49th IEEE Conference on Decision and Control
December 15-17, 2010
Hilton Atlanta Hotel, Atlanta, GA, USA

On the Infinite Horizon Constrained Switched LQR Problem

Maximilian Balandat, Wei Zhang, Alessandro Abate

***Abstract*— This paper studies the Discrete-Time Switched LQR problem over an infinite time horizon subject to polyhedral constraints on state and control input. The overall constrained, infinite-horizon problem is split into two subproblems: (i) an *unconstrained*, infinite-horizon problem and (ii) a *constrained*, finite-horizon one. We derive a stationary suboptimal policy for problem (i) with analytical bounds on its optimality, and develop a formulation of problem (ii) as a Mixed-Integer Quadratic Program. By introducing the concept of a *safe set*, the solutions of the two subproblems are combined to achieve the overall control objective. It is shown that, by proper choice of the design parameters, the error of the overall sub-optimal solution can be made arbitrarily small. The approach is tested through a numerical example.**

I. Introduction

Among the template problems in optimal control, the Linear Quadratic Regulator (LQR) is a fundamental one [1]. The study of this problem in a discrete-time framework in the recent past has witnessed enticing extensions to models subject to hard constraints on the states and control inputs [2], [3], [4], [5]. In addition, an extension of the LQR problem to Discrete-Time Switched Linear Systems (DSLS), referred to as the Discrete-Time Switched LQR (DSLQR) problem, has also been extensively studied [6], [7], [8], [9]. However, one of the main restrictions of this line of work is that it deals exclusively with unconstrained DSLS.

The main contribution of this paper is the development of a general framework for solving the Discrete-Time Constrained Switched LQR (DCSLQR) problem over an infinite horizon. Specifically, we characterize the infinite-horizon *hybrid-control sequence* (continuous and discrete controls) that minimizes a quadratic cost function, subject to polyhedral constraints on the state and on the input. Motivated by the solution of the classical (non-hybrid) constrained LQR problem developed in [4], we split the infinite-horizon DCSLQR problem into the following two related subproblems:

(i) Infinite-Horizon Unconstrained DSLQR: the solution to this problem can be computed efficiently using the numerical relaxation framework from [6], [8] and is a stationary hybrid-control law, characterized by a set of positive semidefinite matrices.

(ii) Constrained, Finite-Time, Optimal Hybrid Control (CFTOHC): this problem can be formulated as a Mixed Integer Quadratic Program (MIQP) and solved efficiently using available optimization algorithms.

The above two subproblems are connected through the so-called *safe set*, that is a set of states for which the solution to the unconstrained DSLQR problem is guaranteed to be always feasible. We show that if the unconstrained DSLS model is stabilizable, then for reasonable constraints a non-trivial safe set always exists. Moreover, we show that in case the constrained system starting from the given intial condition is stablizable, then with a sufficiently large optimization horizon the solution of the CFTOHC problem (ii) can always drive the state trajectory into the safe set, from where the solution of the (unconstrained) DSLQR problem (i) is feasible and optimal. Thus, by concatenating the solutions of the first and the second subproblem, we can obtain the solution to the overall infinite-horizon DCSLQR problem.

Based on the above ideas, an efficient algorithm is developed to solve the infinite-horizon DCSLQR problem with guaranteed suboptimal performance. We formally show that, by proper choice of the design parameters, the suboptimality error can be made arbitrarily small. The performance of the algorithm is demonstrated through a numerical example.

Due to space constraints, the proofs of most of the theorems in this paper will be omitted.

Notation: n, p and M are arbitrary finite positive integers; $\mathbb{Z}^+$ denotes the set of nonnegative integers, $\mathbb{M} \triangleq \{1, \dots, M\}$ is the set of subsystem indices, I_n is the $n \times n$ identity matrix; $\|\cdot\|$ represents the standard Euclidean norm in $\mathbb{R}^n$, and the induced norm over n-dim. matrices; $|\cdot|$ denotes the cardinality of a given set; $\mathcal{A}$ denotes the set of positive semidefinite (p.s.d.) matrices; $2^{\mathcal{A}}$ is the power set of $\mathcal{A}$; $\lambda_{\min}(\cdot)$ and $\lambda_{\max}(\cdot)$ characterize the smallest and the largest eigenvalues, respectively, of a given positive semidefinite (p.s.d.) matrix. The variable z denotes a generic initial state of system (1).

II. Problem Formulation

Consider the DSLS described by:

$$x_{t+1} = A_{v_t} x_t + B_{v_t} u_t \tag{1}$$

subject to the constraints

$$x_t \in \mathcal{X},\, u_t \in \mathcal{U},\; \forall\, t \in \mathbb{Z}^+, \tag{2}$$

where $x_t \in \mathbb{R}^n$ is the continuous state, $u_t \in \mathbb{R}^p$ is the continuous control and $v_t \in \mathbb{M}$ is the discrete control that determines the discrete mode at time t. The sets $\mathcal{X}$ and $\mathcal{U}$ are polyhedra that contain the origin in their interiors. The sequence of pairs $\psi_\infty = \{(u_t, v_t)\}_{t=0}^{\infty}$ is called the *hybrid-control sequence*. For each $i \in \mathbb{M}$, A_i and B_i are constant

M. Balandat is with the Department of Electrical Engineering and Information Technology, Technische Universität Darmstadt, Darmstadt, Germany balandat@stud.tu-darmstadt.de

W. Zhang is with the Department of Electrical Engineering and Computer Science, University of California, Berkeley, CA, USA weizhang@eecs.berkeley.edu

A. Abate is with the Delft Center for Systems and Control, Technische Universiteit Delft, Delft, The Netherlands a.abate@tudelft.nl

matrices of appropriate dimensions and the pair (A_i, B_i) denotes a subsystem.

