

Resource Constrained LQR Control Under Fast Sampling *

Jerome Le Ny
Department of Electrical and
Systems Engineering
University of Pennsylvania
200 South 33rd Street
Philadelphia, PA 19104, USA
jeromel@seas.upenn.edu

Eric Feron
School of Aerospace
Engineering
Georgia Institute of
Technology
270 Ferst Drive
Atlanta, GA 30332, USA
feron@gatech.edu

George J. Pappas
Department of Electrical and
Systems Engineering
University of Pennsylvania
200 South 33rd Street
Philadelphia, PA 19104, USA
pappasg@seas.upenn.edu

ABSTRACT

We investigate a state feedback Linear Quadratic Regulation problem with a constraint on the number of actuation signals that can be updated simultaneously. Such a constraint arises for example in networked and embedded control systems, due to limited communication and computation capabilities. Following recent results on the dual problem of scheduling Kalman filters, we first develop a bound on the achievable performance that can be computed efficiently by semidefinite programming. This bound can be approached arbitrarily closely by an analog periodic controller that can switch between control inputs arbitrarily fast. We then discuss implementation issues on digital platforms, i.e., the discretization of the analog controller in the presence of a relatively fast but finite sampling rate.

Categories and Subject Descriptors

C.3 [Special-Purpose and Application-Based Systems]: real-time and embedded systems; G.1.10 [Mathematics of Computing]: Numerical Analysis—*Applications*; F.2.2 [Analysis of algorithms and problem complexity]: Non-numerical Algorithms and Problems—*Sequencing and Scheduling*

General Terms

Algorithms, Theory

Keywords

Networked and Embedded Control Systems

1. INTRODUCTION

Modern control systems are increasingly implemented on networked and embedded platforms [10]. This trend raises many interesting questions at the interface of control, computing, and communications, see e.g. [2, 21]. In particular, for many control systems it is important to understand the impact of limited computational and communication resources. Taking this aspect into account as much as possible during the control system design phase can help reduce the overall system cost and provides more flexibility at the implementation and system integration phases [3, 15].

In this paper we consider a control problem for a plant with multiple actuation channels, only one of which can be updated at a time, see Fig. 1. Such a constraint arises for example if the controller is implemented on a platform with limited computational power or executing other computation-intensive tasks. In this case, it might be desirable or necessary to divide the control function into subfunctions with shorter execution times that update only a subset of the control signals when executed [15]. A natural question is then to decide which input signals should be updated more frequently. A similar situation occurs if the control signals must be sent to the actuators via a communication network. In this case, the restriction on the number of simultaneous control signal updates is due to the limited communication capacity constraint between the controller and the plant. Here we limit our discussion to the Linear Quadratic Regulation (LQR) problem [1].

Similar problems at the interface of control and scheduling have been considered at least since Meier et al. studied the dual problem of scheduling measurement systems [18]. More recently much research has been done on linear quadratic control problems and other optimal control problems for switched systems, see e.g. [11, 16, 17, 22, 26], and on joint control and scheduling problems, especially in the context of Networked Control Systems (NCS), see e.g. [3, 20, 21]. In this paper, we consider a continuous-time infinite-horizon linear quadratic control problem under state feedback. Lee [16] considers a closely related but more general output feedback infinite-horizon Linear Quadratic Gaussian (LQG) problem in discrete time, and Zhang et al. [26] study the finite-horizon problem, also under state feedback. The importance of the continuous-time infinite-horizon version of the problem, which is the dual of a Kalman filter scheduling problem considered in [14], comes from the fact that it admits

*This work was supported by NSF award No. 0931239

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

HSCC'11, April 12–14, 2011, Chicago, Illinois, USA.
Copyright 2011 ACM 978-1-4503-0629-4/11/04 ...\$10.00.




Figure 1: Plant-Controller configuration with an additional resource constraint, due to the presence of a communication network or because the processor can only update one control input at a time.

a simple solution that requires only solving a Linear Matrix Inequality (LMI) whose size is of the order of the plant dimension. In contrast, the available solutions of discrete-time formulations typically involve searching over the set of possible input sequences [18], whose size grows exponentially with the length of the sequence. Moreover, there is no a priori known bound on the size of the search space for infinite-horizon formulations [16]. The continuous-time approach also has the advantage of considering the inter-sample behavior of closed-loop sampled-data systems. The potential drawback of this method is that the controller obtained is a continuous-time (periodic, fast switching) controller, which must be discretized for most applications. Hence an important goal of this paper is to discuss in some details implementation issues of this analog controller on digital platforms for the control of continuous-time plants.

The paper is organized as follows. Section 2 describes our formulation of the resource-constrained LQR control problem. A solution for this problem consists in designing a feedback control signal together with a scheduling policy specifying which control channel to use at each time. These two subproblems can in fact be treated independently, and section 3 describes the optimal analog controller given a scheduling policy, which is simply obtained from the solution of a standard LQR control problem. In section 4 we show how to compute a bound on the performance achievable by any controller and scheduling policy, by solving a semidefinite program. If we assume that we can implement an analog control law, this bound can be approximated arbitrarily closely by certain periodic switching policies described in section 5, with the performance gap simply controlled by the rate at which the policies are allowed to switch between inputs. The rest of the paper is devoted to the discussion of digital implementations of these analog controllers. Two possible implementations, one assuming possibly time-varying sampling rates and the other assuming a time-triggered platform respectively, are described in section 6, and their performance is evaluated by simulations in section 7.

2. PROBLEM FORMULATION

Consider a linear continuous-time plant with dynamics

$$dx = Ax dt + B_{\sigma(t)}u(t) dt + dw, \quad (1)$$

where $x \in \mathbb{R}^n$ is the n -dimensional state vector, $u \in \mathbb{R}^m$ is the input vector, and $A \in \mathbb{R}^{n \times n}$. Moreover, $\sigma(t) \in \{1, \dots, N\} =: [N]$ is an additional control parameter used to select the matrix $B_{\sigma(t)} \in \{B_1, \dots, B_N\}$, where $B_i \in \mathbb{R}^{n \times m}$ for all $i \in [N]$. The noise process w is a vector Wiener process with zero mean and incremental covariance $W dt$. We assume that we have access to the full state x to design a feedback controller. Also, assume for simplicity that x_0 is deterministic, so that $\Sigma(0) = E[x(0)x(0)]^T = 0$. Finally, assume that the pair $(A, [B_1, \dots, B_N])$ is controllable.

