Proceedings of the American Control Conference
Chicago, Illinois • June 2000

Stabilization of LTI Systems with Communication Constraints[1]

Dimitrios Hristu[2]
Institute for Systems Research,
University of Maryland,
College Park, MD 20742
hristu@isr.umd.edu

Abstract

This work is aimed at exploring the interaction of communication and control in systems whose sensors and actuators are distributed across a shared network. Examples of such systems include groups of autonomous vehicles, MEMS arrays and other network-controlled systems. We generalize recent results concerning the stabilization of LTI systems under limited communication. We seek a stabilizing static output feedback controller whose communication with the underlying plant follows a given periodic pattern. We present an algorithm that allows us to pass to a time-invariant formulation of the problem and use simulated annealing to search for stabilizing feedback gains.

1 Introduction

Recent advances in communications and networking technologies are now enabling the construction of complex systems whose sensors, actuators and computing elements are connected by means of a network. Examples of such systems include groups of vehicles, satellite clusters, smart structures and MEMS arrays, to name a few. The performance of these distributed systems is often limited as much by the lack of time on a shared network of sensors and actuators as it is by lack of computational power. This fact has led to recent efforts towards bringing together aspects of control and communication, under a framework that will lead to a better understanding of control systems with communication constraints [3, 8, 5, 7, 2].

One approach to analyzing the effects of communication on the control of a distributed system, is to introduce a "communication sequence" [3] which allows multiple (sub)systems to share the attention of a centralized controller [5]. Communication sequences quantify the amount of "attention" that the decision-maker pays to each component of a control system. Previous work has addressed the problem of stabilizing a set of LTI systems when the controller can only communicate with one sub-system at a time [6]. In this paper we generalize those results by separating sensing and actuation events without the need for grouping states into sub-systems. In Sec. 4 we present a new "extensification algorithm" which transforms the stabilization problem into an equivalent problem, involving a search for stable elements of a family of matrices [6].

2 A Prototype Computer-Controlled System

Consider an n-dimensional LTI system with input $u \in \mathbb{R}^m$ and output $y \in \mathbb{R}^p$. The system is driven by a digital controller (Fig. 1) that does not have simultaneous access to all inputs/outputs of the control system. In particular:
• The controller sends inputs to and receives measurements from the system every Δ units of time, via a zero-order-hold stage.
• Inputs/outputs are transmitted via a bus which has limited capacity. Specifically, the bus can "carry" at most $b > 0$ signals, with $b < m + p$, $b \in \mathbb{N}^*$.

The capacity of the communication bus is to be

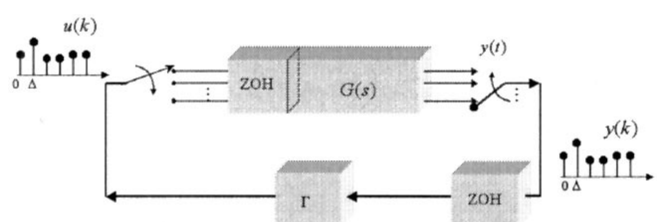

Figure 1: A closed-loop computer-controlled system

split between input and output signals, with b_r channels used for sampling the output of the underlying LTI system and b_w channels used for transmitting control inputs. We will refer to these two groups of channels as the "input" and "output bus." Of course, b_r and b_w may change at any time as long as $b_r + b_w = b$. This represents a rather general

[1]This work was funded by NSF Grant no. EEC 94-02384 and ARO Grant no.DAAG 55-97-1-0144.
[2]This work was done while the author was with the Division of Engineering and Applied Sciences at Harvard University, Cambridge, MA.

0-7803-5519-9/00 $10.00 © 2000 AACC

setting for describing controller-plant communication, allowing for dynamic reconfiguration of the available channels. To simplify the discussion, we will take the number of input and output channels to be constant, with $b_w < m$ and $b_r < p$.