For each $t \in \mathbb{Z}^+$, denote by $\xi_t \triangleq (\mu_t, \nu_t) : \mathbb{R}^n \mapsto \mathbb{R}^p \times \mathbb{M}$ the *(state-feedback) hybrid-control law* of system (1), where $\mu_t : \mathbb{R}^n \mapsto \mathbb{R}^p$ is called the *continuous-control law* and $\nu_t : \mathbb{R}^n \mapsto \mathbb{M}$ is called the *switching-control law*. A sequence of hybrid-control laws constitutes an *infinite horizon hybrid-control policy* $\pi_\infty = \{\xi_0, \xi_1, \dots\}$. A policy is called *stationary* if it consists of the same control law at all time, i.e., $\xi_t = \xi$ for all $t \in \mathbb{Z}^+$. The closed-loop dynamics of a system controlled by a policy π_∞ are given by:

$$x_{t+1} = A_{\nu_t(x_t)} x_t + B_{\nu_t(x_t)} \mu_t(x_t). \tag{3}$$

Denote by $\pi_\infty(z)$ the hybrid-control sequence generated by the policy π_∞ for the initial condition $x_0 = z$, i.e. $\pi_\infty(z) = \{(\mu_t(x_t), \nu_t(x_t))\}_{t\in\mathbb{Z}^+}$. Define the running cost function as:

$$L(x,u,v) = x^T Q_v x + u^T R_v u, \;\; \forall x \in \mathcal{X}, u \in \mathcal{U}, v \in \mathbb{M},$$

where $Q_v = {Q_v}^T \succ 0$ and $R_v = {R_v}^T \succ 0$ be the weighting matrices for the state and continuous-control input in subsystem $v \in \mathbb{M}$. Each hybrid-control sequence $\psi_\infty = \{(u_t, v_t)\}_{t=0}^\infty$ is associated with a quadratic performance index:

$$J_\infty(z, \psi_\infty) = \sum_{t=0}^{\infty} L(x_t, u_t, v_t), \tag{4}$$

where x_t is the closed-loop trajectory controlled by ψ_∞, with initial condition $x_0 = z \in \mathbb{R}^n$. Our objective is to find the optimal hybrid-control sequence that solves the following constrained optimal control problem:

$$\begin{cases} J^*_\infty(z) = \inf\limits_{\substack{v_0, v_1, \dots \\ u_0, u_1, \dots}} J_\infty(z, \psi_\infty) \\ \textbf{s.t. } (1) \text{ and } (2) \text{ with } x_0 = z. \end{cases} \tag{5}$$

We will refer to the above problem as the *Discrete-Time Constrained Switched LQR problem* (DCSLQR). The stated problem is an extension of the classical discrete-time LQR controller synthesis to Switched Linear Systems subject to polyhedral input and state constraints. Because of the complexity of the problem at hand, we do not require to find the optimal policy π^*_∞ for all initial states $z \in \mathcal{X}$, but rather the control sequence ψ^*_∞ for a given initial condition.

Clearly, if the *unconstrained* system (1) is not stabilizable, $J_\infty(x_0, \psi_\infty)$ will be infinite for all possible control sequences ψ_∞ for any $x \neq 0$. Therefore, the stabilizability is a minimum requirement for the well-posedness of problem (5).

Definition 1 (exp. stabilizability): The unconstrained system (1) is called *exponentially stabilizable* if there exists a policy π_∞ and constants $a \geq 1$ and $0 < c < 1$ such that the closed-loop trajectory under the policy π_∞ starting from any initial state $x_0 = z$ satisfies $\|x_t\|^2 \leq a c^t \|z\|^2$, $\forall t \in \mathbb{Z}^+$.

The following assumption is made throughout this paper:

(A1) The unconstrained SLS (1) is exp. stabilizable.

Remark 1: Assumption (A1) holds for most problems of practical interest and can be easily verified. In particular, if one of the unconstrained subsystems is stabilizable, the assumption will be satisfied. Furthermore, even in the case that none of the subsystems is stabilizable, it is still possible for the overall unconstrained switched linear system to be exponentially stabilizable (see Section VI and [8]).

III. The Unconstrained Hybrid Control Problem

In this section, we recall some recent results on the DSLQR problem [9], which can be viewed as a special case of problem (5) with the trivial constraints $\mathcal{X} = \mathbb{R}^n$, $\mathcal{U} = \mathbb{R}^p$. These results are crucial for solving the general DCSLQR problem, as will be further discussed in Section V-B.

A. The Value Function

For each $k \in \mathbb{Z}^+$, define the k-horizon value function as

$$\begin{cases} J^*_{k,uc}(z) = \inf\limits_{\substack{u_t \in \mathbb{R}^p, v_t \in \mathbb{M} \\ 0 \leq t \leq k-1}} \sum_{t=0}^{k-1} L(x_t, u_t, v_t) \\ \textbf{s.t. } \text{equation (1) with } x_0 = z. \end{cases} \tag{6}$$

The objective of DSLQR is to solve the infinite horizon unconstrained optimal control problem, i.e. to find the infinite-horizon value function $J^*_{\infty,uc}(z) = \lim_{k\to\infty} J^*_{k,uc}(z)$.

By a standard result of dynamic programming, the finite-horizon value functions can be computed recursively through the *one-stage value iteration*:

$$J^*_{k+1,uc}(z) = \inf_{u\in\mathbb{R}^p, v\in\mathbb{M}} \{L(z,u,v) + J^*_{k,uc}(A_v z + B_v u)\}.$$

An important feature of the DSLQR problem is that its finite-horizon value function can be characterized analytically. The key idea is to generalize the notion of the well known Difference Riccati Equation (DRE) [1] to the Switched LQR problem. For each subsystem $i \in \mathbb{M}$, define the *Riccati Mapping* $\rho_i : \mathcal{A} \mapsto \mathcal{A}$ as:

$$\rho_i(P) = Q_i + A_i^T P A_i - A_i^T P B_i \left(R_i + B_i^T P B_i\right)^{-1} B_i^T P A_i. \tag{7}$$

Definition 2 (Switched Riccati Mapping): The mapping $\rho_\mathbb{M} : 2^\mathcal{A} \mapsto 2^\mathcal{A}$ defined by:

$$\rho_\mathbb{M}(\mathcal{H}) = \{\rho_i(P),\ i \in \mathbb{M},\ P \in \mathcal{H}\} \tag{8}$$

is called the *Switched Riccati Mapping* (SRM) associated with the infinite horizon value function $J^*_{\infty,uc}(z)$.

