Infinite-Horizon Switched LQR Problems in Discrete Time: A Suboptimal Algorithm With Performance Analysis

Wei Zhang, Jianghai Hu, and Alessandro Abate

***Abstract*— This paper studies the quadratic regulation problem for discrete-time switched linear systems (DSLQR problem) on an infinite time horizon. A general relaxation framework is developed to simplify the computation of the value iterations. Based on this framework, an efficient algorithm is developed to solve the infinite-horizon DSLQR problem with guaranteed closed-loop stability and suboptimal performance. Due to its stability and suboptimal performance guarantees, the proposed algorithm can be used as a general controller synthesis tool for switched linear systems.**

This paper studies an extension of the classical LQR problem to switched linear systems (SLS), which will be referred to as the Discrete-Time Switched LQR (DSLQR) Problem. The goal is to find both the continuous-control and switching-control strategies to minimize a quadratic cost functional over an infinite time horizon. The problem is expected to play a fundamental role in the study of switched and hybrid systems as the classical LQR problem does for linear systems.

In our earlier work [11], we have proved some analytical properties for the finite-horizon DSLQR problem. Due to the discrete nature of the switching control sequence, the exact solution to the DSLQR problem is NP hard even for a finite time horizon. The main contribution of this paper is the development of an efficient algorithm to solve the infinite-horizon DSLQR problem with guaranteed suboptimal performance. The algorithm is based on a relaxation framework that can yield efficient representations of the (approximate) value functions. The key idea is to use convex optimization to identify and remove matrices that are redundant in characterizing the value functions and the corresponding (sub)optimal strategies. This is in line with many existing methods on approximate dynamic programming (ADP) ([2], [7], [8]), which aim at simplifying the computations by finding compact representations of the value functions up to certain numerical relaxation errors. Central to most ADP approaches is the analysis of the evolution of the relaxation errors through value iterations. Most existing error analysis methods for ADP either require a discount factor strictly less than one [2] or assume the infinite-horizon value function and the running cost function jointly satisfy some technical conditions that are in general difficult to verify *a priori* [7].

In this paper, by taking advantage of the particular structure of the DSLQR problem, we develop a relaxation framework that enables error analysis for undiscounted cost functions under easy-to-check assumptions. Our analysis indicates that as long as the SLS is stabilizable, the closed-loop performance of the DSLQR solution can be made arbitrarily close to the optimal one by properly choosing the relaxation parameter. The distance-to-optimality error bound is also derived analytically in terms of the subsystem matrices; and the bound can be evaluated *a priori*. Moreover, the suboptimal DSLQR solution is guaranteed to be a stabilizing controller whenever the SLS is stabilizable. These stability and performance guarantees are of fundamental importance and are often not provided by other existing optimal control strategies for switched/hybrid systems [3], [7], especially on infinite control horizon.

A preliminary version of this work has appeared in an earlier conference paper [9], and has also been successfully applied to study stabilization of SLSs [10]. The main distinction of this paper as compared with [9] lies in the development of a *stationary* suboptimal infinite horizon policy rather than a receding-horizon type of policy proposed in [9]. In addition, the asymptotic stabilizability assumption adopted in this paper is much weaker than the one in [9] that assumes one of the subsystems is stabilizable. These advances are not of theoretical importance, but also dramatically simplifies the design of DSLQR controllers for a wider range of applications.

I. Problem Formulation

Let n, M and p be some positive integers, and let $\mathbb{Z}_+$ denote the set of nonnegative integers. We consider a discrete time switched linear system (SLS) given by:

$$x(t+1) = A_{v(t)}x(t) + B_{v(t)}u(t), \quad t \in \mathbb{Z}_+, \tag{1}$$

where $x(t) \in \mathbb{R}^n$ is the continuous state, $u(t) \in \mathbb{R}^p$ is the continuous control, $v(t) \in \mathbb{M} := \{1, \ldots, M\}$ is the discrete control that determines the discrete mode at time t. The sequence $\{(u(t), v(t))\}_{t=0}^{\infty}$ is called a *hybrid-control sequence*. The hybrid-control action $(u(t), v(t))$ at time t is determined through a state-feedback hybrid-control law, namely, a function $\xi \triangleq (\mu, \nu) : \mathbb{R}^n \to \mathbb{R}^p \times \mathbb{M}$ that maps the continuous state $x(t)$ to a hybrid-control action $\xi(x(t)) = (u(t), v(t))$. Here, $\mu : \mathbb{R}^n \to \mathbb{R}^p$ and $\nu : \mathbb{R}^n \to \mathbb{M}$ are called the *(state-feedback) continuous-control law* and the *switching-control law*, respectively. A sequence of hybrid-control laws $\{\xi_t\}_{t=0}^{\infty}$ constitutes an *infinite-horizon feedback policy* $\pi_\infty = \{\xi_0, \xi_1, \ldots\}$. The closed-loop dynamics driven by a feedback policy $\pi_\infty = \{(\mu_t, \nu_t)\}_{t \in \mathbb{Z}_+}$ is governed by

$$x(t+1) = A_{\nu_t(x(t))}x(t) + B_{\nu_t(x(t))}\mu_t(x(t)), \quad t \in \mathbb{Z}_+. \tag{2}$$

To be more specific, the closed-loop trajectory under π_∞ with initial state $z \in \mathbb{R}^n$ will be denoted by $x(\cdot; z, \pi_\infty)$. Let $L(x, u, v) = x^T Q_v x + u^T R_v u$, $\forall x \in \mathbb{R}^n, u \in \mathbb{R}^p, v \in \mathbb{M}$, be the running cost function, where $Q_v = Q_v^T \succ 0$ and

W. Zhang is with the Department of Electrical and Computer Engineering, The Ohio State University, Columbus OH 43210 Email: zhang@ece.osu.edu

J. Hu is with the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47906. Email: jianghai@purdue.edu

A. Abate is with the Delft Center for Systems and Control, Delft University of Technology, The Netherlands. Email: a.abate@tudelft.nl

$R_v = R_v^T \succ 0$, $v \in \mathbb{M}$, are strictly positive definite matrices of appropriate dimensions. Denote by $\lambda_{\min}(\cdot)$ and $\lambda_{\max}(\cdot)$ the smallest and largest eigenvalues of a symmetric matrix, respectively. Define $\lambda_Q^- = \min_{i\in\mathbb{M}}\{\lambda_{\min}(Q_i)\}$. The strict positive definiteness of $\{Q_i\}_{i\in\mathbb{M}}$ implies that $\lambda_Q^- > 0$. For a given initial state $x(0) = z \in \mathbb{R}^n$, the performance of a feedback policy π_∞ is measured through the following cost function:

$$J(z;\pi_\infty) = \sum_{t=0}^{\infty} L(x(t), \mu_t(x(t)), \nu_t(x(t))). \tag{3}$$

Clearly, the boundedness of $J(\cdot;\pi_\infty)$ requires the stabilizability of system (1). System (1) is called *exponentially stabilizable* if $\exists \pi_\infty, b \geq 1, a \in (0,1)$ such that $\|x(t;z,\pi_\infty)\|^2 \leq ba^t\|z\|^2,\ \forall t \in \mathbb{Z}^+, z \in \mathbb{R}^n$. Such a policy π_∞ is called an *exponentially stabilizing policy*. It has been proved in [5] that exponential stability and asymptotic stability are equivalent for SLSs. Hence, without loss of generality, the following assumption is adopted throughout this paper.

