TECHNICAL RESEARCH REPORT

Interrupt-based feedback control over a shared communication medium

by D. Hristu-Varsakelis, P. R. Kumar

TR 2003-34

Interrupt-based feedback control over a shared communication medium [1]

D. Hristu-Varsakelis
Mechanical Engineering and
Institute for Systems Research
University of Maryland
College Park, MD
hristu@glue.umd.edu

P. R. Kumar
Electrical and Computer Engineering
University of Illinois Urbana-Champaign,
Urbana, IL
prkumar@uiuc.edu

Abstract

This work is a continuation of recent efforts aimed at understanding the interplay of control, communication and computation in systems whose sensors, actuators and computing elements are distributed across a network. We investigate the simultaneous stabilization of a group of linear systems whose feedback loops are closed over an idealized shared medium. The capacity of that medium is constrained so that only a limited number of controller-plant connections can be accommodated at any one time. We introduce a feedback communication policy – inspired by previous work on queuing systems and real-time scheduling – for deciding which system(s) should be admitted into the network and for how long. The use of feedback in making communication decisions results in a set of autonomous dynamical systems which are coupled to one another due to the presence of communication constraints. We give conditions for the stability of the collection under the proposed communication policy and present simulation results that illustrate our ideas.

1 Introduction

The ongoing interest in networked control systems has motivated a host of research on the analysis and control of dynamical systems whose operation is subject to communication constraints (see for example [7, 18], also [10] and references therein). One of the directions that are being pursued focuses on understanding the effects of communication constraints on variants of classical control problems through the use of so-called "communication sequences" [3, 8, 6, 16] which specify the spatial and temporal order in which a system's sensors, actuators and controller exchange data. In many of these works, communication events occur at integer multiples of a common period; a (sometimes periodic) communication sequence is chosen a priori, then an optimal control problem is solved with the communication sequence regarded as a parameter. Such polling-based communication schemes preserve linearity (when the underlying systems are themselves linear) and often make the analysis easier.

One of the challenges in the approach outlined above lies in the difficulties encountered when jointly optimizing the control law and communication sequence; in some cases, the best one can hope for is a useful heuristic (for example [12]). Moreover, specifying a communication sequence a priori may be undesirable for practical reasons, including the need for a timer and memory as well as the lack of a feedback mechanism that would allow for changes in the communication sequence in response to disturbances or other events. For these reasons it is sometimes preferable to use an interrupt-based policy for deciding which parts of a networked dynamical system should interact at a particular time. Such a scheme is appealing because it leads to autonomous dynamics and because there is at least some evidence [17] that suggests that feedback communication policies can outperform their "static" counterparts.

In this paper we investigate the stability of a collection of continuous-time feedback LTI systems that rely on an idealized network of limited capacity in order to close their feedback loops. We show that a sufficient condition for stability of the collection under periodic communication (detailed in [9]) is also sufficient under an interrupt-based communication policy that is inspired by results from real-time scheduling and queuing theory [15]. We note that our model for limited communication leads to systems with switched or hybrid dynamics; in those settings known stability conditions are typically conservative [2, 17] or computationally complex (see [11] for an excellent survey of standard results). Our plan for improving on existing results is based on a simple communication policy that "pays attention" to the system whose state is furthest

[1] This research was supported by ODDR&E MURI01 Grant No. DAAD19-01-1-0465, (Center for Communicating Networked Control Systems, through Boston University) and by NSF CRCD Grant No. EIA 0088081.

from the origin (this will be made precise in Sec. 3).

Section 2 describes the stability problem we are concerned with. We propose a simple interrupt-based strategy (a variant of the "Clear the Largest Buffer" policy introduced in [15]) for deciding on-line which systems should be allowed use of the network. In Section 3 we give sufficient conditions for the stability of a collection of dynamical systems under the proposed communication policy, and improve upon previously established results. In the case of scalar dynamical systems operating under the proposed communication policy, our stability condition is necessary and sufficient. Section 4 contains simulation results that illustrate the main ideas.