Definition 3 (Switched Riccati Sets): The sequence of sets $\{\mathcal{H}_k\}_{k=0}^\infty$ generated iteratively by $\mathcal{H}_{k+1} = \rho_\mathbb{M}(\mathcal{H}_k)$ with $\mathcal{H}_0 = \{0\}$ is called the *Switched Riccati Sets* (SRS) associated with the infinite horizon value function $J^*_{\infty,uc}(z)$.

Starting from the singleton set $\{0\}$, the SRSs evolve according to the SRM. For any finite k, the set $\mathcal{H}_k$ consists of at most M^k p.s.d. matrices. It has been shown [9] that the SRS completely characterize the value function $J^*_{\infty,uc}(z)$.

Theorem 1 ([9]): For each $k \in \mathbb{Z}^+$, the k-horizon value function of the unconstrained DSLQR problem is given by:

$$J^*_{k,uc}(z) = \min_{P\in\mathcal{H}_k} z^T P z, \quad \forall z \in \mathbb{R}^n. \tag{9}$$

B. Efficient Computation of the SRSs

For nontrivial DSLSs with $M > 1$, the cardinality of $\mathcal{H}_k$ grows exponentially with k. However, usually not all the matrices in $\mathcal{H}_k$ contribute to the overall minimum of (9). The main idea therefore is to at each iteration k remove "redundant" matrices from the sets $\mathcal{H}_k$ [9]. Additionally, a numerical relaxation parameter ϵ can be introduced to further reduce the size of the associated SRS, possibly resulting in a suboptimal solution. This motivates the following definitions.

Definition 4 (ϵ-redundancy): For any constant $\epsilon \geq 0$, a matrix $\hat{P} \in \mathcal{H}_k$ is called *ϵ-redundant with respect to $\mathcal{H}_k$* if

$$\min_{P \in \mathcal{H}_k \setminus \hat{P}} z^T P z \leq \min_{P \in \mathcal{H}_k} z^T (P + \epsilon I_n) z, \quad \text{for any } z \in \mathbb{R}^n. \tag{10}$$

Definition 5 (ϵ-equivalent subset):
For any constant $\epsilon \geq 0$, the set $\mathcal{H}_k^\epsilon$ is called an *ϵ-equivalent subset* of $\mathcal{H}_k$ if $\mathcal{H}_k^\epsilon \subseteq \mathcal{H}_k$ and for all $z \in \mathbb{R}^n$,

$$\min_{P \in \mathcal{H}_k} z^T P z \leq \min_{P \in \mathcal{H}_k^\epsilon} z^T P z \leq \min_{P \in \mathcal{H}_k} z^T (P + \epsilon I_n) z. \tag{11}$$

Removing the ϵ-redundant matrices from $\mathcal{H}_k$ will result in an ϵ-equivalent subset $\mathcal{H}_k^\epsilon$. If we replace $\mathcal{H}_k$ with $\mathcal{H}_k^\epsilon$ in equation (9), Definition 5 guarantees that the corresponding value function will deviate from the original one by a term at most equal to $\epsilon \|z\|^2$. To simplify computation, we shall remove as many ϵ-redundant matrices as possible. A convex condition is derived in ([6], Lemma 1) to test whether a matrix in $\mathcal{H}_k$ is ϵ-redundant or not. By removing the ϵ-redundant matrices after each iteration, a relaxed version of the SRSs can be obtained iteratively as:

$$\mathcal{H}_0^\epsilon = \mathcal{H}_0, \quad \text{and} \quad \mathcal{H}_{k+1}^\epsilon = ES_\epsilon\left(\rho_{\mathbb{M}}(\mathcal{H}_k^\epsilon)\right), \tag{12}$$

where $ES_\epsilon(\mathcal{H})$ denotes an algorithm that computes the ϵ-equivalent subset of $\mathcal{H}$ by implementing the test of the condition in [6]. Similar to (9), one can define the *approximate value function* based on the relaxed SRSs as

$$J_{k,uc}^\epsilon(z) = \min_{P \in \mathcal{H}_k^\epsilon} z^T P z, \quad \forall z \in \mathbb{R}^n. \tag{13}$$

Since the set $\mathcal{H}_k^\epsilon$ computed through (12) typically contains much fewer matrices than $\mathcal{H}_k$, the approximate value function $J_{k,uc}^\epsilon$ is usually much easier to compute than $J_{k,uc}^*$. Moreover, by choosing ϵ sufficiently small and k sufficiently large, the function $J_{k,uc}^\epsilon$ can be made arbitrarily close to the infinite-horizon value function $J_{\infty,uc}^*$ [6].

Theorem 2: Under (A1), there always exist constants $\beta < \infty$, $\eta < \infty$ and $\gamma < 1$, all independent of k and ϵ, such that

(i) $\lambda_Q^- \|z\|^2 \leq J_{k,uc}^\epsilon(z) \leq \beta \|z\|^2, \ \forall z \in \mathbb{R}^n, k \in \mathbb{Z}^+$;

(ii) $|J_{k,uc}^\epsilon(z) - J_{\infty,uc}^*(z)| \leq \eta \gamma^k \epsilon \, \|z\|^2, \ \forall z \in \mathbb{R}^n, k \in \mathbb{Z}^+$,

where $\lambda_Q^- = \min_{i \in \mathbb{M}} \{\lambda_{\min}(Q_i)\}$.