Assumption 1: System (1) is exponentially stabilizable.

The goal of this paper is to solve the following optimal control problem.

Problem 1 (DSLQR problem): Under Assumption 1, find an infinite-horizon policy π_∞ that solves: $V^*(z) = \inf_{\pi_\infty} J(z;\pi_\infty)$, $\forall z \in \mathbb{R}^n$.

Problem 1 is a natural extension of the classical LQR problem to SLSs and is thus called the (infinite-horizon) *Discrete-Time Switched LQR (DSLQR) Problem*. The resulting infimum cost V^* will be referred to as the *infinite-horizon* value function of the DSLQR problem.

II. Approximate Value Iteration and Error Propagation

A. Exact Value Iteration

One way to tackle Problem 1 is through dynamic programming. Denote by $\mathcal{G}_+ \triangleq \{g : \mathbb{R}^n \to \mathbb{R}_+ \cup \{+\infty\}\}$ the space of all nonnegative, extended-valued functions on $\mathbb{R}^n$. For an arbitrary control law $\xi = (\mu,\nu) : \mathbb{R}^n \to \mathbb{R}^p \times \mathbb{M}$, the operator $\mathcal{T}_\xi : \mathcal{G}_+ \to \mathcal{G}_+$ is defined by

$$\mathcal{T}_\xi[g](z) = L(z,\mu(z),\nu(z)) + g(A_{\nu(z)}z + B_{\nu(z)}\mu(z)),$$

for all $z \in \mathbb{R}^n$, $g \in \mathcal{G}_+$. Minimizing $\mathcal{T}_\xi$ over ξ yields the *one-stage value iteration operator* $\mathcal{T} : \mathcal{G}_+ \to \mathcal{G}_+$, i.e.,

$$\mathcal{T}[g](z) = \inf_{u\in\mathbb{R}^p, v\in\mathbb{M}} \big\{L(z,u,v) + g(A_v z + B_v u)\big\}.$$

Denote by $\mathcal{T}^k$ the composition of $\mathcal{T}$ with itself k times, i.e., $\mathcal{T}^k = \mathcal{T} \circ \mathcal{T}^{k-1}$ for all $k = 1,2,\ldots$. Let $V_0 \equiv 0$ and define $V_k := \mathcal{T}^k[V_0]$ as the k-horizon value function of Problem 1. An important consequence of Assumption 1 is the exponential convergence of V_k to V^*.

Theorem 1: Assumption 1 implies (i) $\exists \beta < \infty$ such that $\lambda_Q^-\|z\|^2 \leq V^*(z) \leq \beta\|z\|^2$, and (ii) $0 \leq V^*(z) - V_k(z) \leq \alpha_V \gamma_V^k \|z\|^2$, for all $k \in \mathbb{Z}_+$, where $\alpha_V = \frac{\beta^2 - (\lambda_Q^-)^2}{\lambda_Q^-}$ and $\gamma_V = \frac{1}{1+\lambda_Q^-/\beta}$. The constant β can be chosen, in particular, as in (13) in the Appendix.

Proof: See the Appendix. ■

The above theorem indicates that under Assumption 1, V_k is a good approximation of V^* for large k. It turns out that for DSLQR problems, V_k takes a simple analytical form [11]. To see this, let $\mathcal{A}$ be the set of all positive semidefinite (p.s.d.) matrices, and let $\mathcal{F}$ be the set of all finite subsets of $\mathcal{A}$, i.e., $\mathcal{F} := \{\mathcal{H} \subset \mathcal{A} : |\mathcal{H}| < \infty\}$, where $|\cdot|$ denotes the cardinality of a set. Furthermore, we denote by $\rho_i : \mathcal{A} \to \mathcal{A}$ the *Riccati mapping* associated with subsystem $i \in \mathbb{M}$, i.e.,

$$\rho_i(P) = Q_i + A_i^T P A_i - A_i^T P B_i (R_i + B_i^T P B_i)^{-1} B_i^T P A_i. \tag{4}$$

We call the mapping $\rho_{\mathbb{M}} : \mathcal{F} \to \mathcal{F}$ defined by: $\rho_{\mathbb{M}}(\mathcal{H}) = \{\rho_i(P) : i \in \mathbb{M} \text{ and } P \in \mathcal{H}\}$, $\forall \mathcal{H} \in \mathcal{F}$, the *Switched Riccati Mapping* (SRM) associated with Problem 1. The sets $\{\mathcal{H}_k\}_{k\in\mathbb{Z}_+}$ generated iteratively by $\mathcal{H}_{k+1} = \rho_{\mathbb{M}}(\mathcal{H}_k)$ with $\mathcal{H}_0 = \{0\}$ are called the *Switched Riccati Sets* (SRSs), and they characterize the finite-horizon value functions $\{V_k\}_{k\in\mathbb{Z}_+}$.

Theorem 2 ([11]): The k-horizon value function of the DSLQR problem is $V_k(z) = \min_{P\in\mathcal{H}_k} z^T P z$, $\forall k \in \mathbb{Z}_+$.

Remark 1: It is beneficial to view $\mathcal{H}_k$ as a representation of V_k in the space $\mathcal{F}$. From this perspective, $\rho_{\mathbb{M}}$ is essentially a representation of the value iteration operator $\mathcal{T}$ in $\mathcal{F}$.

B. Relaxed Value Iteration

As k increases, the number of matrices in $\mathcal{H}_k$ grows exponentially, making the exact characterization of V_k increasingly expensive. One way to alleviate this computational challenge is to remove matrices in $\mathcal{H}_k$ that are less important in terms of characterizing V_k. To formalize the idea, we introduce a few definitions. For any $\mathcal{H} \in \mathcal{F}$, define $V_{\mathcal{H}}(z) = \min_{P\in\mathcal{H}} z^T P z, \forall z \in \mathbb{R}^n$.

