

Mode Reduction for Markov Jump Systems

ZHE DU  (Student Member, IEEE), LAURA BALZANO  (Senior Member, IEEE),
AND NECMIYE OZAY  (Senior Member, IEEE)

Department of Electrical Engineering and Computer Science, University of Michigan Ann Arbor, Ann Arbor, MI 48109 USA

CORRESPONDING AUTHOR: Zhe Du (e-mail: zhedu@umich.edu)

The work of Zhe Du and Necmiye Ozay was supported by ONR under Grants N00014-18-1-2501 and N00014-21-1-2431. The work of Zhe Du and Laura Balzano was supported in part by AFOSR YIP under Award FA9550-19-1-0026 and in part by NSF CAREER Award under Grant CCF-1845076. The work of Laura Balzano was supported by NSF BIGDATA Award under Grant IIS-1838179.

ABSTRACT Switched systems are capable of modeling processes with underlying dynamics that may change abruptly over time. To achieve accurate modeling in practice, one may need a large number of modes, but this may in turn increase the model complexity drastically. Existing work on reducing system complexity mainly considers state space reduction, whereas reducing the number of modes is less studied. In this work, we consider Markov jump linear systems (MJSs), a special class of switched systems where the active mode switches according to a Markov chain, and several issues associated with its mode complexity. Specifically, inspired by clustering techniques from unsupervised learning, we are able to construct a reduced MJS with fewer modes that approximates the original MJS well under various metrics. Furthermore, both theoretically and empirically, we show how one can use the reduced MJS to analyze stability and design controllers with significant reduction in computational cost while achieving guaranteed accuracy.

INDEX TERMS Clustering, learning for control, Markov jump systems, model/controller reduction.

I. INTRODUCTION

As the control and machine learning communities build tools to model ever more complex dynamical systems, it will become increasingly important to identify redundant aspects of a model and remove them using various unsupervised learning techniques. State dimensionality reduction has long been common in control systems, using principal component analysis and similar techniques. In this paper we consider the setting where switched systems have redundant modes, and we apply clustering – another fundamental unsupervised learning technique – to remove redundancies.

Switched systems generalize time-invariant systems by allowing the dynamics to switch over time. They have been used to model abrupt changes in the environment (e.g. weather and road surfaces), controlled plants (e.g. functioning statuses of different components), disturbances, or even control goals (e.g. cost functions in the optimal control). Switched system models are used in a variety of applications including controlling a Mars rover exploring an unknown heterogeneous terrain, solar power generation, investments in financial markets, and communications with packet losses [1], [2], [3], [4], [5], [6], [7]. However, these benefits are accompanied by new

complexity challenges: the number of modes that is needed to model systems accurately and thoroughly may grow undesirably large. For example, for controlled plants composed of multiple components, if we model each combination of health statuses, e.g. working and faulty, of all components as a mode, then the number of modes grows exponentially with the number of components. Given this rate, there can be an huge amount of modes even with a moderate number of components. Given a system with this many modes, analysis can become computationally intractable. For example, in finite horizon linear quadratic regulator (LQR) problems with horizon T , the total number of controllers to be computed is s^T where s denotes the number of modes. This lack of scalability calls for systematic and theoretically guaranteed ways to reduce the number of modes.

Existing work on (switched) system reduction mainly focuses on reducing the state dimension [8] or constructing finite abstractions for the continuous state space [9]. Reducing the mode complexity, however, is still mainly an uncharted territory.

In this work, we study how one can perform mode reduction for Markov jump linear systems (MJS), a class of switched

systems with individual modes given by state space linear time-invariant (LTI) models and the mode switching governed by Markov chains. Our main contributions are the following:

- We propose a clustering-based method that takes mode dynamics as features and use the estimated clusters to construct a mode-reduced MJS.
- The reduced MJS is proved to well approximate the original MJS under several approximation metrics.
- We show the reduced MJS can be used as a surrogate for the original MJS to analyze stability and design controllers with guaranteed performance and significant reduction in computational cost.
- Compared to the preliminary conference version [10], this paper establishes stronger approximation guarantees (Section V), provides novel stability analysis (Section VI), includes all the proofs that were omitted in [10], and includes a more thorough discussion and survey of related work, including a table that lays out relationships among existing work and our work.

Our work adds a new dimension, i.e., reduction of modes, to the research of switched system reduction. This framework can be generalized to other problems such as robust and optimal control, invariance analysis, partially observed systems, etc. Other than constructing and analyzing the reduced MJS, the technical tools we develop in this work regarding perturbations can be applied to cases when there are model mismatches, e.g. system estimation errors incurred when dynamics are learned in identification or data-driven adaptive control as in [11].

II. RELATED WORK

Depending on the problems of interest and methodologies, the work on reduction for stochastic (switched) systems can be roughly divided into three categories: bisimulation, symbolic abstraction, and order reduction.

Bisimulation: To evaluate the equivalency between two stochastic switched systems, notions of (approximate) probabilistic bisimulation are proposed in [12], [27], [28]. Approximation metrics [29] from different perspectives are developed to compare two systems, e.g. one(multi)-step transition kernels [19] and trajectories [13], [22], [30]. Based on the approximate bisimulation notion in [14], a technique for reducing the state space of labeled Markov chains through state aggregation is proposed in [15]. Unlike existing work that typically defines the notions of (approximate) bisimulation on the state space, we provide an algorithm that constructs a reduced system by aggregating the mode space, which provably approximates the original one. Our work shares the idea of aggregation of Markov chains with [15], but we also seek to recover the best aggregation partition which is otherwise assumed as prior knowledge in [15].

Symbolic Abstraction: Given a system with continuous state space, abstraction [31] considers discretizing the state space and then constructing a finite state symbolic model, which can be used as a surrogate for model verification [32], [33] or controller synthesis [34]. The work on abstraction for

stochastic hybrid systems starts with the autonomous cases. Under uniform discretization, [35] and [19] provide approximation guarantees that depend on the discretization width. An adaptive partition scheme is proposed in [20], which mitigates the curse of dimensionality suffered by uniform sampling. Since the systems under consideration are autonomous, these work mainly serves verification purposes, but falls short toward controller synthesis goals. [9] addressed this by allowing inputs in the systems. The idea of partitioning the continuous state space is similar to our work except that our partition is performed on the mode space, i.e., the discrete state space in hybrid systems, which provides a new yet closely related dimension to the existing abstraction work.

Order Reduction: Another important line of research on system reduction is order reduction [36], where one seeks to reduce the dimension of the state space to satisfy certain criteria. With the help of linear matrix inequalities (LMIs), various methods have been applied for MJS, including \mathcal{H}_∞ reduction [8], balanced truncation [25], and \mathcal{H}_2 reduction [26], etc. Order reduction is also applied to more complex models with time-varying transition probabilities [24].

The reduction of Markov chains, a class of simplified yet fundamental stochastic switched models, has also attracted the learning and statistics communities. Several notions of lumpability are proposed in [37], which coincide with the notion of bisimulation in [27]. Lumpability allows one to reduce the original Markov chain to a smaller scale yet equivalent Markov chain by lumping the Markovian states. Similar research focusing on the equivalence metrics and bounding the difference of transition kernel, can be found in [38], [39], [40] under the name of coarse graining. Compared with the bisimulation work for general stochastic systems, which is mostly conceptual, the restriction to Markov chains allows for practical “low-rank + clustering” methods [41] to uncover the lumpability structure. [16] considers the case when the Markov matrix is estimated from a trajectory, and the approximate lumpability case are studied in [17], [18]. Furthermore [17] studies the reduction of Markov chains that are embedded in switched autoregressive exogenous models, but the overall dynamical models are not reduced. Based on the ideas in [17], our work further extends the reduction to the overall MJS.

A comprehensive comparison of the related work together with our work is listed in Table 1. The entry “exact bisimulation condition” tells whether ideal case sufficient conditions are provided under which a system can be reduced without introducing any model inaccuracy, i.e., they are bisimilar. In practice, when the reduced system is constructed, these principled conditions can help gain more insight into the original system. In practice, these system reduction methods developed under different perspectives can be combined to achieve overall better performance. For example, for the continuous state space, one can apply order reduction followed by finite abstraction (the former can help remove the curse of dimensionality for the latter), and meanwhile our work can further help reduce the discrete mode space.

TABLE 1. Related Work on Reduction for Stochastic Switched Systems

Reference	Model	Reduction Target	Exact Bisimulation Condition	Reduction Method	Approximation Metric
[12]	Controlled Markov Process	State Space Cardinality	Yes	N.A.	Formula Metric
[13], [14]	Labelled Markov Chains (autonomous)			Clustering	Trajectory
[15]	Markov Chains			N.A.	N.A.
[16]–[18]	Switching Stochastic Sys.		State Space Discretization	Trajectory	Transition Kernel
[19]	Stochastic Hybrid System (autonomous)		(Bi)simulation Function	Invariance Probability	Invariance Probability
[20]	Jump Linear Stochastic System	N.A.	\mathcal{H}_∞ Reduction	Trajectory	Trajectory
[21]–[23]	MJS		Balanced Truncation	\mathcal{H}_∞ norm	\mathcal{H}_2 norm
[8], [24]			\mathcal{H}_2 Reduction	\mathcal{H}_2 norm	
[25]					
[26]					
This Work	MJS	Modes	Yes	Clustering	Trajectory

This work is organized as follows: we present the preliminaries and mode reduction problem setup in Section III-A; an clustering-based reduction approach is proposed in Section IV; in Section V, we discuss how the reduced MJS approximates the original one under setups; Section VI and VII respectively show that one can use the reduced MJS as a surrogate to evaluate stability and design LQR controllers for the original MJS; simulation experiments are presented in Section VIII.

III. PRELIMINARIES AND PROBLEM SETUP

For a matrix \mathbf{E} , $\mathbf{E}(i, :)$ denotes the i th row of \mathbf{E} , and $\mathbf{E}(i, j:k)$ denotes the i th row preserving only the j th to k th columns. For any index set A , $\mathbf{E}(i, A)$ denotes the i th row preserving columns given by A . Let $\sigma_i(\mathbf{E})$ ($\lambda_i(\mathbf{E})$) denote its i th largest singular (eigen) value. For any $s \in \mathbb{N}$, we let $[s] := \{1, 2, \dots, s\}$. We say $\Omega_{1:r} := \{\Omega_1, \dots, \Omega_r\}$ is a r -cluster partition of $[s]$ if $\bigcup_{i=1}^r \Omega_i = [s]$, $\Omega_i \cap \Omega_j = \emptyset$ for any $i \neq j$, and $\Omega_i \neq \emptyset$. We let $\Omega_{(i)}$ denote the cluster with i th largest cardinality. For a sequence of variables X_0, X_1, \dots, X_N , let $X_{0:N} := \{X_i\}_{i=0}^N$. Notation \otimes denotes the Kronecker product. Notation \mathbf{I}_n denotes the n -dimensional identity matrix, and $\mathbf{1}_n$ denotes the n -dimensional all-ones vector.

A. PRELIMINARIES

In this work, we consider Markov jump systems (MJSs) with dynamics given by

$$\Sigma := \{\mathbf{x}_{t+1} = \mathbf{A}_{\omega_t} \mathbf{x}_t + \mathbf{B}_{\omega_t} \mathbf{u}_t, \omega_t \sim \text{MarkovChain}(\mathbf{T})\} \quad (1)$$

where $\mathbf{x}_t \in \mathbb{R}^n$ and $\mathbf{u}_t \in \mathbb{R}^p$ denote the state and input at time t . There are s modes parameterized by $\{\mathbf{A}_i, \mathbf{B}_i\}_{i=1}^s$ where $\mathbf{A}_i \in \mathbb{R}^{n \times n}$ and $\mathbf{B}_i \in \mathbb{R}^{n \times p}$ are state and input matrices for mode i . The active mode at time t is indexed by $\omega_t \in [s]$, and the mode switching sequence $\omega_{0:t}$ follows a Markov chain with Markov matrix $\mathbf{T} \in \mathbb{R}^{s \times s}$, i.e., $\mathbb{P}(\omega_{t+1} = j | \omega_t = i) = \mathbf{T}(i, j)$. We assume the Markov chain \mathbf{T} is ergodic. By properties of ergodicity, \mathbf{T} has a unique stationary distribution

$\pi \in \mathbb{R}^s$, and we let π_{\max} and π_{\min} denote the largest and smallest element in π . In the remainder of the paper, we use $\Sigma := \text{MJS}(\mathbf{A}_{1:s}, \mathbf{B}_{1:s}, \mathbf{T})$ to denote the groundtruth MJS in (1) that we want to study, and similarly use notation $\text{MJS}(\cdot, \cdot, \cdot)$ to parameterize any MJS with expressions given in (1). We introduce the following two special types of Markov chains, which are closely tied to the main focus of this work.

Definition 1 (Lumpability and Aggregatability [37]): Markov matrix $\mathbf{T} \in \mathbb{R}^{s \times s}$ is *lumpable* w.r.t. partition $\Omega_{1:r}$ on $[s]$ if for any $k, l \in [r]$, and $i, i' \in \Omega_k$, we have $\sum_{j \in \Omega_l} \mathbf{T}(i, j) = \sum_{j \in \Omega_l} \mathbf{T}(i', j)$. As a special case, it is further *aggregatable* if $\mathbf{T}(i, :) = \mathbf{T}(i', :)$.

Lumpability of a Markov chain coincides with the definition of probabilistic bisimulation in [28], which describes an equivalence relation on $[s]$, i.e., two members are equivalent if they belong to the same cluster. For a Markov chain \mathbf{T} that is lumpable with respect to partition $\Omega_{1:s}$, we use $\zeta_t \in [r]$ to index the active cluster at time t , i.e., $\zeta_t = k$ if and only if $\omega_t \in \Omega_k$, and use $\zeta_{0:t}$ to denote the active cluster sequence.

B. PROBLEM FORMULATION

With the notions of Markov chain lumpability and aggregatability, in this work, we consider reducing the number of modes for Σ under the following two problem settings.

Problem P1 (Lumpable Case): Assume the dynamics of $\Sigma = \text{MJS}(\mathbf{A}_{1:s}, \mathbf{B}_{1:s}, \mathbf{T})$ is known. Suppose there exists a hidden partition $\Omega_{1:r}$ on $[s]$ and $\epsilon_A, \epsilon_B, \epsilon_T \geq 0$ such that

$$\begin{aligned} \sum_{k \in [r]} \sum_{i, i' \in \Omega_k} \|\mathbf{A}_i - \mathbf{A}_{i'}\|_F &\leq \epsilon_A, \\ \sum_{k \in [r]} \sum_{i, i' \in \Omega_k} \|\mathbf{B}_i - \mathbf{B}_{i'}\|_F &\leq \epsilon_B, \\ \sum_{k, l \in [r]} \sum_{i, i' \in \Omega_k} \left| \sum_{j \in \Omega_l} \mathbf{T}(i, j) - \sum_{j \in \Omega_l} \mathbf{T}(i', j) \right| &\leq \epsilon_T. \end{aligned} \quad (3)$$

Then, given $\{\mathbf{A}_i, \mathbf{B}_i\}_{i=1}^s$, \mathbf{T} and r , we seek to estimate the partition $\Omega_{1:r}$ by clustering the modes, construct a reduced

MJS for Σ , and provide guarantees on the behavior difference incurred by the reduction.

Throughout this work, we refer to ϵ_A , ϵ_B , and ϵ_T as *perturbations*. When $\epsilon_T = 0$, by definition \mathbf{T} is lumpable. Thus, in condition (3), one can view \mathbf{T} as *approximately* lumpable. One can think of $\{\omega_t, \mathbf{x}_t\} \in [s] \times \mathbb{R}^n$ as a hybrid state [19]. Then, condition (3) guarantees the existence of an approximate equivalence relation in the discrete domain $[s]$, while condition (2) guarantees this in the continuous domain \mathbb{R}^n . More discussions for the special case when $\epsilon_A = \epsilon_B = \epsilon_T = 0$ follow in the next subsection.

For the aggregatable case, we separately formulate a similar problem in Problem P2.

Problem P2 (Aggregatable Case): In Problem P1, replace (3) with $\sum_{k \in [r]} \sum_{i, i' \in \Omega_k} \|\mathbf{T}(i, :)^\top - \mathbf{T}(i', :)^\top\|_1 \leq \epsilon_T$.

P2 differs from P1 in terms of ϵ_T , which quantifies the violation of lumpability and aggregatability respectively for the matrix \mathbf{T} . Consider a 3-state Markov matrix \mathbf{T} with rows $\mathbf{T}(1, :) = [0.2, 0.4, 0.4]$, $\mathbf{T}(2, :) = [0.7, 0.1, 0.2]$, and $\mathbf{T}(3, :) = [0.7, 0, 0.3]$. Then for the partition $\Omega_1 = \{1\}$, $\Omega_2 = \{2, 3\}$, we obtain $\epsilon_T = 0$ in P1 but $\epsilon_T = 0.4$ in P2. In other words, \mathbf{T} is exactly lumpable but non-aggregatable with a violation level of 0.4. In P2, $\epsilon_T = 0$ only when $\mathbf{T}(i, :) = \mathbf{T}(i', :)$ for all $i, i' \in \Omega_k$, i.e. the rows are equal. On the other hand, in P1, $\epsilon_T = 0$ is possible even if no row equalities exist. Hence, being ϵ_T -aggregatable in P2 is a stronger assumption than being ϵ_T -lumpable in P1. As a result, in Section IV, the clustering guarantee for P2 is stronger and more interpretable than that of P1.

C. EQUIVALENCE BETWEEN MJSs

To compare the original and mode-reduced MJSs as mentioned in P1, we need a notion of equivalency between two MJSs with different numbers of modes. This is provided below via a surjection from modes of the larger MJS to the smaller one, which extends the bijection idea in [8], [22] that can only compare two MJSs with equal amounts of modes.

Definition 2 (Equivalency between MJSs): Consider two MJSs Σ_1 and Σ_2 with the same state and input dimensions n, p , but different number of modes s_1 and s_2 respectively. WLOG, assume $s_1 > s_2$. Let $\{\mathbf{x}_t^{(1)}, \mathbf{u}_t^{(1)}, \omega_t^{(1)}\}$ and $\{\mathbf{x}_t^{(2)}, \mathbf{u}_t^{(2)}, \omega_t^{(2)}\}$ denote their respective state, input, and mode index. Σ_1 and Σ_2 are equivalent if there exists a partition $\Omega_{1:s_2}$ on $[s_1]$ such that Σ_1 and Σ_2 have the same transition kernels, i.e. for any time t , any mode $k, k' \in [s_2]$, any $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^n$, and any $\mathbf{u} \in \mathbb{R}^p$

$$\begin{aligned} & \mathbb{P}\left(\omega_{t+1}^{(1)} \in \Omega_{k'}, \mathbf{x}_{t+1}^{(1)} = \mathbf{x}' \mid \omega_t^{(1)} \in \Omega_k, \mathbf{x}_t^{(1)} = \mathbf{x}, \mathbf{u}_t^{(1)} = \mathbf{u}\right) \\ &= \mathbb{P}\left(\omega_{t+1}^{(2)} = k', \mathbf{x}_{t+1}^{(2)} = \mathbf{x}' \mid \omega_t^{(2)} = k, \mathbf{x}_t^{(2)} = \mathbf{x}, \mathbf{u}_t^{(2)} = \mathbf{u}\right). \end{aligned} \quad (4)$$

The trivial perturbation-free case, i.e., $\epsilon_A, \epsilon_B, \epsilon_T = 0$, provides a sufficient condition that guarantees that an MJS can be reduced to a smaller MJS with equivalency between them.

Definition 3 (Mode-reducibility Condition): If in P1, $\epsilon_A, \epsilon_B, \epsilon_T = 0$, we say Σ is *mode-reducible* with respect to $\Omega_{1:r}$.

If this condition holds for Σ , we can construct a mode-reduced MJS $\check{\Sigma} := \text{MJS}(\check{\mathbf{A}}_{1:r}, \check{\mathbf{B}}_{1:r}, \check{\mathbf{T}})$ such that for any $k, l \in [r]$, any $i \in \Omega_k$, $\check{\mathbf{A}}_k = \mathbf{A}_i$, $\check{\mathbf{B}}_k = \mathbf{B}_i$, and $\check{\mathbf{T}} \in \mathbb{R}^{r \times r}$ with $\check{\mathbf{T}}(k, l) = \sum_{j \in \Omega_l} \mathbf{T}(i, j)$, which is illustrated in Fig. 1. Let $\{\check{\mathbf{x}}_t, \check{\mathbf{u}}_t, \check{\omega}_t\}$ denote the state, input, and mode index for the reduced $\check{\Sigma}$. Then, the following fact shows that $\check{\Sigma}$ and Σ are equivalent according to Definition 2.

