calculated the required Kalman filter gains and the LQ optimal feedback gains. The state evolution of the closed-loop system under optimal control is shown in Fig. 2-(a). Because of the transmission delays, the controller received its first sensor measurement at $k = 2$, so that the first input arrived at the plant at $k = 5$, with the plant running under zero control until that time.

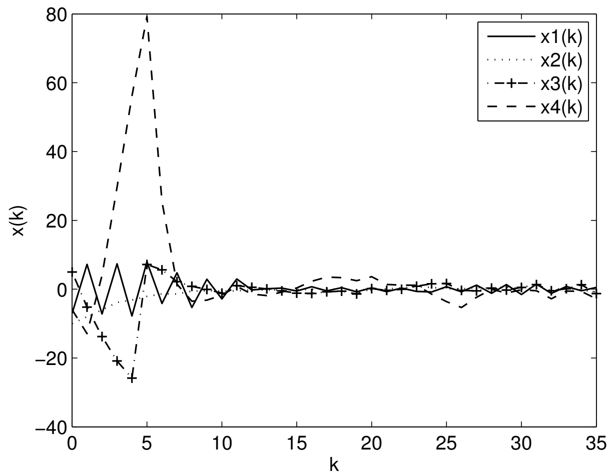

(a)

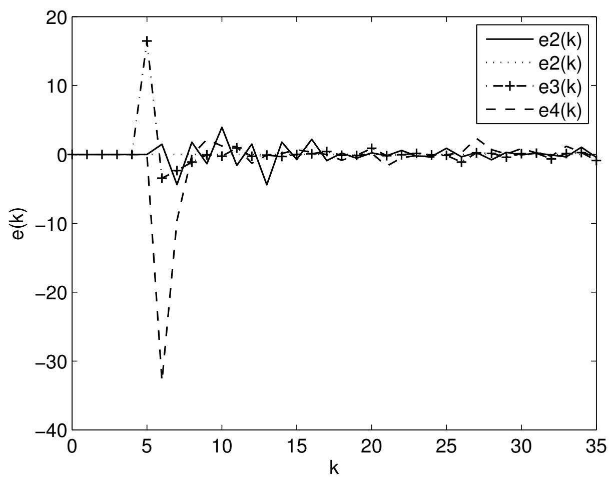

(b)

Fig. 2. (a): State evolution of the closed-loop stochastic NCS (no ZOH) under optimal LQG control. The loop delays were $\Delta_1 = 2,\ \Delta_2 = 3$. (b): Difference between the response in (a) and that of the ZOH-inclusive NCS for the same initial conditions and noise sample paths.

For comparison, we constructed an optimal LQG controller for the ZOH-inclusive periodic NCS (8), with the same parameters, delays, communication sequences, and cost function penalizing the states and the input (in this case, the contents of the ZOH). The plant satisfied condition (9), thus the communication sequences designed for the non-ZOH NCS preserved the structural properties of the ZOH-inclusive model as well, in the sense described in Cor. 2.

The response of the ZOH-inclusive NCS was very similar to that of the simpler model; the difference, $e(k)$, between the state trajectory in Fig. 2-(a) and that of the ZOH-inclusive plant (under the same initial conditions and noise sample paths) is shown in Fig. 2-(b). With the addition of a ZOH however, the average LQG cost was higher. In particular, for 1000 simulations starting from the same initial conditions and using the same noise sample paths in both models over $k \in [0, 5000]$, the average cost was 5.504×10^5 for the non-ZOH model versus 6.07×10^5 for the ZOH-inclusive NCS. It is interesting that the inclusion of a ZOH did not lead to better long-term performance. We would expect the opposite to be true whenever successive input values are correlated, and have observed such cases in our numerical experiments with various plant parameters.

V. Conclusions and Future Work

We discussed the control of NCSs which are subject to medium access constraints and known delays. In that setting, control and communication are coupled because the optimality of a controller depends on the sequence in which medium access is alloted to the plant's sensors and actuators. Solving the problem of simultaneously optimizing control and communication appears to be intractable, and we made use of a "decoupling" technique which restricts communication to periodic sequences that preserve the plant's structural properties.

Forgoing the use of ZOH (i.e., sensors an actuators are ignored while they are "off-line") lowers the complexity of the NCS model and makes the selection of useful communication sequences straightforward, based only on the plant's controllability/observability indices. In particular, it is always possible to design periodic communication sequences that preserve the detectability and stabilizability of the underlying plant in the presence of communication constraints, with an upper bound on their period which is lower than previously established ones. Furthermore, if the plant's parameters satisfy an orthogonality condition, communication sequences designed under the non-ZOH model, are equally effective in the more practical setting where a ZOH *is* included. Interestingly, the inclusion of a ZOH at the plant's input stage does not always result in a lower LQG cost.

Opportunities for future work include extensions of this work to NCSs which are subject to "dropped" data packets and random delays, in addition to the constraints discussed here. It would also be interesting to pursue the question of whether optimal communication sequences are generally periodic, as suggested by the results of [17] and other work on sensor scheduling.

References

[1] D. Hristu-Varsakelis and R. W. Brockett. Experimenting with hybrid control. *IEEE Control Systems Magazine*, 22(1):82–95, Feb. 2002.

[2] D. Hristu-Varsakelis and W. S. Levine. An undergraduate laboratory for networked digital control systems. *IEEE Control Systems Magazine*, 25(1):60–62, Feb. 2004.

[3] H. Zeltwanger. An inside look at the fundamentals of CAN. *Control Engineering*, 42(1):81–7, Jan. 1995.