Let us first fix a time horizon T , and consider the problem of designing a switching signal $\sigma(t)$ together with a control input $u(t)$ that jointly minimize the quadratic cost

$$J_T(\sigma, u) = \frac{1}{T} E \left\{ \int_0^T x^T Q x + u^T R_\sigma u dt + x^T(T) Q_f x(T) \right\}, \quad (2)$$

where $Q \succeq 0$, $Q_f \succ 0$, and $R_i \succ 0$, for all $i \in [N]$. We are mostly concerned with the infinite-horizon version of this problem, where the goal is to design a policy σ and a control law u minimizing

$$J(\sigma, u) = \limsup_{T \rightarrow \infty} J_T(\sigma, u). \quad (3)$$

To study this problem, we assume throughout the paper that $(A, Q^{1/2})$ is observable.

To the switching signal $\sigma(t)$ we associate an N -dimensional signal of binary variables $\beta(t) \in \{0, 1\}^N$ such that

$$\beta_i(t) = 1 \Leftrightarrow \sigma(t) = i.$$

In other words, β is a unit vector with a one at the index σ . The scalars $\beta_i(t)$ corresponding to the signal $\sigma(t)$ are therefore subject to the constraints $\beta_i(t) \in \{0, 1\}$ for all $i \in [N]$ and

$$\sum_{i=1}^N \beta_i(t) \leq 1. \quad (4)$$

Clearly there is a one-to-one correspondence between σ and β and we employ both notations interchangeably. Note in particular that we can rewrite the dynamics equation (1) as

$$dx = Ax dt + \left(\sum_{i=1}^N \beta_i(t) B_i \right) u(t) dt + dw.$$

Example 1. Consider the situation of Fig. 1, where there are N control channels $u = [u_1, \dots, u_N]$ available, each possibly multidimensional, with $u_i \in \mathbb{R}^{m_i}$ and $\sum_{i=1}^N m_i = m$. However, only one channel be used at each time. The control inputs that cannot be updated are set to zero, as in e.g. [25]. In other words, the matrices B_i are of the form

$$B_i = [\mathbf{0}_{n \times m_1} \quad \dots \quad \mathbf{0}_{n \times m_{i-1}} \quad \hat{B}_i \quad \mathbf{0}_{n \times m_{i+1}} \quad \dots \quad \mathbf{0}_{n \times m_N}]$$

where $\hat{B}_i \in \mathbb{R}^{n \times m_i}$.

3. CONTROLLER DESIGN

Given a switching policy σ , the system (1) evolves as a linear time-varying system. Fixing $T < \infty$, the controller design problem is then a standard linear quadratic regulation problem and so the optimal controller is a static feedback controller

$$u_\sigma(t) = L_\sigma(t; T)x, \quad (5)$$

where $L_\sigma(t; T) = -R_\sigma^{-1}(t)B_\sigma^T(t)P_\sigma(t; T)$ and $P_\sigma(t; T)$ satisfies the differential Riccati equation

$$\begin{aligned} \dot{P}_\sigma(t; T) &= -A^T P_\sigma(t; T) - P_\sigma(t; T)A - Q \\ &\quad + P_\sigma(t; T)B_\sigma(t)R_\sigma^{-1}B_\sigma^T(t)P_\sigma(t; T), \quad (6) \\ P_\sigma(T; T) &= Q_f. \end{aligned}$$

Note that (6) can be rewritten

$$\begin{aligned} \dot{P}_\sigma(t; T) &= -A^T P_\sigma(t; T) - P_\sigma(t; T)A - Q \quad (7) \\ &\quad + P_\sigma(t; T) \left(\sum_{i=1}^N \beta_i(t) B_i R_i^{-1} B_i^T \right) P_\sigma(t; T). \end{aligned}$$

Moreover, denoting $\Sigma(t) = E[x(t)x^T(t)]$ the correlation matrix of the state vector, and using $u = L_\sigma(t; T)x$, the cost expression (2) becomes

$$J_T(\sigma) = \frac{1}{T} \left\{ \int_0^T \text{Tr}[(Q + L_\sigma^T(t)R_\sigma L_\sigma(t; T))\Sigma(t)] dt + \text{Tr}[Q_f \Sigma(T)] \right\}.$$

It is useful to further transform these expressions, as follows. First, we have

$$\begin{aligned} A^T P_\sigma + P_\sigma A - P_\sigma B_\sigma R_\sigma^{-1} B_\sigma^T P_\sigma + Q \\ = (A + B_\sigma L_\sigma)^T P_\sigma + P_\sigma (A + B_\sigma L_\sigma) + L_\sigma^T R_\sigma L_\sigma + Q, \end{aligned}$$

which gives, in the Riccati equation (6)

$$\dot{P}_\sigma(t; T) = -(A + B_\sigma L_\sigma)^T P_\sigma - P_\sigma (A + B_\sigma L_\sigma) - L_\sigma^T R_\sigma L_\sigma - Q.$$

We also have a Lyapunov equation describing the dynamics of the covariance matrix $\Sigma(t)$ under the control $u = L_\sigma x$, namely [4]

$$\dot{\Sigma}(t) = (A + B_\sigma L_\sigma)\Sigma(t) + \Sigma(t)(A + B_\sigma L_\sigma)^T + W.$$

Therefore

$$\begin{aligned} \text{Tr}[(Q + L_\sigma^T(t)R_\sigma L_\sigma(t; T))\Sigma(t)] \\ = -\text{Tr}[(\dot{P}_\sigma + (A + B_\sigma L_\sigma)^T P_\sigma + P_\sigma (A + B_\sigma L_\sigma))\Sigma(t)] \\ = -\text{Tr}[\dot{P}_\sigma \Sigma(t) + P_\sigma (\dot{\Sigma} - W)] \\ = \text{Tr}[P_\sigma W] - \text{Tr} \left[\frac{d}{dt} (P_\sigma \Sigma) \right]. \end{aligned}$$