Definition 1 (ϵ-Redundancy): For any $\epsilon > 0$ and $\mathcal{H} \in \mathcal{F}$, a matrix $\bar{P} \in \mathcal{H}$ is called *ϵ-redundant* with respect to $\mathcal{H}$ if $V_{\mathcal{H}\setminus\bar{P}}(z) \leq V_{\mathcal{H}}(z) + \epsilon\|z\|^2$, for any $z \in \mathbb{R}^n$.

Definition 2 (ϵ-ES): For any $\epsilon > 0$ and $\mathcal{H} \in \mathcal{F}$, a subset $\mathcal{H}^\epsilon \subset \mathcal{H}$ is called an *ϵ-Equivalent-Subset (ϵ-ES)* of $\mathcal{H}$ if $V_{\mathcal{H}}(z) \leq V_{\mathcal{H}^\epsilon}(z) \leq V_{\mathcal{H}}(z) + \epsilon\|z\|^2$, for any $z \in \mathbb{R}^n$.

To simplify the computation of V_k at each step k, we shall prune out as many redundant matrices as possible from the corresponding set $\mathcal{H}_k$. However, checking whether a matrix in $\mathcal{H}_k$ is redundant or not is itself a challenging problem. Geometrically, any p.s.d. matrix $\bar{P}$ defines an ellipsoid (possibly degenerate) in $\mathbb{R}^n$: $\{x \in \mathbb{R}^n : x^T\bar{P}x \leq 1\}$. It can be verified that $\bar{P} \in \mathcal{H}_k$ is ϵ-redundant if and only if the ellipsoid corresponding to $P + \epsilon I_n$ is contained in the union of all the ellipsoids corresponding to the matrices in $\mathcal{H}_k \setminus \{\bar{P}\}$. Since the union of ellipsoids is usually not convex, there is no general way to efficiently verify this geometric condition or equivalently the condition given in Definition 1. Nevertheless, a sufficient condition for ϵ-redundancy can be easily derived.

Lemma 1: $\bar{P} \in \mathcal{H}$ is ϵ-redundant if there exist nonnegative constants $\alpha_1,\ldots,\alpha_{|\mathcal{H}|-1}$ such that $\sum_{j=1}^{|\mathcal{H}|-1}\alpha_j = 1$ and $\bar{P} + \epsilon I_n \succeq \sum_{j=1}^{|\mathcal{H}|-1}\alpha_j P^{(j)}$, where $\{P^{(j)}\}_{j=1}^{|\mathcal{H}|-1}$ is an enumeration of $\mathcal{H} \setminus \{\bar{P}\}$.

Proof: Under the condition in the Lemma, for any $z \in \mathbb{R}^n$, we have $z^T(\bar{P} + \epsilon I_n)z \geq \sum_{j=1}^{|\mathcal{H}|-1} z^T\alpha_j P^{(j)}z \geq$

Algorithm 1 $[ES_\epsilon(\mathcal{H})]$

Set $\mathcal{H}^\epsilon = \emptyset$.
for each $P \in \mathcal{H}$ **do**
  if P does NOT satisfy the condition in Lemma 1 with respect to $\mathcal{H}^\epsilon$ **then**
    $\mathcal{H}^\epsilon \leftarrow \mathcal{H}^\epsilon \cup \{P\}$
  end if
end for
Return $\mathcal{H}^\epsilon$

$z^T P^{(j_z)} z$, for some $j_z \in \{1, \dots, |\mathcal{H}|-1\}$. Thus, by definition, $\bar{P}$ is ϵ-redundant in $\mathcal{H}$. ∎

For given $\bar{P}$ and $\mathcal{H}$, the condition in Lemma 1 can be efficiently verified by solving a convex optimization problem.

Lemma 2: The condition in Lemma 1 holds if and only if the solution of the following convex optimization problem

$$\begin{cases} \max\limits_{\{\alpha_1,\dots,\alpha_{|\mathcal{H}|-1}\}} \sum_{j=1}^{|\mathcal{H}|-1} \alpha_j \\ \text{subject to: } \begin{cases} \sum_{j=1}^{|\mathcal{H}|-1} \alpha_j P^{(j)} \preceq \bar{P} + \epsilon I_n \\ \alpha_j \geq 0, j = 1, \dots, |\mathcal{H}|-1 \end{cases} \end{cases} \tag{5}$$

satisfies $\sum_{j=1}^{|\mathcal{H}|-1} \alpha_j \geq 1$.

Proof: Straightforward. ∎

The optimization problem in (5) can be easily solved using various convex optimization algorithms [4, Chapter 11]. Based on the above lemma, an efficient algorithm (Algorithm 1) is developed to compute an ϵ-ES for any given set $\mathcal{H} \in \mathcal{F}$. Qualitatively, the algorithm simply removes all the matrices that satisfy the condition of Lemma 1 and returns the set of remaining matrices. Denote by $ES_\epsilon(\mathcal{H})$ the ϵ-ES returned by Algorithm 1.

To reduce complexity, starting from $\mathcal{H}_0$, we can apply $ES_\epsilon(\cdot)$ after each SRM $\rho_\mathbb{M}$ to obtain an ϵ-ES with fewer matrices before further propagating the set with the next SRM. The SRM $\rho_\mathbb{M}$ followed by the pruning algorithm ES_ϵ can be viewed as a relaxed version of the original SRM.

Definition 3 (ϵ-relaxed SRSs): For $\epsilon > 0$, the composite mapping $ES_\epsilon \circ \rho_\mathbb{M} : \mathcal{F} \to \mathcal{F}$ is called the *ϵ-relaxed SRM* of system (1). The sets $\{\mathcal{H}_k^\epsilon\}_{k \in \mathbb{Z}_+}$ generated iteratively by:

$$\mathcal{H}_0^\epsilon = \{0\} \text{ and } \mathcal{H}_{k+1}^\epsilon = ES_\epsilon \circ \rho_\mathbb{M}(\mathcal{H}_k^\epsilon), \quad k \in \mathbb{Z}_+, \tag{6}$$

are called the *ϵ-relaxed SRSs* associated with Problem 1.

