

Efficient Suboptimal Solutions of Switched LQR Problems

Wei Zhang, Alessandro Abate and Jianghai Hu

Abstract—This paper studies the discrete-time switched LQR (DSLQR) problem using a dynamic programming approach. Efficient algorithms are proposed to solve both the finite-horizon and the infinite-horizon suboptimal DSLQR problems. More importantly, we establish analytical conditions under which the strategies generated by the algorithms are stabilizing and suboptimal. These conditions are derived explicitly in terms of subsystem matrices and are thus easy to verify. The proposed algorithms and the analysis provide a systematical way of solving the DSLQR problem with guaranteed closed-loop stability and suboptimal performance. Simulation results indicate that the proposed algorithms can efficiently solve not only specific but also randomly generated DSLQR problems, making NP-hard problems numerically tractable.

I. INTRODUCTION

This paper studies the discrete-time linear quadratic regulation problem for switched linear systems, hereby referred to as the DSLQR problem. The goal is to develop a computationally appealing algorithm to construct an optimal or suboptimal feedback strategy that minimizes a quadratic cost function. The problem is of fundamental importance both in theory and practice and has challenged researchers for many years. The bottleneck mostly lies in the determination of the optimal switching strategy. Many methods have been proposed to tackle this problem, most of which are in a divide-and-conquer manner. Algorithms for optimizing the switching instants for a fixed mode sequence have been developed for general switched systems in [1] and for switched systems with autonomous dynamics in [2]. Although an algorithm for updating the switching sequence is discussed in [2], finding the best switching sequence is still an NP-hard problem, even for switched linear systems.

This paper studies the (DSLQR) problem based on a dynamic programming (DP) approach. A DP problem can usually be solved by gridding the state space. Such a gridding-based method has been used to study various optimal control problems for switched systems [3], [4], [5]. Its main drawback lies in the exponential growth of the complexity as the state dimension increases. Fortunately, for some simple classes of problems, the value functions may have some nice analytical properties that can be used to simplify the numerical computation. For this reason, the quadratic optimal control problem of switched/piecewise linear systems has

attracted much research attention [6], [7], [8]. In our previous papers [9], [10], we showed that the value function of the DSLQR problem is the pointwise minimum over a finite number of quadratic functions. More importantly, we showed that these quadratic functions are exactly characterized by a set of positive semidefinite matrices which can be obtained using the difference Riccati equation. These properties can dramatically simplify the solution of the DSLQR problem.

The main contribution of this paper is the design and analysis of various efficient algorithms for solving the finite-horizon and infinite-horizon suboptimal DSLQR problems. The key idea is to use convex optimization to identify and remove the matrices that are redundant in terms of characterizing the optimal and suboptimal strategies. This is in line with the approaches of Neuro-dynamic programming ([11]) and approximate dynamic programming ([12]), which both try to simplify the computation by finding a compact representation of the value functions up to certain numerical relaxations. Compared with the previous work, our distinction mostly lies on the analysis of these algorithms. We establish conditions under which the strategies generated by the proposed algorithms are stabilizing and suboptimal. More importantly, these conditions are derived explicitly in terms of subsystem matrices and are very easy to verify. Therefore, the proposed algorithms, together with the analysis, provide a systematic way of solving the DSLQR problem with guaranteed closed-loop stability and suboptimal performance. Simulation results indicate that the proposed algorithms can efficiently solve not only specific but also randomly generated DSLQR problems, making the NP-hard problems numerically tractable.

This paper is organized as follows. The DSLQR problem is formulated in Section II. Some previous results on the value functions of the DSLQR problem are reviewed in Section III. Efficient algorithms are proposed to solve the finite-horizon and infinite-horizon DSLQR problems in Section IV and Section V, respectively. The algorithms are tested through some numerical examples in Section VI. Finally, some concluding remarks are given in Section VII.

II. PROBLEM FORMULATION

Consider the discrete-time switched linear system described by:

$$x(t+1) = A_{v(t)}x(t) + B_{v(t)}u(t), t \in T_N \triangleq \{0, \dots, N-1\}. \quad (1)$$

where $x(t) \in \mathbb{R}^n$ is the continuous state, $v(t) \in \mathbb{M} \triangleq \{1, \dots, M\}$ is the switching control that determines the discrete mode, $u(t) \in \mathbb{R}^p$ is the continuous control and

This work was partially supported by the National Science Foundation under Grant CNS-0643805

W. Zhang and J. Hu are with the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN47906. Email: {zhang70,jianghai@purdue.edu}.

A. Abate is with the Department of Aeronautics and Astronautics, Stanford University, CA 94305. Email: {aabate@stanford.edu}.

T_N is the control horizon with length N (possibly infinite). The sequence of pairs $\{(u(t), v(t))\}_{t=0}^{N-1}$ is called the *hybrid control sequence*. For each $i \in \mathbb{M}$, A_i and B_i are constant matrices of appropriate dimensions and the pair (A_i, B_i) is called a subsystem of (1). This switched linear system is time invariant in the sense that the set of available subsystems $\{(A_i, B_i)\}_{i=1}^M$ is independent of time t . We assume that there is no internal forced switchings, i.e., the system can stay at or switch to any mode at any time instant. At each time $t \in T_N$, denote by $\xi_{t,N} \triangleq (\mu_{t,N}, \nu_{t,N}) : \mathbb{R}^n \rightarrow \mathbb{R}^p \times \mathbb{M}$ the *hybrid-feedback law* of system (1), where $\mu_{t,N} : \mathbb{R}^n \rightarrow \mathbb{R}^p$ is called the *continuous-feedback law* and $\nu_{t,N} : \mathbb{R}^n \rightarrow \mathbb{M}$ is called the *switching-feedback law*. A sequence of hybrid-feedback laws over the horizon T_N constitutes an N -*horizon feedback policy*: $\pi_N \triangleq \{\xi_{0,N}, \xi_{1,N}, \dots, \xi_{N-1,N}\}$. If system (1) is driven by a feedback policy π_N , then the closed-loop dynamics is governed by

$$x(t+1) = A_{\nu_{t,N}(x(t))}x(t) + B_{\nu_{t,N}(x(t))}\mu_{t,N}(x(t)), t \in T_N. \quad (2)$$

For a given initial state $x(0) = z$, the performance of the feedback policy π_N can be measured by the following cost functional:

$$J_{\pi_N}(z) = \psi(x(N)) + \sum_{t=0}^{N-1} L(x(t), \mu_{t,N}(x(t)), \nu_{t,N}(x(t))), \quad (3)$$

where $\psi : \mathbb{R}^n \rightarrow \mathbb{R}^+$ and $L : \mathbb{R}^n \times \mathbb{R}^p \times \mathbb{M} \rightarrow \mathbb{R}^+$ are called the *terminal cost function* and the *running cost function*, respectively, which are assumed to take the following quadratic forms:

$$\psi(x) = x^T Q_f x, \quad L(x, u, v) = x^T Q_v x + u^T R_v u,$$

for $x \in \mathbb{R}^n$, $u \in \mathbb{R}^p$, $v \in \mathbb{M}$, where $Q_f = Q_f^T \succeq 0$ is the terminal state-weighting matrix, and $Q_v = Q_v^T \succeq 0$ and $R_v = R_v^T \succ 0$ are the running weighting matrices for the state and the control, respectively, for subsystem $v \in \mathbb{M}$. When the control horizon N is infinite, the terminal cost will never be incurred and the objective function becomes:

$$J_{\pi_\infty}(z) = \sum_{t=0}^{\infty} L(x(t), \mu_{t,\infty}(x(t)), \nu_{t,\infty}(x(t))). \quad (4)$$

Denote by Π_N the set of all admissible N -horizon policies. The goal of this paper is to solve the following problem.