Consider the discretization of $G(s)$ (with sampling period Δ) and let $x(k+1) = Ax(k)+Bu(k)$, $y(k) = Cx(k)$ be a state-space representation of the resulting discrete-time system. Each element of u retains its value (by virtue of the ZOH stage) until that element is updated by the controller. At the same time, the controller may only receive partial information about the output $y(k)$. One way of negotiating the constraints imposed by the bus, is to choose a sequence of operations for the switches (see Fig. 1) that select which inputs/outputs are to be updated/sampled at a particular time. This is captured in the idea of a "communication sequence" (originally introduced in [3]):

Definition 1 *An* **N-periodic communication sequence** *is an element of*

$$\mathbb{E}_{per}^{m\times N} = \{(\sigma(0), \sigma(1), ..., \sigma(N-1), \sigma(0), ..., \sigma(N-1), ...) : \sigma(i) \in \{0,1\}^m\} \quad (1)$$

for some $m > 0$.

Controller-plant communication follows a periodic pattern, specified by a pair of N-periodic sequences: a "control" sequence $\sigma_w \in \mathbb{E}_{per}^{m\times N}$ will be used to transmit inputs and a "measurement" sequence $\sigma_r \in \mathbb{E}_{per}^{p\times N}$ will provide a pattern for sampling the system output. The entries of $\sigma_w(i)$ ($\sigma_r(i)$) indicate which elements of $u(k)$ ($y(k)$)are to be updated (measured) at the k^{th} time step. We will ignore quantization errors associated with the representation of signal samples in the digital controller and with the transmission of those samples through the communication bus.

Definition 2 *Consider a computer-controlled system $G(z)$ with $b_w < m$ ($b_r < p$) being the dimension of the input (output) communication bus. A pair of communication sequences $\sigma_w \in \mathbb{E}^{m\times N}$, $\sigma_r \in \mathbb{E}^{p\times N}$ is* **admissible** *if:*

- $||\sigma_w(i)||^2 \leq b_w, ||\sigma_r(i)||^2 \leq b_r \quad \forall i = 0, ..., N-1$
- *$Span\{\sigma_w(0), \ldots, \sigma_w(N-1)\} = \mathbb{R}^m$ and $Span\{\sigma_r(0), \ldots, \sigma_r(N-1)\} = \mathbb{R}^p$*

The above conditions require that no more than b_w (b_r) of the system inputs (outputs) be updated (measured) by the controller at each step and that the pair (σ_w, σ_r) allow communication with all inputs and outputs of the linear system at least once every period.

3 Stabilization with Limited Communication

We now focus on the problem of stabilizing a computer-controlled system with communication constraints, using static output feedback.

Problem 1 *Given: a computer-controlled LTI system $G(z)$ with $b_r, b_w \in N^*$ denoting the size of the input and output communication busses and a pair of admissible N-periodic communication sequences $\sigma_r, \sigma_w \in \mathbb{E}_{per}^{p\times N}, \mathbb{E}_{per}^{m\times N}$, find a constant output feedback gain $\Gamma \in \mathbb{R}^{p\times m}$ that stabilizes the closed-loop system.*

In [6], we showed (by providing a so-called "extensification algorithm") that a version of this problem was equivalent to the NP-hard problem:

Problem 2 *Given a collection of matrices $A_i \in \mathbb{R}^{q\times q}, 0 \leq i \leq i_{max}$, and scalars $\gamma_1, ..., \gamma_{i_{max}} \in \mathbb{R}$, find a stable element of the affine subspace*

$$\mathcal{A} = \mathcal{A}_0 + \sum_{i=1}^{i_{max}} \gamma_i \mathcal{A}_i \quad (2)$$

In the next section we present a new, more general algorithm which considers actuation and measurement events separately, requires no a-priori grouping of states into "sub-systems" and arrives at a similar construction for the matrices that span the affine subspace of interest, with $q = (2N^2 - N)n$ and $i_{max} = mp$ in the statement of Problem 2.

4 Extensive Form of a Discrete LTI System

Consider the (discretized) system of Fig. 1, to which a static output feedback controller is attached, subject to limited communication as described in Sec. 2.