Finally, we obtain the following expression for the cost function, using the optimal control law u given σ

$$J_T(\sigma) = \frac{1}{T} \left\{ \int_0^T \text{Tr}[P_\sigma(t; T)W] dt - \text{Tr}[P_\sigma(T; T)\Sigma(T)] + \text{Tr}[P_\sigma(0; T)\Sigma(0)] + \text{Tr}[Q_f \Sigma(T)] \right\},$$

that is,

$$J_T(\sigma) = \frac{1}{T} \int_0^T \text{Tr}[P_\sigma(t; T)W] dt \quad (8)$$

since $P_\sigma(T; T) = Q_f$ and $\Sigma(0) = 0$. The remaining goal is then to obtain a scheduling policy σ minimizing (9), or the infinite-horizon cost

$$J(\sigma) = \limsup_{T \rightarrow \infty} \frac{1}{T} \int_0^T \text{Tr}[P_\sigma(t; T)W] dt. \quad (9)$$

4. PERFORMANCE BOUND

Our approach is to first derive a bound on the performance achievable by any switching policy σ . Let

$$X_\sigma(t; T) = P_\sigma^{-1}(t; T).$$

Then in terms of this new variable, equation (7) becomes

$$\dot{X}_\sigma(t; T) = X_\sigma(t; T)A^T + AX_\sigma(t; T) + X_\sigma(t; T)QX_\sigma(t; T) \quad (10)$$

$$- \left(\sum_{i=1}^N \beta_i(t) B_i R_i^{-1} B_i^T \right), \quad (11)$$

$$X_\sigma(T; T) = Q_f^{-1}.$$

Now define the time averages

$$\tilde{X}_\sigma(T) = \frac{1}{T} \int_0^T X_\sigma(t; T) dt, \quad \tilde{\beta}_i(T) = \frac{1}{T} \int_0^T \beta_i(t) dt.$$

Integrating equation (11) over the interval $[0, T]$, we get

$$\begin{aligned} \frac{X_\sigma(T; T) - X_\sigma(0; T)}{T} &= \tilde{X}_\sigma(T)A^T + A\tilde{X}_\sigma(T) \\ &\quad - \left(\sum_{i=1}^N \tilde{\beta}_i(T) B_i R_i^{-1} B_i^T \right) + \frac{1}{T} \int_0^T X_\sigma(t; T)QX_\sigma(t; T) dt. \end{aligned}$$

Now by Jensen's inequality applied to the matrix convex function $X \rightarrow XQX$ [7, p. 110], we have

$$\begin{aligned} \frac{1}{T} \int_0^T X_\sigma(t; T)QX_\sigma(t; T) dt \\ \succeq \left(\frac{1}{T} \int_0^T X_\sigma(t; T) dt \right) Q \left(\frac{1}{T} \int_0^T X_\sigma(t; T) dt \right). \end{aligned}$$

Together with the fact that $X_\sigma(0; T) \succeq 0$, we get the following convex quadratic inequality

$$\begin{aligned} \frac{Q_f^{-1}}{T} \succeq \tilde{X}_\sigma(T)A^T + A\tilde{X}_\sigma(T) \quad (12) \\ - \left(\sum_{i=1}^N \tilde{\beta}_i(T) B_i R_i^{-1} B_i^T \right) + \tilde{X}_\sigma(T)Q\tilde{X}_\sigma(T). \end{aligned}$$

Also, note that in the cost function

$$\frac{1}{T} \int_0^T \text{Tr}[P_\sigma(t; T)W] dt = \text{Tr} \left[\left(\frac{1}{T} \int_0^T P_\sigma(t; T) dt \right) W \right],$$

and, again by Jensen's inequality and the matrix convexity of $X \rightarrow X^{-1}$ on the positive definite matrices [7, p. 76], we have

$$\left(\frac{1}{T} \int_0^T P_\sigma(t; T) dt \right) = \left(\frac{1}{T} \int_0^T X_\sigma(t; T)^{-1} dt \right) \succeq \tilde{X}_\sigma(T)^{-1}.$$

Finally, we see therefore that for any T , the optimal cost is lower bounded by the quantity

$$\text{Tr}[\tilde{X}_\sigma(T)^{-1}W] \quad (13)$$

where $\tilde{X}_\sigma(T)$ satisfies the constraint (12) and $\tilde{\beta}_i(T)$ is subject to

$$\sum_{i=1}^N \tilde{\beta}_i(T) \leq 1, \quad 0 \leq \tilde{\beta}_i(T), i = 1, \dots, N.$$

A lower bound on the achievable performance for T finite can then be obtained by letting $\tilde{X}_\sigma(T) = X$ in (13), where X is the solution of the following program with variables $X, \{b_i\}_{1 \leq i \leq N}$

$$\min_{X \succ 0, \{b_i\}_{1 \leq i \leq N}} \text{Tr}[X^{-1}W] \quad (14)$$

$$\text{s.t. } \frac{Q_f^{-1}}{T} \succeq XA^T + AX \quad (15)$$

$$- \left(\sum_{i=1}^N b_i B_i R_i^{-1} B_i^T \right) + XQX$$

$$\sum_{i=1}^N b_i \leq 1, \quad 0 \leq b_i, i = 1, \dots, N.$$

For $(A, Q^{1/2})$ observable, this optimization problem also has a solution for $T \rightarrow \infty$ (see [14]), where (15) is replaced by

$$0 \succeq XA^T + AX - \left(\sum_{i=1}^m b_i B_i R_i^{-1} B_i^T \right) + XQX. \quad (16)$$

The resulting optimal value is a bound on the performance achievable by any switching policy for the infinite-horizon problem.

Finally, we can compute this performance bound efficiently by solving a semidefinite program. We consider only the infinite-horizon problem from now on, i.e., with the constraint (16). Introduce the slack variable $Y \succeq X^{-1}$. By taking Schur complements, we then see that the bound can be obtained by solving

$$Z^* = \min_{X, Y, \{b_i\}_{1 \leq i \leq N}} \text{Tr}[YW] \quad (17)$$

$$\text{s.t. } \begin{bmatrix} Y & I \\ I & X \end{bmatrix} \succ 0$$

$$\begin{bmatrix} XA^T + AX - \left(\sum_{i=1}^N b_i B_i R_i^{-1} B_i^T \right) & XQ^{1/2} \\ Q^{1/2} X & -I \end{bmatrix} \preceq 0$$

$$\sum_{i=1}^N b_i \leq 1, \quad 0 \leq b_i, i = 1, \dots, N.$$

5. OPTIMAL CONTINUOUS-TIME POLICIES

In the rest of the paper, we discuss certain switching policies approaching the performance bound (17). Note first that one can evaluate empirically the cost of any policy (say, by simulation) and compare this cost to the performance bound to obtain an indication of the policy performance with respect to an optimal policy. Depending on certain implementation issues, one can approach this bound more or less closely. In this section, we start by assuming essentially no limitation on the signals σ, u . In particular, σ is allowed to switch arbitrarily fast, and u can be implemented as the continuous-time function (5). In this context, there are periodic continuous-time policies that approach the lower bound on achievable performance arbitrarily closely, and hence are essentially optimal. These policies can serve as a benchmark to evaluate the performance of policies satisfying more realistic implementation constraints, such as the ones discussed in section 6.