Fact 1: Suppose Σ is mode-reducible and $\check{\Sigma}$ is constructed as above. Consider the case when Σ and $\check{\Sigma}$ have (i) initial mode distributions satisfy $\mathbb{P}(\omega_0 \in \Omega_k) = \mathbb{P}(\check{\omega}_0 = k)$ for all $k \in [r]$, (ii) the same initial states ($\mathbf{x}_0 = \check{\mathbf{x}}_0$), and (iii) the same input sequences ($\mathbf{u}_{0:t-1} = \check{\mathbf{u}}_{0:t-1}$). Then, these two MJSs have the same mode and state transition kernels, i.e. $\mathbb{P}(\omega_t \in \Omega_k, \mathbf{x}_t = \mathbf{x}) = \mathbb{P}(\check{\omega}_t = k, \check{\mathbf{x}}_t = \mathbf{x})$ for all t , all $k \in [r]$ and $\mathbf{x} \in \mathbb{R}^n$. Particularly, there exists a special type of reduced $\check{\Sigma}$ such that the modes are synchronized: for all t , $\check{\omega}_t = \zeta_t$. In this case, $\check{\mathbf{x}}_t = \mathbf{x}_t$ for all t .

Fact 1 first shows the equivalency between Σ and $\check{\Sigma}$ in terms of the transition kernels, which is then extended to trajectory realizations if certain synchrony exists between $\zeta_{0:t}$ and $\check{\omega}_{0:t}$. The condition $\check{\omega}_t = \zeta_t$ in Fact 1 essentially establishes a coupling between the Markov chains $\omega_{0:t}$ and $\check{\omega}_{0:t}$ such that $\mathbb{P}(\omega_t \in \Omega_k, \check{\omega}_t = k) = \mathbb{P}(\omega_t \in \Omega_k) = \mathbb{P}(\check{\omega}_t = k)$. Establishing coupling between the stochastic systems usually allows for stronger equivalency and approximation result. Similar coupling scheme is implicitly used in [8], [22]; an optimal coupling by minimizing Wasserstein distance is discussed in [13]; and a weaker coupling using the idea of HMM is discussed in [24].

In Definition 2 and Fact 1, one can view $\{\omega_t, \mathbf{x}_t\} \in [s] \times \mathbb{R}^n$ as a hybrid state [19]. \mathbf{T} being lumpable guarantees the existence of an equivalence relation in the discrete domain $[s]$ as in Definition 1, while state/input matrices being the same guarantees this in the continuous domain \mathbb{R}^n .

IV. CLUSTERING-BASED MODE REDUCTION FOR MJS

In this section, we first propose Algorithm 1 to estimate the latent partition $\Omega_{1:r}$ and construct the reduced MJS for Problem P1 and P2, and then provide its theoretical guarantees for partition estimation in Section IV-A.

We treat the estimation of partition $\Omega_{1:r}$ essentially as a mode clustering problem with the dynamics matrices \mathbf{A}_i , \mathbf{B}_i and transition distribution $\mathbf{T}(i, :)$ serving as features for mode i . In Algorithm 1, we first construct the feature matrix Φ from Line 2 to Line 8, with $\Phi(i, :)$ denoting the features of mode i . For the aggregatable case in Problem P2, we simply stack the vectorized \mathbf{A}_i , \mathbf{B}_i and $\mathbf{T}(i, :)$, and use $\alpha_A, \alpha_B, \alpha_T$ to denote their weights respectively. One way to choose these weights is as a normalization, e.g. $\alpha_A \propto 1/\max_i \|\mathbf{A}_i\|$, so that these three features would have the same scales. Though in the aggregatable case P2, similarities among the rows of \mathbf{T} shed light on the groundtruth partition $\Omega_{1:r}$, this is no longer valid in the lumpable case P1 as two modes belonging to the same cluster

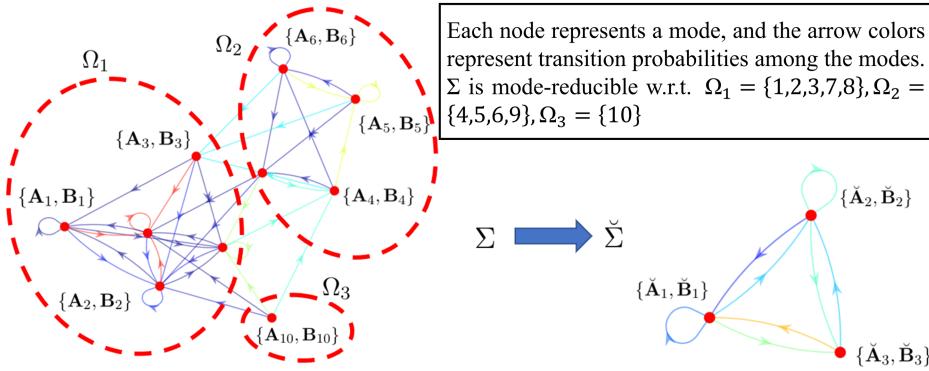


FIGURE 1. Illustration of reduction under mode-reducibility condition.

Algorithm 1: System Reduction for MJS.

Input: $\mathbf{A}_{1:s}$, $\mathbf{B}_{1:s}$, \mathbf{T} , π , r , and non-negative tuning weights $\alpha_{\mathbf{A}}$, $\alpha_{\mathbf{B}}$, $\alpha_{\mathbf{T}}$ that sum to 1

1. Construct feature matrix Φ : $\forall i \in [s]$,
2. **case Problem P2 do**
3. $\Phi(i, :) = [\alpha_{\mathbf{A}} \text{vec}(\mathbf{A}_i)^T, \alpha_{\mathbf{B}} \text{vec}(\mathbf{B}_i)^T, \alpha_{\mathbf{T}} \mathbf{T}(i, :)]$.
4. **case Problem P1 do**
5. $\mathbf{H} = \text{diag}(\pi)^{\frac{1}{2}} \mathbf{T} \text{diag}(\pi)^{-\frac{1}{2}}$ **do**
6. $\mathbf{W}_r \leftarrow$ top r left singular vectors of \mathbf{H}
7. $\mathbf{S}_r = \text{diag}(\pi)^{-\frac{1}{2}} \mathbf{W}_r$
8. $\Phi(i, :) = [\alpha_{\mathbf{A}} \text{vec}(\mathbf{A}_i)^T, \alpha_{\mathbf{B}} \text{vec}(\mathbf{B}_i)^T, \alpha_{\mathbf{T}} \mathbf{S}_r(i, :)]$
9. $\mathbf{U}_r \leftarrow$ top r left singular vectors of Φ
10. Solve k-means problem: $\hat{\Omega}_{1:r}, \hat{\mathbf{c}}_{1:r} = \arg \min_{\hat{\Omega}_{1:r}, \hat{\mathbf{c}}_{1:r}} \sum_{k \in [r]} \sum_{i \in \hat{\Omega}_k} \|\mathbf{U}_r(i, :) - \hat{\mathbf{c}}_k\|^2$
11. Construct $\hat{\Sigma}$, $\forall k, l \in [r]$
 $\hat{\mathbf{A}}_k = \frac{1}{|\hat{\Omega}_k|} \sum_{i \in \hat{\Omega}_k} \mathbf{A}_i$, $\hat{\mathbf{B}}_k = \frac{1}{|\hat{\Omega}_k|} \sum_{i \in \hat{\Omega}_k} \mathbf{B}_i$,
 $\hat{\mathbf{T}}(k, l) = \frac{1}{|\hat{\Omega}_k|} \sum_{i \in \hat{\Omega}_k, j \in \hat{\Omega}_l} \mathbf{T}(i, j)$

Output: $\hat{\Sigma} : \text{MJS}(\hat{\mathbf{A}}_{1:r}, \hat{\mathbf{B}}_{1:r}, \hat{\mathbf{T}})$

can still have different transition probabilities $\mathbf{T}(i, :)$, even if $\epsilon_{\mathbf{T}} = 0$. According to (3), the groundtruth partition $\Omega_{1:r}$ is only embodied in the mode-to-cluster transition probabilities $\sum_{j \in \Omega_l} \mathbf{T}(i, j)$ constructed using the groundtruth partition itself. This leaves us in a “chicken-and-egg” dilemma. To deal with this, from Line 5 to 8, we compute the first r left singular vectors \mathbf{W}_r of matrix $\text{diag}(\pi)^{\frac{1}{2}} \mathbf{T} \text{diag}(\pi)^{-\frac{1}{2}}$, and then weight it by $\text{diag}(\pi)^{-\frac{1}{2}}$ to obtain matrix $\mathbf{S}_r \in \mathbb{R}^{s \times r}$, which is used to construct features in Φ for the lumpable case P1. We will later justify using \mathbf{S}_r as features by showing row similarities in \mathbf{S}_r reflect the partition under certain assumptions.

With the feature matrix Φ , to recover the partition, we resort to k-means: in Line 10, k-means is applied to the first r left singular vector \mathbf{U}_r of Φ . The typical algorithm for k-means is Lloyd’s algorithm, where the cluster centers and partition membership are updated alternately. Based on the solution $\hat{\Omega}_{1:r}$ obtained via k-means, we construct the reduced $\hat{\Sigma}$ by averaging modes within the same estimated cluster. A more

subtle averaging scheme is through the weights provided in the stationary distribution π which describes the frequency of each mode being active in the long run. $\hat{\Sigma}$ generated by this scheme (or any weighted averaging) would have the same performance guarantees as the uniform averaging, which is provided in Section IV-A.

In practice, if we have no good prior knowledge which model of Problem P1 and P2 would yield the best system reduction performance, we can first obtain the partitions for both cases and then pick the one yields smaller hindsight perturbations $\epsilon_{\mathbf{A}}$, $\epsilon_{\mathbf{B}}$, $\epsilon_{\mathbf{T}}$ in Problem P1 and P2. When one picks $\alpha_{\mathbf{A}} = \alpha_{\mathbf{B}} = 0$, i.e only the Markov matrix \mathbf{T} is used to cluster the modes, then our clustering scheme under the aggregatable case P2 is equivalent to [16] which studies clustering for Markov matrices that is estimated from a single trajectory. The lumpable case P1, on the other hand, is based on preliminary analysis in [41].

We note that several aspects of this algorithm we have guarantees for do not directly consider the metrics important to this problem; for example averaging dynamics matrices within the same cluster may not yield the dynamics that gives an optimal fit for prediction or controller design. That said, even for this straightforward approach, Section V provides several strong approximation guarantees. We are hopeful that future generalizations will be able to build on this theory and further improve the control performance of our mode-reduction approach.

A. THEORETICAL GUARANTEES FOR CLUSTERING

In this section, we discuss the clustering performance by comparing the estimated partition $\hat{\Omega}_{1:r}$ and the true $\Omega_{1:r}$. As k-means algorithms are known to have local convergence properties [42], we instead assume for the k-means problem in Algorithm 1, a $(1 + \epsilon)$ approximate solution can be obtained, i.e., $\sum_{k \in [r], i \in \hat{\Omega}_k} \|\mathbf{U}_r(i, :) - \hat{\mathbf{c}}_k\|^2 \leq (1 + \epsilon) \min_{\Omega'_{1:r}, \mathbf{c}'_{1:r}} \sum_{k \in [r], i \in \Omega'_k} \|\mathbf{U}_r(i, :) - \mathbf{c}'_k\|^2$. Many efficient algorithms have been developed that can provide $(1 + \epsilon)$ approximate solutions. For $\epsilon = 1$, a linear time (in terms of r and s) algorithm is provided in [43]. For smaller ϵ , [44] proposes a linear time

algorithm using random sampling; [45] gives a polynomial time algorithm with computational complexity independent of s . We later show how ϵ affects the overall clustering performance.

To evaluate the performance of partition estimation, we define misclustering rate (MR) as $\text{MR}(\hat{\Omega}_{1:r}) = \min_{h \in \mathcal{H}} \sum_{k \in [r]} \frac{|\{i: i \in \Omega_k, i \notin \hat{\Omega}_{h(k)}\}|}{|\Omega_k|}$, where \mathcal{H} is the set of all bijections from $[r]$ to $[r]$ so that the comparison finds the best cluster label matching. The error metric MR counts the total misclustered modes normalized by the cluster sizes, which implies clustering errors occurring in smaller clusters would yield larger MR.

We define the following averaged feature matrix $\bar{\Phi}$ based on the underlying partition $\Omega_{1:r}$: for all $i \in [s]$ (suppose $i \in \Omega_k$ for some $k \in [r]$), $\bar{\Phi}(i, :) = \frac{1}{|\Omega_k|} \sum_{i' \in \Omega_k} \Phi(i', :)$. By construction, there are up to r unique rows in $\bar{\Phi}$, hence $\text{rank}(\bar{\Phi}) \leq r$. We first present the clustering guarantee for Problem P2, i.e., the aggregatable case.

Theorem 1: Consider Problem P2 and Algorithm 1. Suppose $\hat{\Omega}_{1:r}$ is a $(1 + \epsilon)$ k-means solution. Let $\epsilon_{\text{Agg}}^2 = \alpha_A^2 \epsilon_A^2 + \alpha_B^2 \epsilon_B^2 + \alpha_T^2 \epsilon_T^2$. Then, if $\text{rank}(\bar{\Phi}) = r$ and $\epsilon_{\text{Agg}} \leq \frac{\sigma_r(\bar{\Phi}) \sqrt{|\Omega_{(r)}| + |\Omega_{(1)}|}}{8\sqrt{(2+\epsilon)|\Omega_{(1)}|}}$, we have

$$\text{MR}(\hat{\Omega}_{1:r}) \leq 64(2 + \epsilon) \sigma_r(\bar{\Phi})^{-2} \epsilon_{\text{Agg}}^2. \quad (5)$$

Additionally, if $\epsilon_{\text{Agg}} \leq \frac{\sigma_r(\bar{\Phi})}{8\sqrt{(2+\epsilon)|\Omega_{(1)}|}}$, then $\text{MR}(\hat{\Omega}_{1:r}) = 0$.

The key term ϵ_{Agg} measures how modes within the same cluster differ from each other, i.e., inner-cluster distance. On the other hand, the singular value $\sigma_r(\bar{\Phi})$ measures the differences of modes from different clusters, i.e., inter-cluster distance. This is because when modes belonging to different clusters have similar features, their corresponding rows in the averaged feature matrix $\bar{\Phi}$ will also be similar, which could give small $\sigma_r(\bar{\Phi})$. Particularly, if two different clusters share the same features, then $\text{rank}(\bar{\Phi}) = r - 1$ and $\sigma_r(\bar{\Phi}) = 0$. In the theorem, when the inner-cluster distance is small compared to the inter-cluster distance, the misclustering rate can be bounded by their ratio $\epsilon_{\text{Agg}}/\sigma_r(\bar{\Phi})$. By definition of misclustering rate, the smallest nonzero value it can take is given by $\frac{1}{|\Omega_{(1)}|}$. Therefore, whenever the upper bound in (5) is smaller than $\frac{1}{|\Omega_{(1)}|}$, one can guarantee $\text{MR}(\hat{\Omega}_{1:r}) = 0$, which yields the final claim in Theorem 1.

The clustering guarantee for the lumpable case in Problem P1 is more involved than the aggregatable case. We first provide a few more notions and definitions that can help the exposition. We say a Markov matrix \mathbf{T} is *reversible* if there exists a distribution $\pi \in \mathbb{R}^s$ such that $\pi(i)\mathbf{T}(i, j) = \pi(j)\mathbf{T}(j, i)$ for all $i, j \in [s]$. This condition translates to $\mathbf{diag}(\pi)\mathbf{T} = \mathbf{T}^T \mathbf{diag}(\pi)$ when \mathbf{T} is ergodic with stationary distribution π . For a reversible Markov matrix that is also lumpable, we have the following property.

Lemma 1 (Appendix A in [41]): For a reversible Markov matrix \mathbf{T} that is also lumpable with respect to partition $\Omega_{1:r}$, it is diagonalizable with real eigenvalues. Let $\mathbf{S} \in \mathbb{R}^{s \times s}$ denote

an arbitrary eigenvector matrix of \mathbf{T} . Then, there exists an index set $\mathcal{A} \subseteq [s]$ with $|\mathcal{A}| = r$ such that for all $k \in [r]$, for all $i, i' \in \Omega_k$, we have $\mathbf{S}(i, \mathcal{A}) = \mathbf{S}(i', \mathcal{A})$.

We say \mathbf{T} in Lemma 1 has *informative spectrum* if $\mathcal{A} = [r]$ and $|\lambda_r(\mathbf{T})| > |\lambda_{r+1}(\mathbf{T})|$, which implies that the r eigenvectors that carry partition information in Lemma 1 correspond to the r leading eigenvalues. For lumpable Markov matrices, we define the $\epsilon_{\mathbf{T}}$ -neighborhood of \mathbf{T} :

$$\begin{aligned} \mathcal{L}(\mathbf{T}, \Omega_{1:r}, \epsilon_{\mathbf{T}}) := & \left\{ \mathbf{T}_0 \in \mathbb{R}^{s \times s} : \mathbf{T}_0 \text{ is Markovian,} \right. \\ & \left. \|\mathbf{T}_0 - \mathbf{T}\|_{\infty} \leq \epsilon_{\mathbf{T}}, \|\mathbf{T}_0 - \mathbf{T}\|_{\text{F}} \leq \epsilon_{\mathbf{T}}, \right. \\ & \forall k, l \in [r], \forall i \in \Omega_k, \sum_{j \in \Omega_l} \mathbf{T}_0(i, j) = \frac{1}{|\Omega_k|} \sum_{\substack{i' \in \Omega_k \\ j \in \Omega_l}} \mathbf{T}(i', j) \}. \end{aligned} \quad (6)$$

Under the approximate lumpability condition in (3), one can show this neighborhood set is non-empty. Related discussions are provided in Appendix B. To find such a $\mathbf{T}_0 \in \mathcal{L}(\mathbf{T}, \Omega_{1:r}, \epsilon_{\mathbf{T}})$, one only needs to solve a feasibility linear programming problem. Then we provide the clustering guarantee for the lumpable case.

Theorem 2: Consider Problem P1 and Algorithm 1. Let $\gamma_1 := \sum_{i=2}^s \frac{1}{1 - \lambda_i(\mathbf{T})}$, $\gamma_2 := \min\{\sigma_r(\mathbf{H}) - \sigma_{r+1}(\mathbf{H}), 1\}$, $\gamma_3 := \frac{16\gamma_1 \sqrt{r\pi_{\max} \|\mathbf{T}\|_{\text{F}}}}{\gamma_2 \pi_{\min}^2}$, and $\epsilon_{\text{Lmp}}^2 := \alpha_A^2 \epsilon_A^2 + \alpha_B^2 \epsilon_B^2 + \alpha_T^2 \gamma_3^2 \epsilon_{\mathbf{T}}^2$. Assume there exists an ergodic and reversible $\mathbf{T}_0 \in \mathcal{L}(\mathbf{T}, \Omega_{1:r}, \epsilon_{\mathbf{T}})$ with informative spectrum. Suppose $\hat{\Omega}_{1:r}$ is a $(1 + \epsilon)$ k-means solution. Then, if $\text{rank}(\bar{\Phi}) = r$, $\epsilon_{\mathbf{T}} \leq \frac{\pi_{\min}}{\gamma_1}$, $\epsilon_{\text{Lmp}} \leq \frac{\sigma_r(\bar{\Phi}) \sqrt{|\Omega_{(r)}| + |\Omega_{(1)}|}}{8\sqrt{s(2+\epsilon)|\Omega_{(1)}|}}$, we have

$$\text{MR}(\hat{\Omega}_{1:r}) \leq 64(2 + \epsilon) \sigma_r(\bar{\Phi})^{-2} \epsilon_{\text{Lmp}}^2. \quad (7)$$

Additionally, if $\epsilon_{\text{Lmp}} \leq \frac{\sigma_r(\bar{\Phi})}{8\sqrt{(2+\epsilon)|\Omega_{(1)}|}}$, then $\text{MR}(\hat{\Omega}_{1:r}) = 0$.

Theorem 2 for the lumpable case is similar to Theorem 1 for the aggregatable case with an additional γ_3 term. This is a result of using \mathbf{S}_r and \mathbf{T} to construct features in Algorithm 1 for these two cases. γ_3 describes how much the lumpability perturbation $\epsilon_{\mathbf{T}}$ on \mathbf{T} affects the row equalities of its spectrum-related matrix \mathbf{S}_r in Lemma 1. The assumption on the existence of \mathbf{T}_0 with informative spectrum guarantees (i) the partition $\Omega_{1:r}$ information is carried by the leading eigenvectors of \mathbf{T}_0 as introduced in Lemma 1, and (ii) this information can still be preserved in \mathbf{S}_r as long as \mathbf{T} is close to \mathbf{T}_0 . Because of this, Theorem 2 may not hold for arbitrary lumpable \mathbf{T} , but only those close to Markov matrices with informative spectra.