C. Stationary Suboptimal Policy

Denote by $\xi_{k,uc}^\epsilon$ the hybrid-control law generated by the approximate value function $J_{k,uc}^\epsilon$, namely,

$$\begin{aligned} \xi_{k,uc}^\epsilon(z) &= \left(\mu_{k,uc}^\epsilon(z), \nu_{k,uc}^\epsilon(z)\right) \\ &\triangleq \arg\inf_{u \in \mathbb{R}^p, v \in \mathbb{M}} \left\{L(z,u,v) + J_{k,uc}^\epsilon(A_v z + B_v u)\right\}. \end{aligned} \tag{14}$$

Applying this law at each time step yields a stationary policy $\pi_{\infty,uc}^{\epsilon,k} = \{\xi_{k,uc}^\epsilon, \xi_{k,uc}^\epsilon, \ldots\}$. The particular structure (13) of $J_{k,uc}^\epsilon$ enables an analytical characterization of $\xi_{k,uc}^\epsilon$.

Lemma 1: The hybrid-control law $\xi_{k,uc}^\epsilon$ is given by:

$$\begin{aligned} \xi_{k,uc}^\epsilon(z) &\triangleq \left(\mu_{k,uc}^\epsilon(z), \nu_{k,uc}^\epsilon(z)\right) \\ &= \left(-K_{i_k^\epsilon(z)}\left(P_k^\epsilon(z)\right) z \, , \, i_k^\epsilon(z)\right), \\ \text{with} \quad \left(P_k^\epsilon(z), i_k^\epsilon(z)\right) &= \arg\min_{P \in \mathcal{H}_k^\epsilon, i \in \mathbb{M}} z^T \rho_i(P) z, \end{aligned} \tag{15}$$

where $K_\cdot(\cdot)$ is the Kalman gain defined by:

$$K_i(P) \triangleq (R_i + B_i^T P B_i)^{-1} B_i^T P A_i, \ i \in \mathbb{M}, P \in \mathcal{A}. \tag{16}$$

By Theorem 2, we know that $J_{k,uc}^\epsilon$ is a good approximation of $J_{\infty,uc}^\epsilon$ for sufficiently large k and small ϵ. Intuitively, the performance of the policy $\pi_{\infty,uc}^{\epsilon,k}$ generated by $J_{k,uc}^\epsilon$ should also be close to the optimal one. This can indeed be guaranteed if $\pi_{\infty,uc}^{\epsilon,k}$ is exponentially stabilizing.

Theorem 3 ([10]): Suppose for each $P \in \mathcal{H}_k^\epsilon$, there exist nonnegative constants $\alpha_j, j = 1, \ldots, j^*$ and κ_1 such that

$$\sum_{j=1}^{j^*} \alpha_j = 1 \text{ and } P \succ \sum_{j=1}^{j^*} \alpha_j \left(\hat{P}^{(j)} + (\kappa_1 - \kappa_*) I_n\right) \tag{17}$$

where $\{\hat{P}^{(j)}\}_{j=1}^{j^*}$ is an enumeration of the set $\rho_{\mathbb{M}}(\mathcal{H}_k^\epsilon)$, and $\kappa_* = \min_{i \in \mathbb{M}, P \in \mathcal{H}_k^\epsilon} \lambda_{\min}\left(K_i(P)^T R_i K_i(P) + Q_i\right)$, with $K_i(P)$ being defined in (16). Then the closed-loop trajectory x_t driven by $\pi_{\infty,uc}^{\epsilon,k}$ is exponentially stable with

$$\|x_t\|^2 \leq \tfrac{\beta}{\lambda_Q^-} \left(\tfrac{1}{1+\kappa_1/\lambda_Q^-}\right)^t \|z\|^2, \tag{18}$$

where β and λ_Q^- are constants defined in Theorem 2.

Checking condition (17) can be formulated as a LMI feasibility problem and thus may be verified efficiently.

Theorem 4: If $\mathcal{H}_k^\epsilon$ satisfies (17) in Theorem 3, then the cost associated with the policy $\pi_{\infty,uc}^{\epsilon,k}$ is bounded above by:

$$J_\infty(z, \pi_{\infty,uc}^{\epsilon,k}(z)) \leq J_{\infty,uc}^*(z) + \eta \gamma^k \epsilon \left(1 + \tfrac{\beta(\kappa + \lambda_Q^-)}{\kappa \cdot \lambda_Q^-}\right) \|z\|^2,$$

where β, λ_Q^-, η, γ and κ_1 are the constants defined in Theorem 2 and Theorem 3.

Theorem 4 indicates that by further increasing k and reducing ϵ, we can make the performance of this stabilizing policy arbitrarily close to the optimal one. These results allow us to construct a suboptimal policy in a systematical way as described in Algorithm 1, which returns a relaxed SRS that characterizes a suboptimal policy independent of the initial state z of the system.

Both the relaxation algorithm ES_ϵ and the algorithm for checking (17) involve only simple convex optimization programs. Additionally, although the size of the SRSs grows exponentially fast, experience shows that due to the numerical relaxation, the size of the relaxed SRSs $\mathcal{H}_k^\epsilon$ usually grows slowly and saturates at a small number, even in high-dimensional state space (see [9] for details). Therefore, Algorithm 1 can be carried out rather efficiently.

For the remainder of this paper, we shall denote by $\mathcal{H}_k^\epsilon$ the relaxed SRS returned by Algorithm 1, whose corresponding policy $\pi_{\infty,uc}^{\epsilon,k}$ is exponentially stabilizing with a upper bound on the performance given in Theorem 4.

Algorithm 1 (Unconstrained Suboptimal Policy)

Input: ϵ, $\epsilon_{\min}$ and $k_{\max}$

Set $\mathcal{H}_0^\epsilon = \{0\}$.
while $\epsilon > \epsilon_{\min}$ **do**
  for $k = 1$ to $k_{\max}$ **do**
    $\mathcal{H}_k^\epsilon = ES_\epsilon(\rho_{\mathbb{M}}(\mathcal{H}_k^\epsilon))$
    if $\mathcal{H}_k^\epsilon$ satisfies the condition of Theorem 3 **then**
      stop and return $\mathcal{H}_k^\epsilon$ (which characterizes $\pi_{\infty,uc}^{\epsilon,k}$)
    end if
  end for
  $\epsilon = \epsilon/2$
end while

IV. Safe Set and Its Computation

A. The Safe Set

We define the safe set of an unconstrained, infinite-horizon, hybrid-control policy as the set of initial states for which the closed-loop system driven by this policy satisfies the constraints (2) for all $t \geq 0$.