Just as the SRM $\rho_\mathbb{M}$ represents the value iteration operator $\mathcal{T}$ of the DSLQR problem in the space $\mathcal{F}$, the relaxed SRM $ES_\epsilon \circ \rho_\mathbb{M}$ represents the relaxed value iteration operator defined by $\mathcal{R}_\epsilon \circ \mathcal{T}$, where $\mathcal{R}_\epsilon : \mathcal{G}_+ \to \mathcal{G}_+$ will be referred to as the *relaxation operator with parameter* ϵ and is defined as $\mathcal{R}_\epsilon[V_\mathcal{H}](z) = V_{ES_\epsilon(\mathcal{H})}(z)$, $\forall z \in \mathbb{R}^n, \mathcal{H} \in \mathcal{F}$. Therefore, the relaxed value iteration is the exact value iteration followed by a relaxation step.

C. Approximate Value Function and Control Law

We introduce two important functions,

$$V_k^\epsilon = (\mathcal{R}_\epsilon \circ \mathcal{T})^k [V_0], \text{ and } \tilde{V}_{k+1}^\epsilon = \mathcal{T}[V_k^\epsilon], k \in \mathbb{Z}_+. \tag{7}$$

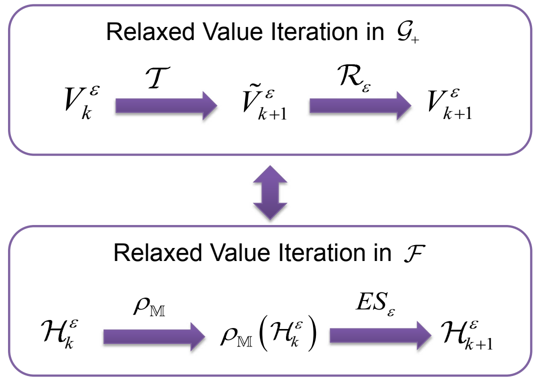

Fig. 1. Representations of the relaxed value iteration operator in $\mathcal{G}_+$ and $\mathcal{F}$.

The function V_k^ϵ is called the *k-horizon approximate value function* of Problem 1. The function $\tilde{V}_k^\epsilon$ is an auxiliary function that is useful in studying the properties of V_k^ϵ. The relationships between V_k^ϵ and $\tilde{V}_k^\epsilon$ and their connections to the relaxed SRSs are illustrated in Fig. 1. According to the lemma below, the error $V_k^\epsilon - V_k$ can be fully controlled by tuning the relaxation parameter ϵ.

Lemma 3: For any $z \in \mathbb{R}^n$ and $\epsilon \geq 0$, the functions defined in (7) satisfy:

1) $\tilde{V}_k^\epsilon(z) \leq V_k^\epsilon(z) \leq \tilde{V}_k^\epsilon(z) + \epsilon \|z\|^2$, for $k \geq 1$;
2) $V_k^\epsilon(z) \leq (1 + \epsilon/\lambda_Q^-) V_k(z)$, for $k \in \mathbb{Z}_+$, where $V_k(z)$ is the exact value function with no relaxation.

Proof: By Definition 2, we have $0 \leq \mathcal{R}_\epsilon[V_\mathcal{H}](z) - V_\mathcal{H}(z) \leq \epsilon \|z\|^2$, $\forall \mathcal{H} \in \mathcal{F}, z \in \mathbb{R}^n$. This together with the fact that $V_k^\epsilon = \mathcal{R}_\epsilon[\tilde{V}_k^\epsilon]$ implies part 1) of this lemma. The second part can be proved by induction. The result clearly holds for $k = 0$ because $V_0^\epsilon = V_0 \equiv 0$. Assuming it is true for some $k \in \mathbb{Z}_+$, we show that it also holds for $k+1$. For a fixed $z \in \mathbb{R}^n$, we have $\tilde{V}_{k+1}^\epsilon(z) = \mathcal{T}[V_k^\epsilon](z) \leq \inf_{u,v}\{L(z,u,v) + (1+\epsilon/\lambda_Q^-) V_k(A_v z + B_v u)\}$. Thus,

$$\begin{aligned} V_{k+1}^\epsilon(z) &= \mathcal{R}_\epsilon \circ \mathcal{T}[V_k^\epsilon](z) \leq \tilde{V}_{k+1}^\epsilon(z) + \epsilon \|z\|^2 \\ &\leq \inf_{u,v} \left\{ (1+\epsilon/\lambda_Q^-) \left[L(z,u,v) + V_k(A_v z + B_v u) \right] \right\} \\ &= (1+\epsilon/\lambda_Q^-) V_{k+1}(z), \end{aligned}$$

where the second step is due to the fact that $L(z,u,v) \geq \lambda_Q^- \|z\|^2$ for all $u \in \mathbb{R}^p$ and $v \in \mathbb{M}$. ∎

Denote by ξ_k^ϵ the hybrid-control law generated by V_k^ϵ, i.e.,

$$\begin{aligned} &\xi_k^\epsilon(z) = (\mu_k(z), \nu_k(z)) \\ &= \arg\min_{(u,v)} \{L(z,u,v) + V_k^\epsilon(A_v z + B_v u)\}, \forall z \in \mathbb{R}^n. \end{aligned} \tag{8}$$

The function V_k^ϵ and the corresponding control law ξ_k^ϵ can be characterized analytically using the relaxed SRS $\mathcal{H}_k^\epsilon$.

Theorem 3: For any $k \in \mathbb{Z}_+$, $z \in \mathbb{R}^n$ and $\epsilon \geq 0$, we have

$$V_k^\epsilon(z) = \min_{P \in \mathcal{H}_k^\epsilon} z^T P z, \quad \tilde{V}_{k+1}^\epsilon(z) = \min_{P \in \rho_\mathbb{M}(\mathcal{H}_k^\epsilon)} z^T P z,$$

$$\text{and } \xi_k^\epsilon(z) = \Big(-K_{i_k^\epsilon(z)}(P_k^\epsilon(z)) z, i_k^\epsilon(z) \Big)$$

where

$$\Big(P_k^\epsilon(z), i_k^\epsilon(z) \Big) = \arg\min_{P \in \mathcal{H}_k^\epsilon, i \in \mathbb{M}} z^T \rho_i(P) z, \tag{9}$$

and $K_i(P)$ is the Kalman gain for subsystem i corresponding to matrix P, i.e.,

$$K_i(P) \triangleq (R_i + B_i^T P B_i)^{-1} B_i^T P A_i. \tag{10}$$

Proof: The result follows easily by solving explicitly the quadratic optimization problem in (8) followed by a standard induction argument. ■

III. Suboptimal Policy With Performance Bound

According to a standard result of dynamic programming [1], the stationary policy $\pi_\infty^* = \{\xi^*, \xi^*, \ldots\}$ is an optimal solution to the infinite-horizon DSLQR problem, where ξ^* is the control law satisfying the Bellman equation: $\mathcal{T}_{\xi^*}[V^*] = V^*$. However, exact characterizations of V^* and ξ^* are often impossible for DSLQR problems. A natural solution is to use the control law ξ_k^ϵ defined in (8) with a sufficiently large k and small ϵ in place of ξ^* to construct a stationary policy $\pi_\infty^{\epsilon,k} = \{\xi_k^\epsilon, \xi_k^\epsilon, \ldots\}$. In the rest of this section, we shall first establish the conditions under which $\pi_\infty^{\epsilon,k}$ is exponentially stabilizing, and then derive a performance bound for such a stabilizing policy.