V. APPROXIMATION GUARANTEES

With perturbations $\epsilon_A, \epsilon_B, \epsilon_{\mathbf{T}}$, the reduced $\hat{\Sigma}$ may not be equivalent to the original Σ as in Fact 1. In this case, if certain approximation guarantees can be established, they can be used in verification tasks such as safety [22] and invariance [20] evaluations. In this section, we show that the reduced system

$\hat{\Sigma}$ can be guaranteed to well approximate the original system Σ under metrics such as transition kernels (distributions) and trajectory realizations. Particularly, these metrics reach 0 when $\epsilon_A, \epsilon_B, \epsilon_T = 0$, i.e., the mode-reducibility condition in Definition 3 holds. We have shown in Theorem 1 and 2 that $\text{MR}(\hat{\Omega}_{1:r}) = 0$ when perturbations $\epsilon_A, \epsilon_B, \epsilon_T$ are small. Hence, in this section together with Section VI and VII, we assume $\Omega_{1:r} = \hat{\Omega}_{1:r}$ for simplicity. In these sections, the theory holds for perturbations $\epsilon_A, \epsilon_B, \epsilon_T$ introduced in either P1 or P2.

We study the approximation in a setup where Σ and $\hat{\Sigma}$ start with the same initial condition and are driven by the same input.

Setup S1 (Initialization-Excitation Setup): Systems Σ and $\hat{\Sigma}$ have (i) initial mode distributions satisfy $\mathbb{P}(\omega_0 \in \hat{\Omega}_k) = \mathbb{P}(\hat{\omega}_0 = k)$ for all $k \in [r]$; (ii) the same initial states, i.e., $\mathbf{x}_0 = \hat{\mathbf{x}}_0$, and (iii) the same inputs $\mathbf{u}_t = \hat{\mathbf{u}}_t$ for all t .

Note that when Σ and $\hat{\Sigma}$ have fixed and shared initial conditions and inputs as in setup S1, we can at most evaluate the difference between \mathbf{x}_t and $\hat{\mathbf{x}}_t$ in terms of their distributions (or, the transition kernels of Σ and $\hat{\Sigma}$). However, the actual realizations of \mathbf{x}_t and $\hat{\mathbf{x}}_t$, i.e., when we only generate a single sample for each, can be very different. This is because \mathbf{x}_t and $\hat{\mathbf{x}}_t$ are driven not only by the input excitation, but also the mode switching sequences $\omega_{0:t-1}$ and $\hat{\omega}_{0:t-1}$; thus, \mathbf{x}_t and $\hat{\mathbf{x}}_t$ are likely to be far away from each other if the realizations of $\omega_{0:t-1}$ and $\hat{\omega}_{0:t-1}$ are different. On the other hand, if the reduced model $\hat{\Sigma}$ is to be used online to predict the future behavior of Σ , and if the mode ω_t can be observed at run-time, we can assume the following and derive stronger relations on the state realization difference $\|\mathbf{x}_t - \hat{\mathbf{x}}_t\|$.

Setup S2 (Mode Synchrony): Mode $\hat{\omega}_t$ of $\hat{\Sigma}$ is synchronous to ω_t of Σ , i.e., for all t , if $\omega_t \in \hat{\Omega}_k$ then $\hat{\omega}_t = k$.

Mode synchrony setup essentially establishes the strongest possible coupling between $\omega_{0:t}$ and $\hat{\omega}_{0:t}$ as discussed in Section III-C. When the mode sequence $\hat{\omega}_{0:t}$ of $\hat{\Sigma}$ is synchronized with that of Σ , this amounts to having $\hat{\Sigma}$ being driven by an external switching signal $\omega_{0:t}$.

In the following, we provide bounds on how close $\hat{\Sigma}$ is to Σ in terms of the following approximation metrics: (i) under the mean-square stability of Σ , the difference $\|\mathbf{x}_t - \hat{\mathbf{x}}_t\|$ in trajectories (Theorem 3); (ii) under uniform stability, the difference in trajectories (Theorem 4 (T1)) and the difference of transition kernels (Theorem 4 (T2)).

A. RESULT WITH MEAN-SQUARE STABILITY

Due to the stochasticity of mode switching, the stability for MJS is typically studied in mean-square sense. We say Σ is mean-square stable (MSS), if there exists Σ_∞ such that $\lim_{t \rightarrow \infty} \mathbb{E}[\mathbf{x}_t \mathbf{x}_t^\top] = \Sigma_\infty$. Define the augmented state matrix $\mathcal{A} \in \mathbb{R}^{sn^2 \times sn^2}$ with its ij -th $n^2 \times n^2$ block given by $[\mathcal{A}]_{ij} := \mathbf{T}(j, i) \cdot \mathbf{A}_j \otimes \mathbf{A}_i$, and let $\rho(\mathcal{A})$ denote the spectral radius of \mathcal{A} . Then, for the autonomous case, i.e., $\mathbf{u}_t = 0$, Σ being MSS is equivalent to $\rho(\mathcal{A}) < 1$. It is known that MSS does not imply each individual mode is stable, and vice versa [46]. For

any $\rho \geq \rho(\mathcal{A})$, let us define $\tau(\mathcal{A}, \rho) := \sup_{k \in \mathbb{N}} \|\mathcal{A}^k\| / \rho^k$. This term compares the convergence of $\|\mathcal{A}^k\|$ and ρ^k . In the remainder of this paper, we use τ as a shorthand notation for $\tau(\mathcal{A}, \rho)$, but keep in mind that τ depends on the choice of the free parameter ρ . It is easy to see (i) $\tau \geq 1$, (ii) $\|\mathcal{A}^k\| \leq \tau \rho^k$, and (iii) by Gelfand's formula, τ is finite for any $\rho > \rho(\mathcal{A})$. When \mathcal{A} is diagonalizable with decomposition $\mathcal{A} = \mathbf{V} \Lambda \mathbf{V}^{-1}$, a generic property of matrices in $\mathbb{R}^{sn^2 \times sn^2}$, τ is also finite for $\rho = \rho(\mathcal{A})$. In this case, we have $\tau \leq \|\mathbf{V}\| \|\mathbf{V}^{-1}\|$. We let $\bar{A} := \max_i \|\mathbf{A}_i\|$, $\bar{B} := \max_i \|\mathbf{B}_i\|$. The following theorem provides an upper bound for $\|\mathbf{x}_t - \hat{\mathbf{x}}_t\|$ under MSS.

Theorem 3: Consider setup S1 and S2 where the shared initial state \mathbf{x}_0 and inputs $\mathbf{u}_{0:t}$ can be arbitrary as long as for all t , \mathbf{u}_t is bounded, i.e., $\|\mathbf{u}_t\| \leq \bar{u}$. Assume Σ is MSS and $\hat{\Omega}_{1:r} = \Omega_{1:r}$ in Algorithm 1. For any $\rho \in [\rho(\mathcal{A}), 1)$ and its corresponding τ , let $\rho_0 := \frac{1+\rho}{2}$. For perturbation, assume $\epsilon_A \leq \min\{\bar{A}, \frac{1-\rho}{6\sqrt{\bar{A}}\|\mathbf{T}\|}\}$ and $\epsilon_B \leq \bar{B}$. Then, $\mathbb{E}[\|\mathbf{x}_t - \hat{\mathbf{x}}_t\|] \leq 4\sqrt{n\sqrt{s}\tau} \epsilon_t^{mss}$ where $\epsilon_t^{mss} := \rho_0^{\frac{t-1}{2}} \sqrt{t\bar{A}\|\mathbf{T}\|} \epsilon_A \|\mathbf{x}_0\| + \sqrt{\bar{B}\bar{u}} \left(\frac{\sqrt{\rho_0}}{(1-\sqrt{\rho_0})^2} \sqrt{\bar{A}\|\mathbf{T}\|} \epsilon_A + \frac{\sqrt{2}}{1-\sqrt{\rho_0}} \sqrt{\epsilon_B} \right)$.

In this theorem, ϵ_t^{mss} is the key element in the upper bound. In its definition, the first term describes the effect of ϵ_A through initial state \mathbf{x}_0 . Since $\rho < 1$ due to Σ being MSS, we know $\rho_0 < 1$, which implies exponential decay. The rest of the terms in ϵ_t^{mss} characterize the effects of ϵ_A and ϵ_B through the inputs. And if there is no input, the trajectory difference $\|\mathbf{x}_t - \hat{\mathbf{x}}_t\|$ converges to 0 exponentially with t . The condition $\epsilon_A \leq \frac{1-\rho}{6\sqrt{\bar{A}}}$ is used to guarantee perturbation ϵ_A is small such that $\hat{\Sigma}$ is still MSS, as otherwise the difference will grow exponentially, and no meaningful results can be established in this case. Conditions $\epsilon_A < \bar{A}$ and $\epsilon_B \leq \bar{B}$ are only used to simplify the expressions, and similar bounds can be established without them.

Fact 1 provides a sanity check for Theorem 3: when $\epsilon_A = \epsilon_B = 0$, we have $\mathbf{x}_t = \hat{\mathbf{x}}_t$. In the autonomous case, i.e., $\mathbf{u}_t = 0$, as a direct corollary of Theorem 3, we can further obtain a probabilistic bound on the difference over an entire trajectory using Markov inequality: with probability at least $1 - \delta$, $\sum_{t=0}^{\infty} \|\mathbf{x}_t - \hat{\mathbf{x}}_t\| \leq \frac{4\sqrt{npt} \|\mathbf{x}_0\| \sqrt{\bar{A}\epsilon_A}}{\delta(1-\sqrt{\rho_0})^2}$.

B. RESULTS WITH UNIFORM STABILITY

MSS in Section V-A is a weak notion of stability in that it only requires stability in expectation while still allowing a set of mode switching sequences that result in explosive \mathbf{x}_t , even a set with nonzero probability. In this section, we consider uniform stability, which guarantees stable \mathbf{x}_t even with an arbitrary switching sequence. Uniform stability allows us to further build approximation results without enforcing mode synchrony as in S2.

We let $\xi(\mathbf{A}_{1:s})$ denote the joint spectral radius of the set of state matrices $\mathbf{A}_{1:s}$ of Σ , i.e., $\xi(\mathbf{A}_{1:s}) := \lim_{k \rightarrow \infty} \max_{\omega_{1:k} \in [s]^k} \|\mathbf{A}_{\omega_1} \cdots \mathbf{A}_{\omega_k}\|^{\frac{1}{k}}$. We say Σ is uniformly (and exponentially)

stable if $\xi(\mathbf{A}_{1:s}) < 1$. For any $\xi \geq \xi(\mathbf{A}_{1:s})$, we define $\kappa(\mathbf{A}_{1:s}, \xi)$ to compare the convergence of $\|\mathbf{A}_{\omega_1} \cdots \mathbf{A}_{\omega_k}\|$ and ξ^k for any mode switching sequence $\omega_{1:k}$: $\kappa(\mathbf{A}_{1:s}, \xi) := \sup_{k \in \mathbb{N}} \max_{\omega_{1:k} \in [s]^k} \|\mathbf{A}_{\omega_1} \cdots \mathbf{A}_{\omega_k}\| / \xi^k$. In the remainder of this paper, we use κ as a shorthand notation for $\kappa(\mathbf{A}_{1:s}, \xi)$. Note that the pair $\{\xi, \kappa\}$ for uniform stability is just the counterpart of $\{\rho, \tau\}$ defined earlier for MSS. And similarly, we have (i) $\kappa \geq 1$, (ii) $\|\mathbf{A}_{\omega_1} \cdots \mathbf{A}_{\omega_k}\| \leq \kappa \xi^k$, and (iii) κ is finite for any $\xi > \xi(\mathbf{A}_{1:s})$. Furthermore, we let $\bar{\mathcal{T}} := \max_{i,j} \mathbf{T}(i, j)$.

We first formally define the transition kernels for Σ and $\hat{\Sigma}$ and their distance. Under fixed initial state \mathbf{x}_0 and input sequence $\mathbf{u}_{0:t-1}$, we define the reachable set $\mathcal{X}_t := \{\mathbf{x}_t : \mathbf{x}_{0:t} \text{ is a solution to (1)} \forall \omega_{0:t-1} \in [s]^t\}$. Then we define the t -step transition kernel as $p_t(\mathbf{x}) := \mathbb{P}(\mathbf{x}_t = \mathbf{x})$ for all $\mathbf{x} \in \mathcal{X}_t$. Note that both \mathcal{X}_t and $p_t(\mathbf{x})$ depend on the choice of the initial state and input sequence as well. We omit this dependency in the notation not only for simplicity but also because the approximation results we provide hold for arbitrary initial state and input sequence. Similarly, for the reduced $\hat{\Sigma}$, we use $\hat{\mathcal{X}}_t$ to denote the reachable set at time t , and for $\hat{\mathbf{x}} \in \hat{\mathcal{X}}_t$, we let $\hat{p}_t(\hat{\mathbf{x}}) := \mathbb{P}(\hat{\mathbf{x}}_t = \hat{\mathbf{x}})$. Then, for $\ell \geq 1$ the ℓ -Wasserstein distance $W_\ell(p_t, \hat{p}_t)$, between distributions p_t and \hat{p}_t is defined as the optimal objective value of the following mass transportation problem:

$$\begin{aligned} \min_{f \geq 0} \quad & \left(\sum_{\mathbf{x} \in \mathcal{X}_t, \hat{\mathbf{x}} \in \hat{\mathcal{X}}_t} f(\mathbf{x}, \hat{\mathbf{x}}) \|\mathbf{x} - \hat{\mathbf{x}}\|^\ell \right)^{1/\ell} \\ \text{s.t.} \quad & \sum_{\mathbf{x} \in \mathcal{X}_t} f(\mathbf{x}, \hat{\mathbf{x}}) = \hat{p}_t(\hat{\mathbf{x}}), \forall \hat{\mathbf{x}} \\ & \sum_{\hat{\mathbf{x}} \in \hat{\mathcal{X}}_t} f(\mathbf{x}, \hat{\mathbf{x}}) = p_t(\mathbf{x}), \forall \mathbf{x}. \end{aligned} \quad (8)$$

The constraints describe the transportation of probability mass distributed according p_t to the support of \hat{p}_t so that the mass after transportation distributes the same as \hat{p}_t . We can view $f(\mathbf{x}, \hat{\mathbf{x}})$ as the mass that is transported from point \mathbf{x} to $\hat{\mathbf{x}}$ and $\|\mathbf{x} - \hat{\mathbf{x}}\|$ as the distance it travels. When $\ell = 1$, the goal is to minimize the total weighted travel distance, and the resulting W_1 is also known as the earth mover's distance. Now we are ready to present our results for the uniform stability assumption.

Theorem 4: Consider setup S1 where the shared initial state \mathbf{x}_0 and inputs $\mathbf{u}_{0:t}$ can be arbitrary as long as for all t , \mathbf{u}_t is bounded, i.e., $\|\mathbf{u}_t\| \leq \bar{u}$. Assume Σ is uniformly stable and $\hat{\Omega}_{1:r} = \Omega_{1:r}$ in Algorithm 1. For any $\xi \in [\xi(\mathbf{A}_{1:s}), 1)$ and its corresponding κ , let $\xi_0 := \frac{1+\xi}{2}$. For perturbation, we assume $\epsilon_A \leq \frac{1-\xi}{2\kappa}$ and $\epsilon_B \leq \bar{B}$. Then, we have the following results.

- T1) Under S2, $\|\mathbf{x}_t - \hat{\mathbf{x}}_t\| \leq \epsilon_t^{us} := t \xi_0^{t-1} \kappa^2 \|\mathbf{x}_0\| \epsilon_A + \frac{2(1+t\xi_0^t)\kappa^2 \bar{B}\bar{u}}{1-\xi_0} \epsilon_A + \frac{\kappa\bar{u}}{1-\xi} \epsilon_B$ almost surely.
- T2) Consider the autonomous case, i.e., $\mathbf{B}_{1:s} = 0$. (S2 is not mandatory.) Then, $W_\ell(p_t, \hat{p}_t) \leq t \xi_0^{t-1} \kappa^2 \|\mathbf{x}_0\| \epsilon_A + 2r^2 t \kappa \|\mathbf{x}_0\| r^t (\kappa \epsilon_A + \xi)^t (\bar{\mathcal{T}} + \epsilon_T)^{(t-2)/\ell} \epsilon_T^{1/\ell}$.

In Theorem 4, condition $\epsilon_A \leq \frac{1-\xi}{2\kappa}$ guarantees the reduced $\hat{\Sigma}$ is uniformly stable with joint spectral radius upper bounded by ξ_0 . The condition $\epsilon_B \leq \bar{B}$ simplifies the expression, and similar results can be obtained when it is relaxed. (T1) upper

bounds the realization difference with the mode synchrony setup. We can see the similarity between the upper bounds ϵ_t^{us} and ϵ_t^{mss} of Theorem 3 under MSS assumption. The fact that uniform stability and MSS upper bound $\|\mathbf{x}_t - \hat{\mathbf{x}}_t\|$ deterministically and in expectation respectively is a manifestation of the difference between these two stability notions for MJS.

In (T2), we bound the Wasserstein distance between p_t and \hat{p}_t . This bound depends on both perturbations ϵ_A and ϵ_T . Let μ and \mathbf{S} denote the mean and covariance for \mathbf{x}_t ; and similarly define $\hat{\mu}$ and $\hat{\mathbf{S}}$ for $\hat{\mathbf{x}}_t$. From [47, Theorem 4], we obtain $\|\mu - \hat{\mu}\|^2 + d(\mathbf{S}, \hat{\mathbf{S}}) \leq W_2(p_t, \hat{p}_t)^2$, where $d(\mathbf{S}, \hat{\mathbf{S}}) := \text{tr}(\mathbf{S} + \hat{\mathbf{S}} - 2(\mathbf{S}^{\frac{1}{2}} \hat{\mathbf{S}} \mathbf{S}^{\frac{1}{2}})^{\frac{1}{2}})$ is a metric between \mathbf{S} and $\hat{\mathbf{S}}$. Hence, by setting $\ell = 2$ in (T2), we also obtain upper bounds for the differences between p_t and \hat{p}_t in terms of their first and second order moments, i.e., $\|\mu - \hat{\mu}\|$ and $d(\mathbf{S}, \hat{\mathbf{S}})$. These metrics can be used to obtain performance bounds in other control problems such as covariance steering [48], [49], [50] and ensemble control [51].

VI. STABILITY ANALYSIS

In this section, we study whether the stability properties of Σ can be deduced from those of $\hat{\Sigma}$. Recall that MSS of Σ depends on $\rho(\mathbf{A})$, the spectral radius of its augmented state matrix \mathbf{A} , and its uniform stability depends on $\xi(\mathbf{A}_{1:s})$, the joint spectral radius of state matrices $\mathbf{A}_{1:s}$. Similarly, for $\hat{\Sigma}$, we define its augmented state matrix $\hat{\mathbf{A}} \in \mathbb{R}^{rn^2 \times rn^2}$ with its ij -th $n^2 \times n^2$ block given by $[\hat{\mathbf{A}}]_{ij} := \hat{\mathbf{T}}(j, i) \cdot \hat{\mathbf{A}}_j \otimes \hat{\mathbf{A}}_i$ and let $\rho(\hat{\mathbf{A}})$ denote its spectral radius; we let $\xi(\hat{\mathbf{A}}_{1:r})$ denote the joint spectral radius of state matrices $\hat{\mathbf{A}}_{1:r}$. With these notations, we want to analyze when $\rho(\hat{\mathbf{A}})$ (or $\xi(\hat{\mathbf{A}}_{1:r})$) can be taken as an approximation for $\rho(\mathbf{A})$ (or $\xi(\mathbf{A}_{1:s})$) since computing or approximating $\rho(\hat{\mathbf{A}})$ and $\xi(\hat{\mathbf{A}}_{1:r})$ may require much less computation compared with $\rho(\mathbf{A})$ and $\xi(\mathbf{A}_{1:s})$ as $\hat{\Sigma}$ has much fewer number of modes than Σ .