Problem 1 (DSLQR problem): For a given initial state $z \in \mathbb{R}^n$ and a possibly infinite positive integer N , find the N -horizon policy $\pi_N \in \Pi_N$ that minimizes $J_{\pi_N}(z)$ subject to (2).

To solve Problem 1, for each time $t \in T_N$, we define the value function $V_{t,N} : \mathbb{R}^n \rightarrow \mathbb{R}$ as:

$$V_{t,N}(z) = \inf_{\substack{u(j) \in \mathbb{R}^p, v(j) \in \mathbb{M} \\ t \leq j \leq N-1}} \left\{ \psi(x(N)) + \sum_{j=t}^{N-1} L(x(j), u(j), v(j)) \right\} \\ \text{subject to eq. (1) with } x(t) = z. \quad (5)$$

The $V_{t,N}(z)$ so defined is the minimum cost-to-go starting from state z at time t . The minimum cost for the DSLQR

problem with an initial condition $x(0) = z$ is simply $V_{0,N}(z)$. Due to the time-invariant nature of the switched system (1), its value function depends only on the number of remaining time steps, i.e.,

$$V_{t,N}(z) = V_{t+m, N+m}(z),$$

for all $z \in \mathbb{R}^n$ and all integers $m \geq -t$. In the rest of this paper, when no ambiguity arises, we will denote by $V_k(z) \triangleq V_{N-k, N}(z)$ and $\xi_k \triangleq \xi_{N-k, N}$ the value function and the hybrid-feedback law, respectively, at time $t = N - k$ when there are k time steps left. With the new notations, the N -horizon policy π_N can also be written as $\pi_N = \{\xi_N, \dots, \xi_1\}$. For any positive integer k , the feedback law ξ_k can be thought of as the first step of a k -horizon policy.

By a standard result of Dynamic Programming [13], for any finite integer N , the value function V_N can be obtained recursively using the one-stage *value iteration*:

$$V_{k+1}(z) = \inf_{u,v} \{L(z, u, v) + V_k(A_v z + B_v u)\}, \forall z \in \mathbb{R}^n,$$

with the initial condition $V_0(z) = \psi(z)$, $\forall z \in \mathbb{R}^n$. Denote by $V_\infty(\cdot)$ the pointwise limit (if it exists) of the sequence of functions $\{V_k(\cdot)\}_{k=0}^\infty$ generated by the value iterations. It is well known ([13]) that even if $V_\infty(z)$ exists, it may not always coincide with the true infinite-horizon value function. To emphasize its substantial difference from the finite-horizon value function, the infinite-horizon value function is specially denoted by $V^*(z)$, i.e., $V^*(z) = \inf_{\pi_\infty \in \Pi_\infty} J_{\pi_\infty}(z)$.

III. THE VALUE FUNCTION

The finite-horizon value functions of the DSLQR problem have been characterized analytically in [9]. The particular structure of these value functions is crucial for the design and analysis of the suboptimal algorithms to be developed in Sections IV and V.

It is well-known that the value function of the classical LQR problem is quadratic and the positive semidefinite (p.s.d.) matrix characterizing this quadratic value function can be obtained recursively through difference Riccati equation (DRE). Let \mathcal{A} be the set of all p.s.d. matrices. Then, the DRE can be viewed as a mapping from \mathcal{A} to \mathcal{A} depending on the system matrices. We call this mapping the *Riccati Mapping* and denote by $\rho_i : \mathcal{A} \rightarrow \mathcal{A}$ the Riccati Mapping of subsystem $i \in \mathbb{M}$, i.e.,

$$\rho_i(P) = Q_i + A_i^T P A_i - A_i^T P B_i (R_i + B_i^T P B_i)^{-1} B_i^T P A_i. \quad (6)$$

To define the value function of the DSLQR problem, switchings are incorporated into the Riccati mapping, resulting in the switched Riccati mapping.

Definition 1: Let $2^{\mathcal{A}}$ be the power set of \mathcal{A} . The mapping $\rho_{\mathbb{M}} : 2^{\mathcal{A}} \rightarrow 2^{\mathcal{A}}$ defined by:

$$\rho_{\mathbb{M}}(\mathcal{H}) = \{\rho_i(P) : \text{for some } i \in \mathbb{M} \text{ and } P \in \mathcal{H}\}$$

is called the *Switched Riccati Mapping* (SRM) of Problem 1.

In words, the SRM maps a **set** of p.s.d. matrices to another **set** of p.s.d. matrices and each matrix in $\rho_{\mathbb{M}}(\mathcal{H})$ is obtained

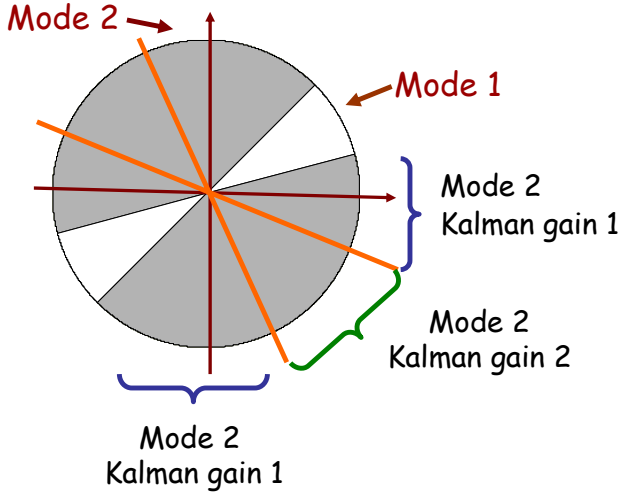


Fig. 1. Typical optimal decision regions of a 2-switched system, where mode 1 is optimal within the white region and mode 2 is optimal within the gray region. The optimal mode region is further divided into smaller conic regions, each of which corresponds to a different Kalman gain.

by taking the classical Riccati mapping of some matrix in \mathcal{H} through some subsystem $i \in \mathbb{M}$.

Definition 2: The sequence of sets $\{\mathcal{H}_k\}_{k=0}^N$ generated iteratively by $\mathcal{H}_{k+1} = \rho_{\mathbb{M}}(\mathcal{H}_k)$ with initial condition $\mathcal{H}_0 = \{Q_f\}$ is called the *Switched Riccati Sets* (SRSs) of Problem 1.

An important fact about the DSLQR problem is that its value functions are completely characterized by the SRSs.

