53rd IEEE Conference on Decision and Control
December 15-17, 2014. Los Angeles, California, USA

Performance Analysis of a Class of Linear Quadratic Regulators for Switched Linear Systems

D. Antunes, W.P.M.H. Heemels

***Abstract*— We analyze a class of suboptimal policies for regulating the state of a switched linear system to zero while minimizing a quadratic cost. Our novel results provide upper and lower bounds on the performance of instances of this class with respect to the optimal policy and other policies of interest. We present a numerical example illustrating the applicability of the results to event-triggered control.**

I. INTRODUCTION

We consider the following switched linear system (SLS)

$$x_{k+1} = A_{\sigma_k} x_k + B_{\sigma_k} u_k, \quad k \in \mathbb{N}_0, \tag{1}$$

where $x_k \in \mathbb{R}^n$ is the state, $u_k \in \mathbb{R}^{n_u}$ is the control input, and $\sigma_k \in \{1, \dots, m\}$ is the switching input at discrete time $k \in \mathbb{N}_0 := \{0, 1, 2, \dots\}$. Our goal is to design a policy for the control and the switching input which regulates the state to zero while minimizing a quadratic infinite-horizon cost.

The control of SLSs arises in various applications such as mixing of fluids [1], HIV treatment [2], DC power conversion [3], event-triggered control [4], multi-agent systems [5], real-time control task scheduling [6], and damping of vibrating structures [7]. Model (1) captures many of these applications and can be used as an approximation to study several others by linearization and discretization of related non-linear and continuous-time models. The problem of minimizing a quadratic cost parallels the well-known LQR problem [8], which is one of the tools per excellence for regulating/stabilizing (1) in the special case $m = 1$, where only the control input is to be designed.

However, obtaining an optimal policy for the control and the switching inputs when $m > 1$ is in general computationally intractable [9] and thus one must settle for suboptimal policies. The work [3] proposes a suboptimal strategy resulting from an iterative relaxed dynamic programming algorithm; a relaxation parameter, specifying an acceptable loss of performance with respect to the optimal strategy, allows to trade performance cost and computational complexity. The work [9] proposes a related value function iterative method, formally establishing that a stabilizing stationary policy can be found after enough iterations; [9] also provides a bound on the cost of such policy with respect to the optimal policy. Nevertheless, for a given relaxation parameter, the computational complexity of the resulting policy can still be large. A different approach is proposed in [4], [10]; assuming the knowledge of a stabilizing base policy, a low complexity policy is derived based on rollout ideas [8, Ch. 6]. Yet, [4], [10] provide no performance guarantees with respect to the optimal policy. For related work see also [11].

In this paper we analyze a class of linear quadratic regulators parameterized by a set of positive semidefinite matrices, which define a piecewise quadratic function $\hat{J}$ approximating the optimal cost. This class of control policies is broad enough to capture the policies in [3], [9] and in [4], [10]. Our aim is to provide low complexity policies with (preferably tight) performance guarantees. To this effect, our novel results provide upper and lower bounds on the cost difference between $\hat{J}$ and the cost of the proposed policy. Choosing $\hat{J}$ as the cost of a given base policy we recover the policies in [4], [10]. We can then use our novel results to assert the gain obtained by a rollout policy over a base policy. On the other hand choosing $\hat{J}$ as a lower bound on the optimal cost, we can estimate the distance of the resulting policy from the optimal.

We illustrate the applicability of the results presented in this paper by considering a switched system arising from an event-triggered control setup proposed in [4]. For this example, we conclude that the bounds provided are reasonably close to the true values obtained by simulation.

The remainder of the paper is organized as follows. Section II presents the proposed class of linear quadratic regulators and Section III provides two general results for analyzing the performance of these regulators. We discuss the implications of these results when $\hat{J}$ is a special upper and lower bound on the optimal cost in Sections IV and V, respectively. Section VI presents a numerical example and Section VII provides concluding remarks.

II. A CLASS OF LINEAR QUADRATIC REGULATORS

Consider the following quadratic cost

$$\sum_{k=0}^{\infty} g(x_k, u_k, \sigma_k), \tag{2}$$

where

$$g(x, u, i) := \begin{bmatrix} x^\intercal & u^\intercal \end{bmatrix} \begin{bmatrix} Q_i & S_i \\ S_i^\intercal & R_i \end{bmatrix} \begin{bmatrix} x \\ u \end{bmatrix}$$

is assumed to be positive semi-definite for every $i \in \mathcal{M} := \{1, \dots, m\}$. We are interested in finding a policy μ, consisting of a function from the state to the control and the

Duarte Antunes and Maurice Heemels are with the Control Systems Technology Group, Department of Mechanical Engineering, Eindhoven University of Technology, the Netherlands. {D.Antunes,M.Heemels}@tue.nl

This work is supported by the Dutch Science Foundation (STW) and the Dutch Organization for Scientific Research (NWO) under the VICI grant "Wireless controls systems: A new frontier in Automation" (No. 11382), by the European 7th Framework Network of Excellence by the project "Highly-complex and networked control systems (HYCON2-257462)".