To begin with, we first construct an intermediate MJS by expanding the reduced $\hat{\Sigma}$: we let $\bar{\Sigma} := \text{MJS}(\bar{\mathbf{A}}_{1:s}, \bar{\mathbf{B}}_{1:s}, \bar{\mathbf{T}})$ such that $\bar{\mathbf{T}} \in \mathcal{L}(\mathbf{T}, \hat{\Omega}_{1:r}, \epsilon_T)$, and for all $i \in [s]$ (suppose $i \in \hat{\Omega}_k$), $\bar{\mathbf{A}}_i = \hat{\mathbf{A}}_k$, $\bar{\mathbf{B}}_i = \hat{\mathbf{B}}_k$. By definition of $\mathcal{L}(\mathbf{T}, \hat{\Omega}_{1:r}, \epsilon_T)$, we can solve for $\bar{\mathbf{T}}$ through a linear programming feasibility problem with constraints given by the definition of $\mathcal{L}(\cdot, \cdot, \cdot)$ in (6). Particularly, if it is the aggregatable case P2, it suffices to let $\bar{\mathbf{T}}(i, :) := |\hat{\Omega}_k|^{-1} \sum_{i \in \hat{\Omega}_k} \mathbf{T}(i, :)$ if $i \in \hat{\Omega}_k$. Note that by construction, $\bar{\Sigma}$ is mode-reducible with respect to $\hat{\Omega}_{1:r}$ and can be reduced to $\hat{\Sigma}$. According to Fact 1, $\bar{\Sigma}$ has the same dynamics as $\hat{\Sigma}$. Since $\bar{\Sigma}$ has the same number of modes as Σ , we can use $\bar{\Sigma}$ as a bridge to compare Σ and $\hat{\Sigma}$. We let $\rho(\bar{\mathbf{A}})$ denote the spectral radius of $\bar{\mathbf{A}} \in \mathbb{R}^{sn^2 \times sn^2}$ whose ij -th $n^2 \times n^2$ block is given by $[\bar{\mathbf{A}}]_{ij} := \bar{\mathbf{T}}(j, i) \cdot \bar{\mathbf{A}}_j \otimes \bar{\mathbf{A}}_i$ and let $\xi(\bar{\mathbf{A}}_{1:s})$ denote the joint spectral radius of $\bar{\mathbf{A}}_{1:s}$. The following preliminary result (proof omitted due to its simplicity) says $\bar{\Sigma}$ and $\hat{\Sigma}$ have the same stability properties.

Lemma 2: For $\hat{\Sigma}$ and $\bar{\Sigma}$, we have $\rho(\hat{\mathbf{A}}) = \rho(\bar{\mathbf{A}})$ and $\xi(\hat{\mathbf{A}}_{1:s}) = \xi(\bar{\mathbf{A}}_{1:s})$.

One implication of Lemma 2 is that if an MJS is mode-reducible, the reduced MJS has the same MSS and uniform

stability as the original MJS in terms of (joint) spectral radius. When Σ is not exactly mode-reducible, Lemma 2 allows us to compare the stability properties of $\hat{\Sigma}$ and Σ via the intermediate expanded $\bar{\Sigma}$ as presented in Theorem 5. For analysis purposes, similar to τ and κ defined for Σ , we define $\bar{\tau} := \sup_{k \in \mathbb{N}} \|\bar{\mathcal{A}}^k\|/\hat{\rho}^k$ for any $\hat{\rho} \geq \rho(\hat{\mathcal{A}})$ and $\bar{\kappa} := \sup_{k \in \mathbb{N}} \max_{\omega_{1:k} \in [r]^k} \|\bar{\mathcal{A}}_{\omega_1} \cdots \bar{\mathcal{A}}_{\omega_k}\|/\hat{\xi}^k$ for any $\hat{\xi} \geq \xi(\hat{\mathcal{A}}_{1:s})$

Theorem 5 (Stability Analysis): Assume $\hat{\Omega}_{1:r} = \Omega_{1:r}$ in Algorithm 1, then Σ and $\hat{\Sigma}$ have the following relations.

T1) (MSS) For any $\rho \geq \rho(\mathcal{A})$ and its corresponding τ , any $\hat{\rho} \geq \rho(\hat{\mathcal{A}})$ and its corresponding $\bar{\tau}$, we have

$$\begin{aligned} \rho(\hat{\mathcal{A}}) - \rho(\mathcal{A}) &\leq \tau \epsilon_\rho + (\rho - \rho(\mathcal{A})) \\ \rho(\mathcal{A}) - \rho(\hat{\mathcal{A}}) &\leq \bar{\tau} \epsilon_\rho + (\hat{\rho} - \rho(\hat{\mathcal{A}})) \end{aligned} \quad (9)$$

where $\epsilon_\rho := \sqrt{s}((2\bar{\mathcal{A}} + \epsilon_A)\epsilon_A + \bar{\mathcal{A}}^2\epsilon_T)$.

T2) (Uniform stability) For any $\xi \geq \xi(\mathcal{A}_{1:s})$ and its corresponding κ , any $\hat{\xi} \geq \xi(\hat{\mathcal{A}}_{1:r})$ and its corresponding $\bar{\tau}$, we have

$$\begin{aligned} \xi(\hat{\mathcal{A}}_{1:r}) - \xi(\mathcal{A}_{1:s}) &\leq \kappa \epsilon_A + (\xi - \xi(\mathcal{A}_{1:s})) \\ \xi(\mathcal{A}_{1:s}) - \xi(\hat{\mathcal{A}}_{1:r}) &\leq \bar{\kappa} \epsilon_A + (\hat{\xi} - \xi(\hat{\mathcal{A}}_{1:r})). \end{aligned} \quad (10)$$

Proof: From Lemma 2, it suffices to prove

$$\rho(\bar{\mathcal{A}}) \leq \tau \epsilon_\rho + \rho, \quad \rho(\mathcal{A}) \leq \bar{\tau} \epsilon_\rho + \hat{\rho}. \quad (11)$$

$$\xi(\bar{\mathcal{A}}_{1:s}) \leq \kappa \epsilon_A + \xi, \quad \xi(\mathcal{A}_{1:s}) \leq \bar{\kappa} \epsilon_A + \hat{\xi}. \quad (12)$$

Since we assume $\hat{\Omega}_{1:r} = \Omega_{1:r}$, then for Σ and $\bar{\Sigma}$, we have $\|\bar{\mathcal{A}}_i - \mathcal{A}_i\| \leq \epsilon_A$, $\|\bar{\mathcal{B}}_i - \mathcal{B}_i\| \leq \epsilon_B$, and $\|\bar{\mathcal{T}} - \mathcal{T}\|_\infty \leq \epsilon_T$. Consider matrix $\bar{\mathcal{A}}$ and \mathcal{A} , we have $[\bar{\mathcal{A}}]_{ij} - [\mathcal{A}]_{ij} = \bar{\mathcal{T}}(j, i)\bar{\mathcal{A}}_j \otimes \bar{\mathcal{A}}_i - \mathcal{T}(j, i)\mathcal{A}_j \otimes \mathcal{A}_i = \bar{\mathcal{T}}(j, i)(\bar{\mathcal{A}}_j \otimes \bar{\mathcal{A}}_i - \mathcal{A}_j \otimes \mathcal{A}_i) + (\bar{\mathcal{T}}(j, i) - \mathcal{T}(j, i))\mathcal{A}_j \otimes \mathcal{A}_i$. Note that $\bar{\mathcal{A}}_j \otimes \bar{\mathcal{A}}_i - \mathcal{A}_j \otimes \mathcal{A}_i = (\bar{\mathcal{A}}_j - \mathcal{A}_j) \otimes \mathcal{A}_i + \mathcal{A}_j \otimes (\bar{\mathcal{A}}_i - \mathcal{A}_i) + (\bar{\mathcal{A}}_j - \mathcal{A}_j) \otimes (\bar{\mathcal{A}}_i - \mathcal{A}_i)$, which gives $\|\bar{\mathcal{A}}_j \otimes \bar{\mathcal{A}}_i - \mathcal{A}_j \otimes \mathcal{A}_i\| \leq (2\bar{\mathcal{A}} + \epsilon_A)\epsilon_A$. Then, we have $\|[\bar{\mathcal{A}}]_{ij} - [\mathcal{A}]_{ij}\| \leq \bar{\mathcal{T}}(j, i)(2\bar{\mathcal{A}} + \epsilon_A)\epsilon_A + \|\bar{\mathcal{T}}(j, i) - \mathcal{T}(j, i)\|\bar{\mathcal{A}}^2$. To simplify the notation, we let $c_1 := (2\bar{\mathcal{A}} + \epsilon_A)\epsilon_A$ and $c_2 := \bar{\mathcal{A}}^2$. By Cauchy-Schwarz inequality, we have $\sum_i \|[\bar{\mathcal{A}}]_{ij} - [\mathcal{A}]_{ij}\|^2 \leq (c_1 \|\bar{\mathcal{T}}(j, :)\| + c_2 \|\bar{\mathcal{T}}(j, :)\|)^2$. Thus, $\|\bar{\mathcal{A}} - \mathcal{A}\| \leq \sqrt{s} \max_j (\sum_i \|[\bar{\mathcal{A}}]_{ij} - [\mathcal{A}]_{ij}\|)^{0.5} \leq \sqrt{s}(c_1 + c_2 \epsilon_T) =: \epsilon_\rho$.

With Corollary 2 in the appendix, we have $\|\bar{\mathcal{A}}^k\| \leq \tau(\tau \epsilon_\rho + \rho)^k$. By Gelfand's formula, $\rho(\bar{\mathcal{A}}) = \limsup_{k \rightarrow \infty} \|\bar{\mathcal{A}}^k\|^{1/k} \leq \tau \epsilon_\rho + \rho$, which shows the left inequality of (11). If we use Corollary 2 the other way, we have $\|\mathcal{A}^k\| \leq \bar{\tau}(\bar{\tau} \epsilon_\rho + \hat{\rho})^k$, which similarly implies $\rho(\mathcal{A}) \leq \bar{\tau} \epsilon_\rho + \hat{\rho}$. With these results, (11) is proved. (12) can be shown similarly by noticing $\|\bar{\mathcal{A}}_i - \mathcal{A}_i\| \leq \epsilon_A$ and then using Lemma 10 in the appendix. ■

Theorem 5 provides upper bounds on $|\rho(\hat{\mathcal{A}}) - \rho(\mathcal{A})|$ and $|\xi(\hat{\mathcal{A}}_{1:r}) - \xi(\mathcal{A}_{1:s})|$. By definition, τ decreases when ρ increases, and the same applies to the pairs $\{\bar{\tau}, \hat{\rho}\}$, $\{\kappa, \xi\}$, and $\{\bar{\kappa}, \hat{\xi}\}$. Hence, for fixed ϵ_ρ and ϵ_A , by tuning the free parameters ρ , $\hat{\rho}$, ξ , and $\hat{\xi}$, one may obtain tighter upper bounds in Theorem 5. When $\rho = \rho(\mathcal{A})$, $\hat{\rho} = \rho(\hat{\mathcal{A}})$, the bound in (T1)

becomes tight at $\epsilon_\rho = 0$ as the upper and lower bounds meet at 0. Note that these results hold for both stable and unstable Σ , and does not require perturbation $\epsilon_A, \epsilon_B, \epsilon_T$ to be small, which is in contrast to approximation results in Theorem 3 and 4.

Now we briefly compare the complexities for computing or approximating $\rho(\mathcal{A})$, $\rho(\hat{\mathcal{A}})$, $\xi(\mathcal{A}_{1:s})$, and $\xi(\hat{\mathcal{A}}_{1:r})$. Since \mathcal{A} has dimension $sn^2 \times sn^2$, the complexity to compute its spectral radius $\rho(\mathcal{A})$ is $\mathcal{O}(s^3 n^6)$, but it only requires $\mathcal{O}(r^3 n^6)$ for $\rho(\hat{\mathcal{A}})$. Computation of the joint spectral radius is in general undecidable [52]. An iterative approach [53] provides an approximation for $\xi(\mathcal{A}_{1:s})$ with computational complexity $\mathcal{O}(s)$, whereas it only requires $\mathcal{O}(r)$ for $\xi(\hat{\mathcal{A}}_{1:r})$.

VII. CONTROLLER DESIGN WITH CASE STUDY ON LQR

When the mode of an MJS can be measured at run-time, one can use *mode-dependent* controllers. A mode-dependent controller is essentially a collection of individual controllers, one per mode, and the deployed controller switches with corresponding modes. Therefore, if we can reduce the modes, that would also reduce the number of controllers in a mode-dependent control. That is, with the reduced $\hat{\Sigma}$, we can design mode-dependent controller $\hat{\mathbf{K}}_{1:r}$ for $\hat{\Sigma}$ and then associate every mode i in Σ with $\hat{\mathbf{K}}_i$ if $i \in \hat{\Omega}_k$. Since $\hat{\Sigma}$ has a smaller scale than Σ , the computational cost may be reduced but the question is how this simplified controller performs on the original system Σ . In this section, we show how this idea can be used for linear quadratic regulator (LQR) for MJS and provide suboptimality guarantees for the reduced controller.

In the infinite horizon MJS LQR problems, given positive definite cost matrices \mathbf{Q} and \mathbf{R} , we define quadratic cumulative cost $J_T = \mathbb{E}[\sum_{t=0}^{T-1} (\mathbf{x}_t^\top \mathbf{Q} \mathbf{x}_t + \mathbf{u}_t^\top \mathbf{R} \mathbf{u}_t) + \mathbf{x}_T^\top \mathbf{Q} \mathbf{x}_T]$. The goal is to design inputs to minimize the infinite time average cost $\limsup_{T \rightarrow \infty} \frac{1}{T} J_T$ under Σ . To ease the exposition, we let $\mathbb{S}_s^+ := \{\mathbf{X}_{1:s} : \forall i \in [s], \mathbf{X}_i \in \mathbb{R}^{n \times n}, \mathbf{X}_i \succeq 0\}$. For $\mathbf{X}_{1:s} \in \mathbb{S}_s^+$, for all $i \in [s]$, define three operators $\varphi_i(\mathbf{X}_{1:s}) := \sum_{j \in [s]} \mathbf{T}(i, j)\mathbf{X}_j$, $\mathcal{K}_i(\mathbf{X}_{1:s}) := -(\mathbf{R} + \mathbf{B}_i^\top \varphi_i(\mathbf{X}_{1:s}) \mathbf{B}_i)^{-1}(\mathbf{B}_i^\top \varphi_i(\mathbf{X}_{1:s}) \mathbf{A}_i)$, and

$$\begin{aligned} \mathcal{R}_i(\mathbf{X}_{1:s}) &:= \mathbf{Q} + \mathbf{A}_i^\top \varphi_i(\mathbf{X}_{1:s}) \mathbf{A}_i - \mathbf{A}_i^\top \varphi_i(\mathbf{X}_{1:s})^\top \mathbf{B}_i \\ &\quad \cdot (\mathbf{R} + \mathbf{B}_i^\top \varphi_i(\mathbf{X}_{1:s}) \mathbf{B}_i)^{-1} \mathbf{B}_i^\top \varphi_i(\mathbf{X}_{1:s}) \mathbf{A}_i. \end{aligned} \quad (13)$$

Then, the solution to the infinite horizon LQR is the following: we first solve for the coupled Riccati equations $\mathbf{P}_i = \mathcal{R}_i(\mathbf{P}_{1:s})$, $\forall i \in [s]$, and then if $\omega_t = i$ at time t , we let input $\mathbf{u}_t = \mathbf{K}_i \mathbf{x}_t$ where $\mathbf{K}_i = \mathcal{K}_i(\mathbf{P}_{1:s})$. The solution existence, uniqueness, and optimality can be guaranteed by the following assumption according to [46].

Assumption A1: Σ is mean-square stabilizable. Cost matrices $\mathbf{Q} \succ 0$ and $\mathbf{R} \succ 0$.

To design controllers with the reduced $\hat{\Sigma}$, we can first compute controller $\hat{\mathbf{K}}_{1:r}$ by solving LQR problem with $\hat{\Sigma}$ as the MJS dynamics. This requires solving r coupled Riccati equations, each of which is parameterized by $\hat{\mathbf{A}}_i, \hat{\mathbf{B}}_i, \hat{\mathbf{T}}(i, :), \mathbf{Q}, \mathbf{R}$. To apply $\hat{\mathbf{K}}_{1:r}$ to the original Σ , we simply let $\mathbf{u}_t = \hat{\mathbf{K}}_k \mathbf{x}_t$ if

$\omega_t = k$. Since the number of coupled Riccati equations is the same as the number of modes, the computational cost for Σ is $\mathcal{O}(s)$ while only $\mathcal{O}(r)$ for $\bar{\Sigma}$, thus the saving is prominent when $r \ll s$.

Next, we analyze the suboptimality when applying controllers computed with $\bar{\Sigma}$. To begin with, similar to the notations for Σ , for $\bar{\Sigma}$ we define $\bar{\varphi}_{1:r}$, $\bar{\mathcal{R}}_{1:r}$, $\bar{\mathcal{K}}_{1:r}$, $\bar{\mathbf{P}}_{1:r}$, and $\hat{\mathbf{K}}_{1:r}$. Particularly, $\bar{\mathbf{P}}_{1:r}$ denotes the Riccati solution such that $\hat{\mathbf{P}}_i = \bar{\mathcal{R}}_i(\bar{\mathbf{P}}_{1:r})$, and $\hat{\mathbf{K}}_{1:r}$ is computed such that $\hat{\mathbf{K}}_i = \bar{\mathcal{K}}_i(\bar{\mathbf{P}}_{1:r})$. We will take the expanded and mode-reducible MJS $\bar{\Sigma}$ constructed with $\bar{\Sigma}$ and $\hat{\Omega}_{1:r}$ in Section VI as a bridge. For $\bar{\Sigma}$, we similarly define $\bar{\varphi}_{1:s}$, $\bar{\mathcal{R}}_{1:s}$, $\bar{\mathcal{K}}_{1:s}$, $\bar{\mathbf{P}}_{1:s}$, and $\bar{\mathbf{K}}_{1:s}$. In terms of LQR solutions, the relation between $\bar{\Sigma}$ and $\bar{\Sigma}$ is given below.

Lemma 3: Assume the Riccati solution $\bar{\mathbf{P}}_{1:s}$ exists and $\bar{\mathbf{P}}_i > 0$ for all i . Then, (i) there exists a unique Riccati solution $\hat{\mathbf{P}}_{1:r}$ in \mathbb{S}_s^+ ; (ii) $\hat{\mathbf{P}}_k = \bar{\mathbf{P}}_i$, $\hat{\mathbf{K}}_k = \bar{\mathbf{K}}_i$ for any $i \in \hat{\Omega}_k$ for any k .

Proof: We consider the Riccati operator iteration defined as follows: $\bar{\mathbf{P}}_i^{(0)} = \mathbf{Q}$, $\bar{\mathbf{P}}_i^{(h+1)} = \bar{\mathcal{R}}_i(\bar{\mathbf{P}}_i^{(h)})$ for all $i \in [s]$, $h \in \mathbb{N}$ and $\hat{\mathbf{P}}_k^{(0)} = \mathbf{Q}$, $\hat{\mathbf{P}}_k^{(h+1)} = \hat{\mathcal{R}}_k(\hat{\mathbf{P}}_k^{(h)})$ for all $k \in [r]$, $h \in \mathbb{N}$. Then, note that by construction, for all $i \in \hat{\Omega}_k$ and all $l \in [r]$, we have $\sum_{j \in \hat{\Omega}_l} \bar{\mathbf{T}}(i, j) = \hat{\mathbf{T}}(k, l)$. Through induction and algebra, it is easy to show that for all $h \in \mathbb{N}$, and for any $i, i' \in \hat{\Omega}_k$ for any k , we have $\bar{\mathbf{P}}_i^{(h)} = \bar{\mathbf{P}}_{i'}^{(h)} = \hat{\mathbf{P}}_k^{(h)}$.

Since $\bar{\mathbf{P}}_i > 0$, by [54, Fact 4], we know $\bar{\mathbf{P}}_{1:s}$ is the unique solution among \mathbb{S}_s^+ , and $\bar{\mathbf{K}}_{1:s}$ stabilizes $\bar{\Sigma}$. According to [46, Proposition A.23], the stabilizability of $\bar{\Sigma}$ and the fact $\mathbf{Q}, \mathbf{R} > 0$ imply $\lim_{h \rightarrow \infty} \bar{\mathbf{P}}_i^{(h)} = \bar{\mathbf{P}}_i$. Combining this convergence result with the Riccati iteration results we just showed, we further have, for any $i, i' \in \hat{\Omega}_k$ and any k , we have $\bar{\mathbf{P}}_i = \bar{\mathbf{P}}_{i'} = \hat{\mathbf{P}}_k$. Then, it is easy to show that $\bar{\mathbf{K}}_i = \bar{\mathbf{K}}_{i'} = \hat{\mathbf{K}}_k$. The uniqueness of $\hat{\mathbf{P}}_{1:r}$ can be shown by contradiction. ■

With this lemma, we have the following suboptimality guarantees in terms of applying controller $\hat{\mathbf{K}}_{1:r}$ to Σ .