Theorem 1 ([9]): The value function for the DSLQR problem at time $N - k$, i.e., with k time steps left, is

$$V_k(z) = \min_{P \in \mathcal{H}_k} z^T P z. \quad (7)$$

Furthermore, for $k = 1, \dots, N$, if we define

$$(P_k^*(z), i_k^*(z)) = \arg \min_{(P \in \mathcal{H}_{k-1}, i \in \mathbb{M})} z^T \rho_i(P) z, \quad (8)$$

then the optimal hybrid-feedback law at time $t = N - k$ is $\xi_k^*(z) = (\mu_k^*(z), \nu_k^*(z))$, where $\mu_k^*(z) = -K_{i_k^*(z)}(P_k^*(z))z$ and $\nu_k^*(z) = i_k^*(z)$. Here, $K_i(P)$ is the Kalman gain for subsystem i with matrix P , i.e.,

$$K_i(P) \triangleq (R_i + B_i^T P B_i)^{-1} B_i^T P A_i. \quad (9)$$

Compared with the classical LQR problem, the value function of the DSLQR problem is no longer a single quadratic function; it becomes the pointwise minimum of a finite number of quadratic functions. At each time step, instead of having a single Kalman gain for the entire state space, the optimal state feedback gain becomes state dependent. Furthermore, the minimizer $(P_k^*(z), i_k^*(z))$ of equation (8) is radially invariant, indicating that at each time step all the points along the same radial direction have the same optimal hybrid-feedback law. These properties are illustrated in Fig. 1 using an example in \mathbb{R}^2 with 2 subsystems. At each time step, the state space is decomposed into several conic *decision*

regions, each of which corresponds to a pair of optimal mode and optimal feedback gain. In addition, all the gray conic regions have the same optimal mode, say mode 2. It is worth mentioning that in a higher dimensional state space, the conic decision regions may become nonconvex and rather complicated. A salient feature of the DSLQR problem is that all these complex decision regions are completely encoded in a finite number of matrices in the SRSs.

IV. SUBOPTIMAL CONTROL IN FINITE HORIZON

It is clear from Theorem 1 that the key for solving the DSLQR problem is the computation of the SRSs $\{\mathcal{H}_k\}_{k=0}^N$. The main challenge lies in the exponential growth of $|\mathcal{H}_k|$ as k increases. Two important facts about the value functions can be used to simplify the computation of $\{\mathcal{H}_k\}_{k=0}^N$. (i) The matrices in \mathcal{H}_k that make no contribution to the minimum of (7) can be directly removed without causing any error. (ii) When suboptimal performance is acceptable, the matrices in \mathcal{H}_k that make only a “small” contribution to the minimum (7) can also be removed. The goal of this subsection is to use these two ideas to simplify the computation of the (sub)-optimal solution and to analyze the impacts of these simplifications on various aspects of the closed-loop system.

A. Redundancy and Equivalent Subsets

To formalize the above idea, we introduce a few definitions.

Definition 3 (Algebraic Redundancy): A matrix $\hat{P} \in \mathcal{H}$ is called (*algebraic*) *redundant* if for any $z \in \mathbb{R}^n$, there exist a matrix $P \in \mathcal{H}$ such that $z^T P z \leq z^T \hat{P} z$.

If $\hat{P} \in \mathcal{H}$ is redundant, then \mathcal{H} and $\mathcal{H} \setminus \{\hat{P}\}$ will define the same value functions. In this sense, these two sets are equivalent.

Definition 4 (Equivalent Subset (ES)): Let \mathcal{H} and $\bar{\mathcal{H}}$ be two sets of p.s.d matrices.

- 1) The set $\bar{\mathcal{H}}$ is called *equivalent* to \mathcal{H} , denoted by $\mathcal{H} \sim \bar{\mathcal{H}}$, if $\min_{P \in \mathcal{H}} z^T P z = \min_{\bar{P} \in \bar{\mathcal{H}}} z^T \bar{P} z, \forall z \in \mathbb{R}^n$.
- 2) $\bar{\mathcal{H}}$ is called an *equivalent subset* of \mathcal{H} if $\bar{\mathcal{H}} \subseteq \mathcal{H}$ and $\bar{\mathcal{H}} \sim \mathcal{H}$.

Some matrices in \mathcal{H}_k are “almost” algebraic redundant in the sense that removing them will only introduce a small error to the value function. We call these matrices numerically redundant.

Definition 5 (Numerical Redundancy): A matrix $\hat{P} \in \mathcal{H}_k$ is called (*numerically*) ϵ -*redundant* with respect to \mathcal{H}_k if

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

Definition 6 (ϵ -ES): The set \mathcal{H}_k^ϵ is called an ϵ -Equivalent-Subset (ϵ -ES) of \mathcal{H}_k if $\mathcal{H}_k^\epsilon \subseteq \mathcal{H}_k$ and for all $z \in \mathbb{R}^n$,

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

Removing the ϵ -redundant matrices may introduce some error for the value function; but the error is no larger than ϵ for $\|z\| \leq 1$. To simplify the computation, for a given tolerance ϵ , we want to prune out as many ϵ -redundant matrices as possible. The following lemma provides a sufficient condition for testing the ϵ -redundancy for a given matrix.

Algorithm 1 [$Algo(\mathcal{H}_k)$]

Specify a tolerance ϵ and set $\mathcal{H}_k^\epsilon = \emptyset$.
for each $P \in \mathcal{H}_k$ **do**
 if P does NOT satisfies the condition in Lemma 1
 with respect to \mathcal{H}_k^ϵ **then**
 $\mathcal{H}_k^\epsilon = \mathcal{H}_k^\epsilon \cup \{P\}$;
 end if
end for
Return \mathcal{H}_k^ϵ .

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

Checking the condition in Lemma 1 is a LMI feasibility problem [14] and can be easily solved using various convex optimization algorithms [15]. To compute an ϵ -ES of \mathcal{H}_k , one only need to remove the matrices in \mathcal{H}_k that satisfy the condition in Lemma 1. The detailed procedure is summarized in Algorithm 1. Denote by $Algo_\epsilon(\mathcal{H}_k)$ the ϵ -ES of \mathcal{H}_k returned by the algorithm. The next step is to combine this algorithm with the SRM. To this end, we define $\{\mathcal{H}_k^\epsilon\}_{k=0}^N$ iteratively as:

$$\mathcal{H}_0^\epsilon = \mathcal{H}_0 \text{ and } \mathcal{H}_{k+1}^\epsilon = Algo_\epsilon(\rho_{\mathbb{M}}(\mathcal{H}_k^\epsilon)), \text{ for } k \in T_N. \quad (10)$$

The iteration (10) computes a sequence of relaxed SRSs $\{\mathcal{H}_k^\epsilon\}_{k=0}^N$. Using the formulas in Theorem 1, these sets $\{\mathcal{H}_k^\epsilon\}_{k=0}^N$ also define a sequence of ‘‘approximate’’ value functions and the corresponding feedback policy. Intuitively speaking, the cost associated with this relaxed policy should be very close to the optimal one if the relaxation parameter ϵ is small. In the next subsection, we shall formally analyze the performance of this policy. Before doing that, we first demonstrate the simplicity of computing such a relaxed policy.