Definition 6 (Safe Set): For an arbitrary infinite-horizon policy $\pi_\infty = \{(\mu_t, \nu_t)\}_{t\in\mathbb{Z}^+}$, the *safe set* $\mathcal{X}(\pi_\infty)$ is:

$$\mathcal{X}(\pi_\infty) = \Big\{x_0 \in \mathbb{R}^n \,\Big|\, x_t \in \mathcal{X},\ \mu_t(x_t) \in \mathcal{U}, \\ x_{t+1} = A_{\nu_t(x_t)}x_t + B_{\nu_t(x_t)}\mu_t(x_t),\ \forall t \in \mathbb{Z}^+\Big\}. \tag{19}$$

From the above definition, given a policy π, if $x_{t_0} \in \mathcal{X}(\pi)$ for some $t_0 \in \mathbb{Z}^+$, then $x_t \in \mathcal{X}(\pi)$ for all $t \geq t_0$. In other words, the safe set is the maximal positive invariant set for the closed-loop system subject to constraints (2), i.e. the *maximal output admissible set* [11]. In the following, denote $\mathcal{X}_{\infty,uc} = \mathcal{X}(\pi_{\infty,uc}^{\epsilon,k})$. Notice that a closed-loop trajectory starting from any point in any arbitrarily shaped subset $\underline{\mathcal{X}}_{\infty,uc} \subset \mathcal{X}_{\infty,uc}$ will stay inside $\mathcal{X}_{\infty,uc}$ for all time and thus never violate the constraints. Since an exact characterization of $\mathcal{X}_{\infty,uc}$ is very hard to obtain, the rest of this subsection is devoted to the computation of a subset $\underline{\mathcal{X}}_{\infty,uc} \subset \mathcal{X}_{\infty,uc}$.

B. Analytical Characterization of the Safe Set

Theorem 5: Under (A1), there exists a constant $r^* > 0$ such that the set

$$\underline{\mathcal{X}}_{\infty,uc} = \{z \in \mathbb{R}^n \mid ||z|| \leq r^*\} \tag{20}$$

is a subset of the safe set $\mathcal{X}_{\infty,uc}$.

Proof: Let $r_0 = \max\{r : \|z\| \leq r \Rightarrow z \in \mathcal{X}\}$. Since $0 \in \text{int}(\mathcal{X})$, we have $r_0 > 0$. By Theorem 3, $\pi_{\infty,uc}^{\epsilon,k}$ is exponentially stabilizing and thus $||x_t|| \leq c_1 ||x_0||,\ \forall t \geq 0$, for some finite positive constant c_1. From (15), we know $||u_t|| = ||K_i(P)x_t|| \leq c_2 ||x_t||,\ \forall\, t \in \mathbb{Z}^+$, where $c_2 = \max_{i\in\mathbb{M}, P\in\mathcal{H}_k^\epsilon} \{\|(K_i(P))\|\}$. Let $\underline{\mathcal{X}}_{\infty,uc}$ be the Euclidean ball $B(r^*)$, with $r^* = \min\left\{\frac{r_0}{c_1}, \frac{r_0}{c_1 c_2}\right\}$, centered at the origin. The values of r_0, c_1 and c_2 are all finite and thus is r^*. It can be easily seen that for any initial state in $\underline{\mathcal{X}}_{\infty,uc}$, the closed-loop trajectory and the corresponding continuous-control sequence will always satisfy constraints (2). ■

Following the above proof, the characterization of the safe subset $\underline{\mathcal{X}}_{\infty,uc}$ requires only the estimation of the three constants r_0, c_1 and c_2. Therefore, the safe subset $\underline{\mathcal{X}}_{\infty,uc}$ in (20) can be easily computed in a state space of arbitrary dimension. However, this approach may be too conservative, resulting in a set $\underline{\mathcal{X}}_{\infty,uc}$ much smaller than the actual safe set $\mathcal{X}_{\infty,uc}$. In the following, we discuss a computational approach which can be used to under-approximate $\mathcal{X}_{\infty,uc}$, especially in lower-dimensional state spaces.

C. Computational Approach via Invariant Set

The most straightforward approach to obtain a safe set $\mathcal{X}_{\infty,uc}$ is to compute the positive invariant set in Definition 6. The computation of invariant sets can be reframed as a reachability problem and is reminiscent of the seminal work in [12], [13]. Different approaches have been developed in the literature to compute reachable sets for dynamical systems, such as polytopic or zonotopic methods, ellipsoidal methods, level-set methods and others [14]. In our case, the main issue preventing the direct use of those computational tools is the implicit form of the control law (15), as its associated *decision regions* (regions in the state space which yield the same minimizing pair (P_k^ϵ, i_k^ϵ) of (15)), are non-convex subsets of second-order cones in general [9].

One immediate approach for computing an invariant set is through gridding the state space, as implemented in Algorithm 2 (for the proposed algorithm, $\mathcal{X}$ is assumed to be bounded). Let $G_{\overline{\mathcal{X}}}$ be the set of all points that constitute a uniform grid with step size δ_{grid} over the smallest hyper-rectangle $\overline{\mathcal{X}}$ in the state space that contains the constraint polyhedron $\mathcal{X}$. Let $\mathcal{G}$ describe the region in the state space covered by the gridpoints g_i in $G_{\overline{\mathcal{X}}}$. A mapped state z is regarded as contained in $\mathcal{G}$, if $\min_{g\in G_{\overline{\mathcal{X}}}} \|z - g\|_\infty \leq \frac{\delta_{grid}}{2}$.