[4] R. W. Brockett. Stabilization of motor networks. In *Proc. of the 34th IEEE Conf. on Decision and Control*, pages 1484–8, 1995.

[5] D. Hristu-Varsakelis and K. Morgansen. Limited communication control. *Systems and Control Letters*, 37(4):193–205, Jul. 1999.

[6] D. Hristu. Generalized inverses for finite-horizon tracking. In *Proc. of the 38th IEEE Conf. on Decision and Control*, volume 2, pages 1397–402, 1999.

[7] H. Rehbinder and M. Sanfridson. Scheduling of a limited communication channel for optimal control. In *Proc. of the* 39^{th} *IEEE Conf. on Decision and Control*, volume 1, pages 1011–16, Dec. 2000.

[8] B. Lincoln and B. Bernhardsson. Efficient pruning of search trees in LQR control of switched linear systems. In *Proc. of the* 39^{th} *IEEE Conf. on Decision and Control*, volume 2, pages 1828–33, Dec. 2000.

[9] D. Hristu-Varsakelis and P. R. Kumar. Interrupt-based feedback control over a shared communication medium. In *Proc. of the IEEE Conf. on Decision and Control*, pages 3223–8, Dec. 2002.

[10] L. A. Montestruque and P. J. Antsaklis. Stability of model-based networked control systems with time-varying transmission times. *IEEE Trans. on Automatic Control*, 49(9):1562–72, Sep. 2004.

[11] L. A. Montestruque and P. J. Antsaklis. Networked Control Systems: A model-based approach. In D. Hristu-Varsakelis and W. S. Levine, editors, *Handbook of Networked and Embedded Control Systems*, pages 601–625. Birkhauser, 2005.

[12] L. Zhang and D. Hristu-Varsakelis. Stabilization of networked control systems: Designing effective communication sequences. In *Proc. of the 16th IFAC World Congress*, Jul. 2005.

[13] L. Zhang and D. Hristu-Varsakelis. Communication and control co-design for networked control systems. *Automatica*, 42(6):953–958, Jun. 2006.

[14] D. Hristu-Varsakelis and L. Zhang. LQG control of networked control systems with access constraints and delays. *International Journal of Control*, to appear 2008.

[15] D. Hristu-Varsakelis. Stabilization of networked control systems with access contstraints and delays. In *Proc. of the 44th IEEE Int'l. Conf. on Decision and Control*, Dec. 2006.

[16] C. Ionete and A. Cela. Structural properties and stabilization of ncs with medium access constraints. In *Proc. of the* 45^{th} *IEEE Conf. on Decision and Control*, pages 1141–7, Dec. 2006.

[17] O. C. Imer and T. Basar. To measure or to control:optimal control of lti systems with scheduled measurements and controls. In *Proc. of the 2006 American Control Conf.*, Jun. 2006.

[18] N. Elia and S. K. Mitter. Stabilization of linear systems with limited information. *IEEE Trans. on Automatic Control*, 46(9):1384–1400, Sept. 2001.

[19] W. S. Wong and R. W. Brockett. Systems with finite bandwidth constraints - part II: Stabilization with limited information feedback. *IEEE Trans. on Automatic Control*, 42(5):1049–1052, 1999.

[20] S. Tatikonda and S. Mitter. Control under communication constraints. *IEEE Trans. Automatic Control*, 49(7):1056–68, Jul. 2004.

[21] G. N. Nair and R. J. Evans. Stabilization with data-rate-limited feedback: tightest attainable bounds. *Systems and Control Letters*, 41(1):49–56, Sept. 2000.

[22] B. Azimi-Sadjadi. Stability of networked control systems in the presence of packet losses. In *Proc., 42nd IEEE Conf. on Decision and Control*, pages 676 – 81, Dec. 2003.

[23] O. C. Imer, S. Yuksel, and T. Basar. Optimal control of dynamical systems over unreliable communication links. In *Proc. of the 6th IFAC Symposium on Nonlinear Control Systems (NOLCOS)*, Sep. 2004.

[24] J. Nilsson. *Real-Time Control Systems with Delays*. PhD thesis, Department of Automatic Control, Lund Institute of Technology, Lund, Sweden., 1998.

[25] B. Lincoln and B. Bernhardsson. Optimal control over networks with long random delays. In *Proc. of the Int'l. Symp. on Mathematical Theory of Networks and Systems*, 2000.

[26] Chung-Yao Kao and Bo Lincoln. Simple stability criteria for systems with time-varying delays. *Automatica*, 40(8):1429–1434, August 2004.

[27] D. Hristu-Varsakelis. Feedback control with communication constraints. In D. Hristu-Varsakelis and W. S. Levine, editors, *Handbook of Networked and Embedded Control Systems*, pages 575–600. Birkhauser, 2005.

[28] L. Zhang. *Access Scheduling and Controller Design in Networked Control Systems*. PhD thesis, Department of Mechanical Engineering, University of Maryland, College Park, USA, 2005.

[29] D. P. Bertsekas. *Dynamic Programming and Optimal Control*. Athena Scientific, Belmont, MA, 2nd edition, 2000.

[30] S. Bittanti, P. Colaneri, and G. De Nicolao. The difference periodic Riccati equation for the periodic prediction problem. *IEEE Trans. on Automatic Control*, 33(8):706–712, Aug. 1988.

[31] L. Zhang and D. Hristu-Varsakelis. LQG control under limited communication. In *Proc. of the 44th IEEE Conf. on Decision and Control*, pages 185–190, Dec. 2005.

[32] R. Luck and A. Ray. An observer-based compensator for distributed delays. *Automatica*, 26(5):903–7, 1990.