Consider the example with the following matrices:

$$A_1 = \begin{bmatrix} 2 & 1 \\ 1 & 1 \end{bmatrix}, A_2 = \begin{bmatrix} 2 & 1 \\ 0 & 0.5 \end{bmatrix}, A_3 = \begin{bmatrix} 3 & 1 \\ 0 & 2 \end{bmatrix}, \\ A_4 = \begin{bmatrix} 3 & 1 \\ 0 & 0.8 \end{bmatrix}, B_1 = \begin{bmatrix} 1 \\ 1 \end{bmatrix}, B_2 = \begin{bmatrix} 1 \\ 2 \end{bmatrix}, B_3 = B_1, \\ B_4 = B_2, Q_i = Q_f = I_2, R_i = 1, i = 1, \dots, 4, \text{ and } N = 20.$$

By Theorem 1, the solution of this problem is completely determined by the SRSs $\{\mathcal{H}_k\}_{k=0}^N$. As shown in Fig. 2, a direct computation of $\{\mathcal{H}_k\}_{k=0}^N$ results in a combinatorial complexity on the order of 10^{12} . However, if we use the relaxed iteration (10) with $\epsilon = 10^{-3}$, eventually \mathcal{H}_N^ϵ only contains 14 matrices. This example shows that the numerical relaxation can dramatically simplify the computation. Our next task is to prove that a small relaxation ϵ will indeed result in a suboptimal policy and will not change some important properties, such as the stability, of the closed-loop system.

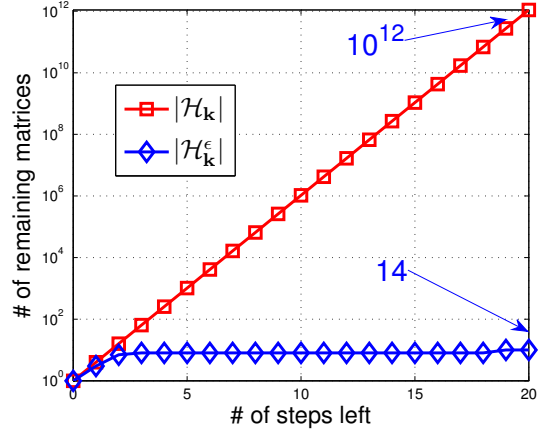


Fig. 2. Evolution of $|\mathcal{H}_k^\epsilon|$

B. Performance Analysis

We first introduce some notations. Let $\|\cdot\|$ be the 2-norm of a given matrix or vector. Let \mathbb{Z}^+ be the set of all nonnegative integers. Denote by $\lambda_{\min}(\cdot)$ and $\lambda_{\max}(\cdot)$ the smallest and the largest eigenvalues of a p.s.d. matrix, respectively. Define $\lambda_Q^- = \min_{i \in \mathbb{M}} \{\lambda_{\min}(Q_i)\}$ and $\lambda_f^+ \triangleq \lambda_{\max}(Q_f)$. For $k = 1, \dots, N$, define $V_k^\epsilon(z) = \min_{P \in \mathcal{H}_k^\epsilon} z^T P z$ and let $\xi_k^\epsilon(\cdot)$ be the feedback law generated by V_{k-1}^ϵ , namely,

$$\xi_k^\epsilon(z) = \arg \min_{(u,v)} \{L(z, u, v) + V_{k-1}^\epsilon(A_v z + B_v u)\}. \quad (11)$$

Following a similar proof of Theorem 1 ([9]), one can easily obtain:

$$\xi_k^\epsilon(z) = \left(-K_{i_k^\epsilon(z)}(P_k^\epsilon(z))z, i_k^\epsilon(z) \right), \quad (12)$$

$$\text{where } \left(P_k^\epsilon(z), i_k^\epsilon(z) \right) = \arg \min_{P \in \mathcal{H}_{k-1}^\epsilon, i \in \mathbb{M}} z^T \rho_i(P)z,$$

where $K_i(\cdot)$ denotes the Kalman gain for a given p.s.d. matrix as defined in (9). Let $\pi_N^\epsilon = \{\xi_N^\epsilon, \dots, \xi_1^\epsilon\}$ be the N -horizon policy generated by $\{V_k^\epsilon\}_{k=0}^{N-1}$. Let $\pi_N^* = \{\xi_N^*, \dots, \xi_1^*\}$ be the optimal policy generated by the *exact* value functions $\{V_k\}_{k=0}^{N-1}$. Typically, π_N^ϵ is much easier to compute than π_N^* because \mathcal{H}_k^ϵ contains much fewer matrices than \mathcal{H}_k . However, the relaxation $Algo_\epsilon(\cdot)$ introduces an error and this error propagates through the iteration (10). Therefore, to take advantage of the simplicity of π_N^ϵ , it must be ensured that $J_{\pi_N^\epsilon}(z)$, namely, the actual cost associated π_N^ϵ , does not deviate too far from the optimal cost $V_N(z)$.

The goal of this subsection is to derive conditions under which the feedback policy π_N^ϵ is stabilizing and suboptimal. A general N -horizon policy π_N is called δ -suboptimal over a set E if for any initial state $x_0 \in E$, the cost under π_N is within the δ -neighborhood of the optimal cost, i.e., $|J_{\pi_N}(x_0) - V_N(x_0)| \leq \delta$. Let $x_{z,N}^*(\cdot)$ be an N -horizon optimal trajectory originating from z at time 0. Similarly, denote by $x_{z,N}^\epsilon(\cdot)$ the N -horizon state trajectory driven by π_N^ϵ with initial condition $x_{z,k}^\epsilon(0) = z$. Define

$$\tilde{V}_{k+1}^\epsilon(z) = \min_{u,v} \{L(z, u, v) + V_k^\epsilon(A_v z + B_v u)\}. \quad (13)$$

Following immediately from (12), we have

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

Therefore, \tilde{V}_{k+1}^ϵ is the function defined based on the set $\rho_{\mathbb{M}}(\mathcal{H}_k^\epsilon)$. As the relaxed SRSs evolve according to (10) from \mathcal{H}_k^ϵ to $\rho_{\mathbb{M}}(\mathcal{H}_k^\epsilon)$, and from $\rho_{\mathbb{M}}(\mathcal{H}_k^\epsilon)$ to $\text{Algo}_\epsilon(\rho_{\mathbb{M}}(\mathcal{H}_k^\epsilon)) = \mathcal{H}_{k+1}^\epsilon$, the corresponding functions defined based on these sets change from V_k^ϵ to \tilde{V}_{k+1}^ϵ , and from \tilde{V}_{k+1}^ϵ to V_{k+1}^ϵ , respectively. Based on this relation and the definition of the ϵ -ES, we have

$$\tilde{V}_{k+1}^\epsilon(z) \leq V_{k+1}^\epsilon(z) \leq \tilde{V}_{k+1}^\epsilon(z) + \epsilon \|z\|^2. \quad (14)$$

We make the following two assumptions in the rest of this paper.

- (A1) At least one subsystem is stabilizable;
- (A2) $Q_i \succ 0, \forall i \in \mathbb{M}$.

Remark 1: The above assumptions are analogous to the stabilizability and detectability conditions commonly adopted in the study of the classical LQR problems. They are not restrictive because randomly generated subsystem matrices will satisfy them with probability 1.

Lemma 2 ([10]): Under assumption (A1), there exists a constant $\beta < \infty$, such that $V_N(z) \leq \beta \|z\|^2$ for all $N \in \mathbb{Z}^+$ and $z \in \mathbb{R}^n$.