Algorithm 2 (Grid-based approx. computation of $\mathcal{X}_{\infty,uc}$)

Input: $\xi_{k,uc}^\epsilon = \{(\mu_{k,uc}^\epsilon, \nu_{k,uc}^\epsilon)\}$, $\mathcal{X}$, $\mathcal{U}$, $G_{\overline{\mathcal{X}}}$

1: set $G_0 = \left\{g_i \in G_{\overline{\mathcal{X}}} \mid g_i \in \mathcal{X},\ \mu_{k,uc}^\epsilon(g_i) \in \mathcal{U}\right\}$
2: $G_{k+1} = \left\{g_i \in G_k \middle| A_{\nu_{k,uc}^\epsilon(g_i)} g_i + B_{\nu_{k,uc}^\epsilon(g_i)} \mu_{k,uc}^\epsilon(g_i) \in \mathcal{G}_k\right\}$
3: **if** $G_{k+1} = G_k$ **then**
4:   **return** $G_{\mathcal{X}_{\infty,uc}} = G_{k+1}$
5: **else**
6:   ***GOTO 2***
7: **end if**

Algorithm 2 was implemented in `MATLAB` and tested for state dimensions $n \leq 4$. In principle, gridding also works in higher dimensions. However, the computational complexity grows exponentially with the state dimension. This "curse of dimensionality" prohibits dense gridding for higher dimensional problems.

V. Solution to a General DCSLQR Problem

The goal of this section is to solve a general DCSLQR problem, i.e. to find an infinite-horizon hybrid-control sequence for a given initial state z to achieve at least suboptimal performance with respect to the cost function $J_\infty^*(z)$.

A. Stabilizable Set

The DCSLQR problem is meaningful only when the given initial state z results in a finite cost $J_\infty^*(z)$. To characterize the set of such initial states, we introduce the following:

Definition 7 (Stabilizable Set): The set defined by

$$\mathcal{S}_\infty = \big\{z \in \mathbb{R}^n \,\big|\, \exists\, \psi_\infty = \{(u_t, v_t)\}_{t\in\mathbb{Z}^+} \text{ such that } x_t \in \mathcal{X},\ u_t \in \mathcal{U} \text{ and } x_t \to 0 \text{ exponentially fast}\big\} \tag{21}$$

is called the *stabilizable set* of system (1) subject to constraints (2), where x_t is the closed-loop trajectory driven by ψ_∞ with initial state $x_0 = z$.

For constrained LQR of linear systems ($M = 1$), it is possible to compute the stabilizable set $\mathcal{S}_\infty$ for compact sets $\mathcal{X}$ and $\mathcal{U}$ [3]. This is achieved by combining multiparametric quadratic programming [2] with reachability analysis. However, the algorithmic approach used by the authors in [3] can not be easily used for DSLS ($M > 1$). See [15] for further details on this problem.

B. DCSLQR formulation as an MIQP

For a given initial state in $\mathcal{S}_\infty$, our strategy in solving the DCSLQR problem is to first drive the system state into the safe subset $\underline{\mathcal{X}}_{\infty,uc}$, and then use the suboptimal infinite-horizon policy $\pi^{\epsilon,k}_{\infty,uc}$ to further regulate the state towards the origin. To this end, we introduce the following *constrained finite-time optimal hybrid control (CFTOHC)* problem:

$$\begin{cases} J_N^*(z;\phi) = \min\limits_{(u_t,v_t)} \left\{ \phi(x_N) + \sum\limits_{t=0}^{N-1} L(x_t,u_t,v_t) \right\}, \\ \textbf{s.t } (1) \text{ and } (2) \text{ with } x_0 = z, \end{cases} \tag{22}$$

with the terminal cost function $\phi : \mathcal{X} \to \mathbb{R}^+$. Denote by $x_{N|0}$ the state at time $t = N$ when system (1) is controlled by the solution of (22). The reason for introducing the above optimization problem is that with a properly chosen terminal cost function $\phi(\cdot)$, the optimal cost J_N^* will coincide with the value function J_∞^* of the DCSLQR problem.

Theorem 6: $J_N^*(z;\phi) = J_\infty^*(z)$ if, for all $z \in \mathcal{X}$:

(i) $\phi(z) = J_\infty^*(z)$ *or*

(ii) $\phi(z) = J_{\infty,uc}^*(z)$ and $x_{N|0} \in \underline{\mathcal{X}}_{\infty,uc}$.

Theorem 6 indicates in particular that solving Problem (5) is equivalent to solving Problem (22) when $\phi = J_{\infty,uc}^*$ and $x_{N|0} \in \underline{\mathcal{X}}_{\infty,uc}$. By Theorem 2, the function $J_{\infty,uc}^*$ can be accurately approximated by $J_{k,uc}^\epsilon$ for large k and small ϵ. Thus, $J_{k,uc}^\epsilon$ serves us as a local CLF inside $\underline{\mathcal{X}}_{\infty,uc}$. With $\phi(z) = J_{k,uc}^\epsilon(z) = \min_{P\in\mathcal{H}_k^\epsilon} z^T P z$, Problem (22) becomes:

$$\begin{cases} J_N^*(z; J_{k,uc}^\epsilon) \\ = \min\limits_{P\in\mathcal{H}_k^\epsilon} \left\{ \min\limits_{(u_t,v_t)} \left[x_N^T P x_N + \sum\limits_{t=0}^{N-1} L(x_t,u_t,v_t) \right] \right\}, \\ \textbf{s.t } (1) \text{ and } (2) \text{ with } x_0 = z. \end{cases} \tag{23}$$

The above formulation is obtained by first substituting $\phi(z) = \min_{P\in\mathcal{H}_k^\epsilon} z^T P z$ into (22) and then changing the order of the two minimizations. The change on the order of the minimizations will not affect the solution because there are only finitely many matrices in $\mathcal{H}_k^\epsilon$. By Theorem 6, $J_N^*(z; J_{k,uc}^\epsilon)$ will be close to $J_\infty^*(z)$ if the controlled terminal state $x_{N|0}$ is in $\underline{\mathcal{X}}_{\infty,uc}$. This can be always guaranteed if N is chosen sufficiently large.