To design optimal continuous-time policies for the infinite-horizon problem (3), we start by solving (17) and thus obtain

a set of optimal parameters $\{b_i\}_{1 \leq i \leq N}$. Let $b_0 = 0$. We then choose a period length ϵ to execute the following policy, for an ϵ -periodic function $t \rightarrow P(t)$ to be described next. We divide each period in subintervals as follows

(i) Use control

$$u = L_i(t)x = -R_i^{-1}B_i^T P(t)x \quad (18)$$

for the interval

$$[\epsilon \sum_{k=0}^{i-1} b_k, \epsilon \sum_{k=0}^i b_k], 1 \leq i \leq N.$$

(ii) If $\sum_{k=0}^N b_k < 1$, then set $u(t) = 0$ for the interval $[\epsilon \sum_{k=0}^N b_k, \epsilon]$.

Hence the policy spends a proportion b_i of each period in mode i , and does not use any actuation signal for a proportion $1 - \sum_{i=1}^N b_i$ of each period (although we typically have $\sum_{i=1}^N b_i = 1$).

We define a periodic switching signal σ such that $\sigma(t) = i$ over the interval $[\epsilon \sum_{k=0}^{i-1} b_k, \epsilon \sum_{k=0}^i b_k]$, and $\sigma(t) = 0$ over the interval $[\epsilon \sum_{k=0}^N b_k, \epsilon]$ in each period, where $\sigma = 0$ signifies that no input signal is sent to the plant. Define also the formal notation $B_0 R_0^{-1} B_0^T := \mathbf{0}_{n \times n}$. The matrix $P(t)$ used in (18) is the unique positive definite stabilizing ϵ -periodic solution of the following periodic Riccati differential equation (PRE)

$$\dot{P}(t) = -A^T P(t) - P(t)A - Q + P(t)B_{\sigma(t)}R_{\sigma(t)}^{-1}B_{\sigma(t)}^T P(t), \quad (19)$$

see [6, 14]. Denote by $Z(\epsilon)$ the value of the infinite-horizon cost $J(\sigma)$ for the ϵ -periodic continuous-time policy described above, and recall that we denoted by Z^* the value of the performance bound (17). The following theorem says that in the limit of arbitrarily fast switching ($\epsilon \rightarrow 0$), these policies perform essentially optimally.

THEOREM 1. *We have $Z(\epsilon) - Z^* = O(\epsilon)$ as $\epsilon \rightarrow 0$. In particular $Z(\epsilon) \rightarrow Z^*$ as $\epsilon \rightarrow 0$.*

The proof of this theorem follows from the results in [14] on scheduling continuous-time Kalman filters by duality. We now turn to evaluating the performance impact of more realistic implementation constraints.

6. DIGITAL IMPLEMENTATION

There are essentially three ways of designing digital controllers for sampled-data systems [5, 8]. We can first discretize the plant dynamics and design the controller in discrete-time, as in much of the recent literature on the optimal control of switched systems [11, 16, 26]. This potentially ignores some behaviors of the closed-loop continuous-time system, such as hidden oscillations [5]. The second approach, adopted here, is to design a continuous-time controller which is then discretized in actual implementations. The last approach is the most challenging and relatively unexplored in the switched systems literature, and consists in modeling the digital implementation explicitly at the continuous-time level [8] prior to control design.

The policies described in the previous section are continuous-time control laws, i.e., they require a continuous update of the control signal u . The function $t \rightarrow P(t)$ in (i) can be

computed a priori, but this requires a numerical integration method to obtain the stabilizing periodic solution of the Riccati equation (19). Moreover, in practice most digital control implementations use piecewise constant input signals. Another issue with the result of theorem 1 is that in actual implementations the period ϵ is a finite positive constant governed by the rate at which the physical system can switch between modes or sample the state. Hence the purpose of this section is to discuss more practical digital implementation schemes.

6.1 Restrictions on the Switching Times: Time-Triggered Policies

Fixing ϵ a priori imposes that the system following (i), (ii) is able to switch mode at times $b_1\epsilon, (b_1 + b_2)\epsilon, (b_1 + b_2 + b_3)\epsilon, \dots$. In fact, the result of theorem 1 remains valid for any continuous-time policy σ that spends a proportion b_i of its period in mode i , not necessarily in a single interval as in (i), which gives additional flexibility to the schedule. Let us call a policy or schedule σ where ϵ can be fixed a priori and the system can guarantee a time $b_i\epsilon$ per period in mode i a policy of type (U) (unrestricted). Typical technological limitations impose a bound of the form

$$\min_{1 \leq i \leq N} \{b_i\epsilon\} \geq \Delta,$$

where Δ is here the minimum dwell-time, i.e., the minimum time that the system must remain in any mode, e.g. due to limited sampling rates or to the implementation overhead of the mode switching mechanism.

In some cases, it can be easier in an embedded control system implementation to further restrict the times at which the mode can change to an a priori fixed set of regularly spaced times $k\Delta, k \geq 0$. This is the case for example if a time-triggered protocol is used for scheduling tasks on an embedded processor [12], which can simplify system integration and verification, see e.g. [20, 24]. Let us call a policy σ satisfying this constraint of fixed computing slots of size Δ a policy of type (TT) (time-triggered). Policies of type (U) and type (TT) give somewhat different discretized models, as discussed in the next subsection. We now turn to the problem of approximating a schedule of type (U) by a more restricted policy of type (TT).