2 A collection of switched dynamical systems

Consider a collection of continuous-time LTI systems

$$\begin{aligned} \dot{x}_i(t) &= A_i x_i(t) + B_i u_i(t); \quad i = 1, \ldots, N \\ x_i(t) &\in \mathbb{R}^n, u_i(t) \in \mathbb{R}^m \end{aligned} \tag{1}$$

whose open loop dynamics are unstable ($Re\{\lambda(A_i)\} > 0, \quad i = 1, \ldots, N$). Each system communicates with a remotely located controller over an idealized shared network, according to the static state feedback law[1] $u_i(t) = K_i x_i(t)$ (see Fig. 1), with the gains K_i designed

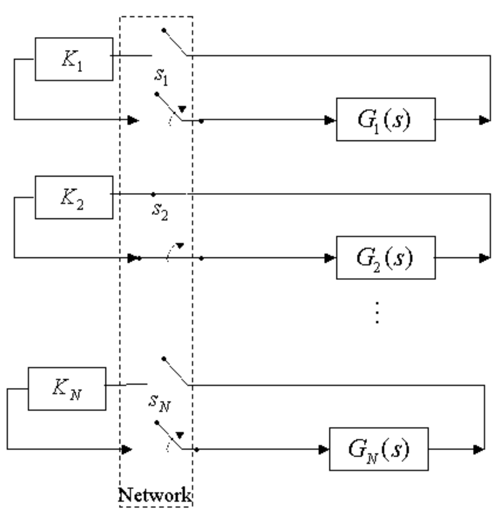

Figure 1: A collection of networked control systems $G_i(s) = I(sI - A_i)^{-1} B_i$ driven by static feedback controllers K_i via a network. Only k of N switches s_i can be closed at any one time.

a priori so that $Re\{\lambda(A_i + B_i K_i)\} < 0, \quad i = 1, \ldots, N$.

We assume that the "shared network" is an idealized communication medium which provides connectivity between a system and its controller in a discrete sense (*on* or *off*). We will not consider the effects of network layers, delays or complications due to packed-based communication. Controller-plant communication is limited in the sense that a maximum of $k < N$ plants may close their feedback loops at any one time. This could be because:

- each system has its own controller, with communication taking place over a medium which can only accommodate a maximum of k "users" at a time, or
- all systems are stabilized by a single centralized controller which can perform a limited amount of computation per unit time, or has only enough outputs to communicate with k out of N systems.

We are interested in finding a **feedback-based** *sequence for establishing and terminating communication between each system and its controller, in a way that stabilizes all systems in the collection.* By "feedback-based" we mean a communication sequence that is a function of the system states.

It should be noted that our switch-based model (like those in [8, 6]) captures computational as well as communication constraints. When the network is busy, the controller may have to wait to receive data from its sensor suite. On the other hand, even if bandwidth is virtually "unlimited", a centralized controller may be unable to process data from all N sensors suites simultaneously when it comes to evaluating the corresponding feedback laws. In this paper we will always refer to "communication constraints", but it should be understood that our model can capture both types of constraints and in fact the two cases are indistinguishable as far as the evolution of the system is concerned.

We use $A_i^o \triangleq A_i$ and $A_i^c \triangleq A_i + B_i K_i$ to denote the open and closed loop dynamics of the collection, and write the dynamics of the resulting switched system as:

$$\dot{x}_i(t) = A_{s_i(t)} x_i(t); \quad i = 1, \ldots, N \tag{2}$$

where $A_{s_i(t)} \in \{A_i^o, A_i^c\}$, and $s_i(t) \in \{0, 1\}$ are piecewise constant functions that indicate when the i^{th} loop is closed ($s_i(t) = 1$). For the systems (2), we represent the communication sequence [3, 8, 9] by a piecewise constant $\sigma : \mathbb{R}^+ \to \{0, 1\}^N$ with

$$\sigma(t) = [s_1(t) \; s_2(t) \; \cdots \; s_N(t)]^T$$

and express the communication constraint as $\sum_1^N \sigma_i(t) \leq k$. Because we have assumed that the A_i^c are stable by choice of K_i, there exist quadratic Lyapunov functions $V_i^c(x_i) = x_i^T P_i x_i$, $P_i = P_i^T > 0$

[1] Here we have assumed that state is available, although our results apply in the case of output feedback as well.

such that for $Q_i = Q_i^T > 0$,

$$(A_i^c)^T P_i + P_i A_i^c = -Q_i < 0 \quad (3)$$
$$\dot{V}_i^c(x) \le \lambda_i V_i^c(x), \quad t \in [jT, jT + \Delta s) \quad (4)$$

for some $\lambda_i \in \mathbb{R}^-$, $j = 0, 1, \ldots$. The following then holds:

Theorem 1 (from [9]) *Consider the collection of networked LTI systems described in (2) and assume that at most k out of N systems are allowed to close their feedback loops at any one time. For $i = 1, \ldots, N$, let $V_i^c(x_i) = x_i^T P_i x_i$, $P_i = P_i^T > 0$ be Lyapunov functions for the closed-loop systems satisfying $(A_i^c)^T P_i + P_i A_i^c < \lambda_i P_i < 0$ when communication is available (feedback loop closed) and $(A_i^o)^T P_i + P_i A_i^o < \mu_i P_i$ otherwise (for some $\lambda_i < 0$, $\mu_i > 0$). Then, for any $T > 0$, there exists a T-periodic communication sequence that stabilizes all N systems if*

$$\sum_{i=1}^{N} \frac{\mu_i}{\mu_i - \lambda_i} < k. \quad (5)$$

The utility of the bound in (5) depends of course on how closely the estimates λ_i and μ_i approximate the decay/growth rates of the open and closed loop dynamics of the systems. There is a set of optimal quadratic Lyapunov functions (i.e. those P_i that will result the tightest possible λ_i and μ_i), which can be computed by solving the following problem:

Problem 1 *Given A_i^c, A_i^o, $i = 1, \ldots, N$, minimize $c \triangleq -\mu_i/\lambda_i > 0$ over all $P_i = P_i^T > 0$, $\lambda_i < 0$, subject to:*

$$\begin{aligned} (A_i^c)^T P_i + P_i A_i^c &< \lambda_i P_i \\ (A_i^o)^T P_i + P_i A_i^o &< -c\lambda_i P_i \quad (6) \\ P_i > 0, \quad \lambda_i < 0, \quad c &> 0 \end{aligned}$$

The above inequalities are linear in each of the variables P_i, λ_i, c and Problem 1 can be solved using bisection on c and solving the BMI problem ([1, 5]) that results (see [9] for details). In general, (5) is conservative even if the P_i are optimal according to Problem 1.

3 A feedback-based communication policy

Instead of specifying the switching functions $s_i(t)$ in advance for the systems (2), we would like for the $s_i(t)$ to depend only on the states $x_i(t)$. In the following, we will assume without loss of generality that $k = 1$, i.e. the network can only accommodate one feedback loop at a time (the case $k > 1$ follows easily). For $k = 1$, we can simplify matters by setting $s_i(t) = 0$ for all indices $i \in \{1, 2, \ldots, N\}$ except one which we denote by $i_*(t)$, corresponding to the unique system whose feedback loop is closed at t. We now define the following policy for choosing $i_*(t)$ (and therefore the switching functions $s_i(t)$):

Definition 1 (CLS-ϵ) *Let $i_*(t)$ denote the index of the system whose feedback loop is closed at time t.*

- *1. Let t_0 denote the current time. Set $i_*(t_0) = argmax(\|x_i(t_0)\|)$.*
- *2. When $\|x_{i_*}(t)\| = \epsilon > 0$, repeat from step 1.*

This policy, which seeks to "Contain the Largest State" (abrv. CLS-ϵ), can be viewed as the analog of the "Clear the Largest Buffer" policy, originally introduced in the study of distributed manufacturing systems [15]. In the following, we show how those results relate to our problem as we investigate the behavior of the systems described by Eq. 2 under variations of the CLS-ϵ policy.

Remark 1: At first glance one might think that there is a tradeoff between making ϵ too small or too large. When ϵ is large, performance will of course suffer since the states remain large. If ϵ is too small, then one may think that again the performance will be low because the uncontrolled states will tend to become very large while the controller is trying to make a particular state very small (smaller than ϵ). However, that is not true as we shall see.

3.1 Systems with scalar dynamics

Consider the case where $A_i^c, A_i^o \in \mathbb{R}$ for all $i = 1, \ldots, N$.

Theorem 2 *Consider the collection of networked LTI systems described in (2) with $A_{s_i(t)} \in \{A_i^o, A_i^c\}$, $A_i^o > 0$, $A_i^c < 0$. Assume that at most $k = 1$ out of N systems [2] are allowed to close their feedback loops at any one time and that the binary-valued $s_i(t)$ are determined by CLS-ϵ for any fixed $\epsilon > 0$. Then, all $|x_i(t)|$ will approach ϵ if and only if*

$$p \triangleq \sum_{i=0}^{N} \frac{A_i^o}{A_i^o - A_i^c} < 1 \quad (7)$$

Furthermore, if $p > 1$ then there exists no stabilizing communication sequence.