Two important inequalities that are frequently used throughout the subsequent discussions are given in the following lemma.

Lemma 3: Under assumptions (A1) and (A2), for any integer $N \geq 0$, we have

$$\begin{aligned} V_N(z) &\leq V_N^\epsilon(z) \leq V_N(z) + \epsilon \eta \|z\|^2 & (15) \\ \text{and } \tilde{V}_N^\epsilon(z) &\leq V_N(z) + \epsilon(\eta - 1) \|z\|^2, & (16) \end{aligned}$$

where $\eta = \frac{1 + (\beta/\lambda_Q^- - 1)\gamma}{1 - \gamma}$.

Proof: See [16]. ■

It has been proved in our previous paper [10] that under assumptions (A1) and (A2), the optimal trajectory $x_{z,N}^*(\cdot)$ is exponentially stable. Intuitively speaking, this property should also hold for $x_{z,N}^\epsilon(\cdot)$ when ϵ is sufficiently small. We now derive an upper bound of ϵ that guarantees the stability of $x_{z,N}^\epsilon(\cdot)$. The following lemma is the key in deriving this upper bound.

Lemma 4: Under assumptions (A1) and (A2), the trajectory $x_{z,N}^\epsilon(\cdot)$ satisfies

$$\begin{aligned} \|x_{z,N}^\epsilon(t)\|^2 &\leq \left(\gamma + \frac{\epsilon\gamma\eta}{\beta}\right)^t \left(\frac{\beta + \epsilon\eta}{\lambda_Q^-}\right) \|z\|^2, \text{ for } t < N, \\ \text{and } \|x_{z,N}^\epsilon(N)\|^2 &\leq \left(\gamma + \frac{\epsilon\gamma\eta}{\beta}\right)^{N-1} \left(\frac{\zeta^2(\beta + \epsilon\eta)}{\lambda_Q^-}\right) \|z\|^2. \end{aligned}$$

where β is the constant defined in Lemma 2,

$$\gamma = \frac{1}{1 + \lambda_Q^-/\beta} \text{ and } \zeta = \max_{i \in \mathbb{M}} \|A_i - B_i K_i(Q_f)\|. \quad (17)$$

Proof: In this proof, we denote $x_{z,N}^\epsilon(\cdot)$ by $\hat{x}(\cdot)$ and assume the corresponding hybrid control sequence is

$(\hat{u}(\cdot), \hat{v}(\cdot))$. By (11), (13–16) and Lemma 2, for each $t = 1, \dots, N$, we have

$$\begin{aligned} &V_{N-(t-1)}^\epsilon(\hat{x}(t-1)) - V_{N-t}^\epsilon(\hat{x}(t)) \\ &\geq \tilde{V}_{N-(t-1)}^\epsilon(\hat{x}(t-1)) - V_{N-t}^\epsilon(\hat{x}(t)) \\ &= L(\hat{x}(t-1), \hat{u}(t-1), \hat{v}(t-1)) \geq \lambda_Q^- \|\hat{x}(t-1)\|^2 \\ &\geq \frac{\lambda_Q^-}{\beta} V_{N-(t-1)}^\epsilon(\hat{x}(t-1)) \\ &\geq \frac{\lambda_Q^-}{\beta} \left(V_{N-(t-1)}^\epsilon(\hat{x}(t-1)) - \epsilon \eta \|\hat{x}(t-1)\|^2 \right) \\ &\geq \frac{\lambda_Q^-}{\beta} V_{N-t}^\epsilon(\hat{x}(t)) - \frac{\lambda_Q^- \epsilon \eta}{\beta} \|\hat{x}(t-1)\|^2. \end{aligned}$$

Therefore, for $t = 1, \dots, N$,

$$\begin{aligned} V_{N-t}^\epsilon(\hat{x}(t)) &\leq \gamma \left[V_{N-(t-1)}^\epsilon(\hat{x}(t-1)) + \frac{\lambda_Q^- \epsilon \eta}{\beta} \|\hat{x}(t-1)\|^2 \right] \\ &\leq \left[\gamma \left(1 + \frac{\epsilon \eta}{\beta} \right) \right] V_{N-(t-1)}^\epsilon(\hat{x}(t-1)) \\ &\leq \left(\gamma + \frac{\epsilon \eta}{\beta} \right)^t V_N^\epsilon(z) \leq \left(\gamma + \frac{\epsilon \eta}{\beta} \right)^t (\beta + \epsilon \eta) \|z\|^2. \end{aligned}$$

Here, the second inequality follows from the fact that $V_k^\epsilon(z) \geq \lambda_Q^- \|z\|^2$ for $k \geq 0$. Using this fact again yields

$$\|\hat{x}(t)\|^2 \leq \left(\gamma + \frac{\epsilon \eta}{\beta} \right)^t \left(\frac{\beta + \epsilon \eta}{\lambda_Q^-} \right) \|z\|^2, \text{ for } t < N. \quad (18)$$

For $t = N$, according to (12), we have that $\hat{x}(N) = (A_i - B_i K_i(Q_f)) \cdot \hat{x}(N-1)$ for some $i \in \mathbb{M}$. Therefore, $\|\hat{x}(N)\|^2 \leq \zeta^2 \|\hat{x}(N-1)\|^2$, where ζ is defined in (17). The desired result follows from (18). ■

With Lemma 4, the following theorem follows immediately.

Theorem 2: Under (A1) and (A2), if $\epsilon < \frac{(1-\gamma)\beta}{\gamma\eta}$, the policy π_N^ϵ is stabilizing.

We now derive an upper bound for the actual cost associated with the policy π_N^ϵ .

Theorem 3: Under assumptions (A1) and (A2), $J_{\pi_N^\epsilon}(z) \leq V_N(z) + \epsilon(\eta - 1) \|z\|^2$ for any $z \in \mathbb{R}^n$ and $N \geq 0$.

Proof: Let $\hat{x}(\cdot)$ and $(\hat{u}(\cdot), \hat{v}(\cdot))$ be the same as in the proof of Lemma 4. By (11) and (13), we have $L(\hat{x}(t), \hat{u}(t), \hat{v}(t)) = \tilde{V}_{N-t}^\epsilon(\hat{x}(t)) - V_{N-(t+1)}^\epsilon(\hat{x}(t+1))$ for each $t = 0, \dots, N-1$. Therefore,

$$\begin{aligned} J_{\pi_N^\epsilon}(z) &= \sum_{t=0}^{N-1} L(\hat{x}(t), \hat{u}(t), \hat{v}(t)) + \psi(\hat{x}(N)) \\ &= \sum_{t=0}^{N-1} [\tilde{V}_{N-t}^\epsilon(\hat{x}(t)) - V_{N-(t+1)}^\epsilon(\hat{x}(t+1))] + \psi(\hat{x}(N)) \\ &= \tilde{V}_N^\epsilon(z) + \sum_{t=1}^{N-1} \left\{ [\tilde{V}_{N-t}^\epsilon(\hat{x}(t)) - V_{N-t}^\epsilon(\hat{x}(t))] \right. \\ &\quad \left. + [\psi(\hat{x}(N)) - V_0^\epsilon(\hat{x}(N))] \right\}. \end{aligned}$$

Since by definition $\psi(z) = V_0^\epsilon(z)$ and $\tilde{V}_{N-t}^\epsilon(z) \leq V_{N-t}^\epsilon(z)$ for any $z \in \mathbb{R}^n$ and $t = 1, \dots, N-1$, we have

$$J_{\pi_N^\epsilon}(z) \leq \tilde{V}_N^\epsilon(z) \leq V_N(z) + \epsilon(\eta - 1) \|z\|^2. \quad \blacksquare$$

Remark 2: Notice that the error function $\epsilon(\eta - 1) \|z\|^2$ does not grow with respect to the horizon N . This property plays a crucial role in deriving the suboptimal policies for the infinite-horizon DSLQR problems.