978-1-4673-6088-3/14/$31.00 ©2014 IEEE

switching inputs, such that (2) is minimized when

$$(u_k, \sigma_k) = \mu(x_k). \tag{3}$$

Finding an optimal μ leads in general to a combinatorial problem and thus is computationally intractable [9].

Here, we consider a class of suboptimal policies parameterized by a set $\mathcal{T}$ of n_T positive semi-definite matrices. Given $\mathcal{T}$ we define

$$\hat{J}(x) = \min_{T \in \mathcal{T}} x^\intercal T x, \tag{4}$$

and

$$\bar{J}_\mu(x) := \min_{u \in \mathbb{R}^{n_u}, i \in \mathcal{M}} g(x, u, i) + \hat{J}(A_i x + B_i u). \tag{5}$$

For each state $x \in \mathbb{R}^n$, $\mu(x)$ is defined based on one of the (possibly many) control and switching inputs that achieve the minimum in (5), and is characterized as follows. For a square matrix $T \in \mathbb{R}^n$, and $i \in \mathcal{M}$, let

$$G_i(T) := -(B_i^\intercal T B_i + R_i)^\dagger (B_i^\intercal T A_i + S_i^\intercal),$$

where the symbol † denotes the pseudo-inverse, and

$$F_i(T) := A_i^\intercal T A_i + Q_i - G_i(T)^\intercal (B_i^\intercal T B_i + R_i) G_i(T).$$

Moreover, for $n_P := m n_T$ and $\mathcal{P} := \{1, \dots, n_P\}$, let

$$\{P_j | j \in \mathcal{P}\}, \quad \{K_j | j \in \mathcal{P}\}, \tag{6}$$

be indexations of the sets

$$\{F_i(T) | T \in \mathcal{T}, i \in \mathcal{M}\}, \quad \{G_i(T) | T \in \mathcal{T}, i \in \mathcal{M}\},$$

respectively. Each value $j \in \mathcal{P}$ corresponds to a unique $T \in \mathcal{T}$ and to a unique $i = \pi(j) \in \mathcal{M}$, where

$$\pi : \mathcal{P} \to \mathcal{M} \tag{7}$$

is a map characterizing the latter correspondence. Then the minimum in (5) is achieved by $(u, i) = \mu(x)$, for

$$\mu(x) = (\bar{u}(x), \bar{\sigma}(x)), \tag{8}$$

where

$$\bar{\sigma}(x) = \pi(\iota(x)), \quad \bar{u}(x) = K_{\iota(x)} x, \tag{9}$$

and

$$\iota(x) = \min \operatorname*{argmin}_{j \in \mathcal{P}} \; x^\intercal P_j x. \tag{10}$$

Note that in (10) we have arbitrated that the smallest index is selected if for a given $x \in \mathbb{R}^n$ the minimum of $x^\intercal P_j x$ is achieved by more than one index $j \in \mathcal{P}$. Note also that (9) is in general a choice among the control inputs that achieve the minimum in (5). One such case of interest is when $B_i = 0$ for every $i \in \mathcal{M}$, i.e., only a policy for the switching input is to be designed.

The function $\hat{J}$ can be interpreted as an *approximation* to the optimal cost (2) and the associated policy μ is then an approximate dynamic programming policy [8, Ch. 6]. Several examples will be given below (see Sections IV, V).

III. General Results for Performance analysis

Let $J_\mu(x_0)$ denote the cost (2) of policy μ when (3), (8) is applied to (1) initialized at x_0. Our first result relates this cost to $\hat{J}(x_0)$ through the function $\bar{J}_\mu(x)$. Note that $\bar{J}_\mu(x) = x^\intercal P_{\iota(x)} x$. We assume that μ is stabilizing, i.e., $\|x_k\| \to 0$ as $k \to \infty$ when (3), (8) is applied to (1) and discuss this assumption for specific choices of $\hat{J}$ in the sequel (see Sections IV, V).

Theorem 1: Let $\{x_k\}_{k \geq 0}$ denote the solution to (1) for an initial state x_0 and for a stabilizing μ taking the form (3), (8). Then

$$J_\mu(x_0) = \hat{J}(x_0) + \sum_{k=0}^{\infty} (\bar{J}_\mu(x_k) - \hat{J}(x_k)).$$

□

We can interpret each term in the summation as a gain if

$$\bar{J}_\mu(x_k) - \hat{J}(x_k) \leq 0, \tag{11}$$

or as a loss if

$$\bar{J}_\mu(x_k) - \hat{J}(x_k) \geq 0 \tag{12}$$

and computing these gains and losses along the trajectories of the SLS controlled by μ provides the cost difference between the estimate $\hat{J}$ and the true cost J_μ of policy μ. We will present in the sequel policies which guarantee that either (11) or (12) holds for every x_k, which result from making $\hat{J}$ either an upper or a lower bound on the optimal cost, respectively.