Theorem 7: For every initial condition $x_0 = z \in \mathcal{S}_\infty$, there exists a finite $\hat{N}(z)$ such that for all $N \geq \hat{N}(z)$, the terminal state $x_{N|0}$ of the closed-loop system controlled by the solution of Problem (23) resides inside $\underline{\mathcal{X}}_{\infty,uc}$.

For the case of linear systems ($M{=}1$) and compact sets $\mathcal{X}$ and $\mathcal{U}$, it is possible to compute $\hat{N}$ for all $z \in \mathcal{S}_\infty$ beforehand [3]. However, the issues with the computation of the stabilizable set mentioned in Section V-A also prohibit this for SLS. In the following section, we therefore employ a straightforward approach inspired by [4], [5], which minimizes the computational complexity of (22).

A main contribution of this work is that we are able to characterize the hybrid-control policy $\pi^{\epsilon,k}_{\infty,uc}$ by the set of p.s.d. matrices $\mathcal{H}_k^\epsilon$. This allows us to cast problem (23) as one single augmented Mixed-Integer Quadratic Program (MIQP) for a given initial state z and prediction horizon N. The obtained MIQP can then be solved efficiently using available optimization software.

C. Overall Algorithm for DCSLQR

Theorem 7 guarantees that as N increases, the controlled terminal state $x_{N|0}$ associated with (23) eventually enters the safe subset $\underline{\mathcal{X}}_{\infty,uc}$. In this case, a suboptimal infinite-horizon control sequence for Problem (5) is given by

$$\psi_\infty = \{(\hat{u}_0, \hat{v}_0), \ldots, (\hat{u}_{N-1}, \hat{v}_{N-1}), \pi^{\epsilon,k}_{\infty,uc}(x_{N|0})\}, \tag{24}$$

where $\{(\hat{u}_t, \hat{v}_t)\}_{0\le t<N}$ denotes the solution to the optimization problem (23) and $\pi^{\epsilon,k}_{\infty,uc}(x_{N|0})$ denotes the infinite-horizon hybrid-control sequence generated by the policy $\pi^{\epsilon,k}_{\infty,uc}$ with initial state $x_{N|0}$. A general procedure for solving Problem (5) with initial condition $x_0 = z \in \underline{\mathcal{S}}_\infty$ is summarized in Algorithm 3. The returned control sequence is guaranteed to be suboptimal in the sense that by choosing k sufficiently large and ϵ sufficiently small, its performance can be made arbitrarily close to the optimal one.

Algorithm 3 (Solution of DCSLQR Problem (5))

Input: $x_0 = z \in \underline{\mathcal{S}}_\infty$, method for solving MIQPs

1: Compute $\pi^{\epsilon,k}_{\infty,uc}$ and relaxed SRS $\mathcal{H}_k^\epsilon$ using Algorithm 1
2: Compute $\underline{\mathcal{X}}_{\infty,uc}$ using Theorem (5) or Algorithm 2
3: Set $N = 1$
4: Solve problem (23) with time horizon N
5: **if** $x_{N|0} \in \underline{\mathcal{X}}_{\infty,uc}$ **then**
6: Stop and return the control sequence as defined in (24)
7: **else**
8: Set N=N+1 and go to step 4
9: **end if**

Theorem 8: For any $\delta > 0$ and $z \in \underline{\mathcal{S}}_\infty$, there exists an $k < \infty$ and $\epsilon > 0$ such that the control sequence ψ_∞ returned by Algorithm 3 satisfies

$$J_\infty(z, \psi_\infty) \leq J_\infty^*(z) + \delta.$$

VI. Numerical Example

Consider the following DSLS with two modes:

$$A_1 = \begin{bmatrix} 2 & 0 \\ 0 & 2 \end{bmatrix}, A_2 = \begin{bmatrix} 1.5 & 1 \\ 0 & 1.5 \end{bmatrix}, B_1 = \begin{bmatrix} 1 \\ 2 \end{bmatrix}, B_2 = \begin{bmatrix} 1 \\ 0 \end{bmatrix},$$

$$Q_i = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, R_i = 1, \qquad i = 1, 2,$$

subject to the constraints:

$$\mathcal{X}: \begin{bmatrix} -0.95 & 0.3 \\ 0 & -1 \\ 0.95 & -0.3 \\ 0 & 1 \end{bmatrix} x(t) \leq \begin{bmatrix} 0.85 \\ 1.25 \\ 0.85 \\ 1.25 \end{bmatrix}, \ \mathcal{U}: \begin{bmatrix} 1 \\ -1 \end{bmatrix} u(t) \leq \begin{bmatrix} 0.75 \\ 0.75 \end{bmatrix}.$$

Notice that none of the two subsystems is stabilizable (see Remark 1). The DSLS can therefore only be stabilized by appropriate switching between the two subsystems. The computation of the set $\mathcal{H}_k^\epsilon$ was performed for a time horizon of $k=30$ steps and a numerical relaxation parameter $\epsilon = 10^{-3}$, which led to a cardinality of $|\mathcal{H}_k^\epsilon| = 16$, as opposed to using $2^{30} \approx 10^9$ matrices for characterizing the cost function. It is obvious that, because of the exponential growth of the number of elements in the SRSs, obtaining $\mathcal{H}_k$ without computing equivalent subsets would be computationally prohibitive even for the simple problem at hand.

For an initial condition $x_0 = [0 \ \ 1]^T$, the optimal discrete modes for $0 \leq t \leq 7$ are $\{v_{opt,uc}\} = \{1, 2, 2, 1, 2, 1, 2, 1\}$ and $\{v_{opt}\} = \{1, 2, 1, 2, 1, 2, 1, 2\}$, respectively. The alternating discrete switching sequence that occurs in both the unconstrained and constrained case validates the previous claim that switching between the discrete modes is the only possible control strategy that can stabilize this specific hybrid system. The computation time (excluding the computation of $\mathcal{H}_k^\epsilon$) for this problem on a 3 GHz Intel Core2 CPU was 156 ms using the `CPLEX` solver, where we used `YALMIP` [16] to conveniently parse the optimization problem.