Proof: Without loss of generality, assume that $x_i(t) > 0$. The state equations (2) imply that

$$\ln x_i(t) = A_i^o(t) + \ln x_i(t_0) \qquad \textit{open loop}$$

$$\ln x_i(t) = A_i^c(t) + \ln x_i(t_0) \qquad \textit{closed loop}$$

where t_0 denotes the last switching time. Define

$$V(x) \triangleq \sum_{1}^{N} \frac{\ln(x_i)}{A_i^o - A_i^c} \quad (8)$$

[2] For $k > 1$, replace the right hand side of (7) by k.

and suppose that at time t we are choosing to close the loop of system j. Then,

$$\begin{aligned}\frac{dV}{dt} &= \sum_1^N \frac{1}{x_i(t)} \frac{\dot{x}_i}{A_i^o - A_i^c} \\ &= \frac{A_j^c}{A_j^o - A_j^c} + \sum_{i \neq k} \frac{A_i^o}{A_i^o - A_i^c} \\ &= \frac{A_j^c - A_j^o}{A_j^o - A_j^c} + \sum_1^N \frac{A_i^o}{A_i^o - A_i^c} = -1 + p \qquad (9)\end{aligned}$$

We see that $V < 0$ iff $p < 1$. The condition $p < 1$ is clearly necessary because V is strictly increasing with respect to the x_i. Suppose now that $p < 1$ so that $V(t)$ is strictly decreasing. Because our communication policy requires opening a feedback loop when the associated state reaches $\epsilon > 0$, $V(t)$ will be bounded below by $V(x) = \sum_1^N \frac{\ln(\epsilon)}{A_i^o - A_i^c}$; we conclude that $V(t)$ has a lower limit and so do the $|x_i|$. Finally, suppose that $p > 1$ but that there exists some communication sequence which stabilizes all systems. Then there is a future time T at which all x_i have decreased compared to their values at $t = t_0$. Let $0 < \Delta_i < T$ be the total amount of time that was allocated to the i^{th} system over the interval $[t_0, t_0 + T]$, then $x_i(t_0 + T) = e^{A_i^c \Delta_i} e^{A_i^o (T - \Delta_i)} x_i(t_0) < x_i(t_0)$ for all $i = 1, ..., N$. The last equation requires that $\Delta_i > \frac{T A_i^o}{A_i^o - A_i^c}$ and because $\sum_1^N \Delta_i = T$ we have $\sum_1^N \frac{A_i^o}{A_i^o - A_i^c} = p < 1$, a contradiction. □

From a queuing-theoretic viewpoint, the Lyapunov-like function V is analogous to that used in [15] and attempts to capture the amount of "work" the controller must do in order to bring the states x_i near the origin. Also, from Theorem 2 we see that CLS-ϵ can be used to drive the systems to the origin by gradually decreasing the value of ϵ:

Corollary 1 *If the collection of systems in Eq. 2 is such that $p = \sum_{i=0}^N \frac{A_i^o}{A_i^o - A_i^c} < k$ with $A_i^o, A_i^c \in \mathbb{R}$, then the collection is stable under CLS-ϵ if ϵ varies with time along a sequence that converges asymptotically to zero (e.g. $\epsilon_n = 1/n$).*

Remark 2: Under CLS-ϵ, the switching rate is not bounded. It is possible however to slightly modify the switching policy so that the switching rate is bounded above by $\frac{1}{\tau}$. The "minimum waiting" time $\tau > 0$ will be the analog of the so-called "setup time" in [15]. In that case (which will not be discussed here due to space constraints) the states will remain bounded.

Corollary 2 *If the systems in (2) are all governed by the same set of dynamics $A_{s_i(t)} \in \{A^o, A^c\}$ for all $i = 1, ..., N$ and $A^o, A^c \in \mathbb{R}$, then*

- *As $t \to \infty$, $|x_j| \to \epsilon$ for all $j = 1, ..., N$ if and only if $-\frac{A^o}{A^c} < \frac{1}{N-1}$.*
- *If $-\frac{A^o}{A^c} \geq \frac{1}{N-1}$ then there is no switching policy that stabilizes all systems*
- *The switching policy $i(t)$ will be "round robin".*