For policies of type (TT), we use the following periodic schedules. First, we approximate the parameters $\{b_i\}_{1 \leq i \leq N}$ optimal for (17) by

$$b_i \approx \frac{l_i}{l}, \quad l_i, l \in \mathbb{N}, \quad (20)$$

where l is an admissible length for the period of the overall schedule. For example, if l is of the form $2^p, p \in \mathbb{N}$, then (20) represents a partial binary expansion of the number b_i . In general, increasing l provides a better approximation of the optimal parameters $b_{i, 1 \leq i \leq N}$ but increases the memory requirements of the implementation. The rounding procedure (20) might have to be adjusted slightly in general, to make sure that the constraint $\sum_{i=1}^N l_i \leq l$ is always enforced, corresponding to $\sum_{i=1}^N b_i \leq 1$. We then design a cycle of l slots of length Δ , such that in each cycle, mode i is used in l_i time slots. For resource constrained applications as described in example 1, it is best to spread the slots dedicated to each mode as much as possible within each cycle, to avoid channels operating in an open-loop manner for too long. In

contrast, in order to minimize the number of mode switches, one can schedule the slots for each given mode consecutively. In the simulations presented in section 7 we choose the l_i positions of the slots for each mode randomly among the l slots of a cycle. Note that the length $l\Delta$ of the period of a schedule (TT) can be much longer than the length ϵ of a period of a schedule (U), which must only satisfy $\epsilon \geq \frac{\Delta}{\min_{1 \leq i \leq N} b_i}$. In general for schedules of type (TT), there is a trade-off between reducing the length of the schedule and obtaining a good approximation of the parameters $\{b_i\}_{1 \leq i \leq N}$.

6.2 Controller Discretization

Let us now fix the parameter Δ , representing dwell-time for schedules (U) and time-slot length for schedules (TT). Let us also assume that the continuous-time control law (18) must be implemented on a digital controller, which can therefore only update the control signal u at discrete times (we ignore quantization effects in this paper). We assume that the controller also samples the state at these times and we neglect the computation interval between the sampling time and the control update time. We assume however that for both types of policies Δ is a lower bound on the inter-sampling times. Note that these conventions simply fix a possible choice of implementation constraints for the rest of the discussion, and other scenarios could be considered, e.g. where the inter-sampling times can be shorter than Δ . Finally, a zero-order hold is assumed to be present between the digital controller and the plant, so that the plant sees a piecewise-constant control signal.

Denoting the sampling and computation times $t_k, k \geq 0$, we have $t_k = k\Delta$ for policies of type (TT). For policies of type (U), Let us assume for simplicity of notation that $b_1 = \min_{1 \leq i \leq N} b_i$. We then make the choice $\epsilon = \frac{\Delta}{\min_{1 \leq i \leq N} b_i} = \Delta/b_1$, so that only one sample per period is used in the mode with shortest length. For $i \geq 2$, we divide the length $b_i\epsilon = b_i\Delta/b_1$ of mode i into $\lambda_i = \lfloor b_i/b_1 \rfloor$ blocks of equal length δ_i , and sample and compute the control input for mode i at times $(b_1 + \dots + b_i + k\delta_i)\epsilon, 0 \leq k \leq \lambda_i$, within each period. Note that the time interval between two successive samples in mode i is

$$\delta_i = \frac{b_i\epsilon}{\lfloor \frac{b_i}{b_1} \rfloor} \geq b_1\epsilon = \Delta,$$

hence the constraint on inter-sampling times is satisfied. The sampling times over the whole period for schedules of types (U) are not exactly periodic, which allows us ignore the approximation issue in (20) but could be more complicated to implement. On the other hand, this device allows us to obtain schedules of type (U) of shorter length.

Next, we discuss the time-discretization of the continuous-time function $L_\sigma(t)$ in (18). For simplicity of notation, we discuss this process only for the policies of type (TT). The discretization process for policies of type (U) is similar, except for the slightly different sampling periods in different modes. Hence let us denote by $x_k = x(k\Delta), k \geq 0$, the sample sequence measured by the controller following a schedule (TT). We construct l gain matrices $L_{k,\Delta}, k = 0, \dots, l-1$, which are stored in the controller memory to compute the discrete-time periodic control sequence $u_k, k \geq 0$. With the zero-order hold assumption, the signal $u(t)$ at the input of the plant is piecewise constant equal to u_k on the interval $[k\Delta, (k+1)\Delta)$. We use the notation $\sigma(k) := \sigma(t_k) = \sigma(k\Delta)$

to represent the mode used at time k by the schedule σ . Because of the l -periodicity of the schedule, we have $\sigma(k) = \sigma(k \bmod l)$.

Rather than performing say a simple first-order Euler integration scheme for the continuous-time Riccati equation (6), we use the discrete-time Riccati equation describing the optimal solution for the discrete-time LQR problem. Since we emphasize the connection between continuous-time and discrete-time models here, it is best to use incremental models of discrete-time systems [9, 19]. Incremental models also offer improved numerical properties for a small sampling period Δ [23], which is our focus in this paper.

Hence define

$$A_\Delta = \frac{e^{A\Delta} - I}{\Delta}, \quad B_{i,\Delta} = \frac{1}{\Delta} \int_0^\Delta e^{A\tau} d\tau B_i, \quad 1 \leq i \leq N.$$

Note that $A_\Delta \rightarrow A$ and $B_{i,\Delta} \rightarrow B_i$ as $\Delta \rightarrow 0$. Moreover, the continuous-time cost (2) for the infinite-horizon problem is approximated by the Riemann sum

$$J(\sigma, u) \approx \lim_{T \rightarrow \infty} \frac{1}{T} E \left[\sum_{k=0}^{\lceil T/\Delta \rceil} x_k^T (\Delta Q) x_k + u_k^T (\Delta R_{\sigma(k)}) u_k \right]. \quad (21)$$