The next result provides upper and lower bounds on the cost difference between $J_\mu(x_0)$ and $\hat{J}(x_0)$, exploiting the piecewise quadratic form of these functions. Let

$$\Phi_j := A_{\pi(j)} + B_{\pi(j)} K_j, \quad j \in \mathcal{P}.$$

Theorem 2: Suppose that μ is stabilizing. If there exist a matrix U and a non-negative scalar $\lambda < 1$ such that, for every $x \in \mathbb{R}^n$,

$$\min_{i \in \mathcal{P}} x^\intercal P_i x \leq \min_{T \in \mathcal{T}} x^\intercal T x + x^\intercal U x, \tag{13}$$

and, for every $x \in \mathbb{R}^n$,

$$x^\intercal \Phi_{\iota(x)}^\intercal U \Phi_{\iota(x)} x \leq \lambda x^\intercal U x, \tag{14}$$

then

$$J_\mu(x_0) \leq \hat{J}(x_0) + \frac{1}{1-\lambda} x_0^\intercal U x_0. \tag{15}$$

Moreover, if there exist a matrix W and a non-negative scalar $\gamma < 1$ such that, for every $x \in \mathbb{R}^n$,

$$\min_{i \in \mathcal{P}} x^\intercal P_i x \geq \min_{T \in \mathcal{T}} x^\intercal T x + x^\intercal W x, \tag{16}$$

and, for every $x \in \mathbb{R}^n$,

$$x^\intercal \Phi_{\iota(x)}^\intercal W \Phi_{\iota(x)} x \geq \gamma x^\intercal W x, \tag{17}$$

then

$$J_\mu(x_0) \geq \hat{J}(x_0) + \frac{1}{1-\gamma} x_0^\intercal W x_0. \tag{18}$$

□

This result has several consequences when $\hat{J}$ is a lower or an upper bound on the optimal cost. We discuss this next, also providing a numerical method in terms of LMIs to obtain the bounds (15), (18) for these special cases.

IV. ROLLOUT

We present a simple rollout policy in Section IV-A and in Section IV-B we discuss variants referred to as lifted policies. We show how our results can be used to analyze the performance of these policies in Sections IV-C and IV-D.

A. One-step policy improvement and stability

Suppose that we know a stabilizing policy characterized by a fixed sequence of schedules $b = (b_0, b_1, \dots)$, i.e., $\sigma_k = b_k$, $k \in \mathbb{N}_0$, *independent* of the state, and a linear feedback control input policy $u_k = K_k x_k, \quad k \in \mathbb{N}_0$. Note that once b is fixed the optimal control policy (associated with b) is a natural policy taking this linear feedback form (see [8]). For concreteness, here we consider such optimal control policy, although we could consider more general linear feedback policies. We refer to this policy as the base policy and note that in many problem of interest for SLSs there exist natural choices of base policies, e.g., in event-triggered control, a natural base policy is optimal periodic control [4]; in mixing of fluids, periodic mixing [1]; in protocol design for wireless control, Round-Robin protocols [10], etc.

The cost (2) when this *base policy* is taken for (1) is a positive semi-definite quadratic function $x_0^\intercal P_b x_0$ (cf. [8]). We then pick (4) as

$$\hat{J}(x_0) = \min_{b \in \mathcal{B}} x_0^\intercal P_b x_0, \tag{19}$$

where $\mathcal{B}$ denotes a set of sequences each characterizing a base policy. We can see (19) as a special base policy consisting of the minimum of base policies [8] and the associated policy (8) can then be seen as a rollout policy using a one-step policy improvement [8]. In this case one can conclude that

$$\bar{J}_\mu(x_0) = \min_{c \in \mathcal{C}} x_0^\intercal P_c x_0,$$

where

$$\mathcal{C} := \{(i, b) | i \in \mathcal{M}, b \in \mathcal{B}\} \tag{20}$$

is determined by the choice of $\mathcal{B}$. The set $\mathcal{B}$ is assumed to be such that

$$\mathcal{B} \subseteq \mathcal{C}. \tag{21}$$

For example, if $m = 2$, $\mathcal{B} = \{(1, 2, 1, 2, 1, \dots)\}$ does not satisfy (21) but $\mathcal{B} = \{(2, 1, 2, 1, 2, \dots), (1, 2, 1, 2, 1, \dots)\}$ does. Note that this guarantees that (11) holds, i.e., at each iteration there is always a cost improvement over the base policy, and from this one can show that the cost (2) of the rollout policy $J_\mu(x_0)$ is less that the corresponding base policy [8, Ch. 6].

From this observation we conclude that the cost of the rollout policy is bounded, which implies that the rollout policy is stabilizing under mild conditions (see [10], [8]). One such condition is to assume that the running cost g is (strictly) positive definite, since in this case if the cost is bounded the state must converge to zero as time approaches infinity. A similar reasoning can be used to establish stability for the variants of this policy which we consider next.