Corollary 1: Under the same conditions as in Theorem 2, π_N^ϵ is δ -suboptimal over the unit ball if $\epsilon \leq \frac{\delta}{\eta-1}$.

Based on our analysis, a δ -suboptimal policy can be obtained using the following Algorithm.

Algorithm 2 (Suboptimal Control in Finite Horizon)

- 1) **Initialization:** Specify an error tolerance δ . Let $\epsilon = \frac{\delta}{\eta-1}$ and set $\mathcal{H}_0^c = Q_f$
- 2) **Relaxed ES Iteration:** Perform iteration (10) over the whole horizon N .
- 3) **Suboptimal Strategy:** The suboptimal N -horizon policy $\pi_N^\epsilon = \{\xi_N^\epsilon(x), \dots, \xi_1^\epsilon(x)\}$ is given by:

$$\xi_k^\epsilon(x) = \left(-K_{i_k^\epsilon(x)}(P_k^\epsilon(x))x, i_k^\epsilon(x) \right),$$

$$\text{where } \left(P_k^\epsilon(x), i_k^\epsilon(x) \right) = \arg \min_{P \in \mathcal{H}_{k-1}^c, i \in \mathbb{M}} x^T \rho_i(P)x.$$

V. EXTENSION TO LARGE OR INFINITE HORIZON

It is natural to solve the infinite-horizon case in a divide-and-conquer manner, namely, by applying Algorithm 2 to a reasonably large size of subhorizon, m , and then extending the obtained optimal strategy periodically. We now show that, by choosing proper m and ϵ , such a periodic policy can achieve an arbitrary suboptimal performance. Let $\hat{\pi}_m^\epsilon = (\hat{\xi}_m^\epsilon, \dots, \hat{\xi}_1^\epsilon)$ be the m -horizon policy returned by Algorithm 2 with $Q_f = 0$. It follows from Theorem 3 that

$$J_{\hat{\pi}_m^\epsilon}(z) \leq V_m^0(z) + \epsilon(\eta-1)\|z\|^2 \leq V^*(z) + \epsilon(\eta-1)\|z\|^2, \quad (19)$$

where $V_m^0(z)$ denotes the m -horizon value function with $Q_f = 0$. For $m \geq 2$, let $\pi_\infty^{\epsilon, m}$ be the periodic extension of the first $m-1$ terms of $\hat{\pi}_m^\epsilon$, i.e.¹,

$$\pi_\infty^{\epsilon, m} = \{\hat{\xi}_m^\epsilon, \dots, \hat{\xi}_2^\epsilon, \hat{\xi}_m^\epsilon, \dots, \hat{\xi}_2^\epsilon, \dots\}. \quad (20)$$

We first establish conditions under which the specially constructed policy $\pi_\infty^{\epsilon, m}$ is stabilizing.

Theorem 4: Under assumptions (A1) and (A2), if $\epsilon < \frac{(1-\gamma)\beta}{\gamma\eta}$ and $m > \frac{\ln \lambda_Q - \ln(\beta + \epsilon\eta)}{\ln(\beta\gamma + \epsilon\gamma\eta) - \ln \beta} + 1$, then $\pi_\infty^{\epsilon, m}$ is exponentially stabilizing.

Proof: Denote by $\hat{x}(\cdot)$ the trajectory associated with the policy $\pi_\infty^{\epsilon, m}$ with initial condition $\hat{x}(0) = z$. Let

$$c_m = \left(\gamma + \frac{\epsilon\gamma\eta}{\beta} \right)^{m-1} \left(\frac{\beta + \epsilon\eta}{\lambda_Q} \right). \quad (21)$$

It can be easily verified that under our assumptions, c_m is strictly smaller than 1. By Lemma (4), we have $\|\hat{x}(k(m-1))\|^2 \leq c_m \|\hat{x}((k-1)(m-1))\|^2$ for all $k \geq 1$. Thus, $\|\hat{x}(\cdot)\|^2$ must decrease by a factor of $c_m < 1$ in every $m-1$ steps. It follows that the policy $\pi_\infty^{\epsilon, m}$ is exponentially stabilizing. ■

We now derive a bound for the error between the actual cost $J_{\pi_\infty^{\epsilon, m}}(z)$ and the optimal cost $V^*(z)$.

¹As can be seen from Lemma 4, by using only the first $m-1$ terms of $\hat{\pi}_m^\epsilon$ in constructing $\pi_\infty^{\epsilon, m}$, we can obtain a better bound for the convergence of the closed-loop trajectory.

Theorem 5: Under the same conditions as in Theorem 4, we have

$$V^*(z) \leq J_{\pi_\infty^{\epsilon, m}}(z) \leq V^*(z) + \frac{c_m\beta + \epsilon(\eta-1)}{1-c_m} \|z\|^2, \quad (22)$$

where c_m is defined in (21).

Proof: Obviously, $V^*(z) \leq J_{\pi_\infty^{\epsilon, m}}(z)$ as $\pi_\infty^{\epsilon, m}$ is an infinite-horizon policy. Let $\hat{x}(\cdot)$ be the system trajectory generated by the policy $\pi_\infty^{\epsilon, m}$ starting from $\hat{x}(0) = z$. Define $z_i = \hat{x}(i \cdot (m-1))$ for $i = 0, 1, \dots$. Let $\tilde{\pi} \triangleq \{\hat{\xi}_m^\epsilon, \dots, \hat{\xi}_2^\epsilon\}$ be the first $m-1$ terms of $\hat{\pi}_m^\epsilon$. Then by (19),

$$\begin{aligned} J_{\pi_\infty^{\epsilon, m}}(z) &= \sum_{i=0}^{\infty} J_{\tilde{\pi}}(z_i) \leq \sum_{i=0}^{\infty} J_{\hat{\pi}_m^\epsilon}(z_i) \\ &\leq \sum_{i=0}^{\infty} [V^*(z_i) + \epsilon(\eta-1)\|z_i\|^2]. \end{aligned}$$

By Lemma (4), $\|z_i\|^2 \leq c_m^i \|z\|^2$, where $c_m < 1$ is defined in (21). Therefore, $J_{\pi_\infty^{\epsilon, m}}(z) \leq V^*(z) + \frac{c_m\beta + \epsilon(\eta-1)}{1-c_m} \|z\|^2$ for any initial state z . ■

With the above result, we can easily derive a lower bound for m that guarantees the δ -suboptimality of $\pi_\infty^{\epsilon, m}$ for an arbitrary $\delta > 0$.