Then consider the Difference Periodic Riccati Equation (DPRE) in incremental form (see e.g. [23]), associated to the discrete-time LQR problem (21)

$$\begin{aligned} \frac{P_k - P_{k+1}}{\Delta} &= A_\Delta^T P_{k+1} + P_{k+1} A_\Delta - P_{k+1} B_{\sigma(k),\Delta} \times \\ &\quad (\Delta B_{\sigma(k),\Delta}^T P_{k+1} B_{\sigma(k),\Delta} + R_{\sigma(k)})^{-1} B_{\sigma(k),\Delta}^T P_{k+1} \\ &\quad + \gamma(\sigma(k), \Delta), \end{aligned} \quad (22)$$

where

$$\begin{aligned} \gamma(\sigma(k), \Delta) &= \Delta \left\{ A_\Delta^T P_{k+1} A_\Delta - A_\Delta^T P_{k+1} B_{\sigma(k),\Delta} \right. \\ &\quad \times (\Delta B_{\sigma(k),\Delta}^T P_{k+1} B_{\sigma(k),\Delta} + R_{\sigma(k)})^{-1} B_{\sigma(k),\Delta}^T P_{k+1} \\ &\quad - P_{k+1} B_{\sigma(k),\Delta} (\Delta B_{\sigma(k),\Delta}^T P_{k+1} B_{\sigma(k),\Delta} + R_{\sigma(k)})^{-1} \\ &\quad \times B_{\sigma(k),\Delta}^T P_{k+1} A_\Delta \\ &\quad \left. - \Delta A_\Delta^T P_{k+1} B_{\sigma(k),\Delta} (\Delta B_{\sigma(k),\Delta}^T P_{k+1} B_{\sigma(k),\Delta} + R_{\sigma(k)})^{-1} \right. \\ &\quad \left. \times B_{\sigma(k),\Delta}^T P_{k+1} A_\Delta \right\}. \end{aligned}$$

Note in particular that $\gamma(\sigma(k), \Delta) = O(\Delta)$, so that (22) can be seen as an approximation of the continuous-time dynamics (19) as $\Delta \rightarrow 0$.

Under our controllability and observability assumptions, the DPRE has a unique stabilizing l -periodic solution [6, 14], which is then used to approximate the continuous-time solution $t \rightarrow P(t)$ in (18) at the sampling times. We precompute the matrices P_0, \dots, P_{l-1} defining this periodic solution. Then we define the gain matrices as

$$\begin{aligned} L_{k,\Delta} &= -(\Delta B_{\sigma(k),\Delta} P_{k+1} B_{\sigma(k),\Delta} + R_{\sigma(k)})^{-1} \\ &\quad \times B_{\sigma(k),\Delta}^T P_{k+1} (I + \Delta A_\Delta), \quad k = 0, \dots, l-1, \end{aligned}$$

with $P_l = P_0$. Note again that $L_{i,\Delta} \rightarrow L_i$ as $\Delta \rightarrow 0$, where L_i was defined in (18). Then at period $k \geq 0$, the controller computes

$$u_k = \Delta L_{(k \bmod l),\Delta} x_k,$$

recalling that we neglect the time it takes to compute the matrix-vector product $L_{(k \bmod l),\Delta} x_k$. The dynamics of the

closed-loop system at the sampling times are then

$$x_{k+1} - x_k = (A_\Delta + L_{\sigma(k \bmod l),\Delta}) x_k \Delta + w_k, \quad (23)$$

approximating the continuous-time closed loop dynamics, with $\{w_k\}_{k \geq 0}$ a zero-mean Gaussian white noise with covariance

$$E[w_i w_j^T] = \int_0^\Delta e^{A\tau} W e^{A^T \tau} d\tau.$$

Note that the discretization process for policies of types (U) produces a discrete-time system with (periodic) time-varying dynamics instead of the time invariant system (23), due to the unequal sampling period in different modes.

7. SIMULATION RESULTS

In this section, we present simulation results illustrating the performance of the discretized versions of the optimal continuous-time policies. First, we consider the impact of increasing the dwell time or time-slot length Δ , which is desirable to reduce implementation costs. In particular, we evaluate the region of values Δ where the performance bound (17) is close to the performance achieved by the discrete periodic controllers. Consider the following matrices

$$\begin{aligned} A_1 &= \begin{bmatrix} -1 & 1 & 0 & 0 & 0 \\ 0 & -1 & 1 & 0 & 0 \\ 0 & 0 & -1 & 1 & 0 \\ 0 & 0 & 0 & -1 & 1 \end{bmatrix}, \quad B = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \\ A_2 &= \begin{bmatrix} 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 1 \end{bmatrix}, \quad W = 0.1 I_5. \end{aligned}$$

Here and throughout this section, Q and $R_i, 1 \leq i \leq N$, are taken to be identity matrices. We study the behavior of the policies for the controllable stable system (A_1, B) and unstable system (A_2, B) , assuming a scenario as described in example 1, where each control channel selects one of the three columns of B .

Solving the semidefinite program (17) provides the optimal parameters $b_1 \approx 0.54, b_2 \approx 0.44, b_3 \approx 0.02$ for (A_1, B) and $b_1 \approx 0.74, b_2 \approx 0.12, b_3 \approx 0.14$ for the unstable system (A_2, B) . This results in a schedule of type (U) for the stable system where 27 equidistant samples are taken in the first mode, 22 in the second mode, one in the third mode, before repeating the cycle. For the unstable system, we take 6 samples in the first mode, one sample in mode 2, one sample in mode 3 and repeat the cycle (recall that for schedules (U) the sampling periods in different modes can be different). For both systems, we approximate the optimal values b_1, b_2, b_3 to 2 digits of precision in the (TT) schedules, i.e., the length of the schedule is set to $l = 100$. Hence we obtain $l_1 = 54, l_2 = 44, l_3 = 2$ for the stable system and $l_1 = 74, l_2 = 12, l_3 = 14$ for the unstable system. The positions of the slots for each mode are chosen randomly. The performance curves are shown on Fig. 2 for (A_1, B) and Fig. 3 for (A_2, B) . The achieved average performance for each value of Δ is evaluated via Monte-Carlo simulations, with the cost averaged over 10^5 samples. We also show the performance of the unconstrained controller obtained by first discretizing the dynamics and then designing the optimal discrete-time LQR controller.