B. Lifted policies

We pick $h \in \mathbb{N}_{\geq 2}$ and write (1) as

$$\underline{x}_{\ell+1} = \underline{A}_{\underline{\sigma}_\ell} \underline{x}_\ell + \underline{B}_{\underline{\sigma}_\ell} \underline{u}_\ell, \quad \ell \in \mathbb{N}_0, \tag{22}$$

where $\underline{x}_\ell = x_{\ell h}$, $\underline{u}_\ell = [u_{\ell h}^\intercal \dots u_{(\ell+1)h-1}^\intercal]^\intercal$, and $\underline{\sigma}_k \in \mathcal{M}_h := \{1, \dots, m^h\}$, is such that each $\underline{\sigma}_k = \underline{i} \in \mathcal{M}_h$ corresponds to a unique $(\sigma_0, \dots, \sigma_{h-1}) \in \mathcal{M}^h := \mathcal{M} \times \dots \times \mathcal{M}$. Note that (22) is the lifted system of (1) over h time steps. Moreover, (2) can be written as

$$\sum_{\ell=0}^{\infty} \underline{g}(\underline{x}_\ell, \underline{u}_\ell, \underline{\sigma}_\ell) \tag{23}$$

for a quadratic function positive semi-definite $\underline{g}$. The expressions for $\underline{A}_{\underline{i}}$, $\underline{B}_{\underline{i}}$, $\underline{i} \in \mathcal{M}_h$, and $\underline{g}$ are omitted for brevity. We define a policy for this lifted system as in (8)-(10), i.e.,

$$\underline{\mu}(\underline{x}) = (\underline{u}(x), \underline{i}(x)) \tag{24}$$

where for $\underline{x} \in \mathbb{R}^n$, $\underline{u}(\underline{x}) \in \mathbb{R}^{n_u h}$, $\underline{i}(\underline{x}) \in \mathcal{M}_h$ are control and switching input values attaining the minimum of $\underline{g}(\underline{x}, \underline{u}, \underline{i}) + \hat{J}(\underline{A}_{\underline{i}} \underline{x} + \underline{B}_{\underline{i}} \underline{u})$ and $\hat{J}$ is given by a general (4), which for rollout policies takes the form (19). We then define the *lifted policy* for the *original* system (1) by applying the following control and schedules for $k \in \{\ell h, \dots, (\ell+1)h - 1\}$

$$(u_k, \sigma_k) = (\underline{v}_{k-\ell h}, \underline{s}_{k-\ell h}) \tag{25}$$

where

$$(\underline{v}, \underline{i}_s) = \underline{\mu}(x_{h\ell}), \tag{26}$$

and $\underline{s}$ is the element in $\mathcal{M}^h$ corresponding to $\underline{i}_s \in \mathcal{M}_h$. Note that (26) is only computed at times $k = \ell h$. Then, to analyze a lifted policy for (1), one can apply the results of Section III considering (22) driven by (24), and the cost (23).

C. Lower-bounding the gain

We now show how to use the first part of Theorem 1 to lower-bound the performance improvement of the rollout policy over the corresponding base policy. This is clearly interesting in the applications mentioned before (event-triggered control, mixing of fluids, etc) as it gives a guaranteed cost gain over the base policies.

We start by considering one-step policy improvement policies. Since Assumption (21) guarantees that (11) holds we can assume in (13) that $U = -\underline{U}$ for some positive semi-definite $\underline{U}$. We then rewrite conditions (13) and (14) as requiring that for every $x \in \mathbb{R}^n$

$$\min_{i \in \mathcal{P}} x^\intercal P_i x \leq \min_{T \in \mathcal{T}} x^\intercal T x - x^\intercal \underline{U} x, \tag{27}$$

and

$$x^\intercal \Phi_{\iota(x)}^\intercal \underline{U} \Phi_{\iota(x)} x \geq \lambda x^\intercal \underline{U} x, \tag{28}$$

respectively, where

$$\{P_i | i \in \mathcal{P}\} = \{P_c | c \in \mathcal{C}\}, \qquad \mathcal{T} = \{P_b | b \in \mathcal{B}\}, \tag{29}$$

and the cost *gain* of a rollout policy (8) over the cost (2) of the base policy (19) in (15) is then given by

$$\frac{1}{1-\lambda} x_0^\intercal \underline{U} x_0. \tag{30}$$

Note that (28) is always met for $\lambda = 0$ as $\underline{U} \geq 0$, in which case (15) is a known bound to compare base and rollout policies, see [8, p. 338]. Condition (28) imposes that the gain (11) at time $k+1$ must be always at least λ times the gain at time k for every time k and state x_k. This is only possible if the gain $y^\intercal \underline{U} y$, $y = \Phi_{\iota(x)} x$ is never zero for any state x, which typically amounts to requiring $\underline{U} > 0$ in (27). In turn, if (27) is achieved with $\underline{U} > 0$ then there exists a strictly positive gain for every non-zero state x, which is in general too much to expect from the single-step rollout policy. In fact, this would imply that for every non-zero x the choice in (10) for (29), would correspond to a sequence in $\mathcal{C}$ not contained in $\mathcal{B}$.