Proof: The first two statements follow immediately from Theorem 2. When a system (say x_j) reaches $|x_j| = \epsilon$, its norm is the smallest in the collection. Because the growth rates of all systems are the same, the switching sequence $i(t)$ will return to j precisely after it has taken on all other $N - 1$ indices. This holds for every $j = 1, 2, ..., N$, implying that $i(t)$ is "round robin". □

3.2 The multivariable case

If the systems of Eq. 2 are multivariable, then it is possible that CLS may fail to stabilize the collection but that there are other communication sequences that result in stability. In fact, there are well-known examples of switched systems for which there exists a stabilizing switching sequence even when A^c and A^o are both unstable [2]. This suggests that unlike the scalar case, there may be no necessary condition for stability based solely on the eigenvalues of the systems. We can however obtain sufficient conditions for stability if we are willing to use a **modified CLS-ϵ** policy that makes switching decisions based not on the norms $\|x_i\|$ but rather on the exponential curves that bound the Lyapunov functions from Theorem 1:

$$f_i(t) = \begin{cases} e^{\lambda_i t} V_i(x_i(t_n)) \\ e^{\mu_i t} V_i(x_i(t_n)) \end{cases} \qquad (10)$$

for $t \geq t_n$ where t_n denotes the last switching time, $t_0 = 0$ and λ_i, μ_i are as in Theorem 1 (possibly optimized by solving an instance of Problem 1). We will refer to the $f_i(t)$ in (10) as the *envelope functions*. Because $V_i(t) \leq f_i(t)$ for all i and t and by virtue of our results for the scalar case, we have that

Lemma 1 *Under the assumptions of Th. 1, the modified CLS-ϵ policy (switching on the envelope functions $f_i(t)$ instead of $\|x_i(t)\|$), stabilizes the collection (2) if (5) holds.*

We note that CLS-ϵ is not the only interrupt-based policy that leads to stability under the sufficient condition of Eq. 5. In particular, we have in mind certain queuing and scheduling problems where one seeks to minimize the expected value of queue lengths $Q_k(t)$ given their arrival rates μ_k (possibly Poisson) and the limited processing capacity of a server. In that case, the "μ-c rule" ([13, 4, 14] and references therein) optimizes $\sum c_k \mathcal{E}(Q_k(t))$ by admitting into the server those processes with the highest value of $\mu_k c_k$ value where c_k is the relative cost associated with the k^{th} queue. In Theorem 2 and Lemma 1, the $\ln|x_i|$ (resp. $\ln(V_i)$) play a role similar to that of queue lengths in a scheduling problem with N processes and k servers; the μc rule suggests closing the feedback loops of those systems with the k highest values of $(\mu_i - \lambda_i)c_i$ whose states have not yet reached ϵ in norm.

In general, we expect that the modified CLS-ϵ policy of Lemma 1 will be conservative because it makes switching decisions based on the "envelope functions" that bound the Lyapunov functions V_i. Moreover, it requires that the controller keep track of the envelope functions (10) as well as time. This raises the question of whether the CLS-ϵ policy – this time with *switching decisions based on the Lyapunov functions $V_i(x_i(t))$ themselves as opposed to their bounds $f_i(t)$* – stabilizes a multivariable collection (2) under a condition similar to (5). The complication that arises here is due to the fact that unlike the case of scalar dynamics (or that of Lemma 1 for that matter), the trajectories of the $V_i(x_i(t))$ are not pure exponentials and in fact may not be monotonic between switching times; therefore the state whose Lyapunov function is largest at a given switching time t may not always correspond to the system whose envelope function is largest at t.

Theorem 3 *The collection of systems in Eq. 2 will be stable under the interrupt-based communication policy obtained by replacing $\|x_i(t)\|$ by $V_i(x_i(t))$ in the CLS-ϵ algorithm, if $p = \sum_{i=1}^{N} \frac{\mu_i}{\mu_i - \lambda_i} < k$, where λ_i and μ_i are obtained by solving Prob. 1.*