4.1 Constrained Measurements

Assume for now that $b_w = m$ so that the controller can transmit an entire input vector at once to the LTI system. Only some of the elements of $y(k)$ are received by the controller at each step (as dictated by $\sigma_r \in \mathbb{E}_{per}^{p\times N}$), with all remaining elements holding their last-known values. More precisely,

$$u(k) = \Gamma y_l(k) \quad (3)$$

where $y_l(k)$ is the output vector composed of the most up-to-date measurements available to the controller at the k^{th} step. Typically, we expect that $y_l(k) \neq y(k)$. We write:

$$y_l(k) = diag\Big(\sigma_r(k)\Big)y(k) + \Big(I - diag(\sigma_r(k))\Big)y_l(k-1) \quad (4)$$

where for a vector $x \in \mathbb{R}^n$, $diag(x)$ is an $n \times n$ matrix with the elements of x along its diagonal, all other entries being zero. By iteratively applying Eq. 4 for a number of steps equal to the communication period N, we obtain:

$$y_l(k) = \sum_{i=0}^{N-1} D_R(k,i)Cx(k-i) \quad (5)$$

where

$$D_R(k,i) \triangleq \begin{cases} diag(\sigma_r(k)) & i=0 \\ diag(\sigma_r(k-i)) \prod_{j=0}^{i-1} M_R(k,j) & i>0 \end{cases} \tag{6}$$

$$M_R(k,j) \triangleq I - diag\left(\sigma_r(k-j)\right) \tag{7}$$

Note that $D_R(k,i)$ are diagonal $p \times p$ matrices. The j^{th} diagonal element of $D_R(k,i)$ is 1 if the j^{th} output was last read at the $(k-i)^{th}$ step and is 0 otherwise. We observe that if the communication sequence σ_r is admissible then each of the p elements of the output $y(k)$ will be read at least once every N steps so that the summation in Eq. 5 terminates after at most N terms.

4.2 Constrained Control

We now consider controller-plant communication over the input bus. At the k^{th} step, only the inputs specified by the non-zero entries of $\sigma_w(k)$ are updated, with all other inputs retaining their previous values:

$$u(k) = diag\big(\sigma_w(k)\Gamma y_l(k) + \Big(I - diag\big(\sigma_w(k)\big)\Big)u(k-1) \tag{8}$$

By iterating backwards for a full period (N steps) and assuming that the communication sequence σ_w is admissible, we obtain:

$$u(k) = \sum_{i=0}^{N-1} D_W(k,i)\Gamma y_l(k-i) \tag{9}$$

where D_W (and M_W) are obtained from Eq. 6, 7 simply by replacing σ_r with σ_w. In this case $D_W(k,i)$ is an $m \times m$ diagonal matrix.

4.3 Combining Communication Constraints

Substituting Eq. 5 into Eq. 9, we obtain:

$$Bu(k) \;=\; \sum_{i=0}^{2N-2} F_{ki}x(k-i) \tag{10}$$

where

$$F_{ki} \triangleq B \sum_{j=min(i,N-1)}^{\lfloor \frac{i}{N} \rfloor (i-N-1)} D_W(k,j)\Gamma D_R(k-j,i-j)C \tag{11}$$

It follows that the closed-loop dynamics of the computer-controlled system are given by:

$$x(k+1) = Ax(k) + \sum_{i=0}^{2N-2} F_{ki}x(k-i) \tag{12}$$

Define $Comp(p)$ to be the companion form associated with an n^{th}-degree polynomial $p(s) = \sum_0^n p_i s^i$:

$$Comp(p) \triangleq \begin{bmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & 0 \\ 0 & \cdots & \cdots & 0 & 1 \\ p_n & p_{n-1} & \cdots & p_1 & p_0 \end{bmatrix} \tag{13}$$

If we now use the F_{kj} (Eq. 11) to define the matrix polynomials

$$f_k(s) = A + \sum_{i=0}^{2N-2} F_{kj}s^i \tag{14}$$

then the closed-loop dynamics of Eq. 12 can be expressed in first-order form:

$$\chi(k+1) = Comp(f_k)\chi(k) \tag{15}$$

where $\chi = [x^T_{(k-2N+3)} \;\cdots\; x^T_{(k)}\; x^T_{(k+1)}]^T \in \mathbb{R}^{(2N-1)n}$.