On the other hand, if we consider lifted policies the set $\mathcal{C}$, described by (20) with $\mathcal{M}$ replaced by the much larger set $\mathcal{M}^h$, will generally become sufficiently rich as h increases so that this latter condition can be met. For example if $m = 2$ and $\mathcal{B} = \{b\}$ then $\mathcal{C} = \{(1, b), (2, b)\}$ for a one-step improvement policy, but for a lifted policy with $h = 2$ we have $\mathcal{C} = \{(1, 1, b), (2, 1, b), (1, 2, b), (2, 2, b)\}$. By considering lifted policies we mean that we analyze conditions (27), (28) for the lifted switched system (22)-(24), i.e., the matrices P_i, Φ_i should be replaced by $\underline{P}_i \in \underline{\mathcal{P}}$ and

$$\underline{\Phi}_j := \underline{A}_{\underline{\pi}(j)} + \underline{B}_{\underline{\pi}(j)} \underline{K}_j, \quad j \in \underline{\mathcal{P}},$$

where $\underline{\mathcal{P}}$, $\underline{K}_j$, and $\underline{\pi}$ are defined as in (6), (7), for (22), (23). We will use this reference to lifted policies several times below with this meaning.

Note that there is a trade-off in maximizing the gain (30) by augmenting h for lifted policies. In fact, augmenting h leads to larger $\underline{U}$, in the sense that $\underline{U} > \gamma I$ for larger γ, but reduces the largest λ that satisfies (28) since $\underline{x}_\ell = x_{\ell h}$ will converge to zero faster for larger h. Moreover, since condition (28) is not easy to test, we will need to test it for every i (see (33) below, or the relaxed version (34)) creating many more constraints (exponentially in h) upon testing (28).

To mitigate the latter issue, we propose two methods. The first is to prune the set $\mathcal{P}$ since there may exist redundant matrices in this set that never correspond to the minimum in (10), but still impose a constraint while testing (28) using (33) or (34) below. A matrix $P_j \in \mathcal{P}$ can be pruned if the ellipsoid $\{x | x^\intercal P_j x \leq 1\}$ is covered by the ellipsoids corresponding to the other elements in $\mathcal{P}$. A sufficient condition to prune P_j (cf. [9], [3]) is the existence of non-negative scalars α_ℓ, $\ell \in \mathcal{P}$, adding up to one, such that

$$\sum_{\ell \in \mathcal{P} \setminus \{j\}} \alpha_\ell P_\ell \leq P_j \tag{31}$$

Inspired by [9], one can add ϵI, for a small $\epsilon > 0$, to the right-hand side of (31), leading to more pruned matrices in $\mathcal{P}$, at the expense of less gain guarantees. The second method, which we shall follow in the numerical example, is to compute $\ell^{(\kappa)} = \text{argmin}_{j \in \mathcal{P}} x^{(\kappa)\intercal} P_j x^{(\kappa)}$ for a representative set of $x^{(\kappa)}$ and consider only the switching sequence associated with $P_{\ell^{(\kappa)}}$. Such representative set can be chosen randomly or by simulating the trajectory of (1) driven by the rollout policy for given initial conditions. Note that to guarantee stability of the rollout policy using the arguments mentioned before one should make sure that (21) still holds after pruning by either method.

1) Numerical method to maximize the gain: Using a similar reasoning to the one that lead to (31), we conclude that a *sufficient* condition to test (27) is to assert the existence of non-negative α_i, $i \in \mathcal{P}$, adding up to one, such that

$$\sum_{i \in \mathcal{P}} \alpha_i P_i \leq T - \underline{U}, \quad \text{for all } T \in \mathcal{T}. \tag{32}$$

Requiring

$$\Phi_i^\intercal \underline{U} \Phi_i \geq \lambda \underline{U}, \tag{33}$$

for every $i \in \mathcal{P}$ is sufficient to guarantee (28). However, we can use a relaxation (S-Procedure) to obtain a less restrictive sufficient condition for (28). In fact, one can see that (28) is met if there exist non-negative β_{ji} such that

$$\Phi_i^\intercal \underline{U} \Phi_i + \sum_{j \in \mathcal{P}} \beta_{ji} (P_i - P_j) \geq \lambda \underline{U}, \quad \text{for all } i \in \mathcal{P}. \tag{34}$$

Then to maximize (30) for a particular x_0 we can pick a dense grid of points λ in the interval $[0, 1)$ and for each λ solve the LMI problem:

Problem 1 $\max \ x_0^\intercal \underline{U} x_0$
s.t. $\underline{U} > 0$, (32), (33).

On the other hand if we are interested in maximizing a lower bound on the gains obtained for every initial condition we can consider:

Problem 2 $\max \ \xi$
s.t. $\underline{U} > \xi I, \xi > 0$, (32), (33).