An approximation of the maximal positive invariant set $\mathcal{X}_{\infty,uc}$ for a gridpoint distance of $\delta_{grid} = 10^{-3}$ and the optimal state trajectories are depicted in Figure 1, the associated optimal continous control actions in Figure 2. As can be seen from Figure 1, the safe set $\mathcal{X}_{\infty,uc}$ in this example is rather large and takes up a substantial part of the feasible set $\mathcal{X}$. Consequently, the value of $N = 3$ is rather small in this case. State trajectory and control inputs clearly satisfy the constraints over the whole simulation horizon.

Acknowledgments

Thanks to Francesco Borrelli, Jianghai Hu, and Johan Löfberg, for suggestions and comments on parts of this work.

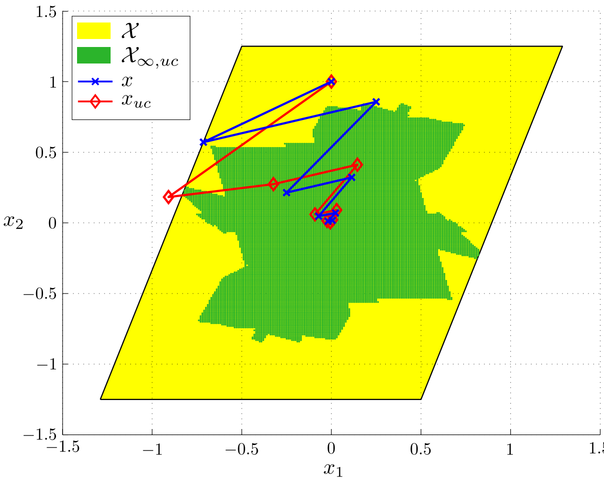

Fig. 1. Feasible region $\mathcal{X}$, gridding-based approximation of the safe set $\mathcal{X}_{\infty,uc}$ and closed-loop state trajectories (constrained and unconstrained)

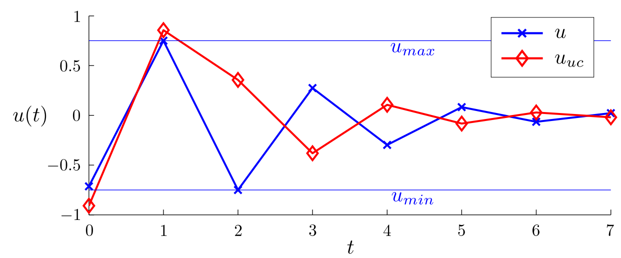

Fig. 2. Continuous optimal control inputs (constrained and unconstrained)

References

[1] E. Sontag, *Mathematical Control Theory: Deterministic Finite Dimensional Systems*, 2nd ed. Springer Verlag, 1998.

[2] A. Bemporad, M. Morari, V. Dua, and E. Pistikopoulos, "The Explicit Linear Quadratic Regulator for Constrained Systems," *Automatica*, vol. 38, no. 1, pp. 3–20, Jan 2002.

[3] P. Grieder, F. Borrelli, F. Torrisi, and M. Morari, "Computation of the Constrained Infinite Time Linear Quadratic Regulator," *Automatica*, vol. 40, no. 4, pp. 701–708, Apr 2004.

[4] P. O. M. Scokaert and J. B. Rawlings, "Constrained Linear Quadratic Regulation," *IEEE Transactions on Automatic Control*, vol. 43, no. 8, pp. 1163–1169, Aug 1998.

[5] M. Sznaier and M. J. Damborg, "Suboptimal Control of Linear Systems with State and Control Inequality Constraints," in *Proc. 26th Conf. on Decision and Control*, Dec 1987, pp. 761–762.

[6] W. Zhang, A. Abate, and J. Hu, "Efficient Suboptimal Solutions of Switched LQR Problems," in *Proceedings of the 28th American Control Conference*, St. Louis, MO, Jun 2009.

[7] B. Lincoln and B. Bernhardsson, "LQR optimization of linear system switching," *IEEE Transactions on Automatic Control*, vol. 47, no. 10, pp. 1701–1705, Oct 2002.

[8] W. Zhang, A. Abate, J. Hu, and M. P. Vitus, "Exponential Stabilization of Discrete-Time Switched Linear Systems," *Automatica*, vol. 45, no. 11, pp. 2526–2536, Nov 2009.

[9] W. Zhang, J. Hu, and A. Abate, "On the value function of the discrete-time switched LQR problem," *IEEE Transactions on Automatic Control*, vol. 54, no. 11, pp. 2669–2674, Nov 2009.

[10] W. Zhang, "Controller synthesis for switched systems using approximate dynamic programming," Ph.D. dissertation, ECE Department, Purdue University, West Lafayette, IN, 2009.

[11] E. G. Gilbert and K. T. Tan, "Linear Systems with State and Control Constraints: The Theory and Application of Maximal Output Admissible Sets," *IEEE Transactions on Automatic Control*, vol. 36, no. 9, pp. 1008–1020, Sep 1991.

[12] D. Bertsekas, "Infinite-time reachability of state-space regions by using feedback control," *IEEE Transactions on Automatic Control*, vol. 17, no. 5, pp. 604–612, Oct 1972.

[13] J. Glover and F. Schweppe, "Control of linear dynamic systems with set constrained disturbances," *IEEE Transactions on Automatic Control*, vol. 16, no. 5, pp. 411–423, Oct 1971.

[14] A. A. Kurzhanskiy and P. Varaiya, "Computation of Reach Sets for Dynamical Systems," CHESS - University of California, Berkeley, Tech. Rep., Sep 2008.

[15] F. Borrelli, *Constrained Optimal Control of Linear and Hybrid Systems*, ser. Lecture Notes in Control and Information Sciences. Springer Verlag, 2003, vol. 290.

[16] J. Löfberg, "YALMIP: A Toolbox for Modeling and Optimization in MATLAB," in *Proceedings of the CACSD Conference*, Taipei, Taiwan, 2004.