A. Stabilizing Condition with Efficient Test

A sufficient condition for $\pi_\infty^{\epsilon,k}$ to be exponentially stabilizing is that the corresponding closed-loop system admits an exponential Lyapunov function. We briefly reccall that a function $g \in \mathcal{G}_+$ is called an *Exponential Lyapunov Function (ELF)* of the closed-loop system (2) if there exist positive finite constants $\kappa_1, \kappa_2, \kappa_3$ such that

$$\begin{cases} \kappa_1 \|z\|^2 \le g(z) \le \kappa_2 \|z\|^2, \forall z \in \mathbb{R}^n \\ g(x(t)) - g(x(t+1)) \ge \kappa_3 \|x(t)\|^2, \forall t \in \mathbb{Z}_+, \end{cases}$$

where $x(t)$ is an arbitrary trajectory of system (2).

Theorem 4 (Lyapunov Theorem): If there exists a policy π_∞ under which system (2) has an ELF with parameters κ_1, κ_2 and κ_3, then the closed-loop trajectory $x(t; z, \pi_\infty)$ of system (2) under π_∞ is exponentially stable and satisfies

$$\|x(t; z, \pi_\infty)\|^2 \le \frac{\kappa_2}{\kappa_1} \left(\frac{1}{1 + \kappa_3/\kappa_2} \right)^t \|z\|^2, \forall t \in \mathbb{Z}_+, z \in \mathbb{R}^n.$$

The above theorem follows easily from standard Lyapunov theory [6, Thm 4.1] and its proof is omitted here. We now show that the approximate value function V_k^ϵ will be an ELF of system (2) under the policy $\pi_\infty^{\epsilon,k}$ for sufficiently large k and sufficiently small ϵ.

Theorem 5: Fix an arbitrary $z \in \mathbb{R}^n$, and let $\hat{x}(t) = x(t; z, \pi_\infty^{\epsilon,k})$. Under Assumption 1, there exist constants $\kappa_3 > 0$, $\hat{k} \in \mathbb{Z}_+$, $\hat{\epsilon} > 0$ such that for any $k \ge \hat{k}$, $\epsilon \le \hat{\epsilon}$,

1) V_k^ϵ is an ELF of system (2) satisfying

$$\lambda_Q^- \|z\|^2 \le V_k^\epsilon(z) \le \beta(1 + \epsilon/\lambda_Q^-)\|z\|^2 \tag{11a}$$
$$V_k^\epsilon(\hat{x}(t)) - V_k^\epsilon(\hat{x}(t+1)) \ge \kappa_3 \|\hat{x}(t)\|^2; \tag{11b}$$

2) the policy $\pi_\infty^{\epsilon,k}$ is exponentially stabilizing, namely, $\|\hat{x}(t)\|^2 \le \alpha_x \gamma_x^t \|z\|^2$, with $\alpha_x = \frac{\beta(1+\epsilon/\lambda_Q^-)}{\lambda_Q^-}$ and $\gamma_x = \frac{\beta(1+\epsilon/\lambda_Q^-)}{\beta(1+\epsilon/\lambda_Q^-)+\kappa_3}$.

Here $\beta \in (0, \infty)$ is the constant given in Theorem 1.

Proof: 1) Fix an arbitrary $z \in \mathbb{R}^n$. Inequality (11a) follows directly from Lemma 3 and part 1) of Theorem 1. To show inequality (11b), let $\hat{x}(t) := x(t; z, \pi_\infty^{\epsilon,k})$ and let $(\hat{u}(t), \hat{v}(t))$ be the corresponding hybrid-control sequence. By (7) and (8), we know that for all $t \in \mathbb{Z}_+$, $\tilde{V}_{k+1}^\epsilon(\hat{x}(t)) - V_k^\epsilon(\hat{x}(t+1)) = \mathcal{T}[V_k^\epsilon](\hat{x}(t)) - V_k^\epsilon(\hat{x}(t+1)) \ge \lambda_Q^- \|\hat{x}(t)\|^2$. Furthermore, by Lemma 3, Theorem 1 and the fact that $V_k(z) \ge \lambda_Q^- \|z\|^2$ for $k \ge 1$, we have, for $k \ge 1$,

$$\begin{aligned}
&\tilde{V}_{k+1}^\epsilon(\hat{x}(t)) \le V_{k+1}^\epsilon(\hat{x}(t)) \le (1 + \epsilon/\lambda_Q^-) V_{k+1}(\hat{x}(t)) \\
&\le (1+\epsilon/\lambda_Q^-) V^*(\hat{x}(t)) \le (1+\epsilon/\lambda_Q^-)(1+\alpha_V \gamma_V^k/\lambda_Q^-) V_k(\hat{x}(t)) \\
&\le (1+\epsilon/\lambda_Q^-)\left(1 + \frac{\alpha_V \gamma_V^k}{\lambda_Q^-}\right) V_k^\epsilon(\hat{x}(t)) \le V_k^\epsilon(\hat{x}(t)) + c_{k,\epsilon}\|\hat{x}(t)\|^2,
\end{aligned}$$

where $c_{k,\epsilon}$ is some constant that can be made arbitrarily small by choosing large k and small ϵ. Therefore, $V_k^\epsilon(\hat{x}(t)) - V_k^\epsilon(\hat{x}(t+1)) \ge \tilde{V}_{k+1}^\epsilon(\hat{x}(t)) - V_k^\epsilon(\hat{x}(t+1)) - c_{k,\epsilon}\|\hat{x}(t)\|^2 \ge (\lambda_Q^- - c_{k,\epsilon})\|\hat{x}(t)\|^2$, which implies (11b).

2) This follows directly from part 1) and Theorem 4. ■

This theorem indicates that as we increase k and reduce ϵ, the function V_k^ϵ eventually becomes an ELF of system (2) with an associated stabilizing policy $\pi_\infty^{\epsilon,k}$. To test whether V_k^ϵ is an ELF, one shall verify condition (11b). By taking advantage of the piecewise quadratic structure of V_k^ϵ, this verification process can be greatly simplified as follows.