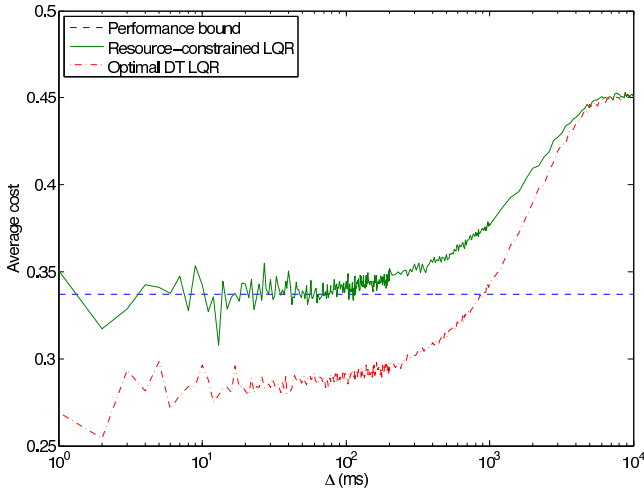


Figure 2: Performance degradation as Δ increases and comparison with the continuous-time performance bound for a stable open-loop system. Note the logarithmic scale on the x-axis (Δ varies from 1ms to 10s). There is no significant performance difference observed between schedules of type (U) and (TT) , and only the cost of the later one is shown.

We can see for these systems that for reasonably small values of Δ (approximately $\Delta \leq 100$ ms for (A_1, B) and $\Delta \leq 10$ ms for (A_2, B)), the performance of the digital controllers of section 6.2 matches the performance bound (17) (up to the noise in the simulation results), hence these controllers and schedules perform essentially optimally. We observe in general that the bound is tight for an interesting range of values of Δ , and typically quite informative for much of the interval of sampling times that is of practical interest given the system dynamics. There is no significant difference between the performance of the schedules of type (U) and (TT) for small values of Δ . Increasing Δ further, we then observe a relatively rapid performance degradation. Nonetheless, the schedules of type (U) show a much better performance than the long schedules of types (TT) in this range for unstable systems (see Fig. 3). Naturally, the value of Δ for which the performance degradation starts to become noticeable depends, among other things, on the degree of stability of the open-loop system (the location of the eigenvalue of A with maximum real part).

Note that the scheduling sequences obtained from the continuous time analysis are not necessarily optimal in general for the problem with large sampling periods Δ , where other effects becomes significant and allow us to distinguish between the performance of schedules (U) and (TT) for example. This can be seen by looking at the unstable system (A_2, B) , and comparing the schedule (TT) to the simple Round-Robin policy with $l = 3$, $l_1 = l_2 = l_3 = 1$, see Fig. 3. For $\Delta > 100$ ms, the Round-Robin policy starts to show a better performance, and diverges at a much smaller rate than for the (TT) schedule (but not the (U) schedule) as Δ continues to increase.

Let us investigate in more details the relationship between the range of values Δ where the performance of the digital implementation matches the continuous-time bound and the degree of stability of the open-loop system, measured here

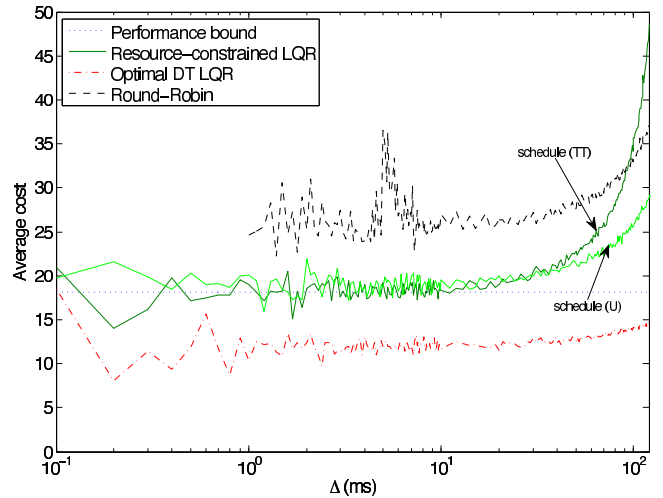


Figure 3: Performance degradation as Δ increases and comparison with the continuous-time performance bound for an unstable open-loop system. The Round-Robin policy simply cycles between the 3 modes by spending one slot in each mode in each cycle.

by the maximum real part of the eigenvalues of A . For this purpose, we fix the matrix B to be the 3×2 matrix

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

and generate 3×3 matrices A randomly with entries generated according to a standard normal distribution. For each such matrix, we evaluate empirically the threshold Δ at which the cost becomes greater than 1.2 times the performance bound. Let us denote Δ_{th} this threshold. Define, for a matrix A

$$s(A) := \max\{\text{Re}(\lambda) \mid \lambda \text{ eigenvalue of } A\}.$$

Fig. 4 shows the variation of Δ_{th} with $s(A)$ for a number of such randomly generated systems, for schedules of type (TT) . The thresholds were approximately determined using Monte-Carlo simulations and a Robbins-Monro procedure [13]. As in the previous example, we see that the region where the performance of the digital controller approximately matches the performance bound (17), here within a tolerance of 20%, extends to fairly large values of slot length, with Δ in the tens of milliseconds allowed even for fairly unstable open-loop dynamics. Better results could in fact be expected with schedules of type (U) .

Finally, Fig. 5 shows an example of regulation result, using the discretized periodic controller and a schedule of type (TT) . The system dynamics is described by a random 10×10 matrix A and a random 10×7 matrix B , and as in example 1 each column of B is associated to a different control channel (i.e. $m_i = 1, i = 1, \dots, 7$). The regulation responses are compared to a design based on implementing the optimal discrete-time linear-quadratic regulator for the system (A, B) , i.e., assuming all input channels can be used simultaneously. As expected, in the presence of the control constraint the response is found to be more sluggish and the




Figure 4: Variation of the threshold value Δ_{th} associated with 20% performance decrease with respect to the performance bound (17) for random linear systems with a three-dimensional state space. On the x -axis we represent the degree of stability of the open loop dynamics, measured by the maximum real part of the eigenvalues of A .

noise perturbing the dynamics cannot be filtered as well as with the standard LQR controller.

8. CONCLUSION

We have discussed a linear quadratic control problem under scheduling constraints. It is found that working directly with the continuous-time formulation of the problem allows us to obtain a simple performance bound, which is optimal assuming an analog controller and in the limit of infinitely fast switching rates. Relaxing these assumptions, the bound remains relevant for the characterization of the performance of digital versions of the analog controller under realistic sampling rates. It is interesting to note that for control problems concerning sampled-data systems subject to resource utilization constraints, it seems to be often easier to work directly in continuous-time rather than discretize the dynamics first. This is related to the well-known possibility to “convexify” the set of control inputs in continuous-time optimal control by fast switching. More generally however, continuous-time modeling is also often a richer modeling framework. For example, from the discussion of the schedules of type (U) in section 6, it emerges that time-varying sampling rates can be useful, a fact that would not be noticed by using discrete-time techniques assuming from the start a fixed sampling rate.