The (linear) system of Eq. 15 is N-periodic in k, and describes the state evolution of the original computer-controlled system under output feedback and periodic communication. We have essentially "extensified" the state vector to include past values up to two communication periods.

It is a fact that every discrete-time periodic system can be expressed as a time-invariant system of higher dimension [4]. Applied to Eq. 15, this fact yields a system of order $(2N^2 - N)n$ which we call the "**extensive form**" of the original system of Problem 1:

$$\mathcal{X}_e(k+1) = \mathcal{A}\mathcal{X}_e(k) \tag{16}$$

where $\mathcal{X}_e(k) \in \mathbb{R}^{(2N^2-N)n}$ and

$$\mathcal{A} = \begin{bmatrix} 0 & \cdots & 0 & 0 & Comp(f_0) \\ Comp(f_1) & 0 & \cdots & 0 & 0 \\ 0 & Comp(f_2) & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \cdots & 0 \\ 0 & \cdots & 0 & Comp(f_{N-1}) & 0 \end{bmatrix} \tag{17}$$

By construction, stability of the extensified system (Eq. 16) is equivalent to the stability of the original system. Moreover, each of the matrices $Comp(f_k)$ are affine in the entries of Γ. By choosing a basis for $\mathbb{R}^{m \times p}$, we can express Γ as $\Gamma = \sum_{i=0}^{mp} \gamma_i E_i$ where E_i is an $m \times p$ matrix whose $(\lfloor \frac{i-1}{p} \rfloor + 1, (i-1)\ mod\ p + 1)^{th}$ entry is "1", with all other entries being zero. In the basis of the $\{E_i\}$ we can write $\mathcal{A}$ as an element of the affine subspace

$$\mathcal{A} = \mathcal{A}_0 + \sum_{i=0}^{mp} \gamma_i \mathcal{A}_i \tag{18}$$

where each of the $\mathcal{A}_i$ are obtained by substituting E_i for Γ in Eq. 11.

In summary, we have given a procedure for converting an output feedback stabilization problem involving LTI systems under limited communication, into a search problem involving a finite collection of $(2N^2 - N)n$-dimensional matrices. These matrices are obtained from the parameters of the original LTI system together with a pair of admissible communication sequences.

5 An Extensification Example

Consider the scalar system

$$x(k+1) = ax(k) + u(k); \qquad y(k) = x(k) \tag{19}$$

We assume that the controller communicates with the above system over a bus of width $b = 1$ according to the following pair of 2-periodic communication sequences:

$$\sigma_r = (1,0,1,0,\ldots), \quad \sigma_w = (0,1,0,1,\ldots) \tag{20}$$

so that the communication channel is used both for measuring the output and for transmitting the control, in an alternating fashion. Clearly, the above sequences are admissible. We want to stabilize the system using a control law of the form

$$u(k) = \gamma y(k - d(\sigma_r, \sigma_w, k)) \tag{21}$$

where $d(\sigma_r, \sigma_w, k)$ is a delay that depends on the communication sequences and the current step k.