Lemma 4: Inequality (11b) holds for some constant $\kappa_3 > 0$ if for each $P \in \mathcal{H}_k^\epsilon$ there exist nonnegative constants $\alpha_j, j = 1, \ldots, j^*$, such that

$$\sum_{j=1}^{j^*} \alpha_j = 1 \quad \text{and} \quad P \succeq \sum_{j=1}^{j^*} \alpha_j \left(\hat{P}^{(j)} + (\kappa_3 - \kappa_*) I_n \right), \tag{12}$$

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 the Kalman gain defined in (10).

Proof: Recall that $\tilde{V}_{k+1}^\epsilon(z) = \min_{P \in \rho_{\mathbb{M}}(\mathcal{H}_k^\epsilon)} z^T P z$, for all $z \in \mathbb{R}^n$. Clearly, condition (12) implies that

$$V_k^\epsilon(z) - \tilde{V}_{k+1}^\epsilon(z) \ge (\kappa_3 - \kappa_*)\|z\|^2, \quad \forall z \in \mathbb{R}^n.$$

Now let $z \in \mathbb{R}^n$ be arbitrary but fixed. Denote by $(\hat{P}, \hat{i})$ the minimizer in (9) for this fixed z. Suppose that the system starts from z at time 0 and is driven by the policy $\pi_\infty^{\epsilon,k}$. Let $\hat{u} = -K_{\hat{i}}(\hat{P})z$ and $\hat{x}_1 = A_{\hat{i}} z + B_{\hat{i}} \hat{u}$ be the continuous control at time 0 and the state at time $t = 1$, respectively. Plugging equations (4) and (10) into $\hat{u}$, we have $V_k^\epsilon(\hat{x}_1) = \min_{P \in \mathcal{H}_k^\epsilon} \left[\hat{x}_1^T \cdot P \cdot \hat{x}_1\right] \le \hat{x}_1^T \cdot \hat{P} \cdot \hat{x}_1 = z^T \rho_{\hat{i}}(\hat{P}) z - \hat{u}^T R_{\hat{i}} \hat{u} - z^T Q_{\hat{i}} z \le z^T \rho_{\hat{i}}(\hat{P}) z - \kappa_* \|z\|^2 = \tilde{V}_{k+1}^\epsilon(z) - \kappa_* \|z\|^2$. which implies $V_k^\epsilon(z) - V_k^\epsilon(\hat{x}_1) \ge V_k^\epsilon(z) - \tilde{V}_{k+1}^\epsilon(z) + \kappa_* \|z\|^2 \ge \kappa_3 \|z\|^2$. ■

Similar to Lemma 1, condition (12) can be verified by solving a convex optimization problem as in Lemma 2.

B. Performance Bound for $\pi_\infty^{\epsilon,k}$

Whenever $\pi_\infty^{\epsilon,k}$ is exponentially stabilizing, the cost $J(\cdot; \pi_\infty^{\epsilon,k})$ will be bounded from above. We now derive an analytical expression for this bound.

Theorem 6: If V_k^ϵ satisfies the condition in Theorem 5, then the cost associated with $\pi_\infty^{\epsilon,k}$ is bounded from above: $J(z;\pi_\infty^{\epsilon,k}) \le \left(1+\eta(\pi_\infty^{\epsilon,k})\right)V^*(z)$, where $\eta(\pi_\infty^{\epsilon,k}) = \left(\frac{\epsilon\beta}{\lambda_Q^-}+\alpha_V\gamma_V^k\right)\frac{\alpha_x}{(1-\gamma_x)\lambda_Q^-}$. Here β, α_V, γ_V, α_x and γ_x are the constants defined in Theorems 1 and 5.

Proof: Fix an arbitrary $z\in\mathbb{R}^n$. Let $\hat{x}(t) = x(t;z,\pi_\infty^{\epsilon,k})$ for $t\in\mathbb{Z}_+$, and let $(\hat{u}(\cdot),\hat{v}(\cdot))$ be the corresponding hybrid-control sequence. Then,

$$\begin{aligned}J(z;\pi_\infty^{\epsilon,k}) &= \sum_{t=0}^{\infty} L(\hat{x}(t),\hat{u}(t),\hat{v}(t))\\ &= \sum_{t=0}^{\infty}\left[\tilde{V}_{k+1}^\epsilon(\hat{x}(t)) - V_k^\epsilon(\hat{x}(t+1))\right]\\ &= \tilde{V}_{k+1}^\epsilon(z) + \sum_{t=1}^{\infty}\left[\tilde{V}_{k+1}^\epsilon(\hat{x}(t)) - V_k^\epsilon(\hat{x}(t))\right],\end{aligned}$$

where the last step follows since the stability of the trajectory $\hat{x}(t)$ implies that $V_k^\epsilon(x(t))\to 0$ as $t\to\infty$. Using Lemma 3, Theorems 1 and 5, and noticing the monotonicity of the value functions, yields

$$\begin{aligned}&J(z;\pi_\infty^{\epsilon,k}) = \tilde{V}_{k+1}^\epsilon(z) + \sum_{t=1}^{\infty}\left[\tilde{V}_{k+1}^\epsilon(\hat{x}(t)) - V_k^\epsilon(\hat{x}(t))\right]\\ &\le \left(1+\frac{\epsilon}{\lambda_Q^-}\right)V^*(z) + \sum_{t=1}^{\infty}\left[(1+\epsilon/\lambda_Q^-)V_{k+1}(\hat{x}(t)) - V_k(\hat{x}(t))\right]\\ &\le V^*(z) + \frac{\epsilon\beta}{\lambda_Q^-}\|z\|^2 + \sum_{t=1}^{\infty}\left[V_{k+1}(\hat{x}(t)) - V_k(\hat{x}(t))\right]\\ &\quad + \sum_{t=1}^{\infty}\frac{\epsilon\beta}{\lambda_Q^-}\|\hat{x}(t)\|^2\\ &\le V^*(z) + \left(\frac{\epsilon\beta}{\lambda_Q^-}+\alpha_V\gamma_V^k\right)\frac{\alpha_x}{1-\gamma_x}\|z\|^2 \le (1+\eta(\pi_\infty^{\epsilon,k}))V^*(z).\end{aligned}$$

■

Remark 2: The error bound $\eta(\pi_\infty^{\epsilon,k})$ is derived based on the worst case scenario and thus may be conservative. The conservative bound is of theoretical importance as it indicates that the error decays linearly as ϵ decreases, and exponentially as k increases. Moreover, the error can be made arbitrarily small by choosing ϵ small enough and k large enough.