Proof: Let $x_i(t)$, $i = 1, \dots, N$ be the states associated with the N systems in our collection and V_i the associated Lyapunov functions obtained by solving Prob. 1. Define $V = \sum_1^N \frac{\ln(V_i(x_i))}{\mu_i - \lambda_i}$ and notice that $\dot{V} \leq -k + p < 0$ when any k loops are closed. Suppose that one of the states x_m is unstable under a CLS-ϵ policy that decides which loop(s) to close based on the largest $V_i(x_i(t))$. Then, if t_j, $j = 0, 1, ...$ denote the times a which the m^{th} loop is closed, we can find an interval $[t_n, t_{n+1}]$ whose length $T = t_{n+1} - t_n$ is arbitrarily large. Moreover, at all other switching times in (t_n, t_{n+1}), there was some state x_q other than x_m for which $V_q(x_q)$ was even larger than $V_m(x_m)$. Because T is arbitrarily large, so is $V_m(x_m(t))$ and therefore so is $V_q(x_q(t))$ for some $t \in (t_n, t_{n+1})$. Proceeding in this manner we see that if there were two states that were unbounded, there must have also been a third one, and so on, and that all states must be unbounded, which contradicts $\dot{V} < 0$. $\square$

We note that it is possible – for appropriate choices of dynamics (2) – to achieve equality in $\dot{V}_i(x_i) \leq \lambda_i V_i(x_i)$ (as well as in the corresponding open loop equation) and thus the condition (5) cannot be relaxed. [3]

[3] One can also consider a feedback communication policy where switching decisions are made based on the values of $V_i(x_i)$ over an interval, e.g. $i^* = argmax_{t \in [t_0, t_0+T]} V(x_i(t))$, where T is chosen large enough. Such a policy will always choose the state corresponding to the largest envelope function and thus result in bounded trajectories, with T acting as a minimum waiting time.

4 Examples

The following examples illustrate the performance of the CLS-ϵ policy on a group of linear systems, sharing a network that can accommodate one system at a time. We first simulated a trio of scalar systems, with open and closed loop dynamics (Eq. 2) characterized by: $A_1^c = -2$, $A_1^o = 4$, $A_2^c = -4$, $A_2^o = 2/3$, and $A_3^c = -6$, $A_3^o = 1$. From Eq. 7, we obtained $p = 0.952$, which implies that the systems are stabilizable under the CLS-ϵ policy. Figure 2 shows the evolution of $|x_i(t)|$, $i = 1, 2, 3$ under CLS-ϵ with $\epsilon = 1$. For a set of three second-order systems, with

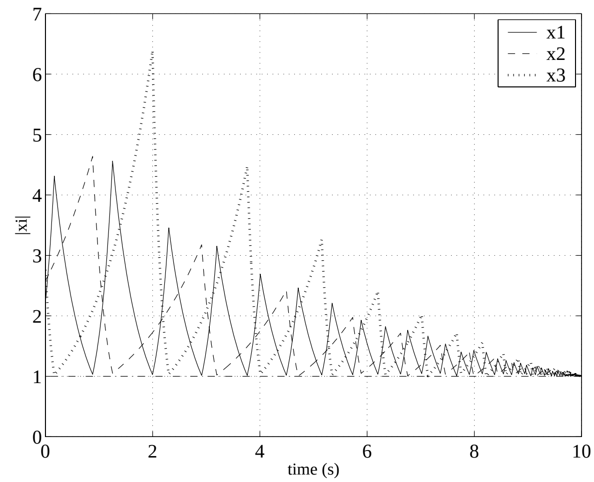

Figure 2: Scalar case - State evolution under CLS, $\epsilon = 1$

$$A_1^c = \begin{bmatrix} -5 & 5 \\ 10 & -6 \end{bmatrix}, \quad A_1^o = \begin{bmatrix} 3 & 4 \\ -3 & 1 \end{bmatrix},$$
$$A_2^c = \begin{bmatrix} -3 & 0 \\ 1 & -3 \end{bmatrix}, \quad A_2^o = \begin{bmatrix} 2 & -1 \\ 2 & \frac{1}{2} \end{bmatrix},$$
$$A_3^c = \begin{bmatrix} -5 & \frac{1}{2} \\ -\frac{1}{2} & -4 \end{bmatrix}, \quad A_3^o = \begin{bmatrix} 1 & 2 \\ -1 & 1 \end{bmatrix}$$

we performed the minimization in Problem 1 for the optimal Lyapunov functions $V_i = x_i^T P_i x_i$ and found that they are be bounded by exponentials with rates $\lambda_1^c = -10.95$, $\lambda_2^c = -5.15$, $\lambda_3^c = -7.90$ and $\mu_1^o = 8.34$, $\mu_2^o = 2.51$, $\mu_3^o = 2.06$, resulting in $p = 0.96 < 1$ (Eq. 5). The evolution of the V_i (starting from random initial conditions, under the same network constraint $k = 1$ and CLS-1 policy that switches based on the values of the V_i at the times when $V_j(t) = \epsilon$ for some j) is shown in Figure 3.