The maximum lower bound on the gain is then obtained by plotting (30) as a function of λ, where $\underline{U}$ results from the solutions to these LMI problems. If these LMI problems are unfeasible even for $\lambda = 0$ then condition (32) is not satisfied meaning that $\underline{U}$ cannot be picked as (strictly) positive definite.

D. Upper-bounding the gain

The second part of Theorem 2 can be used to assert how far is the base policy from a corresponding rollout policy; e.g. in the context of ETC [4] it is useful to guarantee that a periodic control strategy performs already well enough with respect to a rollout (even-triggered) policy.

In this case we can assume in (16) that $W = -\underline{W}$ for a positive semi-definite matrix $\underline{W}$, and rewrite (16) and (17) as requiring that for every $x \in \mathbb{R}^n$,

$$\min_{i \in \mathcal{P}} x^\intercal P_i x \geq \min_{T \in \mathcal{T}} x^\intercal T x - x^\intercal \underline{W} x, \tag{35}$$

and

$$x^\intercal \Phi_{\iota(x)}^\intercal \underline{W} \Phi_{\iota(x)} x \leq \gamma x^\intercal \underline{W} x. \tag{36}$$

We are interested in *minimizing* the upper bound on the cost difference between the rollout and base policies

$$\frac{1}{1-\gamma} x_0^\intercal \underline{W} x_0. \tag{37}$$

Here (36) imposes the most stringent condition, as there may not be a constant $\gamma < 1$ satisfying (36) whereas (35) is always met for sufficiently large $\underline{W}$. Note, however, that for lifted policies we can pick h large enough to reduce γ satisfying (36) and achieve $\gamma < 1$, although this in general leads to larger $\underline{W}$ that satisfies the equivalent version of (35) for the lifted problem. Hence, a trade-off is also present here.

In case $\mathcal{T} = \{T\}$, i.e., $\mathcal{T}$ contains only one matrix, which is the situation in many problems of interest, we can test (35) with the following *equivalent* condition

$$P_i > T - \underline{W}, \quad \text{for all } i \in \mathcal{P}.$$

Moreover one can use a similar relaxation as in (34) to test (36). Numerical methods to minimize the loss function (37) can be derived using the same ideas as in Section IV-C, involving a search over the parameter γ.

V. Policies resulting from lower bounds on the optimal cost

In many applications of SLSs a lower bound on the optimal cost (2) is known. Prime examples are applications in which (1), (2) result from the discretization of a continuous-time model for which an optimal input is known, but cannot be applied due to some constraints in the problem captured by a switching variable. This is the case, e.g., in networked control where the controller communicates with the process via a resource constrained network. As a result the optimal continuous-time control cannot be applied and a switching policy must be designed to orchestrate transmissions, but clearly the optimal continuous-time cost is a lower bound on the cost achieved by any switching policy (cf. Section VI).

Here we consider that such a lower bound is a simple quadratic function, $x^\intercal Tx$, such that

$$x^\intercal Tx \leq J^*(x), \text{ for all } x \in \mathbb{R}^n,$$

where $J^*(x)$ is the optimal cost (2) over the class of feedback policies (3) for an initial state $x \in \mathbb{R}^n$. From standard arguments in dynamic programming [8] it follows that (12) must hold for every state x_k. Our framework would also allow us to consider more general pointwise minimum of quadratic functions, but we consider this special case for simplicity.

In this section, we analyze (8) when $\hat{J}$ is given by such quadratic lower bound. Contrarily to rollout policies, testing the first premise of Theorem 2 that μ is stabilizing may be challenging. We address this in Section V-A. In Section V-B we use the first part of Theorem 2 to bound the distance $J_\mu(x) - x^\intercal Tx$, which is also a bound to the distance of the cost of μ from the optimal since

$$J_\mu(x) - J^*(x) \leq J_\mu(x) - x^\intercal Tx, \text{ for all } x \in \mathbb{R}^n.$$

A. Stability

A natural candidate for a Lyapunov function to test if μ is stabilizing is $x^\intercal Tx$, especially in the mentioned cases where (1), (2) result from the discretization of a continuous-time model. That is, if T is positive definite and if

$$x^\intercal \Phi_{\iota(x)}^\intercal T \Phi_{\iota(x)} x - x^\intercal Tx < 0 \tag{38}$$

for every $x \in \mathbb{R}^n$ then μ is stabilizing. One can conclude that a sufficient condition to test this is the existence of non-negative β_{ij} such that

$$\Phi_i^\intercal T \Phi_i + \sum_{j \in \mathcal{P}} \beta_{ij}(P_j - P_i) - T < 0, \quad \text{forall } i \in \mathcal{P}.$$

Considering lifted policies defined as in Section IV-B, but using the lower bound $\hat{J}(x) = x^\intercal Tx$ as a cost estimate, one can show that (38) holds for large enough h, provided that the SLS is stabilizable, i.e., for every x_0 there exist a control and a switching input sequence that drives the state asymptotically to zero.