The problem considered here is the dual of the Kalman filtering problem with scheduling constraints studied in [14]. Future work will consider the LQG problem under output feedback and constraints on the measurement and control signals.

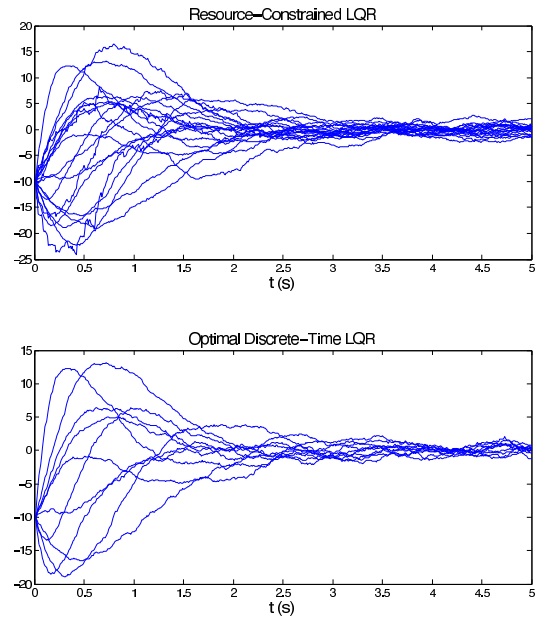


Figure 5: Sample trajectories (plotted at the sample points $k\Delta, k \geq 0$) for a random system with 10 states and 7 inputs, for the digital implementation with schedule (TT) and $\Delta = 10$ ms. One of the 7 inputs can be used in each slot. We also show the response obtained with the standard discrete-time LQR controller (i.e., LQR design after discretizing the dynamics).

9. REFERENCES

- [1] B. D. O. Anderson and J. Moore. *Optimal Control: Linear Quadratic Methods*. Dover, 2007. Republication of the 1990 edition.
- [2] P. J. Antsaklis and J. Bailleul. Special issue on networked control systems technology. *Proceedings of the IEEE*, 95(1):9–28, 2007.
- [3] K.-E. Arzen, A. Cervin, and D. Henriksson. Implementation-aware embedded control systems. In D. Hristu-Varsakelis and W. Levine, editors, *Handbook of Networked and Embedded Control Systems*, 2005.
- [4] K. J. Åström. *Introduction to Stochastic Control Theory*. Dover, 2006. Republication of the 1970 edition.
- [5] K. J. Åström and B. Wittenmark. *Computer-Controlled Systems: Theory and Design*. Prentice Hall, 3rd edition, 1997.
- [6] S. Bittanti, P. Colaneri, and G. D. Nicolao. The periodic Riccati equation. In S. Bittanti, A. J. Laub, and J. C. Willems, editors, *The Riccati equation*. Springer-Verlag, 1991.
- [7] S. P. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2006.
- [8] T. Chen and B. Francis. *Optimal Sampled-Data Control Systems*. Springer, 1995.
- [9] G. C. Goodwin, J. I. Yuz, J. C. Agüero, and M. Cea. Sampling and sampled-data models. In *Proceedings of the American Control Conference*, pages 1–20, Baltimore, MD, 2010.
- [10] D. Hristu-Varsakelis and W. S. Levine, editors.

- Handbook of Networked and Embedded Control Systems*. Springer-Verlag, 2005.
- [11] D. Hristu-Varsakelis and L. Zhang. LQG control of networked control systems with access constraints and delays. *International Journal of Control*, 81(8):1266–1280, August 2008.
- [12] H. Kopetz. *Real-Time Systems: Design Principles for Distributed Embedded Applications*. Kluwer Academic Publishers, 1997.
- [13] H. J. Kushner and G. G. Yin. *Stochastic Approximation and Recursive Algorithms and Applications*. Springer, 2nd edition, 2003.
- [14] J. Le Ny, E. Feron, and M. Dahleh. Scheduling continuous-time Kalman filters. *IEEE Transactions on Automatic Control*, 2011. To appear.
- [15] J. Le Ny and G. J. Pappas. Robustness analysis for the certification of digital controller implementations. In *Proceedings of the First International Conference on Cyber-Physical Systems (ICCPs)*, Stockholm, Sweden, April 2010.
- [16] J.-W. Lee. Infinite-horizon joint LQG synthesis of switching and feedback in discrete time. *IEEE Transactions on Automatic Control*, 54(8):1945–1951, 2009.
- [17] B. Lincoln and B. Bernhardsson. LQR optimization of linear system switching. *IEEE Transactions on Automatic Control*, 47(10):1701–1705, October 2002.
- [18] L. Meier, J. Peschon, and R. Dressler. Optimal control of measurement systems. *IEEE Transactions on Automatic Control*, 12(5):528–536, 1967.
- [19] R. Middleton and G. C. Goodwin. *Digital Control and Estimation. A Unified Approach*. Prentice Hall, Englewood Cliffs, New Jersey, 1990.
- [20] T. Nghiem, G. J. Pappas, R. Alur, and A. Girard. Time-triggered implementations of dynamic controllers. *ACM Transactions in Embedded Computing Systems*, 2010. In press.
- [21] J. Nilsson. *Real-Time Control Systems with Delays*. PhD thesis, Dept. Automatic Control, Lund Institute of Technology, Lund, Sweden, January 1998.
- [22] S. Sager. *Numerical Methods for Mixed-Integer Optimal Control Problems*. PhD thesis, Universität Heidelberg, 2006.
- [23] M. Salgado, R. Middleton, and G. Goodwin. Connection between continuous and discrete Riccati equations with applications to Kalman filtering. *Control Theory and Applications, IEE Proceedings D*, 135(1):28–34, 1988.
- [24] G. Weiss and R. Alur. Automata based interfaces for control and scheduling. In *Proceedings of the 10th International Conference on Hybrid Systems: Computation and Control*, 2007.
- [25] L. Zhang and D. Hristu-Varsakelis. Communication and control co-design for networked control systems. *Automatica*, 42(953-958), 2006.
- [26] W. Zhang, J. Hu, and J. Lian. Quadratic optimal control of switched linear stochastic systems. *Systems and Control Letters*, 2010. submitted.