Theorem 6 (LQR Suboptimality): Assume A1 holds for Σ , and Σ has additive Gaussian noise $\mathcal{N}(0, \sigma_w^2 \mathbf{I}_n)$ that is independent of the mode switching. Let J^* and \hat{J} respectively denote the infinite time average cost incurred by the optimal controller $\mathbf{K}_{1:s}$ and controller $\hat{\mathbf{K}}_{1:r}$ (at time t , $\mathbf{u}_t = \hat{\mathbf{K}}_k \mathbf{x}_t$ if $\omega_t \in \hat{\Omega}_k$). Then, there exists constants $\bar{\epsilon}_{\mathbf{A}, \mathbf{B}}$, $\bar{\epsilon}_{\mathbf{T}}$, $C_{\mathbf{A}, \mathbf{B}}$, and $C_{\mathbf{T}}$, such that when $\max\{\epsilon_{\mathbf{A}}, \epsilon_{\mathbf{B}}\} \leq \bar{\epsilon}_{\mathbf{A}, \mathbf{B}}$ and $\epsilon_{\mathbf{T}} \leq \bar{\epsilon}_{\mathbf{T}}$,

$$\hat{J} - J^* \leq \sigma_w^2 (C_{\mathbf{A}, \mathbf{B}} \max\{\epsilon_{\mathbf{A}}, \epsilon_{\mathbf{B}}\} + C_{\mathbf{T}} \epsilon_{\mathbf{T}})^2 \quad (14)$$

Let J_∞^* and \hat{J}_∞ denote the infinite time cumulative cost incurred by $\mathbf{K}_{1:s}$ and $\hat{\mathbf{K}}_{1:r}$ respectively. Then, when $\sigma_w = 0$, $\epsilon_{\mathbf{T}} \leq \bar{\epsilon}_{\mathbf{T}}$, and $\max\{\epsilon_{\mathbf{A}}, \epsilon_{\mathbf{B}}\} \leq \bar{\epsilon}_{\mathbf{A}, \mathbf{B}}$,

$$\hat{J}_\infty - J_\infty^* \leq (C'_{\mathbf{A}, \mathbf{B}} \max\{\epsilon_{\mathbf{A}}, \epsilon_{\mathbf{B}}\} + C'_{\mathbf{T}} \epsilon_{\mathbf{T}}) \|\mathbf{x}_0\|^2, \quad (15)$$

for some constants $C'_{\mathbf{A}, \mathbf{B}}$ and $C'_{\mathbf{T}}$.

Proof: We will use Lemma 3 and $\bar{\Sigma}$ and $\bar{\mathbf{K}}_{1:s}$ as a bridge to compare $\hat{\mathbf{K}}_{1:r}$ and $\mathbf{K}_{1:s}$. First, we prove (14). Comparing $\bar{\Sigma}$ and Σ , one can see $\|\bar{\mathbf{A}}_i - \mathbf{A}_i\| \leq \epsilon_{\mathbf{A}}$, $\|\bar{\mathbf{B}}_i - \mathbf{B}_i\| \leq \epsilon_{\mathbf{B}}$, and $\|\bar{\mathbf{T}} - \mathbf{T}\|_\infty \leq \epsilon_{\mathbf{T}}$. Then, from [54, Theorem 6] we know when $\max\{\epsilon_{\mathbf{A}}, \epsilon_{\mathbf{B}}\} \leq \bar{\epsilon}_{\mathbf{A}, \mathbf{B}}$ and $\epsilon_{\mathbf{T}} \leq \bar{\epsilon}_{\mathbf{T}}$ for some constants $\bar{\epsilon}_{\mathbf{A}, \mathbf{B}}$ and $\bar{\epsilon}_{\mathbf{T}}$, the Riccati solution $\bar{\mathbf{P}}_{1:s}$ uniquely exists among \mathbb{S}_s^+ and are positive definite, and the cost \bar{J} when applying $\bar{\mathbf{K}}_{1:s}$ to $\bar{\Sigma}$ has suboptimality $\bar{J} - J^* \leq \sigma_w^2 (C_{\mathbf{A}, \mathbf{B}} \max\{\epsilon_{\mathbf{A}}, \epsilon_{\mathbf{B}}\} + C_{\mathbf{T}} \epsilon_{\mathbf{T}})$ for some constants $C_{\mathbf{A}, \mathbf{B}}$ and $C_{\mathbf{T}}$. Using Lemma 3, we know $\hat{\mathbf{P}}_{1:r}$ uniquely exists \mathbb{S}_r^+ , and $\hat{\mathbf{K}}_k = \bar{\mathbf{K}}_i$ for any i belonging to any $\hat{\Omega}_k$, which implies applying $\bar{\mathbf{K}}_{1:s}$ is equivalent to applying $\hat{\mathbf{K}}_{1:r}$, as in the theorem statement. Thus $\bar{J} = \hat{J}$, and $\hat{J} - J^* = \bar{J} - J^* \leq \sigma_w^2 (C_{\mathbf{A}, \mathbf{B}} \max\{\epsilon_{\mathbf{A}}, \epsilon_{\mathbf{B}}\} + C_{\mathbf{T}} \epsilon_{\mathbf{T}})$.

Next, we prove (15). Similar as above, we let \bar{J}_∞ denote the cumulative cost when applying $\bar{\mathbf{K}}_{1:s}$ to Σ , then we have $\hat{J}_\infty = \bar{J}_\infty$. From the proof of [46, Theorem 4.5], we have $\bar{J}_\infty - J_\infty^* = \sum_{t=0}^{\infty} \mathbb{E}[\|\mathbf{M}_{\omega_t}(\bar{\mathbf{K}}_{\omega_t} - \bar{\mathbf{K}}_{\omega_t}^*) \mathbf{x}_t\|^2]$ where $\mathbf{M}_i = \mathbf{R} + \mathbf{B}_i^\top \varphi_i(\bar{\mathbf{P}}_{1:s}) \mathbf{B}_i$ and \mathbf{x}_t is driven by controller $\bar{\mathbf{K}}_{1:s}$. From [54, Theorem 6], we know when $\max\{\epsilon_{\mathbf{A}}, \epsilon_{\mathbf{B}}\} \leq \bar{\epsilon}_{\mathbf{A}, \mathbf{B}}$ and $\epsilon_{\mathbf{T}} \leq \bar{\epsilon}_{\mathbf{T}}$, then $\bar{\mathbf{K}}_{1:s}$ is a stabilizing controller and $\|\bar{\mathbf{K}}_{1:s} - \bar{\mathbf{K}}_{1:s}^*\| \leq C_{\mathbf{A}, \mathbf{B}}^{\mathbf{K}} \max\{\epsilon_{\mathbf{A}}, \epsilon_{\mathbf{B}}\} + C_{\mathbf{T}}^{\mathbf{K}} \epsilon_{\mathbf{T}}$ for some constants $C_{\mathbf{A}, \mathbf{B}}^{\mathbf{K}}$ and $C_{\mathbf{T}}^{\mathbf{K}}$. Following [11, Lemma 10], we know $\sum_{t=0}^{\infty} \mathbb{E}[\|\mathbf{x}_t\|^2] \leq C_{\mathbf{x}} \|\mathbf{x}_0\|^2$ for some constant $C_{\mathbf{x}}$. Combining these results, we have $\bar{J}_\infty - J_\infty^* \leq \|\mathbf{M}_{1:s}\| \|\bar{\mathbf{K}}_{1:s} - \bar{\mathbf{K}}_{1:s}^*\| \sum_{t=0}^{\infty} \mathbb{E}[\|\mathbf{x}_t\|^2] \leq \|\mathbf{M}_{1:s}\| C_{\mathbf{x}} (C_{\mathbf{A}, \mathbf{B}}^{\mathbf{K}} \max\{\epsilon_{\mathbf{A}}, \epsilon_{\mathbf{B}}\} + C_{\mathbf{T}}^{\mathbf{K}} \epsilon_{\mathbf{T}}) \|\mathbf{x}_0\|^2$. ■

In Theorem 6, constants $\bar{\epsilon}_{\mathbf{A}, \mathbf{B}}$, $\bar{\epsilon}_{\mathbf{T}}$, $C_{\mathbf{A}, \mathbf{B}}$, $C_{\mathbf{T}}$, $C'_{\mathbf{A}, \mathbf{B}}$, and $C'_{\mathbf{T}}$ only depend on the original MJS Σ and cost matrices \mathbf{Q} and \mathbf{R} , and their exact expressions can be obtained following the proof and corresponding references. As a sanity check, when there is no perturbation, i.e., mode-reducible case, then we have $\hat{J} = J^*$ and $\hat{J}_\infty = J_\infty^*$, which can also be implied from Lemma 3. For the reduced MJS $\bar{\Sigma}$, its Riccati solution $\bar{\mathbf{P}}_{1:r}$ and thus controllers $\hat{\mathbf{K}}_{1:r}$ are guaranteed to exist when perturbation $\epsilon_{\mathbf{A}}, \epsilon_{\mathbf{B}}$, and $\epsilon_{\mathbf{T}}$ are small enough as required in Theorem 6. The additive noise in Theorem 6 means the MJS dynamics is given by $\mathbf{x}_{t+1} = \mathbf{A}_{\omega_t} \mathbf{x}_t + \mathbf{B}_{\omega_t} \mathbf{u}_t + \mathbf{w}_t$ where $\mathbf{w}_t \sim \mathcal{N}(0, \sigma_w^2 \mathbf{I}_n)$. In the noisy case, both \hat{J}_∞ and J_∞^* are infinite, so the cumulative suboptimality $\hat{J}_\infty - J_\infty^*$ is only studied for the noise-free case as in (15). On the other hand, in the noise-free case, we have not only $J^* = \hat{J}$ as implied by (14), but also $J^* = \hat{J} = 0$ as long as $\hat{\mathbf{K}}_{1:r}$ is stabilizing.

VIII. NUMERICAL EXPERIMENTS

In this section, we present synthetic experiments to evaluate the main results in the paper. We evaluate the clustering performance of Algorithm 1 and the LQR controller designed with the reduced MJS $\bar{\Sigma}$ as discussed in Section VII. All the experiments are performed using MATLAB R2020a on a laptop with Xeon E3-1505M CPU. We use the `kmeans()` function from the Statistics and Machine Learning Toolbox in MATLAB for the k-means problem in Algorithm 1.

A. CLUSTERING EVALUATION

We consider the uniform partition $\Omega_{1:r}$, i.e. $|\Omega_i| = \bar{s} := s/r$ for any i . The system Σ is randomly generated according to P1 or P2 with desired levels of perturbation $\epsilon_{\mathbf{A}}, \epsilon_{\mathbf{B}}, \epsilon_{\mathbf{T}}$ so that in (2) each summand $\|\mathbf{A}_i - \mathbf{A}_{i'}\| \leq \epsilon_{\mathbf{A}}/(r\bar{s}^2)$. The same applies to $\mathbf{B}_{1:s}$ and \mathbf{T} . Specifically, we first randomly generate a small scale MJS $\bar{\Sigma} = \text{MJS}(\bar{\mathbf{A}}_{1:r}, \bar{\mathbf{B}}_{1:r}, \bar{\mathbf{T}})$: we sample each matrix element in $\bar{\mathbf{A}}_k$ and $\bar{\mathbf{B}}_k$ from standard Gaussian distributions and then scale the matrices so that each

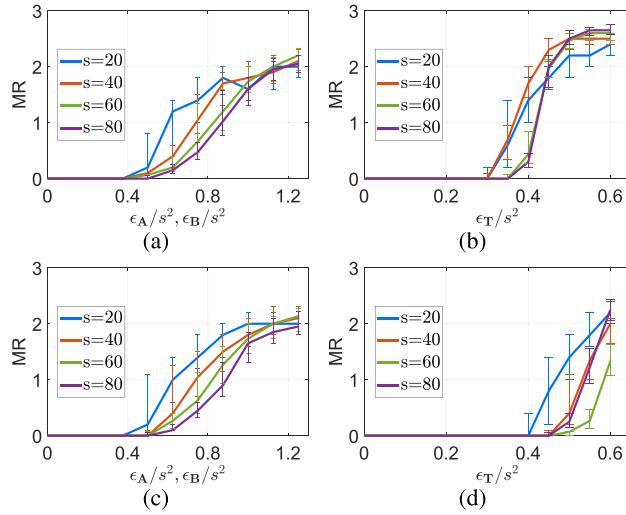


FIGURE 2. MR (median, first and third quartiles) vs number of modes s vs perturbation level. First row: aggregatable case P2. Second row: lumpable case P1. First column: $\epsilon_T = 0.5s^2$, $\alpha_A \propto 1/\max_i \|\mathbf{A}_i\|$, $\alpha_B \propto 1/\max_i \|\mathbf{B}_i\|$, $\alpha_T \propto 0.01/\|\mathbf{T}\|$. Second column: $\epsilon_A, \epsilon_B = 0.5s^2$, $\alpha_A \propto 0.01/\max_i \|\mathbf{A}_i\|$, $\alpha_B \propto 0.01/\max_i \|\mathbf{B}_i\|$, $\alpha_T \propto 1/\|\mathbf{T}\|$.

$\|\check{\mathbf{A}}_k\| = 0.5$ and $\|\check{\mathbf{B}}_k\| = 1$ unless otherwise mentioned; and each $\check{\mathbf{T}}(i, :)$ is sampled from the flat Dirichlet distribution. Then, we generate Σ by augmenting $\check{\Sigma}$. For every mode $i \in \Omega_k$, we let $\mathbf{A}_i = \check{\mathbf{A}}_k + \mathbf{E}_i$ and $\mathbf{B}_i = \check{\mathbf{B}}_k + \mathbf{F}_i$ where we sample elements in \mathbf{E}_i and \mathbf{F}_i from standard Gaussian and then scale them so that $\|\mathbf{E}_i\|_F = \frac{\epsilon_A}{2rs^2}$ and $\|\mathbf{F}_i\|_F = \frac{\epsilon_B}{2rs^2}$. The generation of \mathbf{T} is a bit involved. For the aggregatable case P2, we first generate a Markov matrix $\check{\mathbf{T}} \in \mathbb{R}^{s \times s}$ such that for every $i \in \Omega_k$, $\check{\mathbf{T}}(i, \Omega_l) = \mathbf{a}_{k,l} \check{\mathbf{T}}(k, l)$ where $\mathbf{a}_{k,l} \in \mathbb{R}^{1 \times |\Omega_l|}$ is sampled from the flat Dirichlet distribution; then we let $\mathbf{T}(i, :) = (1 - \frac{\epsilon_T}{2rs^2})\check{\mathbf{T}}(i, :) + \frac{\epsilon_T}{2rs^2}\mathbf{b}_i$ where $\mathbf{b}_i \in \mathbb{R}^{1 \times s}$ is again sampled from the flat Dirichlet distribution. The same steps are used to generate \mathbf{T} for the lumpable case P1 except that $\check{\mathbf{T}}(i, \Omega_l) = \mathbf{a}_{i,l} \check{\mathbf{T}}(k, l)$. Following these steps, Σ satisfies the perturbation conditions in P1 and P2.

To evaluate Algorithm 1, we fix $n = 5$, $p = 3$, $r = 4$ and record the misclustering rate (MR) defined in Section IV-A over 100 runs. Fig. 2 presents the clustering performances under different number of modes s and perturbations ϵ_A , ϵ_B and ϵ_T . In the plots, we normalized the perturbation on the x-axis by s^2 so that the trends under different s can be better visualized. This also follows from the experiment setup: each summand in (2) has $\|\mathbf{A}_i - \mathbf{A}_{i'}\| \leq \mathcal{O}(\epsilon_A/s^2)$. It is clear that the clustering performance degrades with increasing s and perturbations. We can also observe that when the perturbation is small, there are no misclustered modes.

B. LQR CONTROLLER DESIGN

Then, we implement the idea of designing LQR controllers for the original Σ through the reduced $\hat{\Sigma}$ as discussed in Section VII. We let $r = 4$, $n = 10$, $p = 5$, and the system dynamics is generated the same way as the previous section. The process noise variance is $\sigma_w^2 = 0.1$, and the initial

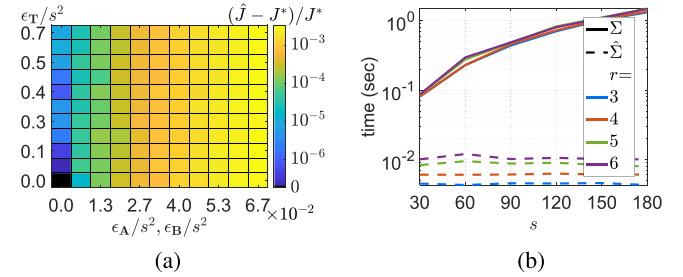


FIGURE 3. (a) LQR suboptimality (median) vs perturbations; (b) LQR computation time (median) for the original MJS Σ and the reduced MJS $\hat{\Sigma}$ with different number of modes and clusters. (We omit the quartiles as they are very close to the median.)

TABLE 2. Suboptimality vs Computation Time vs Selected Number of Modes in $\hat{\Sigma}$

\hat{r}	10	20	30	40	50
$\frac{\hat{J} - J^*}{J^*}$	$1.2e-1$	$3.8e-2$	$4.1e-7$	$6.9e-5$	$8.9e-4$
Time (sec)	$3.7e-2$	$6.8e-2$	$1.4e-1$	$2.2e-1$	$3.2e-1$
\hat{r}	60	70	80	90	100
$\frac{\hat{J} - J^*}{J^*}$	$7.1e-3$	$1.6e-2$	$2.1e-2$	$1.5e-2$	0
Time (sec)	$6.8e-1$	$8.4e-1$	$1.0e0$	$1.1e0$	$1.2e0$

state is $\mathbf{x}_0 = \mathbf{1}$. Fig. 3(a) shows the suboptimality against perturbations for $s = 100$. As one would expect, the suboptimality increases with the perturbation levels and is 0 when there is no perturbation. The trend on ϵ_T is evident when ϵ_A and ϵ_B are small but imperceptible for larger values of ϵ_A and ϵ_B . Fig. 3(b) shows the time to compute controllers via Riccati iterations using Σ and $\hat{\Sigma}$ as a function of s . The computation terminates when the controller difference between two consecutive iterations falls below 10^{-12} . We see when s is large, Σ needs significantly more time than $\hat{\Sigma}$.

Next, we consider a more practical scenario where one has no knowledge of the true number of cluster r , and replace it in Algorithm 1 with a hyper-parameter \hat{r} as the number of modes in $\hat{\Sigma}$. We fix $s = 100$ and $r = 30$, and the rest of the experiment setup is the same as Fig. 3. We record the suboptimality and computation time under different choices of \hat{r} in Table 2. When increases \hat{r} , the suboptimality achieves the minimum when $\hat{r} = r = 30$ and then gradually increases until $\hat{r} = s$, i.e. no reduction is performed at all. This comes as a bit of surprise as one would expect no worse performance when using more clusters than needed. Further investigation suggests that when $\hat{r} > r$, misclustering occurs more frequently than the case of $\hat{r} = r$, which is likely to account for the performance degradation. In practice, to find the best \hat{r} , one could try multiple \hat{r} in Algorithm 1, plug in the resulting partitions into P1 and P2, and select the one that gives the smallest perturbation $\epsilon_A, \epsilon_B, \epsilon_T$.

C. TRAJECTORY APPROXIMATION

In this section, we evaluate the trajectory approximation results from Section V. Let $\theta = \pi/16$, $\check{\mathbf{A}}_1 = [[\cos(\theta),$

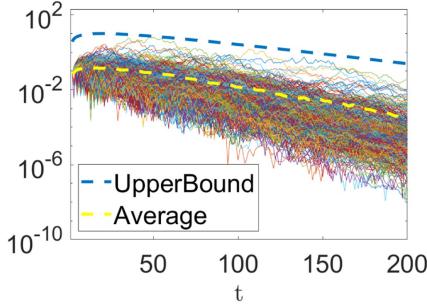


FIGURE 4. Trajectory difference and the upper bound.

$\sin(\theta)]^T, [-\sin(\theta), \cos(\theta)]^T]^T$, $\check{A}_2 = [[0.8, 0]^T, [0, 0.8]^T]^T$, and $\check{A}_3 = [[1.2, 0]^T, [0, 1.2]^T]^T$. Then, we construct an autonomous MJS Σ with 6 modes: for $k = \{1, 2, 3\}$, $A_{2k-1} = \check{A}_k + [[0.1, 0]^T, [0, 0.1]^T]^T$ and $A_{2k} = \check{A}_k - [[0.1, 0]^T, [0, 0.1]^T]^T$. The uniform partition $\{\{1, 2\}, \{3, 4\}, \{5, 6\}\}$ gives $\epsilon_A = 0.6\sqrt{2}$ according to P1. Define T such that for all i , $T(i, j) = 0.2$ if $j \in \{1, 2, 3, 4\}$ and $T(i, j) = 0.1$ if $j \in \{5, 6\}$. By relevant definitions in Section V, the constructed Σ is MSS but not uniformly stable.