Corollary 2: Suppose that the conditions in Theorem 4 hold. For any $\delta > 0$, if we further have $\epsilon < \frac{\delta}{\eta-1}$ and

$$m > m_{\infty}^{\delta, \epsilon} \triangleq \frac{\ln[\delta - \epsilon(\eta-1)]\lambda_Q - \ln(\beta + \delta)(\beta + \epsilon\eta)}{\ln(\beta\gamma + \epsilon\gamma\eta) - \ln \beta} + 1, \quad (23)$$

then the policy $\pi_\infty^{\epsilon, m}$ is δ -suboptimal over the unit ball.

For a given tolerance δ on the optimal cost, we only need to perform $m_{\infty}^{\delta, \epsilon}$ steps of the approximate value iterations (10). The obtained value functions $\{V_k^\epsilon(z)\}_{k=0}^{m_{\infty}^{\delta, \epsilon}}$ characterize the $m_{\infty}^{\delta, \epsilon}$ -horizon feedback policy $\hat{\pi}_m^\epsilon$ which can be used periodically to construct an infinite-horizon policy $\pi_\infty^{\epsilon, m}$. By Corollary 2, such a periodic policy is guaranteed to be δ -suboptimal within the unit ball. This idea also applies when the horizon is large but finite. Denote by $[\pi_\infty^{\epsilon, m}]_N$ the N -horizon truncation of the policy $\pi_\infty^{\epsilon, m}$, i.e., $[\pi_\infty^{\epsilon, m}]_N(t) = \pi_\infty^{\epsilon, m}(t)$ for $t = 0, \dots, N-1$. Similar performance bound as in Theorem 5 can be derived for $[\pi_\infty^{\epsilon, m}]_N(t)$.

Theorem 6: Under the same conditions as in Theorem 4, for any $N \geq m$, we have

$$J_{[\pi_\infty^{\epsilon, m}]_N}(z) \leq V_N(z) + \left[\frac{c_m\beta + \epsilon(\eta-1)}{1-c_m} + \lambda_f^+ c_m^{N_m} \right] \|z\|^2,$$

where c_m is defined in (21) and $N_m = \lfloor N/(m-1) \rfloor$.

Proof: Denote by $\hat{x}(\cdot)$ the closed-loop trajectory generated by the policy $[\pi_\infty^{\epsilon, m}]_N$. Let $\tilde{\pi}$ and z_i be the same as in the proof of Theorem 5. Then by (19),

$$\begin{aligned} J_{[\pi_\infty^{\epsilon, m}]_N}(z) - \psi(\hat{x}(N)) &\leq \sum_{i=0}^{N_m+1} J_{\tilde{\pi}}(z_i) \\ &\leq \sum_{i=0}^{N_m+1} J_{\hat{\pi}_m^\epsilon}(z_i) \leq \sum_{i=0}^{N_m+1} [V_m^0(z_i) + \epsilon(\eta-1)\|z_i\|^2] \end{aligned}$$

Notice that $V_m^0(z) \leq V_N(z)$, $V_m^0(z_i) \leq V^*(z_i)$ and $V^*(z_i) \leq \beta \|z_i\|^2 \leq \beta c_m^i \|z\|^2$, by adding some small positive terms, we have

$$J_{[\pi_\infty^{\epsilon,m}]_N}(z) - \psi(\hat{x}(N)) \leq V_N(z) + \sum_{i=1}^{\infty} \beta c_m^i \|z\|^2 + \sum_{i=0}^{\infty} \epsilon(\eta-1)c_m \|z\|^2. \quad (24)$$

By our hypotheses, we have $c_m < 1$. Thus, $J_{[\pi_\infty^{\epsilon,m}]_N}(z) - \psi(\hat{x}(N)) \leq V_N(z) + \frac{c_m \beta + \epsilon(\eta-1)}{1-c_m} \|z\|^2$. Considering $\psi(\hat{x}(N)) \leq \lambda_f^+ \|\hat{x}(N)\|^2 \leq \lambda_f^+ c_m^N \|z\|^2$, the desired result is proved. ■

Corollary 3: Suppose the conditions in Theorem 4 hold. For any $\delta > 0$, if we further have $\epsilon < \frac{\delta}{\eta-1}$ and

$$N \geq m > m_N^{\delta,\epsilon} \triangleq \frac{\ln[\delta - \epsilon(\eta-1)]\lambda_Q^- - \ln(\beta + \delta + \sigma_f^+)(\beta + \epsilon\eta)}{\ln(\beta\gamma + \epsilon\eta) - \ln\beta} + 1, \quad (25)$$

then the N -horizon policy $[\pi_\infty^{\epsilon,m}]_N$ is δ -suboptimal over the unit all.

Motivated by the above analysis, for a large or infinite N , a δ -suboptimal N -horizon policy can be obtained as follows. First, find the largest ϵ that satisfies all the conditions in Corollary 2. Second, let $m = m_\infty^{\delta,\epsilon}$ or $m = m_N^{\delta,\epsilon}$ depending on whether N is infinite or not. Third, compute the m -horizon suboptimal policy $\hat{\pi}_m^\epsilon$ using Algorithm 2 with $Q_f = 0$. Finally, use $\hat{\pi}_m^\epsilon$ to construct $\pi_\infty^{\epsilon,m}$ based on (20) and keep the first N steps of $\pi_\infty^{\epsilon,m}$ to obtain an N -horizon policy $[\pi_\infty^{\epsilon,m}]_N$. By Corollary 2 or 3, $[\pi_\infty^{\epsilon,m}]_N$ is guaranteed to be δ -suboptimal over the unit ball. The above procedure of constructing the suboptimal control policy is summarized in Algorithm 3. Note that in this procedure, we have assumed that $N > m$. If this is not the case, we should still use Algorithm 2 to carry out the approximate iterations (10) for the whole horizon N .

Algorithm 3 (Large or infinite Horizon Suboptimal Control)

- 1) **Initialization:** Specify an error tolerance δ . Let $\epsilon = \max\{\frac{\delta}{\eta-1}, \frac{\beta(1-\gamma)}{\gamma\eta}\}$.
- 2) **# of iterations steps:** If $N = \infty$, let $m = m_\infty^{\delta,\epsilon}$; otherwise, let $m = m_N^{\delta,\epsilon}$. If $N \leq m$, stop and turn to Algorithm 2.
- 3) **m-horizon Policy:** Calculate the m -horizon suboptimal policy $\hat{\pi}_m^\epsilon$ using Algorithm 2 with $Q_f = 0$.
- 4) **Horizon Extension:** Construct $\pi_\infty^{\epsilon,m}$ from $\hat{\pi}_m^\epsilon$ using (20) and keep its first N terms to obtain $[\pi_\infty^{\epsilon,m}]_N$.

Remark 3: The analytical bounds $m_\infty^{\delta,\epsilon}$ and $m_N^{\delta,\epsilon}$ derived in (23) and (25) may be conservative for some applications. An alternative approach is to start from a smaller value for m in Step 2) of Algorithm 3 and gradually increase its value until the performance saturates. Our analysis guarantees that this tentative procedure can eventually reach any prespecified suboptimal performance by gradually increasing m .