B. Upper-bounding the loss

We can now assume that U is positive semi-definite in (13)-(15) since, as argued before, (12) holds. Using a similar reasoning to the one that lead to (31) we conclude that (13) holds if there exist non-negative scalars α_i, $i \in \mathcal{P}$, adding up to one, such that

$$\sum_{i \in \mathcal{P}} \alpha_i P_i \leq T + U. \tag{39}$$

Moreover, (14) holds if there exist non-negative scalars β_{ji} such that

$$\Phi_i^\intercal U \Phi_i + \sum_{j \in \mathcal{P}} \beta_{ji}(P_j - P_i) \leq \lambda U, \quad \forall i \in \mathcal{P}. \tag{40}$$

Then to minimize the cost loss $\frac{1}{1-\lambda} x_0^\intercal U x_0$ of policy μ with respect to $x_0^\intercal T x_0$, for a particular initial condition x_0, we can solve:

Problem 3 $\min \; x_0^\intercal U x_0$
s.t. $U > 0, \; (39), \; (40),$

for each λ. This is also an upper bound on the loss of μ with respect to the optimal policy. The tighter bound is obtained by a line search over λ. If we are interested in bounding this loss for every initial condition we can consider

Problem 4 $\min \; \xi$
s.t. $0 < U < \xi I, \xi > 0, (39), \; (40).$

If these problems are unfeasible then stability is not guaranteed for the Lyapunov function candidate $x^\intercal Tx$. As mentioned before by considering lifted policies this will be the case, and one can also see that as h increases one can pick smaller λ to satisfy (40). Yet, as h increases, U satisfying (40) becomes larger, and thus there is a trade-off in minimizing the loss.

VI. Numerical Example

Consider the following linear process resulting from the linearization of the model of an inverted pendulum

$$\frac{d}{dt}\begin{bmatrix}\theta(t)\\ \dot{\theta}(t)\end{bmatrix} = \begin{bmatrix}\dot{\theta}(t)\\ \gamma_m \theta + T_m(t)\end{bmatrix}, \tag{41}$$

where $\theta(t)$ is the displacement angle, $\dot{\theta}(t)$ is the angular velocity, and $T_m(t)$ is the torque input at time $t \in \mathbb{R}_{\geq 0}$, and γ_m is a positive constant. The actuators are connected to a controller collocated with the sensors via a communication network. The controller can sample the state sensors, providing full state measurements periodically at times $t_k := k\tau$, $k \in \mathbb{N}_0$, for some sampling period $\tau > 0$. However, to save communication resources, the controller sends control values to the actuators only at a subset of times t_k, $k \in \mathbb{N}_0$, determined by the switching input $\{\sigma_k | k \in \mathbb{N}_0\}$, with $\sigma_k = 1$ when a transmission occurs and $\sigma_k = 2$ otherwise. The control objective is to minimize the following quadratic cost

$$\int_0^\infty \theta(t)^2 + \dot{\theta}(t)^2 + r_c T_m^2(t) dt. \tag{42}$$

Note that in the absence of a network the optimal control input would be the solution to the well known LQR problem, given by $T_m(t) = K_C x_C(t)$, where $x_C(t) = [\theta(t)\, \dot{\theta}(t)]^\intercal$ leading to a cost $x_C(0)^\intercal P_C x_C(0)$. For the numerical values $\gamma_m = 3$ and $r_c = 0.1$ this yields $K_C = [-7.3589 \;\; -4.9717]$ and

$$P_C = \begin{bmatrix} 2.1671 & 0.7359 \\ 0.7359 & 0.4972 \end{bmatrix}.$$

At the actuators side, the control is set to zero if there is no transmission at time t_k and follows a linear state feedback law with gain K_C otherwise, i.e.,

$$T_m(t) = \begin{cases} 0, \text{ if } \sigma_k = 2, \\ K_C x_k, \text{ if } \sigma_k = 1, \end{cases} \quad t \in [t_k, t_{t_k+1}), \tag{43}$$

where $x_k := x_C(t_k)$, $k \in \mathbb{N}_0$. Writing the equations for x_k we obtain a SLSs (1) with two modes $m = 2$ (transmit or not transmit) and $B_i = 0$, $i \in \{1, 2\}$. After exactly discretizing (41), (42) for $\tau = 0.1$ and taking into account (43) we obtain

$$A_0 = \begin{bmatrix} 1.0150 & 0.1005 \\ 0.3015 & 1.0150 \end{bmatrix}, \quad A_1 = \begin{bmatrix} 0.9782 & 0.0756 \\ -0.4381 & 0.5154 \end{bmatrix},$$

and

$$Q_0 = \begin{bmatrix} 0.1040 & 0.0202 \\ 0.0202 & 0.1013 \end{bmatrix}, \quad Q_1 = \begin{bmatrix} 0.6465 & 0.3553 \\ 0.3553 & 0.3065 \end{bmatrix}.$$

We are interested in computing a policy for $\{\sigma_k | k \in \mathbb{N}_0\}$ that minimizes the quadratic cost subject to a transmission rate constraint to save communication resources. Here we consider that the controller can only transmit on average at half of the rate it can measure the state, i.e., $\frac{1}{2\tau}$. To achieve this we use a similar scheme to [4] considering a lifted policy (in the sense of Section IV-B) that incorporates this constraint. We pick $h = 6$ and note that there are 20 scheduling options in $\{1, 2\}^6$ that satisfy the transmission constraint, e.g, $(1, 2, 1, 2, 1, 2)$ and $(2, 1, 1, 2, 2, 1)$.