We fix the initial state $x_0 = [1, 1]^T$, generate 500 independent trajectories for states x_t and \hat{x}_t , and record the difference $\|x_t - \hat{x}_t\|$. In Fig. 4, each thin solid line represents the difference, in log-scale, for each trajectory, the yellow dashed line shows their average, and the blue dashed line depicts the upper bound in Theorem 3. Throughout the time horizon, though not very tight, the theoretical upper bound stays above the averaged difference. Note that, for a given δ , by Markov inequality, shifting the upper bound in the plot upward by $\log(\delta)$ would give a bound on the individual error trajectories with probability $1 - \delta$. As seen in the figure, even the non-shifted version serves a good bound for individual error trajectories.

IX. CONCLUSION AND FUTURE WORK

In this work, we propose a clustering-based method to reduce the number of modes in an MJS. The reduced MJS provably well approximates the original MJS in terms of trajectory, transition kernels, stability, and controller optimality. One future direction could be the generalization of the offline controller design scheme in Section VII to settings where controllers need to be computed in the runtime, such as model predictive control and adaptive control. In these problems, the savings of computation time would be even more prominent. Another potential future direction could be the extension of the fully observed MJS in this work to partially observed MJS, i.e., the state x_t is observed through $y_t = C_{\omega_t} x_t$ for some mode-dependent output matrices $C_{1:s}$. As a side note, the similarity between MJS and Markov decision processes (MDP) hints that the framework and principles developed for MJS in this work may also help the complexity reduction of MDP and reinforcement learning problems.

APPENDIX A

AGGREGATABLE CLUSTERING – PROOF FOR THEOREM 1

We first provide several supporting lemmas. The first one is regarding the perturbation of the left singular vector space.

Lemma 4 (Singular Vectors Perturbation Bound): Consider two arbitrary matrices $\check{\Phi}, \Phi \in \mathbb{R}^{s \times r}$. Let $\bar{U}, U \in \mathbb{R}^{s \times r}$ respectively denote the top- r left singular vectors of $\check{\Phi}$ and Φ with $\bar{U}^T \bar{U} = U^T U = I_r$. Then

$$\min_{O \in \mathcal{O}(r)} \|\bar{U}O - U\|_F \leq \frac{2\sqrt{2}\|\check{\Phi} - \Phi\|_F}{\sigma_r(\check{\Phi}) - \sigma_{r+1}(\check{\Phi})}, \quad (16)$$

where $\mathcal{O}(r)$ denotes the set of all $r \times r$ orthonormal matrices.

This result can be seen simply by combining Lemma 10 and Lemma 11 in [17], where Lemma 10 requires a trivial generalization from spectral norm to the Frobenius norm. The next result says if a matrix has certain rows being identical, its singular vectors share the same identity pattern.

Lemma 5 (Lemma 12 in [17]): Consider a matrix $\check{\Phi} \in \mathbb{R}^{s \times r}$ and a partition $\Omega_{1:r}$ on $[s]$ such that for any $i, i' \in \Omega_k$, $\check{\Phi}(i, :) = \check{\Phi}(i', :)$. Assume $\text{rank}(\check{\Phi}) = r$. Let $\bar{U} \in \mathbb{R}^{s \times r}$ denote the top- r left singular vectors of $\check{\Phi}$ with $\bar{U}^T \bar{U} = I_r$. Then for any $i \in \Omega_k$ and $j \in \Omega_l$, $\|\bar{U}(i, :) - \bar{U}(j, :)\| = (\frac{1}{|\Omega_k|} + \frac{1}{|\Omega_l|})^{0.5}$ if $k \neq l$ and 0 if $k = l$.

The next lemma provides a preliminary result on the performance of k-means when it is applied to a data matrix with feature dimension same as the number of clusters.

Lemma 6 (Lemma 5.3 in [55]): Consider two arbitrary matrices $\bar{U}, U \in \mathbb{R}^{s \times r}$ with $\Delta_U := \|\bar{U} - U\|_F$. Suppose there exists a partition $\Omega_{1:r}$ on $[s]$ such that for any $i, i' \in \Omega_k$, $\bar{U}(i, :) = \bar{U}(i', :)$. Define the inter-cluster distance for cluster k as $\delta_k := \min_{l \in [r] \setminus k} \min_{i \in \Omega_k, j \in \Omega_l} \|\bar{U}(i, :) - \bar{U}(j, :)\|$. Let $\{\hat{\Omega}_{1:r}, \hat{c}_{1:r}\}$ be a $(1 + \epsilon)$ solution to the k-means problem on the rows of U . Then, when $\Delta_U \leq \frac{\min_k \sqrt{|\Omega_k|} \delta_k}{\sqrt{8(2+\epsilon)}}$, we have

$$\min_{h \in \mathcal{H}} \sum_{k \in [r]} |\{i : i \in \Omega_k, i \notin \hat{\Omega}_{h(k)}\}| \cdot \delta_k^2 \leq 8(2 + \epsilon) \Delta_U^2, \quad (17)$$

where \mathcal{H} is the set of all bijections from $[r]$ to $[r]$.

By combining Lemma 4, 5, and 6, we obtain guarantee on the performance of k-means when it is applied to the left singular vectors of the data matrix, which is the key lemma we will use to show Theorem 1 and Theorem 2.

Lemma 7 (Approximate k-means error bound): Consider two arbitrary matrices $\check{\Phi}, \Phi \in \mathbb{R}^{s \times r}$ with $\Delta_\Phi := \|\check{\Phi} - \Phi\|_F$. Suppose there exists a partition $\Omega_{1:r}$ on $[s]$ such that for any $i, i' \in \Omega_k$, $\check{\Phi}(i, :) = \check{\Phi}(i', :)$. Assume $\text{rank}(\check{\Phi}) = r$. Let $U \in \mathbb{R}^{s \times r}$ denote the top- r left singular vectors of Φ with $U^T U = I_r$. Let $\{\hat{\Omega}_{1:r}, \hat{c}_{1:r}\}$ be a $(1 + \epsilon)$ solution to the k-means problem on clustering the rows of U . Then, when $\Delta_\Phi \leq \frac{\sigma_r(\check{\Phi}) \sqrt{|\Omega_{(r)}| + |\Omega_{(1)}|}}{8\sqrt{(2+\epsilon)|\Omega_{(1)}|}}$, we have $\text{MR}(\hat{\Omega}_{1:r}) \leq \frac{64(2+\epsilon)}{\sigma_r(\check{\Phi})^2} \Delta_\Phi^2$.

Proof: Let $\bar{U} \in \mathbb{R}^{s \times r}$ denote the top- r left singular vectors of $\check{\Phi}$ with $\bar{U}^T \bar{U} = I_r$. Then, Lemma 4 implies that exists $O^* \in \mathcal{O}(r)$ such that $\|\bar{U}O^* - U\|_F \leq \frac{2\sqrt{2}\Delta_\Phi}{\sigma_r(\check{\Phi})}$. Note that $\|[\bar{U}O^*](i, :) - [\bar{U}O^*](j, :)\| = \|(\bar{U}(i, :) - \bar{U}(j, :))O^*\| = \|\bar{U}(i, :) - \bar{U}(j, :)\|$. By Lemma 5, we know for any $i \in \Omega_k$,

$j \in \Omega_l$, $\|[\bar{\mathbf{U}}\mathbf{O}^*](i, :) - [\bar{\mathbf{U}}\mathbf{O}^*](j, :)\| = \sqrt{\frac{1}{|\Omega_k|} + \frac{1}{|\Omega_l|}}$ if $k \neq l$ and 0 if $k = l$. Then, for any $k \in [r]$, let $\delta_k := \min_{l \in [r] \setminus k} \min_{i \in \Omega_k, j \in \Omega_l} \|[\bar{\mathbf{U}}\mathbf{O}^*](i, :) - [\bar{\mathbf{U}}\mathbf{O}^*](j, :)\|$, we see $\delta_k \geq \sqrt{\frac{1}{|\Omega_k|} + \frac{1}{|\Omega_{(1)}|}}$.

Note that when $\Delta_\Phi \leq \frac{\sigma_r(\bar{\Phi})\sqrt{|\Omega_{(r)}| + |\Omega_{(1)}|}}{8\sqrt{(2+\epsilon)|\Omega_{(1)}|}}$, one can check $\|\bar{\mathbf{U}}\mathbf{O}^* - \mathbf{U}\|_F \leq \frac{\min_k \sqrt{|\Omega_k|} \delta_k}{\sqrt{8(2+\epsilon)}}$. Then, by Lemma 6, we have $\text{MR}(\hat{\Omega}_{1:r}) = \min_{h \in \mathcal{H}} \sum_{k=1}^r |\{i : i \in \Omega_k, i \notin \hat{\Omega}_{h(k)}\}| \frac{1}{|\Omega_k|} \leq \min_{h \in \mathcal{H}} \sum_{k=1}^r |\{i : i \in \Omega_k, i \notin \hat{\Omega}_{h(k)}\}| \delta_k^2 \leq 64(2+\epsilon) \sigma_r(\bar{\Phi})^{-2} \Delta_\Phi^2$. ■

Main Proof for Theorem 1: Consider Φ in Algorithm 1 Line 3 and its averaged version $\bar{\Phi}$ defined in Section IV-A. Then, by definition, we have $\|\bar{\Phi} - \Phi\|_F^2 = \sum_{k \in [r]} \sum_{i \in \Omega_k} \|\bar{\Phi}(i, :) - \Phi(i, :)\|_F^2 = \alpha_{\mathbf{T}}^2 \cdot \sum_{k \in [r]} \sum_{i \in \Omega_k} \|\mathbf{T}(i, :) - |\Omega_k|^{-1} \cdot \sum_{i' \in \Omega_k} \mathbf{T}(i', :)\|^2 + \alpha_{\mathbf{A}}^2 \cdot \sum_{k \in [r]} \sum_{i \in \Omega_k} \|\mathbf{A}_i - |\Omega_k|^{-1} \cdot \sum_{i' \in \Omega_k} \mathbf{A}_{i'}\|_F^2 + \alpha_{\mathbf{B}}^2 \cdot \sum_{k \in [r]} \sum_{i \in \Omega_k} \|\mathbf{B}_i - |\Omega_k|^{-1} \sum_{i' \in \Omega_k} \mathbf{B}_{i'}\|_F^2$. By the definitions of $\epsilon_{\mathbf{A}}, \epsilon_{\mathbf{B}}, \epsilon_{\mathbf{T}}$ in Problem P2, triangle inequality, and Cauchy-Schwarz inequality, we have $\|\bar{\Phi} - \Phi\|_F \leq \epsilon_{\text{Agg}}$ where $\epsilon_{\text{Agg}} := \sqrt{\alpha_{\mathbf{A}}^2 \epsilon_{\mathbf{A}}^2 + \alpha_{\mathbf{B}}^2 \epsilon_{\mathbf{B}}^2 + \alpha_{\mathbf{T}}^2 \epsilon_{\mathbf{T}}^2}$. By construction, in matrix $\bar{\Phi}$, rows that belong to the same cluster are identical, thus we can apply Lemma 7 to $\{\Phi, \bar{\Phi}\}$ and obtain that when $\epsilon_{\text{Agg}} \leq \frac{\sigma_r(\bar{\Phi})\sqrt{|\Omega_{(r)}| + |\Omega_{(1)}|}}{8\sqrt{(2+\epsilon)|\Omega_{(1)}|}}$, we have $\text{MR}(\hat{\Omega}_{1:r}) \leq 64(2+\epsilon) \sigma_r(\bar{\Phi})^{-2} \epsilon_{\text{Agg}}^2$. ■

APPENDIX B

LUMPABLE CLUSTERING – PROOF FOR THEOREM 2

We first provide a supporting result regarding the perturbation of stationary distribution of Markov chains.

Lemma 8 (Section 3.6 in [56]): For two Markov matrices $\mathbf{T}, \mathbf{T}_0 \in \mathbb{R}^{s \times s}$ and their stationary distributions $\pi, \pi_0 \in \mathbb{R}^s$, we have $\|\pi - \pi_0\|_1 \leq \gamma_1 \|\mathbf{T} - \mathbf{T}_0\|_\infty$, where $\gamma_1 := \sum_{i=2}^s \frac{1}{1 - \lambda_i(\mathbf{T})}$.

When the difference $\|\mathbf{T} - \mathbf{T}_0\|$ is small, we can further have the following corollary.

Corollary 1: In Lemma 8, let $\pi_{\min} := \min_i \pi(i)$, $\pi_{\max} := \max_i \pi(i)$. Suppose $\|\mathbf{T} - \mathbf{T}_0\|_\infty \leq \frac{\pi_{\min}}{\gamma_1}$, then we have

$$\max_i |\pi(i) - \pi_0(i)| \leq \frac{\pi_{\min}}{2}, \quad (18)$$

$$\min_i \pi_0(i) \geq \frac{\pi_{\min}}{2}, \quad \max_i \pi_0(i) \leq \pi_{\max} + \frac{\pi_{\min}}{2} \quad (19)$$

$$\max_i |\pi(i)^{-\frac{1}{2}} - \pi_0(i)^{-\frac{1}{2}}| \leq (\sqrt{2} - 1) \gamma_1 \pi_{\min}^{-\frac{3}{2}} \|\mathbf{T} - \mathbf{T}_0\|_\infty \quad (20)$$

$$\max_i |\pi(i)^{\frac{1}{2}} - \pi_0(i)^{\frac{1}{2}}| \leq (1 - \frac{\sqrt{2}}{2}) \gamma_1 \pi_{\min}^{-\frac{1}{2}} \|\mathbf{T} - \mathbf{T}_0\|_\infty. \quad (21)$$

Proof: Since $\mathbf{1}^\top \pi = \mathbf{1}^\top \pi_0 = 1$, we have $\max_i |\pi(i) - \pi_0(i)| \leq \frac{1}{2} \|\pi - \pi_0\|_1 \leq \frac{\gamma_1}{2} \|\mathbf{T} - \mathbf{T}_0\|_\infty \leq \frac{\pi_{\min}}{2}$. Then using triangle inequality, we can show (18) and (19). Note that the LHS of (20) is equivalent to $\max_i \frac{|\pi_0(i) - \pi(i)|}{\sqrt{\pi(i)\pi_0(i)}(\sqrt{\pi(i)} + \sqrt{\pi_0(i)})}$, then plugging in (19) gives (20). And (21) follows similarly. ■

When the lumpability perturbation $\epsilon_{\mathbf{T}} \neq 0$, matrix \mathbf{S}_r in Algorithm 1 Line 7 no longer has the row identity pattern as discussed in Lemma 1. The next result measures this effect.

Lemma 9: Consider an ergodic Markov matrix $\mathbf{T} \in \mathbb{R}^{s \times s}$ with stationary distribution π and a partition $\Omega_{1:r}$ such that it is approximately lumpable as in (3) with perturbation $\epsilon_{\mathbf{T}}$. Consider the neighborhood of \mathbf{T} given by $\mathcal{L}(\mathbf{T}, \Omega_{1:r}, \epsilon_{\mathbf{T}})$ defined in (6). Assume there exists an ergodic and reversible $\mathbf{T}_0 \in \mathcal{L}(\mathbf{T}, \Omega_{1:r}, \epsilon_{\mathbf{T}})$ that has informative spectrum. Construct $\mathbf{S}_r \in \mathbb{R}^{s \times r}$ with \mathbf{T} and π as in Algorithm 1 Line 7. Construct $\bar{\mathbf{S}}_r \in \mathbb{R}^{s \times r}$ such that for any $i \in [s]$ (suppose $i \in \Omega_k$), $\bar{\mathbf{S}}_r(i, :) = \frac{1}{|\Omega_k|} \sum_{i' \in \Omega_k} \mathbf{S}_r(i', :)$. Let $\pi_{\min} := \max_i \pi(i)$, $\pi_{\max} := \min_i \pi(i)$, $\gamma_1 := \sum_{i=2}^s \frac{1}{1 - \lambda_i(\mathbf{T})}$, $\gamma_2 := \min\{\sigma_r(\mathbf{H}) - \sigma_{r+1}(\mathbf{H}), 1\}$, and $\gamma_3 := \frac{16\gamma_1\sqrt{r}\pi_{\max} \|\mathbf{T}\|_F}{\gamma_2 \pi_{\min}^2}$ where \mathbf{H} is defined in Algorithm 1. Then, when perturbation $\epsilon_{\mathbf{T}} \leq \frac{\pi_{\min}}{\gamma_1}$, we have $\|\mathbf{S}_r - \bar{\mathbf{S}}_r\|_F \leq \gamma_3 \epsilon_{\mathbf{T}}$.

Proof: We will start with analyzing \mathbf{T}_0 and use it as a bridge to prove the claim. Let $\pi_0 \in \mathbb{R}^s$ denote the stationary distribution of \mathbf{T}_0 . Since \mathbf{T}_0 is ergodic, we know π_0 is strictly positive. By definition of reversibility, we know $\text{diag}(\pi_0)\mathbf{T}_0 = \mathbf{T}_0^\top \text{diag}(\pi_0)$, and this further gives $\text{diag}(\pi_0)^{\frac{1}{2}} \mathbf{T}_0 \text{diag}(\pi_0)^{-\frac{1}{2}} = \text{diag}(\pi_0)^{-\frac{1}{2}} \mathbf{T}_0^\top \text{diag}(\pi_0)^{\frac{1}{2}}$. Let $\mathbf{H}_0 := \text{diag}(\pi_0)^{\frac{1}{2}} \mathbf{T}_0 \text{diag}(\pi_0)^{-\frac{1}{2}}$, then we see \mathbf{H}_0 is symmetric. Let $\mathbf{W}_{0,r} \in \mathbb{R}^{s \times r}$ denote the top r left singular vectors of \mathbf{H}_0 , by spectral theorem, we know the columns of $\mathbf{W}_{0,r}$ also serve as the top r eigenvectors of \mathbf{H}_0 . Let $\mathbf{S}_{0,r} := \text{diag}(\pi_0)^{-\frac{1}{2}} \mathbf{W}_{0,r}$, by definition of \mathbf{H}_0 , it is easy to see that the columns of $\mathbf{S}_{0,r}$ are also the top r eigenvectors of \mathbf{T}_0 . Then, by Lemma 1 and the definition of informative spectrum, for any $i, i' \in \Omega_k$, we have $\mathbf{S}_{0,r}(i, :) = \mathbf{S}_{0,r}(i', :)$.