Fig. 3. Convergence of the Value function of Example VI-A.

VI. NUMERICAL EXAMPLES

A. Example 1

First consider a simple infinite-horizon DSLQR problem with the following 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 3 is applied to solve this DSLQR problem. With $\delta = 10^{-3}$, the upper bound of the required number of iterations is $m_\infty^{\delta,\epsilon} = 51$, while as observed in the simulation, the value function already converges in 6 steps. Since $V_k(z)$ is homogeneous and symmetric, in Fig. 3, we plot the evolution of the value functions on the upper half of the unit circle, i.e. the points of the form $z(\theta) = [\cos(\theta), \sin(\theta)]^T$ with $\theta \in [0, \pi]$. The number of matrices in \mathcal{H}_k^ϵ at each step k is listed in Table I. It can be seen that $|\mathcal{H}_k^\epsilon|$ is indeed very small and stays at the maximum value 5 as opposed to growing exponentially as k increases.

TABLE I

$|\mathcal{H}_k^\epsilon|$ FOR EXAMPLE VI-A

k	1	2	3	4	5	6
$ \mathcal{H}_k^\epsilon $	2	4	5	5	5	5

The policy $\pi_\infty^{\epsilon,6}$ can be constructed using (20) based on the sets $\{\mathcal{H}_k^\epsilon\}_{k=1}^5$. The the Euclidean norms of three closed-loop trajectories starting from the same initial state $x(0) = [1, 1]^T$ driven by the policy $\pi_\infty^{\epsilon,6}$ and the two LQR controllers designed for the two individual subsystems are shown in Figure 4. It is clear that $\pi_\infty^{\epsilon,6}$ is stabilizing and results in a better convergence speed of the closed-loop trajectory than the two LQR controllers do.

B. Random Examples

To further demonstrate its effectiveness, Algorithm 3 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.




Fig. 4. Euclidean norms of the closed-loop trajectories with initial condition $x(0) = [1, 1]^T$ driven by different feedback policies.

For both sets, the control horizon N is infinite and $\delta = 10^{-3}$. All of these problems are successfully solved by Algorithm 3 and the distributions of the complexity, namely, the maximum numbers of matrices required for characterizing the suboptimal policy, are plotted in Fig. 5. It can be seen from the figure that all 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 four-dimensional problems need about 40 matrices and some of them may need more than 100 matrices. The complexity of Algorithm 3 depends heavily on the state dimension. In a higher dimensional state space, a larger δ is usually needed in order to retain a high computational speed.

VII. CONCLUSION

The value functions and the optimal strategies of the DSLQR problem can be exactly characterized by the SRSs \mathcal{H}_k , whose size grows exponentially fast. However, with some small relaxation, many matrices in \mathcal{H}_k becomes redundant in terms of characterizing the suboptimal strategies. Efficient algorithms are developed to prune out these redundant matrices and to compute the suboptimal strategies. Analytical conditions under which the strategies generated by the algorithms are stabilizing and suboptimal are derived explicitly in terms of subsystem matrices. The proposed algorithm and the analysis provide a systematical way of solving the finite-horizon and infinite-horizon DSLQR problem with guaranteed closed-loop stability and suboptimality. The results of this paper can be used to study many other problems of the switched linear systems, such as the switched Kalman filtering problem, the switched LQG problem, and the switched receding horizon control problem, etc. All of these will be our future research directions.

REFERENCES

- [1] X. Xu and P.J. Antsaklis. Optimal control of switched systems based on parameterization of the switching instants. *IEEE Transactions on Automatic Control*, 49(1):2–16, 2004.
- [2] M. Egerstedt, Y. Wardi, and F. Delmotte. Optimal control of switching times in switched dynamical systems. *IEEE Transactions on Automatic Control*, 51(1):110–115, 2006.
- [3] C. Seatzu, D. Corona, A. Giua, and A. Bemporad. Optimal control of continuous-time switched affine systems. *IEEE Transactions on Automatic Control*, 51:726–741, May 2006.

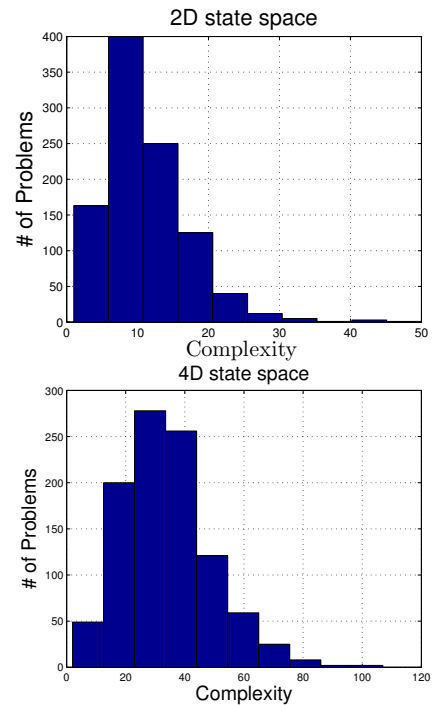


Fig. 5. Complexity distributions of the random examples.

- [4] Sven Hedlund and Anders Rantzer. Convex dynamic programming for hybrid systems. *IEEE Transactions on Automatic Control*, 47(9):1536–1540, sep 2002.
- [5] S. Hedlund and A. Rantzer. Optimal control of hybrid system. In *Proceedings of the IEEE Conference on Decision and Control*, volume 4, pages 3972–3977, Phoenix, AZ, December 1999.
- [6] A. Rantzer and M. Johansson. Piecewise linear quadratic optimal control. *IEEE Transactions on Automatic Control*, Apr. 2000.
- [7] A. Bemporad, A. Giua, and C. Seatzu. Synthesis of state-feedback optimal controllers for continuous-time switched linear systems. In *Proceedings of the IEEE Conference on Decision and Control*, pages 3182–3187, 2002.
- [8] 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, October 2005.
- [9] Wei Zhang and Jianghai Hu. On Optimal Quadratic Regulation for Discrete-Time Switched Linear Systems. In *International Workshop on Hybrid Systems: Computation and Control*, pages 584–597, St Louis, MO, USA, 2008.
- [10] Wei Zhang and Jianghai Hu. On the value functions of the optimal quadratic regulation problem for discrete-time switched linear systems. In *IEEE Conference on Decision and Control*, Cancun, Mexico, December 2008.
- [11] D. P. Bertsekas and J. Tsitsiklis. *Neuro-Dynamic Programming*. Athena Scientific, Sep. 1996.
- [12] B. Lincoln and A. Rantzer. Relaxing dynamic programming. *IEEE Transactions on Automatic Control*, 51(8):1249–1260, August 2006.
- [13] D.P. Bertsekas. *Dynamic Programming and Optimal Control*, volume 2. Athena Scientific, 2 edition, 2001.
- [14] S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan. *Linear Matrix Inequalities in System and Control Theory*. SIAM, Philadelphia, PA, Jun 1994.
- [15] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, New York, NY, USA, 2004.
- [16] W. Zhang and J. Hu. On discrete-time switched LQR problems. Technical report, Purdue University, 2009. TR ECE 09-3.