We consider a policy obtained by considering the lower bound $x_0^\intercal P_C x_0$ as a cost estimate in (4), (8). The costs obtained by simulation and by the bounds obtained by solving Problem 3 are summarized in the next table for three initial conditions

$$x_0^{(1)} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \quad x_0^{(2)} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \quad x_0^{(3)} = \frac{2}{\sqrt{13}} \begin{bmatrix} 1 \\ -\frac{3}{2} \end{bmatrix}.$$

We also show the values obtained by the optimal continuous-time policy and a periodic switching policy for comparison.

Ini. Cond.	J_μ	Problem 3	$x_0^\intercal P_C x0$	Periodic
$x_0^{(1)}$	3.1836	3.2015	2.1671	3.6032
$x_0^{(2)}$	3.0697	3.1822	2.0680	3.6618
$x_0^{(3)}$	0.3625	0.3925	0.3317	0.3720

To obtain these values we restricted the options in $\{1, 2\}^6$ choosing a representative set of states and checking which scheduling options correspond to the choice in (24). With this method we restricted the scheduling options to $(1, 2, 1, 2, 1, 2)$, $(2, 1, 1, 1, 2, 2)$, $(1, 1, 1, 2, 2, 2)$, $(2, 1, 1, 2, 1, 2)$, $(1, 1, 2, 1, 2, 2)$. Note that the guaranteed bounds are reasonably tight to the values obtained by simulation. Considering Problem 4, we obtain $\lambda = 0.53$ and

$$W = \begin{bmatrix} 0.5151 & 0.1617 \\ 0.1617 & 0.2552 \end{bmatrix}$$

which allows us to conclude that

$$J_\rho(x_0) \le x_0^\intercal P_C x_0 + \frac{1}{1-\lambda} x_0^\intercal W x_0 = x_0^\intercal \begin{bmatrix} 3.2630 & 1.0799 \\ 1.0799 & 1.0401 \end{bmatrix} x_0.$$

VII. Concluding Remarks

In this paper we proposed and analyzed a class of suboptimal Linear Quadratic Regulators for SLSs. A numerical example illustrated that the provided performance bounds are reasonably tight to the performance of the suboptimal regulators obtained via simulation.

References

[1] V. Dolk, "Mixing and switching," Master's thesis, Eindhoven University of Technology (TU/e), 2013.

[2] E. Hernandez-Vargas, P. Colaneri, R. Middleton, and F. Blanchini, "Discrete-time control for switched positive systems with application to mitigating viral escape," *International Journal of Robust and Nonlinear Control*, vol. 21, no. 10, pp. 1093–1111, 2011.

[3] B. Lincoln and A. Rantzer, "Relaxing dynamic programming," *IEEE Transactions on Automatic Control*, vol. 51, no. 8, pp. 1249–1260, Aug 2006.

[4] D. Antunes and W. P. M. H. Heemels, "Rollout event-triggered control: Beyond periodic control performance," *IEEE Transactions on Automatic Control*, 2014, accepted for the special issue on Cyber Physical Systems.

[5] R. Olfati-Saber and R. Murray, "Consensus problems in networks of agents with switching topology and time-delays," *IEEE Transactions on Automatic Control*, vol. 49, no. 9, pp. 1520–1533, Sept 2004.

[6] A. Cervin and P. Alriksson, "Optimal on-line scheduling of multiple control tasks: A case study," in *Real-Time Systems, 2006. 18th Euromicro Conference on*, 2006, pp. 10 pp.–150.

[7] F. Blanchini, D. Casagrande, P. Gardonio, and S. Miani, "On optimal damping of vibrating structures," pp. 10 268–10 273, Sep 2011.

[8] D. P. Bertsekas, *Dynamic Programming and Optimal Control*. Athena Scientific, 2005.

[9] W. Zhang, J. Hu, and A. Abate, "Infinite-horizon switched lqr problems in discrete time: A suboptimal algorithm with performance analysis," *IEEE Transactions on Automatic Control*, vol. 57, no. 7, pp. 1815–1821, July 2012.

[10] D. Antunes, W. P. M. H. Heemels, J. Hespanha, and C. Silvestre, "Scheduling measurements and controls over networks- part I: Rollout strategies for protocol design," in *American Control Conference (ACC), 2012*, June 2012, pp. 2036–2041.

[11] D. Gorges, M. Izak, and S. Liu, "Optimal control and scheduling of switched systems," *IEEE Transactions on Automatic Control*, vol. 56, no. 1, pp. 135–140, Jan 2011.