5 Conclusions

We explored the stability of LTI systems whose feedback loops are closed through a shared medium that admits a limited number of "users" at any one time. This

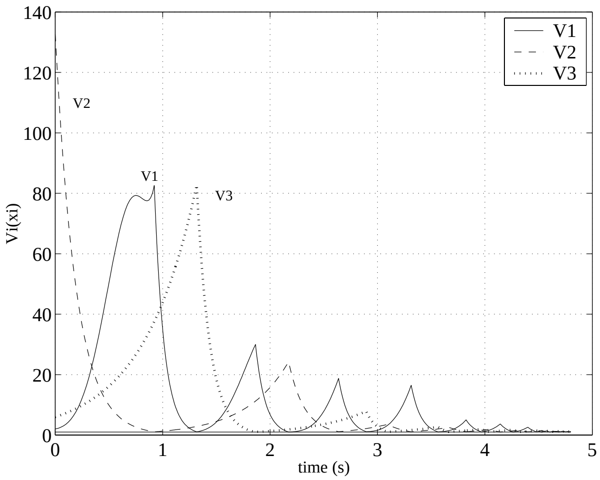

Figure 3: Multivariable case - Lyapunov function evolution under CLS, $\epsilon = 1$

work is part of an effort to understand control-oriented networks, emphasizing questions of control-theoretic interest (in this case stability) as opposed to information flow. We proposed an interrupt-based communication policy (termed **CLS-**ϵ) that "closes the loop" of the system(s) whose state is in some sense furthest from the origin. When that state (or an associated Lyapunov function) is brought within an ϵ-ball from the origin, the corresponding loop is opened and the selection process is repeated. We exposed some connections of the stabilization problem to queuing theory and showed that a sufficient condition for the existence of a stabilizing communication sequence [9] is also sufficient under the CLS-ϵ strategy. If the dynamical systems in question are scalar, our condition is also necessary. In the case of multivariable systems our sufficient condition is "tight", in the sense that if it is violated then there are dynamical systems that cannot be stabilized under any communication policy. Opportunities for future work include finding conditions under which the original CLS-ϵ policy (switching on $\|x_i(t)\|$) stabilizes a collection of systems, computing bounds for the state norms and investigating the effects of communication delays.

References

[1] S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan. *Linear Matrix Inequalities in System and Control Theory*, volume 15 of *Studies in Applied Mathematics*. SIAM, Philadelphia, PA, 1994.

[2] M. S. Branicky. Multiple Lyapunov functions and other analysis tools for switched and hybrid systems. *IEEE Trans. Automatic Control*, 43(4):475–82, Apr. 1998.

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

[4] C. Buyukkoc and P. Varayia and J. Walrand. The μc-rule revisited. *Adv. Appl. Prob.*, vol.17, pp. 237–38, 1985.

[5] P. Gahinet and A. Nemirovskii. *LMI Lab: A package for Manipulating and Solving LMIs.* INRIA, 1993.

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

[7] D. Hristu. *Optimal Control with Limited Communication.* PhD thesis, Harvard University, Division of Engineering and Applied Sciences, 1999.

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

[9] D. Hristu-Varsakelis. Feedback control systems as users of a shared network: Communication sequences that guarantee stability. In *IEEE Conf. on Decision and Control*, Dec. 2000.

[10] IEEE Control Systems Society. *IEEE Control Systems Magazine, special issue on Networks and Control*, Feb. 2001.

[11] D. Liberzon and A. S. Morse. Basic problems in stability and design of switched systems. *IEEE Control Systems Magazine*, 19(5):59–70, Oct. 1999.

[12] 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*, vol. 2, pp. 1828–33, Dec. 2000.

[13] I. Meilijson and U. Yechiali. On optimal right of way policies at a single server station when insertion of idle times is permitted. *Stoch. Proc. Appl.*, v.6, pp. 25–32, 1977.

[14] P. Nain and D. Towsley. Optimal scheduling in a machine with stochastic varying processing rate. *IEEE Trans. Automatic Control*, 39(9), pp. 1853–1855, 1994.

[15] J. Perkins and P. R. Kumar. Stable distributed, real-time scheduling of flexible manufacturing/assembly/disassembly systems. *IEEE Trans. Automatic Control*, 34(2):139–48, Feb. 1989.

[16] 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*, vol. 1, pp. 1011–16, Dec. 2000.

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

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