Following the procedure outlined in Sec.4, we obtain a 3-dimensional periodic system $\chi(k+1) = Comp(f_k)\chi(k)$ with

$$f_k(s) = \begin{cases} a + \gamma s^2 & k \text{ even} \\ a + \gamma s & k \text{ odd} \end{cases} \tag{22}$$

The corresponding companion forms are:

$$Comp(f_k) = \begin{cases} \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ \gamma & 0 & a \end{bmatrix} & k \text{ even} \\ \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & \gamma & a \end{bmatrix} & k \text{ odd} \end{cases} \tag{23}$$

and the extensive form is given by the 6-dimensional LTI system:

$$\mathcal{X}_e(k+1) = \left[\begin{array}{ccc|ccc} & & & 0 & 1 & 0 \\ & 0 & & 0 & 0 & 1 \\ & & & \gamma & 0 & a \\ \hline 0 & 1 & 0 & & & \\ 0 & 0 & 1 & & 0 & \\ 0 & \gamma & a & & & \end{array}\right] \mathcal{X}_e(k) \tag{24}$$

which is to be made stable by choice of γ. The coefficients matrices $\mathcal{A}_0$ and $\mathcal{A}_1$ (Eq. 18) can now be read off from Eq. 24.

6 Finding a Set of Stabilizing Gains

It has been shown that a version of Problem 2 is NP-hard [1] thus in general, one cannot hope to obtain an analytic solution for the stabilizing gains γ_i. One possibility is to pose an optimization problem by asking that the eigenvalues of $\mathcal{A} = \mathcal{A}_0 + \sum_i \gamma_i \mathcal{A}_i$ be enclosed in a circle with the smallest possible radius. This suggests minimizing the spectral radius of the closed-loop system

$$\rho = ||\lambda_{max}(\mathcal{A})|| \tag{25}$$

where $\lambda_{max}(\mathcal{A})$ denotes the largest-magnitude eigenvalue of $\mathcal{A}$. To negotiate the large number of local minima that are expected, we applied simulated annealing on the gains γ_i. Our algorithm numerically computes the gradient $\partial\rho/\partial\gamma_i$ and then lets the γ_i flow along that gradient, adding a white-noise term dw with a gain $g(t)$ that decays to zero:

$$d\gamma_i = \frac{\partial \rho}{\partial \gamma_i} dt + g(t) dw. \tag{26}$$

The "cooling schedule" $g(t)$ should go to zero as $t \to \infty$, but it should do so at a slow enough rate for the spectral radius to approach the global minimum.

7 Simulation Results

Consider the fourth-order, two-input, two-output LTI system:

$$\begin{aligned} x(k+1) &= \begin{bmatrix} 1 & 3/4 & 1/2 & 0 \\ 1/4 & 3 & 1/3 & -1/3 \\ 1/6 & 0 & -1/2 & -3/7 \\ 0 & -1 & 2/5 & 0 \end{bmatrix} x(k) + \begin{bmatrix} 0 & 0 \\ 1 & 0 \\ 0 & 0 \\ 0 & 1 \end{bmatrix} u \\ y(k) &= \begin{bmatrix} 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \end{bmatrix} x(k) \end{aligned} \tag{27}$$

The open-loop system has a spectral radius of $||\lambda_{max}|| = 3.2$ which is also the spectral radius of the extensive form, for $\gamma_i = 0, \quad i = 1, \ldots, 4$. We want to stabilize this LTI system using static output feedback, given that the communication bus can carry two signals to/from any of the inputs or outputs (i.e. $b = 2$). In the following, we investigate the performance of the simulated annealing algorithm for two different communication patterns.

7.1 Control using Non-uniform Attention

We selected a period-four pair of communication sequences,

$$\sigma_w = \sigma_r = \left(\begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \cdots \right) \tag{28}$$

that devote three cycles to the pair (u_1, y_1) for every one cycle allocated to (u_2, y_2). The above sequences were chosen after some experimentation and by noticing that the upper-left 2×2 block of the dynamics for the state evolution (Eq. 27) has a larger spectral radius than the lower-right block (when the coupling between the two blocks is removed). As a result, communicating more often with the (u_1, y_1) pair may lead to better performance.