Recall in Algorithm 1, \mathbf{W}_r denotes the top r left singular vectors of $\mathbf{H} := \text{diag}(\pi)^{\frac{1}{2}} \mathbf{T} \text{diag}(\pi)^{-\frac{1}{2}}$ and let $\mathbf{S}_r = \text{diag}(\pi)^{-\frac{1}{2}} \mathbf{W}_r$. Let $\mathbf{O}^* := \min_{\mathbf{O} \in \mathcal{O}(r)} \|\mathbf{W}_{0,r}\mathbf{O} - \mathbf{W}_r\|_F$, where $\mathcal{O}(r)$ is the set of all $r \times r$ orthonormal matrices. Then, for any $i, i' \in \Omega_k$, we have $[\mathbf{S}_{0,r}\mathbf{O}^*](i, :) = [\mathbf{S}_{0,r}\mathbf{O}^*](i', :)$. Using this, for any $i \in [s]$ (suppose $i \in \Omega_k$), we have

$$\begin{aligned} & \mathbf{S}_r(i, :) - \bar{\mathbf{S}}_r(i, :) \\ &= \frac{|\Omega_k| - 1}{|\Omega_k|} \mathbf{S}_r(i, :) - \frac{1}{|\Omega_k|} \sum_{i': i' \in \Omega_k, i' \neq i} \mathbf{S}_r(i', :) \\ &\leq \frac{|\Omega_k| - 1}{|\Omega_k|} (\mathbf{S}_r(i, :) - [\mathbf{S}_{0,r}\mathbf{O}^*](i, :)) \\ &\quad + \frac{1}{|\Omega_k|} \sum_{i': i' \in \Omega_k, i' \neq i} ([\mathbf{S}_{0,r}\mathbf{O}^*](i', :) - \mathbf{S}_r(i', :)). \end{aligned} \quad (22)$$

WLOG, assume $\{1, \dots, |\Omega_1|\} = \Omega_1$, $\{|\Omega_1| + 1, \dots, |\Omega_1| + |\Omega_2|\} = \Omega_2, \dots$ and define block diagonal matrices $\mathbf{D}, \mathbf{P} \in \mathbb{R}^{s \times s}$ both with r diagonal blocks such that their k -th diagonal blocks $[\mathbf{D}]_k, [\mathbf{P}]_k \in \mathbb{R}^{|\Omega_k| \times |\Omega_k|}$ are given by

$$[\mathbf{D}]_k = \frac{|\Omega_k| - 1}{|\Omega_k|} \mathbf{I}_{|\Omega_k|}, \quad [\mathbf{P}]_k = \frac{1}{|\Omega_k|} (\mathbf{1}_{|\Omega_k|} \mathbf{1}_{|\Omega_k|}^\top - \mathbf{I}_{|\Omega_k|}). \quad (23)$$

Then, stacking (22) for all i , one can verify that $\mathbf{S}_r - \bar{\mathbf{S}}_r = \mathbf{D}(\mathbf{S}_r - \mathbf{S}_{0,r}\mathbf{O}^*) + \mathbf{P}(\mathbf{S}_{0,r}\mathbf{O}^* - \mathbf{S}_r)$. Note that for an arbitrary matrix $\mathbf{E} \in \mathbb{R}^{s \times s}$, we have $\|\mathbf{PE}\|_F^2 = \text{tr}(\mathbf{P}^T \mathbf{PE} \mathbf{E}^T) \leq \text{tr}(\mathbf{D}^T \mathbf{DE} \mathbf{E}^T) = \|\mathbf{DE}\|_F^2$ where the inequality holds since for each diagonal block we have $[\mathbf{P}]_k^T [\mathbf{P}]_k \preceq [\mathbf{D}]_k^T [\mathbf{D}]_k$. Therefore, $\|\mathbf{S}_r - \bar{\mathbf{S}}_r\|_F \leq 2\|\mathbf{D}(\mathbf{S}_r - \mathbf{S}_{0,r}\mathbf{O}^*)\|_F \leq 2\max_k \frac{|\Omega_k|-1}{|\Omega_k|} \|\mathbf{S}_r - \mathbf{S}_{0,r}\mathbf{O}^*\|_F \leq 2\|\mathbf{S}_r - \mathbf{S}_{0,r}\mathbf{O}^*\|_F$. To complete the proof, it suffices to study $\|\mathbf{S}_r - \mathbf{S}_{0,r}\mathbf{O}^*\|_F$.

$$\begin{aligned} & \|\mathbf{S}_r - \mathbf{S}_{0,r}\mathbf{O}^*\|_F \\ &= \|\mathbf{diag}(\pi)^{-\frac{1}{2}}(\mathbf{W}_r - \mathbf{W}_{0,r}\mathbf{O}^*)\|_F \\ & \quad + (\mathbf{diag}(\pi)^{-\frac{1}{2}} - \mathbf{diag}(\pi_0)^{-\frac{1}{2}})\mathbf{W}_{0,r}\mathbf{O}^*\|_F \\ & \leq \pi_{\min}^{-0.5} \|\mathbf{W}_r - \mathbf{W}_{0,r}\mathbf{O}^*\|_F + \sqrt{r} \max_i |\pi(i)^{-\frac{1}{2}} - \pi_0(i)^{-\frac{1}{2}}|. \end{aligned} \quad (24)$$

According to Lemma 4, we know $\|\mathbf{W}_r - \mathbf{W}_{0,r}\mathbf{O}^*\|_F \leq \frac{2\sqrt{2}}{\sigma_r(\mathbf{H}) - \sigma_{r+1}(\mathbf{H})} \|\mathbf{H} - \mathbf{H}_0\|_F$. This together with the upper bound for $\max_i |\pi(i)^{-\frac{1}{2}} - \pi_0(i)^{-\frac{1}{2}}|$ in (20) gives

$$\begin{aligned} \|\mathbf{S}_r - \mathbf{S}_{0,r}\mathbf{O}^*\|_F & \leq \frac{2\sqrt{2}}{(\sigma_r(\mathbf{H}) - \sigma_{r+1}(\mathbf{H}))\pi_{\min}^{0.5}} \|\mathbf{H} - \mathbf{H}_0\|_F \\ & \quad + \frac{(\sqrt{2}-1)\gamma_1\sqrt{r}}{\pi_{\min}^{1.5}} \epsilon_{\mathbf{T}}. \end{aligned} \quad (25)$$

By the definitions of \mathbf{H} and \mathbf{H}_0 , we have $\|\mathbf{H} - \mathbf{H}_0\|_F \leq \|(\mathbf{diag}(\pi)^{\frac{1}{2}} - \mathbf{diag}(\pi_0)^{\frac{1}{2}}) \cdot \mathbf{T} \cdot \mathbf{diag}(\pi)^{-\frac{1}{2}}\|_F + \|\mathbf{diag}(\pi_0)^{\frac{1}{2}} \cdot \mathbf{T} \cdot (\mathbf{diag}(\pi)^{-\frac{1}{2}} - \mathbf{diag}(\pi_0)^{-\frac{1}{2}})\|_F + \|\mathbf{diag}(\pi_0)^{\frac{1}{2}} \cdot (\mathbf{T} - \mathbf{T}_0) \cdot \mathbf{diag}(\pi_0)^{-\frac{1}{2}}\|_F$. Applying Corollary 1 gives $\|\mathbf{H} - \mathbf{H}_0\|_F \leq 2.56\gamma_1\pi_{\max}^{0.5}\pi_{\min}^{-1.5}\|\mathbf{T}\|_{\mathbf{F}}\epsilon_{\mathbf{T}}$. Plugging this into (25), we have $\|\mathbf{S}_r - \mathbf{S}_{0,r}\mathbf{O}^*\|_F \leq \frac{8\gamma_1\sqrt{r}\sqrt{\pi_{\max}}\|\mathbf{T}\|_F}{\gamma_2\pi_{\min}^2}\epsilon_{\mathbf{T}}$, where $\gamma_2 := \min\{\sigma_r(\mathbf{H}) - \sigma_{r+1}(\mathbf{H}), 1\}$. This concludes the proof as we showed that $\|\mathbf{S}_r - \bar{\mathbf{S}}_r\|_F \leq 2\|\mathbf{S}_r - \mathbf{S}_{0,r}\mathbf{O}^*\|_F$. ■

Main Proof for Theorem 2: Consider Φ in Algorithm 1 Line 3 and its averaged version $\bar{\Phi}$ defined in Section IV-A. Then, by definition, we have $\|\bar{\Phi} - \Phi\|_F^2 = \alpha_{\mathbf{T}}^2 \cdot \|\mathbf{S}_r - \bar{\mathbf{S}}_r\|_F^2 + \alpha_{\mathbf{A}}^2 \cdot \sum_{k \in [r]} \sum_{i \in \Omega_k} \|\mathbf{A}_i - |\Omega_k|^{-1}\|$. $\sum_{i' \in \Omega_k} \mathbf{A}_{i'}\|_F^2 + \alpha_{\mathbf{B}}^2 \cdot \sum_{k \in [r]} \sum_{i \in \Omega_k} \|\mathbf{B}_i - |\Omega_k|^{-1} \sum_{i' \in \Omega_k} \mathbf{B}_{i'}\|_F^2$, where $\bar{\mathbf{S}}_r$ is defined in Lemma 9. By Lemma 9 and the definitions of $\epsilon_{\mathbf{A}}$ and $\epsilon_{\mathbf{B}}$ in Problem P1, we have $\|\bar{\Phi} - \Phi\|_F \leq \epsilon_{Lmp}$ where $\epsilon_{Lmp} := \sqrt{\alpha_{\mathbf{A}}^2\epsilon_{\mathbf{A}}^2 + \alpha_{\mathbf{B}}^2\epsilon_{\mathbf{B}}^2 + \alpha_{\mathbf{T}}^2\gamma_3^2\epsilon_{\mathbf{T}}^2}$. By construction, in $\bar{\Phi}$, rows that belong to the same cluster have the same rows, thus we can apply Lemma 7 to $\{\bar{\Phi}, \Phi\}$ and obtain that when $\epsilon_{Lmp} \leq \frac{\sigma_r(\bar{\Phi})\sqrt{|\Omega_{(r)}| + |\Omega_{(1)}|}}{8\sqrt{(2+\epsilon)|\Omega_{(1)}|}}$, we have $\text{MR}(\hat{\Omega}_{1:r}) \leq 64(2+\epsilon)\sigma_r(\bar{\Phi})^{-2}\epsilon_{Lmp}^2$. ■

A NON-EMPTINESS OF $\mathcal{L}(\mathbf{T}, \Omega_{1:r}, \epsilon_{\mathbf{T}})$

Note that both Lemma 9 and Theorem 2 require the set $\mathcal{L}(\mathbf{T}, \Omega_{1:r}, \epsilon_{\mathbf{T}})$, a neighborhood of \mathbf{T} . Now, we show it is non-empty under the approximate lumpability condition (3).

Let $\mathbf{T}_0 := \mathbf{T} + \Delta$ for $\Delta \in \mathcal{D}$ where

$$\begin{aligned} \mathcal{D} := \left\{ \Delta \in \mathbb{R}^{s \times s} : \forall k, l \in [r], \forall i \in \Omega_k, \right. \\ \left. -\mathbf{T}(i, j) \leq \Delta(i, j) \leq 1 - \mathbf{T}(i, j) \quad \forall j \in [s], \right. \\ \left. \sum_{j \in \Omega_l} \Delta(i, j) = -\sum_{j \in \Omega_l} \mathbf{T}(i, j) + |\Omega_k|^{-1} \sum_{\substack{i' \in \Omega_k \\ j \in \Omega_l}} \mathbf{T}(i', j), \right. \\ \left. \|\Delta\|_F \leq \epsilon_{\mathbf{T}}, \quad \|\Delta\|_{\infty} \leq \epsilon_{\mathbf{T}}. \right\} \end{aligned} \quad (26)$$

Then, we see to show there exists $\mathbf{T}_0 \in \mathcal{L}(\mathbf{T}, \Omega_{1:r}, \epsilon_{\mathbf{T}})$, i.e. $\mathcal{L}(\mathbf{T}, \Omega_{1:r}, \epsilon_{\mathbf{T}})$ is non-empty, it is equivalent to show there exists $\Delta \in \mathcal{D}$.

Note that (26) gives that for all $i \in \Omega_k, l \in [r]$, $-\sum_{j \in \Omega_l} \mathbf{T}(i, j) \leq \sum_{j \in \Omega_l} \Delta(i, j) \leq |\Omega_l| - \sum_{j \in \Omega_l} \mathbf{T}(i, j)$. This together with (26) and (27) imply that there exists Δ satisfying both (26) and (27) such that among its elements $\{\Delta(i, j)\}_{j \in \Omega_l}$, the nonzero ones have the same signs as the RHS of (27). Then, for all $i \in \Omega_k, l \in [r]$, we have $\sum_{j \in \Omega_l} |\Delta(i, j)| = |\sum_{j \in \Omega_l} \Delta(i, j)| = |\text{RHS of (27)}| \leq |\Omega_k|^{-1} \sum_{i' \in \Omega_k} |\sum_{j \in \Omega_l} \mathbf{T}(i, j) - \sum_{j \in \Omega_l} \mathbf{T}(i', j)|$. This further gives $\|\Delta\|_F \leq \sum_{k, l \in [r]} \sum_{i \in \Omega_k, j \in \Omega_l} |\Delta(i, j)| \leq |\Omega_k|^{-1} \sum_{k, l \in [r]} \sum_{i, i' \in \Omega_k} |\sum_{j \in \Omega_l} \mathbf{T}(i, j) - \sum_{j \in \Omega_l} \mathbf{T}(i', j)| \leq |\Omega_k|^{-1} \epsilon_{\mathbf{T}}$, where the last inequality follows from (3). These steps also show $\|\Delta\|_{\infty} \leq \epsilon_{\mathbf{T}}$. We have shown $\Delta \in \mathcal{D}$, i.e. \mathcal{D} is non-empty, and so is $\mathcal{L}(\mathbf{T}, \Omega_{1:r}, \epsilon_{\mathbf{T}})$.

APPENDIX C

APPROXIMATION WITH MSS – PROOF FOR THEOREM 3

We first provide several supporting results regarding the perturbation of matrix product.

Lemma 10: Consider two sets of matrices $\mathbf{A}_1, \dots, \mathbf{A}_s$ and $\hat{\mathbf{A}}_1, \dots, \hat{\mathbf{A}}_s$ with $\|\mathbf{A}_i - \hat{\mathbf{A}}_i\| \leq \epsilon_{\mathbf{A}}$ for all $i \in [s]$. Assume there exists a pair $\{\xi, \kappa\}$ such that for all $t \in \mathbb{N}$, $\max_{\sigma_{1:t} \in [s]^t} \|\mathbf{A}_{\sigma_1} \cdots \mathbf{A}_{\sigma_t}\|^{\frac{1}{t}} \leq \kappa \cdot \xi^t$. Then, for all t and any sequence $\sigma_{1:t} \in [s]^t$, we have (i) $\|\prod_{h=1}^t \hat{\mathbf{A}}_{\sigma_h}\| \leq \kappa(\kappa\epsilon_{\mathbf{A}} + \xi)^t$; (ii) $\|\prod_{h=1}^t \hat{\mathbf{A}}_{\sigma_h} - \prod_{h=1}^t \mathbf{A}_{\sigma_h}\| \leq \kappa^2 t(\kappa\epsilon_{\mathbf{A}} + \xi)^{t-1} \epsilon_{\mathbf{A}}$.

Proof: Let $\mathbf{E}_i := \hat{\mathbf{A}}_i - \mathbf{A}_i$, then we see $\|\mathbf{E}_i\| \leq \epsilon_{\mathbf{A}}$ and $\prod_{h=1}^t \hat{\mathbf{A}}_{\sigma_h} = \prod_{h=1}^t (\mathbf{A}_{\sigma_h} + \mathbf{E}_{\sigma_h})$. In the expansion of $\prod_{h=1}^t (\mathbf{A}_{\sigma_h} + \mathbf{E}_{\sigma_h})$, for each $i = 0, 1, \dots, t$, there are $\binom{t}{i}$ terms, each of which is a product where \mathbf{E}_{σ_h} has degree i and \mathbf{A}_{σ_h} has degree $t-i$. We let $\mathbf{F}_{i,j}$ with $i = 0, 1, \dots, t$ and $j \in [\binom{t}{i}]$ to index these expansion terms. Note that $\|\mathbf{F}_{i,j}\| \leq \kappa^{i+1} \xi^{t-i} \epsilon_{\mathbf{A}}^i$. Then, we have $\|\prod_{h=1}^t \hat{\mathbf{A}}_{\sigma_h}\| \leq \sum_{i=0}^t \sum_{j \in [\binom{t}{i}]} \|\mathbf{F}_{i,j}\| \leq \sum_{i=0}^t \binom{t}{i} \kappa^{i+1} \xi^{t-i} \epsilon_{\mathbf{A}}^i \leq \kappa(\kappa\epsilon_{\mathbf{A}} + \xi)^t$.

Similarly, $\|\prod_{h=1}^t \hat{\mathbf{A}}_{\sigma_h} - \prod_{h=1}^t \mathbf{A}_{\sigma_h}\| = \|\sum_{i=1}^t \sum_{j \in [\binom{t}{i}]} \mathbf{F}_{i,j}\| \leq \sum_{i=0}^t \sum_{j \in [\binom{t}{i}]} \|\mathbf{F}_{i,j}\| - \|\mathbf{F}_{0,1}\| \leq \kappa(\kappa\epsilon_{\mathbf{A}} + \xi)^t - \kappa \xi^t \leq \kappa^2 t(\kappa\epsilon_{\mathbf{A}} + \xi)^{t-1} \epsilon_{\mathbf{A}}$, where the last line follows from the fact that for function $f(x) := x^t$ and $x, a \geq 0$, $f(x) \geq f(x+a) - a \cdot f'(x+a)$. ■

Based on Lemma 10, we have the following corollaries, which will be used in different settings in later derivations.

Corollary 2: Consider two matrices \mathcal{A} and $\bar{\mathcal{A}}$ with $\|\mathcal{A} - \bar{\mathcal{A}}\| \leq \epsilon_{\mathcal{A}}$. Suppose there exists a pair $\{\rho, \tau\}$ such that for all $k \in \mathbb{N}$, $\|\mathcal{A}^k\| \leq \tau \rho^k$. Then, we have $\|\bar{\mathcal{A}}^t\| \leq \tau(\tau \epsilon_{\mathcal{A}} + \rho)^t$ and $\|\bar{\mathcal{A}}^t - \mathcal{A}^t\| \leq \tau^2 t(\tau \epsilon_{\mathcal{A}} + \rho)^{t-1} \epsilon_{\mathcal{A}}$.

Corollary 3: Consider two sets of scalars a_1, \dots, a_s and $\hat{a}_1, \dots, \hat{a}_s$ with $|a_i - \hat{a}_i| < \epsilon_a$ and $|a_i| < \bar{a}$ for all $i \in [s]$. Then, for all t and any sequence $\sigma_{1:t} \in [s]^t$, we have $|\prod_{h=1}^t \hat{a}_{\sigma_h}| \leq (\epsilon_a + \bar{a})^t$ and $|\prod_{h=1}^t \hat{a}_{\sigma_h} - \prod_{h=1}^t a_{\sigma_h}| \leq t(\epsilon_a + \bar{a})^{t-1} \epsilon_a$.

The next result considers the evolution of state \mathbf{x}_t in the mean-square sense for autonomous MJSs.

Lemma 11 (Lemma 9 in [11]): Consider MJS($\mathbf{A}_{1:s}, 0, \mathbf{T}$) and define matrix $\mathcal{A} \in \mathbb{R}^{sn^2 \times sn^2}$ with its ij -th $n^2 \times n^2$ block given by $[\mathcal{A}]_{ij} := \mathbf{T}(j, i) \cdot \mathbf{A}_j \otimes \mathbf{A}_i$. Let $\Sigma_t^{(i)} := \mathbb{E}[\mathbf{x}_t \mathbf{x}_t^\top \mathbf{1}_{\{\omega_t=i\}}]$ and $\mathbf{s}_t := [\text{vec}(\Sigma_t^{(1)})^\top, \dots, \text{vec}(\Sigma_t^{(s)})^\top]^\top$. Then, $\mathbf{s}_t = \mathcal{A}^t \mathbf{s}_0$.

Recall in Section V, for Σ , i.e., MJS($\mathbf{A}_{1:s}, \mathbf{B}_{1:s}, \mathbf{T}$), we define the augmented state matrix $\mathcal{A} \in \mathbb{R}^{sn^2 \times sn^2}$ with its ij -th $n^2 \times n^2$ block given by $[\mathcal{A}]_{ij} := \mathbf{T}(j, i) \cdot \mathbf{A}_j \otimes \mathbf{A}_i$; and for any $\rho \geq \rho(\mathcal{A})$ and all $k \in \mathbb{N}$, we have $\|\mathcal{A}^k\| \leq \tau \rho^k$. The next lemma is regarding the augmentation of two MJS with the same \mathbf{A} matrix.

Lemma 12: Construct matrix $\check{\mathcal{A}} \in \mathbb{R}^{4sn^2 \times 4sn^2}$ with its ij -th $4n^2 \times 4n^2$ block given by $[\check{\mathcal{A}}]_{ij} := \mathbf{T}(j, i) \cdot \begin{bmatrix} \mathbf{A}_j & \\ & \mathbf{A}_i \end{bmatrix} \otimes \begin{bmatrix} \mathbf{A}_j & \\ & \mathbf{A}_i \end{bmatrix}$. Then, for all $k \in \mathbb{N}$, $\|\check{\mathcal{A}}^k\| \leq \tau \rho^k$.

To see this result, first notice that there exists a permutation matrix \mathbf{P} such that $\mathbf{P} \check{\mathcal{A}} \mathbf{P}^\top = \mathbf{I}_4 \otimes \mathcal{A}$, where \mathbf{I}_4 denotes the 4×4 identity matrix. This gives $\|\check{\mathcal{A}}^k\| = \|\mathcal{A}^k\|$ and shows the claim.

To prove the result in Theorem 3, we first consider the simplified autonomous case but with potentially different initial states \mathbf{x}_0 and $\hat{\mathbf{x}}_0$.

Proposition 1: Consider the setup in Theorem 3 except that $\mathbf{u}_t = 0$ for all t , and \mathbf{x}_0 and $\hat{\mathbf{x}}_0$ can be different such that $\|\mathbf{x}_0 - \hat{\mathbf{x}}_0\| \leq \epsilon_0$ for some $\epsilon_0 \geq 0$. For perturbation, assume $\epsilon_{\mathbf{A}} \leq \min\{\bar{A}, \frac{1-\rho}{6\tau\bar{A}\|\mathbf{T}\|}\}$. Then, $\mathbb{E}[\|\mathbf{x}_t - \hat{\mathbf{x}}_t\|] \leq 4\sqrt{n\sqrt{s}\tau\rho_0^{\frac{t-1}{2}}}(\|\mathbf{x}_0\|\sqrt{t\bar{A}\|\mathbf{T}\|\epsilon_{\mathbf{A}}} + \sqrt{(\|\mathbf{x}_0\| + \epsilon_0)\epsilon_0})$.