C. Overall Algorithm

If system (1) is exponentially stabilizable, then by Theorem 5, it can always be stabilized by $\pi_\infty^{\epsilon,k}$ for sufficiently large k and sufficiently small ϵ. Such a policy can be found by checking condition (11), or more efficiently, by verifying condition (12) through the solution of a convex optimization problem. If the policy $\pi_\infty^{\epsilon,k}$ is exponentially stabilizing for some k and ϵ, its corresponding cost $J(z;\pi_\infty^{\epsilon,k})$ is bounded from above for all initial states $z\in\mathbb{R}^n$. Theorem 6 implies that further increasing k and decreasing ϵ will eventually result in a policy with any desired suboptimal performance.

Since the control law ξ_k^ϵ is completely characterized by the relaxed SRS $\mathcal{H}_k^\epsilon$, a suboptimal policy of the form $\pi_\infty^{\epsilon,k}$ can be obtained through the relaxed SRM. The basic idea is to evolve $\mathcal{H}_k^\epsilon$ according to the relaxed SRM (6) and stop when the obtained $\mathcal{H}_k^\epsilon$ verifies condition (12). The resulting policy π_∞^ϵ is guaranteed to be suboptimal with a relative error upper bound $\eta(\pi_\infty^{\epsilon,k})$. The detailed solution procedure is given in Algorithm 2. By choosing a proper parameter pair $(\epsilon,k_{\max})$ in the algorithm, the returned policy $\pi_\infty^{\epsilon,k}$ is exponentially stabilizing and can achieve any desired suboptimal performance under Assumption 1. It is also worth mentioning that with the returned set $\mathcal{H}_k^\epsilon$, the closed-loop control sequences driven by $\pi_\infty^{\epsilon,k}$ can also be easily determined using (9). For example, if the state at time t is $x(t)$, then the hybrid-control action at this time step is: $u(t) = -K_{\hat{i}}(\hat{P})x(t)$ and $v(t)=\hat{i}$, where $(\hat{P},\hat{i}) = \arg\min_{P\in\mathcal{H}_k^\epsilon, i\in\mathbb{M}} x(t)^T\rho_i(P)x(t)$.

Algorithm 2 Infinite-Horizon Suboptimal Policy

Choose $\epsilon>0$ and $k_{\max}\in\mathbb{Z}_+$, and set $k=0$ and $\mathcal{H}_0^\epsilon=\{0\}$
for $k=1$ to $k_{\max}$ **do**
  $\mathcal{H}_k^\epsilon \leftarrow ES_\epsilon(\rho_{\mathbb{M}}(\mathcal{H}_k^\epsilon))$
  if $\mathcal{H}_k^\epsilon$ satisfies (12) **then**
    stop and return $\mathcal{H}_k^\epsilon$ that characterizes $\pi_\infty^{\epsilon,k}$ with a relative error bound $\eta(\pi_\infty^{\epsilon,k})$.
  end if
end for

IV. Numerical Examples

A. A Simple 2D Example

Consider a simple infinite-horizon DSLQR problem with two second-order subsystems:

$$A_1=\begin{bmatrix}2&1\\0&1\end{bmatrix}, B_1=\begin{bmatrix}1\\1\end{bmatrix}, A_2=\begin{bmatrix}2&1\\0&0.5\end{bmatrix}, B_2=\begin{bmatrix}1\\2\end{bmatrix}.$$

Suppose that the state and control weights are $Q_1=Q_2=I_2$ and $R_1=R_2=1$, respectively. Both subsystems are unstable but controllable. Algorithm 2 with $\epsilon=10^{-4}$ is applied to solve this DSLQR problem. After 5 steps, we have

$$\mathcal{H}_5^\epsilon = \left\{\begin{bmatrix}6.064&1.205\\1.205&1.905\end{bmatrix},\begin{bmatrix}9.084&3.233\\3.233&2.347\end{bmatrix},\begin{bmatrix}5.107&1.266\\1.266&1.935\end{bmatrix},\begin{bmatrix}7.216&2.560\\2.560&2.106\end{bmatrix}\right\}.$$

We observe that $\mathcal{H}_5^\epsilon$ contains four matrices and it verifies condition (12) with $\kappa_3=0.996$. Therefore, we can stop the iteration now with a stabilizing policy $\pi_\infty^{\epsilon,5}$ characterized by $\mathcal{H}_5^\epsilon$. The relative error bound of this policy is computed to be $\eta(\pi_\infty^{\epsilon,5})=0.13$. The actual performance of the policy should be much better than indicated by this conservative bound. The bound can be improved by carrying out more iterations. For example, after 3 more steps, we will have

$$\mathcal{H}_8^\epsilon = \left\{\begin{bmatrix}6.065&1.206\\1.206&1.905\end{bmatrix},\begin{bmatrix}9.087&3.235\\3.235&2.348\end{bmatrix},\begin{bmatrix}5.108&1.266\\1.266&1.935\end{bmatrix},\begin{bmatrix}7.219&2.561\\2.561&2.107\end{bmatrix}\right\}.$$

Again, the set $\mathcal{H}_8^\epsilon$ still contains only four matrices and all of them are very close to the ones in $\mathcal{H}_5^\epsilon$. The set $\mathcal{H}_8^\epsilon$ verifies condition (12) with $\kappa_3=0.9962$. However, with a larger

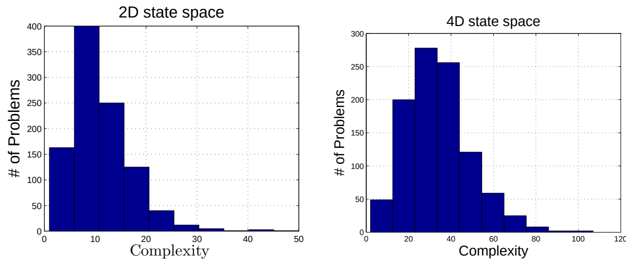

Fig. 2. Complexity distributions of the random examples.

$k = 8$, the conservative bound reduces to $\eta(\pi_\infty^{\epsilon;8}) = 0.009$, which is less than one percent. For this particular example, the complexity of Algorithm 2 as measured by $|\mathcal{H}_k^\epsilon|$, is indeed very small and stays at the maximum value 4 as opposed to growing exponentially.