The matrices composing the extensive form were computed and simulated annealing was performed on the four elements of the feedback matrix Γ. The simulated annealing algorithm was stopped after 5000 steps (approximately 15 minutes on a $500MHz$ PC). The cooling schedule was given by

$$g(i) = \frac{30}{log((1+i))^{(0.8+i/3000)}} \qquad i = 1, \ldots, 4800 \tag{29}$$

i.e. logarithmic decay to a level of 0.11, followed by a linear decay to zero in an additional 200 steps. Equation 29 and its parameters were chosen after some numerical experimentation. Choosing "good" cooling schedules for the stabilization problem considered here remains an open problem. Simulated annealing stabilized the closed-loop system, reducing the spectral radius to 0.795. The final closed-loop eigenvalues are shown in Fig. 2. with the evolution of the spectral

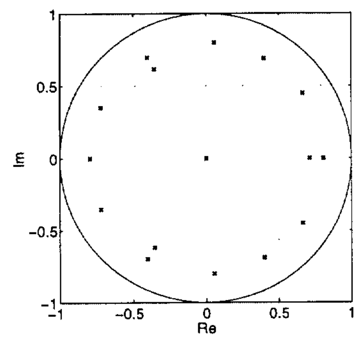

Figure 2: Closed-loop eigenvalues with non-uniform attention, $\lambda_{max} = 0.795$.

radius of the system shown in Fig. 3.

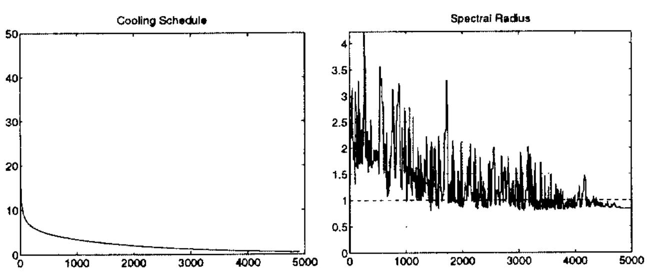

Figure 3: Cooling schedule and evolution of spectral radius (non-uniform attention).

7.2 Control with Uniform Attention

Next, we chose communication sequences corresponding to "uniform attention"

$$\sigma_w = \sigma_r = \left(\begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \cdots \right) \tag{30}$$

with the same cooling schedule as in the in the non-uniform attention case. This time, simulated annealing did not lead to stabilizing gains, even after repeated trials with slower-decaying cooling schedules. The spectral radius of the closed-loop system reached a minimum level of 2.7.

8 Conclusions and Future Work

We discussed the stabilization of LTI systems which operate under limited communication. Our approach is based on the use of periodic communication sequences which direct the flow of control and measurement signals between controller and plant. It is expected (and verified in numerical experiments) that some communication sequences allow more effective control than others. We gave a general "extensification" algorithm for converting the stabilization problem to a time-invariant form and proposed an optimization method for finding stabilizing feedback gains.

There are several issues which require further study, including methods for finding "good" communication sequences, cooling schedules and stopping criteria for simulated annealing. It would be interesting to develop models which provide for more general (perhaps interrupt-based) communication protocols. The design of state observers with limited communication seems to be of importance, especially in light of the NP-hardness result regarding the stabilization problem.

9 Acknowledgment

The author would like to thank Prof. R. W. Brockett for his helpful discussions on the subject.

References

[1] V. Blondell and J. Tsitsiklis. NP hardness of some linear control design problems. *SIAM Journal on Control and Optimization*, 35(6):2118–2127, Nov. 1997.

[2] V. S. Borkar and S. K. Mitter. LQG control with communication constraints. Technical Report LIDS-P-2326, MIT, 1995.

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

[4] Jon H. Davis. Stability conditions derived from spectral theory: discrete systems with periodic feedback. *SIAM Journal of Control*, 10(1):1–13, Feb. 1972.

[5] D. Hristu. Generalized inverses for finite-horizon tracking. In *Proc. of the 38 IEEE Conf. on Decision and Control*, Dec. 1999.

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

[7] G. Walsh, H. Ye, and L. Bushnell. Stability analysis of networked control systems. In *Proc. of the American Control Conf.*, pages 2876–80, Jun. 1999.

[8] 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, May 1999.