B. Complexity Statistics of Randomly Generated Examples

To further demonstrate its effectiveness, Algorithm 2 is tested by two sets of randomly generated DSLQR problems. The first set consists of 1000 two-dimensional DSLQR problems with 10 subsystems. The second set consists of 1000 four-dimensional DSLQR problems with 4 subsystems. For both setups, the control horizon is infinite and $\epsilon = 10^{-3}$. All the instances of these problems are successfully solved by Algorithm 2 and the distributions of the complexity, namely, the number of matrices in $\mathcal{H}_k^\epsilon$ returned by the algorithm, are plotted in Fig. 2. It can be seen from the figure that all of the solutions of the two-dimensional problems require less than 50 matrices and a majority of them only need less than 15 matrices. However, a majority of the solutions of the four-dimensional problems need about 40 matrices and some of them may need more than 100 matrices. In general, the complexity of Algorithm 2 increases with the state dimension because randomly-generated higher-dimensional matrices are less likely to be redundant. In a higher dimensional state space, a larger relaxation ϵ is usually needed to retain the same computational complexity.

V. Conclusion

We have developed an efficient algorithm to solve the infinite-horizon DSLQR problem with guaranteed suboptimal performance. The proposed algorithm together with its analysis provides a general controller synthesis framework for switched linear systems with quadratic performance index. A preliminary version of the proposed algorithm has already been successfully applied to stabilize switched linear systems [10]. Since the LQR controller is widely used to achieve local stability for nonlinear systems, our DSLQR solution can be similarly used to locally stabilize switched nonlinear systems. In addition, we also envision the algorithm and its analysis to be useful in quadratic estimation, robust control, and model predictive control problems of switched linear systems or even general hybrid systems.

Appendix

Denote by $I_B^+ \subset \mathbb{M}$ the set of indices of nonzero B matrices, i.e., $I_B^+ \triangleq \{i \in \mathbb{M} : B_i \neq 0\}$. Define $\lambda_Q^+ = \max_{i\in\mathbb{M}}\{\lambda_{\max}(Q_i)\}$, $\lambda_R^+ = \max_{i\in\mathbb{M}}\{\lambda_{\max}(R_i)\}$, and $\sigma_A^+ = \max_{i\in\mathbb{M}}\left\{\sqrt{\lambda_{\max}(A_i^T A_i)}\right\}$. Let $\sigma_{\min}^+(\cdot)$ be the smallest *positive* singular value of a nonzero matrix. If $I_B^+ \neq \emptyset$, let $\hat{\sigma}_B = \min_{i\in I_B^+}\{\sigma_{\min}^+(B_i)\}$. Let π_∞ be an exponentially stabilizing policy, and let $b \geq 1$ and $a \in (0,1)$ be the constants such that $\|x(t;z,\pi_\infty)\|^2 \leq b a^t \|z\|^2$. Define

$$\beta = \begin{cases} \frac{b\lambda_Q^+}{1-a}, & \text{if } I_B^+ = \emptyset \\ \left(\lambda_Q^+ + \lambda_R^+ \frac{2[a+(\sigma_A^+)^2]}{\hat{\sigma}_B^2}\right)\cdot \dfrac{b}{1-a}, & \text{otherwise;} \end{cases} \tag{13}$$

Part (i) of Theorem 1 can be found in [10, Lemma 4]. To prove Part (ii), we fix an arbitrary $z \in \mathbb{R}^n$ and integer $k > 1$, let $x_k^*(\cdot)$ be the optimal k-horizon trajectory starting from z and let $(u_k^*(\cdot), v_k^*(\cdot))$ be the corresponding optimal hybrid-control sequence. Then, for $t = 0, \ldots, k-1$, we have $V_{k-t}(x_k^*(t)) - V_{k-(t+1)}(x_k^*(t+1)) = L(x_k^*(t), u_k^*(t), v_k^*(t)) \geq \lambda_Q^-\|x_k^*(t)\|^2 \geq \frac{\lambda_Q^-}{\beta}V_{k-t}(x_k^*(t)) \geq \frac{\lambda_Q^-}{\beta}V_{k-(t+1)}(x_k^*(t+1))$. Since $\lambda_Q^-\|x_k^*(t)\|^2 \leq V_{k-t}(x_k^*(t)) \leq \beta\|x_k^*(t)\|^2$, for $t = 0, \ldots, k-1$, the above inequality implies

$$\|x_k^*(k-1)\|^2 \leq \frac{\beta}{\lambda_Q^-}\left(\frac{1}{1+\lambda_Q^-/\beta}\right)^{k-1}\|z\|^2. \tag{14}$$

Note that we can not obtain a similar inequality for $x_k^*(k)$ as $V_0(x_k^*(k)) \equiv 0$. By the optimality of V^*, we have $V^*(z) \leq \sum_{t=0}^{k-2} L(x_k^*(t), u_k^*(t), v_k^*(t)) + V^*(x_k^*(k-1)) = V_k(z) - L(x_k^*(k-1), u_k^*(k-1), v_k^*(k-1)) + V^*(x_k^*(k-1)) \leq V_k(z) + (\beta - \lambda_Q^-)\|x_k^*(k-1)\|^2$. The desired result then follows easily by plugging (14) into the above inequality.

References

[1] D. P. Bertsekas. *Dynamic Programming and Optimal Control*, volume 2. Athena Scientific, 2nd edition, 2001.

[2] D. P. Bertsekas and J. Tsitsiklis. *Neuro-Dynamic Programming*. Athena Scientific, Sep. 1996.

[3] F. Borrelli, M. Baotic, A. Bemporad, and M. Morari. Dynamic programming for constrained optimal control of discrete-time linear hybrid systems. *Automatica*, 41:1709–1721, Oct. 2005.

[4] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, New York, NY, USA, 2004.

[5] J. Hu, J. Shen, and W. Zhang. Generating functions of switched linear systems: Analysis, computation, and stability applications. *IEEE Transaction on Automatic Control*, 56(5):1059 – 1074, May 2011.

[6] H. K. Khalil. *Nonlinear Systems*. Prentice Hall, 2002.

[7] B. Lincoln and A. Rantzer. Relaxing dynamic programming. *IEEE Transactions on Automatic Control*, 51(8):1249–1260, Aug. 2006.

[8] W. B. Powell. *Approximate Dynamic Programming: Solving the Curses of Dimensionality (Wiley Series in Probability and Statistics)*. Wiley-Interscience, 2007.

[9] W. Zhang, A. Abate, and J. Hu. Efficient suboptimal solutions of switched LQR problems. In *Proceedings of the American Control Conference*, St. Louis, MO, Jun. 2009.

[10] W. Zhang, A. Abate, J. Hu, and M. Vitus. Exponential stabilization of discrete-time switched linear systems. *Automatica*, 45(11):2526–2536, Nov. 2009.

[11] W. Zhang, J. Hu, and A. Abate. On the value functions of the discrete-time switched LQR problem. *IEEE Transactions on Automatic Control*, 54(11):2669–2674, Nov. 2009.