Proof: First, we construct two autonomous switched systems:

$$\check{\Pi} := \begin{cases} \check{\mathbf{x}}_{t+1} = \check{\mathbf{A}}_{\check{\omega}_t} \check{\mathbf{x}}_t, \\ \check{\omega}_t = \omega_t, \end{cases} \quad \bar{\Pi} := \begin{cases} \bar{\mathbf{x}}_{t+1} = \bar{\mathbf{A}}_{\bar{\omega}_t} \bar{\mathbf{x}}_t \\ \bar{\omega}_t = \omega_t, \end{cases} \quad (28)$$

where for $i \in [s]$ (suppose $i \in \Omega_k$), $\check{\mathbf{A}}_i := \begin{bmatrix} \mathbf{A}_i & \\ & \hat{\mathbf{A}}_k \end{bmatrix}$, $\bar{\mathbf{A}}_i := \begin{bmatrix} \mathbf{A}_i & \\ & \hat{\mathbf{A}}_k \end{bmatrix}$. Since ω_t of Σ follows Markov chain \mathbf{T} , systems $\check{\Pi}$ and $\bar{\Pi}$ can be viewed as MJS($\check{\mathbf{A}}_{1:s}, 0, \check{\mathbf{T}}$) and MJS($\bar{\mathbf{A}}_{1:s}, 0, \bar{\mathbf{T}}$) respectively with $\check{\mathbf{T}} = \bar{\mathbf{T}} = \mathbf{T}$. We then define observations for $\check{\Pi}$ and $\bar{\Pi}$: $\check{\mathbf{y}}_t = \check{\mathbf{C}} \check{\mathbf{x}}_t$ and $\bar{\mathbf{y}}_t = \bar{\mathbf{C}} \bar{\mathbf{x}}_t$ where $\check{\mathbf{C}} = \bar{\mathbf{C}} = [\mathbf{I}_n, -\mathbf{I}_n]$. We set their initial states as $\check{\mathbf{x}}_0 = [\mathbf{x}_0^\top, \hat{\mathbf{x}}_0^\top]^\top$, $\bar{\mathbf{x}}_0 = [\mathbf{x}_0^\top, \hat{\mathbf{x}}_0^\top]^\top$ where \mathbf{x}_0 and $\hat{\mathbf{x}}_0$ are the initial states of Σ and $\hat{\Sigma}$ respectively.

By construction, we have, for all t , $\check{\mathbf{x}}_t = [\mathbf{x}_t^\top, \hat{\mathbf{x}}_t^\top]^\top$ and $\bar{\mathbf{x}}_t = [\mathbf{x}_t^\top, \hat{\mathbf{x}}_t^\top]^\top$, thus $\check{\mathbf{y}}_t = 0$ and $\bar{\mathbf{y}}_t = \mathbf{x}_t - \hat{\mathbf{x}}_t$. Define $\check{\Sigma}_t := \mathbb{E}[\check{\mathbf{x}}_t \check{\mathbf{x}}_t^\top]$ and $\bar{\Sigma}_t := \mathbb{E}[\bar{\mathbf{x}}_t \bar{\mathbf{x}}_t^\top]$, then we have $\mathbb{E}[\|\mathbf{x}_t - \hat{\mathbf{x}}_t\|^2] = \mathbb{E}[\check{\mathbf{y}}_t \check{\mathbf{y}}_t^\top] = \mathbb{E}[\bar{\mathbf{y}}_t \bar{\mathbf{y}}_t^\top] - \mathbb{E}[\check{\mathbf{y}}_t \bar{\mathbf{y}}_t^\top] = \text{tr}(\check{\mathbf{C}}^\top \check{\mathbf{C}} \check{\Sigma}_t) - \text{tr}(\check{\mathbf{C}}^\top \check{\mathbf{C}} (\bar{\Sigma}_t - \check{\Sigma}_t))$. Since $\check{\mathbf{C}}^\top \check{\mathbf{C}} \succeq 0$, we

further have

$$\mathbb{E}[\|\mathbf{x}_t - \hat{\mathbf{x}}_t\|^2] \leq \text{tr}(\check{\mathbf{C}}^\top \check{\mathbf{C}}) \|\check{\Sigma}_t - \check{\Sigma}_t\| = 2n \|\check{\Sigma}_t - \check{\Sigma}_t\|. \quad (29)$$

Let $\check{\Sigma}_t^{(i)} := \mathbb{E}[\check{\mathbf{x}}_t \check{\mathbf{x}}_t^\top \mathbf{1}_{\{\check{\omega}_t=i\}}]$, $\bar{\Sigma}_t^{(i)} := \mathbb{E}[\bar{\mathbf{x}}_t \bar{\mathbf{x}}_t^\top \mathbf{1}_{\{\bar{\omega}_t=i\}}]$, $\check{\mathbf{s}}_t := [\text{vec}(\check{\Sigma}_t^{(1)})^\top, \dots, \text{vec}(\check{\Sigma}_t^{(s)})^\top]^\top$ and $\bar{\mathbf{s}}_t := [\text{vec}(\bar{\Sigma}_t^{(1)})^\top, \dots, \text{vec}(\bar{\Sigma}_t^{(s)})^\top]^\top$. Note that $\text{vec}(\check{\Sigma}_t) = [\mathbf{I}_{4n^2}, \dots, \mathbf{I}_{4n^2}] \check{\mathbf{s}}_t$ and $\text{vec}(\bar{\Sigma}_t) = [\mathbf{I}_{4n^2}, \dots, \mathbf{I}_{4n^2}] \bar{\mathbf{s}}_t$, thus we have $\|\check{\Sigma}_t - \bar{\Sigma}_t\| \leq \|\check{\Sigma}_t - \bar{\Sigma}_t\|_F = \|\text{vec}(\check{\Sigma}_t - \bar{\Sigma}_t)\| \leq \sqrt{s} \|\check{\mathbf{s}}_t - \bar{\mathbf{s}}_t\|$. Plugging this into (29), we have

$$\mathbb{E}[\|\mathbf{x}_t - \hat{\mathbf{x}}_t\|^2] \leq 2n\sqrt{s} \|\check{\mathbf{s}}_t - \bar{\mathbf{s}}_t\|. \quad (30)$$

By Lemma 11, we have $\check{\mathbf{s}}_t = \check{\mathcal{A}}^t \check{\mathbf{s}}_0$ and $\bar{\mathbf{s}}_t = \bar{\mathcal{A}}^t \bar{\mathbf{s}}_0$, where $\check{\mathcal{A}} \in \mathbb{R}^{4sn^2 \times 4sn^2}$ is constructed such that its ij -th $4n^2 \times 4n^2$ block given by $[\check{\mathcal{A}}]_{ij} = \check{\mathbf{T}}(j, i) \check{\mathbf{A}}_j \otimes \check{\mathbf{A}}_i$, and $\bar{\mathcal{A}}$ is constructed similarly. By triangle inequality, we further have

$$\mathbb{E}[\|\mathbf{x}_t - \hat{\mathbf{x}}_t\|^2] \leq 2n\sqrt{s}(\|\check{\mathcal{A}}^t - \bar{\mathcal{A}}^t\| \|\check{\mathbf{s}}_0\| + \|\bar{\mathcal{A}}^t\| \|\check{\mathbf{s}}_0 - \bar{\mathbf{s}}_0\|) \quad (31)$$

To bound $\mathbb{E}[\|\mathbf{x}_t - \hat{\mathbf{x}}_t\|^2]$, we seek to bound the terms on the RHS individually. Since $\check{\mathbf{s}}_0 = [\text{vec}(\check{\mathbf{x}}_0 \check{\mathbf{x}}_0^\top)^\top \cdot \mathbb{P}(\omega_t = 1), \dots, \text{vec}(\check{\mathbf{x}}_0 \check{\mathbf{x}}_0^\top)^\top \cdot \mathbb{P}(\omega_t = s)]^\top$, we have $\|\check{\mathbf{s}}_0\| = \|\check{\mathbf{x}}_0 \check{\mathbf{x}}_0^\top\|_F \cdot (\sum_{i \in [s]} \mathbb{P}(\omega_t = i)^2)^{\frac{1}{2}} \leq \|\check{\mathbf{x}}_0 \check{\mathbf{x}}_0^\top\|_F = 2\|\mathbf{x}_0\|^2$. Similarly, we have $\|\check{\mathbf{s}}_0 - \bar{\mathbf{s}}_0\| \leq \|\check{\mathbf{x}}_0 \check{\mathbf{x}}_0^\top - \bar{\mathbf{x}}_0 \check{\mathbf{x}}_0^\top\|_F \leq \|\check{\mathbf{x}}_0(\bar{\mathbf{x}}_0 - \check{\mathbf{x}}_0)^\top\|_F + \|(\bar{\mathbf{x}}_0 - \check{\mathbf{x}}_0)\check{\mathbf{x}}_0^\top\|_F \leq \sqrt{2}(\sqrt{3}\|\mathbf{x}_0\| + \epsilon_0)\epsilon_0$.

To bound $\|\check{\mathcal{A}}^t\|$ and $\|\bar{\mathcal{A}}^t - \check{\mathcal{A}}^t\|$, we first evaluate $\|\check{\mathcal{A}} - \bar{\mathcal{A}}\|$. Define $\Delta_i := \check{\mathbf{A}}_i \otimes \check{\mathbf{A}}_i - \bar{\mathbf{A}}_i \otimes \bar{\mathbf{A}}_i$ for all i , and block diagonal matrix $\Delta \in \mathbb{R}^{sn^2 \times sn^2}$ such that the i th $n^2 \times n^2$ block is given by Δ_i . Then one can verify that $\check{\mathcal{A}} - \bar{\mathcal{A}} = (\mathbf{T} \otimes \mathbf{I}_{n^2})\Delta$, which gives $\|\check{\mathcal{A}} - \bar{\mathcal{A}}\| \leq \|\mathbf{T}\| \max_i \|\Delta_i\|$. For $\|\Delta_i\|$, we have $\|\Delta_i\| \leq \|\check{\mathbf{A}}_i\| \|\check{\mathbf{A}}_i - \bar{\mathbf{A}}_i\| + \|\check{\mathbf{A}}_i - \bar{\mathbf{A}}_i\| \|\bar{\mathbf{A}}_i\|$. It is easy to see $\|\check{\mathbf{A}}_i\| \leq \bar{A}$, $\|\check{\mathbf{A}}_i - \bar{\mathbf{A}}_i\| \leq \epsilon_{\mathbf{A}}$, and $\|\bar{\mathbf{A}}_i\| \leq \bar{A} + \epsilon_{\mathbf{A}} \leq 2\bar{A}$. These give $\|\check{\mathcal{A}} - \bar{\mathcal{A}}\| \leq 3\bar{A}\|\mathbf{T}\|\epsilon_{\mathbf{A}}$. From Lemma 12, we know for all $k \in \mathbb{N}$, $\|\check{\mathcal{A}}^k\| \leq \tau \rho^k$. Then, according to Corollary 2, we have $\|\check{\mathcal{A}}^t\| \leq \tau(3\tau\bar{A}\|\mathbf{T}\|\epsilon_{\mathbf{A}} + \rho)^t \leq \tau\rho_0^t$ and $\|\bar{\mathcal{A}}^t - \check{\mathcal{A}}^t\| \leq 3t(3\tau\bar{A}\|\mathbf{T}\|\epsilon_{\mathbf{A}} + \rho)^{t-1} \tau^2 \bar{A}\|\mathbf{T}\|\epsilon_{\mathbf{A}} \leq 3t\rho_0^{t-1} \tau^2 \bar{A}\|\mathbf{T}\|\epsilon_{\mathbf{A}}$, where the premise $\epsilon_{\mathbf{A}} \leq \frac{1-\rho}{6\tau\bar{A}\|\mathbf{T}\|}$ and notation $\rho_0 := \frac{1+\rho}{2}$ are used.

Finally, plugging in the bounds we just derived for each term on the RHS of (31) back, we have $\mathbb{E}[\|\mathbf{x}_t - \hat{\mathbf{x}}_t\|] \leq \sqrt{\mathbb{E}[\|\mathbf{x}_t - \hat{\mathbf{x}}_t\|^2]} \leq 4\sqrt{n\sqrt{s}\tau\rho_0^{\frac{t-1}{2}}}(\|\mathbf{x}_0\|\sqrt{t\bar{A}\|\mathbf{T}\|\epsilon_{\mathbf{A}}} + \sqrt{(\|\mathbf{x}_0\| + \epsilon_0)\epsilon_0})$, which concludes the proof. \blacksquare

Main Proof for Theorem 3: We first decompose \mathbf{x}_t in terms of the contribution from \mathbf{x}_0 , $\mathbf{u}_{0:t-1}$: define $\mathbf{x}_t^{(0')} := (\prod_{h=0}^{t-1} \mathbf{A}_{\omega_h}) \mathbf{x}_0$, for $l = 0, \dots, t-2$, define $\mathbf{x}_t^{(l)} := (\prod_{h=l+1}^{t-1} \mathbf{A}_{\omega_h}) \mathbf{B}_{\omega_l} \mathbf{u}_l$, and $\mathbf{x}_t^{(t-1)} := \mathbf{B}_{\omega_{t-1}} \mathbf{u}_{t-1}$. Then it is easy to see $\mathbf{x}_t = \mathbf{x}_t^{(0')} + \sum_{l=0}^{t-1} \mathbf{x}_t^{(l)}$. Similarly, we define $\hat{\mathbf{x}}_t^{(0')}$ and $\hat{\mathbf{x}}_t^{(l)}$ for $\hat{\mathbf{x}}_t$ such that $\hat{\mathbf{x}}_t = \hat{\mathbf{x}}_t^{(0')} + \sum_{l=0}^{t-1} \hat{\mathbf{x}}_t^{(l)}$. According to Proposition 1, we have

$$\mathbb{E}[\|\mathbf{x}_t^{(0')} - \hat{\mathbf{x}}_t^{(0')}\|] \leq 4\sqrt{n\sqrt{s}\tau\rho_0^{\frac{t-1}{2}}} \sqrt{t\bar{A}\|\mathbf{T}\|\epsilon_{\mathbf{A}}} \|\mathbf{x}_0\|. \quad (32)$$

Note that $\mathbf{x}_t^{(l)}$ and $\hat{\mathbf{x}}_t^{(l)}$ can be viewed as the states at time $t-l$ with respective initial states $\mathbf{B}_{\omega_l} \mathbf{u}_l$ and $\hat{\mathbf{B}}_{\omega_l} \mathbf{u}_l$ and zero inputs.

Therefore, applying Proposition 1 again, we have

$$\mathbb{E}[\|\mathbf{x}_t^{(l)} - \hat{\mathbf{x}}_t^{(l)}\|] \leq 4\sqrt{n\sqrt{s}\tau\rho_0^{\frac{t-l-1}{2}}\sqrt{\bar{B}\bar{u}}} \cdot \left(\sqrt{(t-l-1)\bar{A}\|\mathbf{T}\|\epsilon_A} + \sqrt{2\epsilon_B}\right), \quad (33)$$

where the premise $\epsilon_B \leq \bar{B}$ is applied. With (32) and (33), we have $\mathbb{E}[\|\mathbf{x}_t - \hat{\mathbf{x}}_t\|] \leq \mathbb{E}[\|\mathbf{x}_t^{(0')} - \hat{\mathbf{x}}_t^{(0')}\|] + \sum_{l=0}^{t-1} \mathbb{E}[\|\mathbf{x}_t^{(l)} - \hat{\mathbf{x}}_t^{(l)}\|] \leq 4\sqrt{n\sqrt{s}\tau\rho_0^{\frac{t-1}{2}}\sqrt{t\bar{A}\|\mathbf{T}\|\epsilon_A}}\|\mathbf{x}_0\| + 4\sqrt{n\sqrt{s}\bar{B}\bar{u}}\left(\frac{\sqrt{\rho_0}}{(1-\sqrt{\rho_0})^2}\sqrt{\bar{A}\|\mathbf{T}\|\epsilon_A} + \frac{\sqrt{2}}{1-\sqrt{\rho_0}}\sqrt{\epsilon_B}\right)$, which concludes the proof. \blacksquare

APPENDIX D APPROXIMATION WITH UNIF. STABILITY – PROOF FOR THEOREM 4

Proof for Theorem 4 (T1): \mathbf{x}_t and $\hat{\mathbf{x}}_t$ can be decomposed as: $\mathbf{x}_t = (\prod_{h=0}^{t-1} \mathbf{A}_{\omega_h})\mathbf{x}_0 + \sum_{t'=0}^{t-2} (\prod_{h=t'+1}^{t-1} \mathbf{A}_{\omega_h})\mathbf{B}_{\omega_{t'}}\mathbf{u}_{t'} + \mathbf{B}_{\omega_{t-1}}$, $\mathbf{u}_{t-1}, \hat{\mathbf{x}}_t = (\prod_{h=0}^{t-1} \hat{\mathbf{A}}_{\hat{\omega}_h})\hat{\mathbf{x}}_0 + \sum_{t'=0}^{t-2} (\prod_{h=t'+1}^{t-1} \hat{\mathbf{A}}_{\hat{\omega}_h})\hat{\mathbf{B}}_{\hat{\omega}_{t'}}\hat{\mathbf{u}}_{t'} + \hat{\mathbf{B}}_{\hat{\omega}_{t-1}}\hat{\mathbf{u}}_{t-1}$. Since in Algorithm 1 we let $\hat{\mathbf{A}}_k = |\hat{\Omega}_k|^{-1} \sum_{i \in \hat{\Omega}_k} \mathbf{A}_i$, and the premise gives $\hat{\Omega}_{1:r} = \Omega_{1:r}$, we have $\|\hat{\mathbf{A}}_k - \mathbf{A}_i\| \leq \epsilon_A$ for all $i \in \Omega_k$. Based on the mode synchrony Setup S2, i.e., $\omega_t \in \Omega_{\hat{\omega}_t}$, we further have $\|\hat{\mathbf{A}}_{\hat{\omega}_t} - \mathbf{A}_{\omega_t}\| \leq \epsilon_A$. Similarly, we obtain $\|\hat{\mathbf{B}}_{\hat{\omega}_t} - \mathbf{B}_{\omega_t}\| \leq \epsilon_B$. Then, by Lemma 10 : (i) $\|\prod_{h=t'+1}^{t-1} \hat{\mathbf{A}}_{\hat{\omega}_h}\| \leq \kappa(\kappa\epsilon_A + \xi)^{t-t'-1}$ and (ii) $\|\prod_{h=t'+1}^{t-1} \mathbf{A}_{\omega_h} - \prod_{h=t'+1}^{t-1} \hat{\mathbf{A}}_{\hat{\omega}_h}\| \leq \kappa^2(t-t'-1)(\kappa\epsilon_A + \xi)^{t-t'-2}\epsilon_A$.

With (i) and (ii), and the fact that $\prod_{h=t'+1}^{t-1} \mathbf{A}_{\omega_h} \mathbf{B}_{\omega_{t'}} - \prod_{h=t'+1}^{t-1} \hat{\mathbf{A}}_{\hat{\omega}_h} \hat{\mathbf{B}}_{\hat{\omega}_{t'}} = (\prod_{h=t'+1}^{t-1} \mathbf{A}_{\omega_h} - \prod_{h=t'+1}^{t-1} \hat{\mathbf{A}}_{\hat{\omega}_h})\hat{\mathbf{B}}_{\hat{\omega}_{t'}} - (\prod_{h=t'+1}^{t-1} \mathbf{A}_{\omega_h})(\hat{\mathbf{B}}_{\hat{\omega}_{t'}} - \mathbf{B}_{\omega_{t'}})$, we have

$$\|\prod_{h=t'+1}^{t-1} \mathbf{A}_{\omega_h} \mathbf{B}_{\omega_{t'}} - \prod_{h=t'+1}^{t-1} \hat{\mathbf{A}}_{\hat{\omega}_h} \hat{\mathbf{B}}_{\hat{\omega}_{t'}}\| \leq \kappa\xi^{t-t'-1}\epsilon_B + \kappa^2(t-t'-1)(\kappa\epsilon_A + \xi)^{t-t'-2}(\bar{B} + \epsilon_B)\epsilon_A. \quad (34)$$

According to Setup S1, Σ and $\hat{\Sigma}$ have the same initial states and inputs. Then, applying triangle inequality to the difference $\|\mathbf{x}_t - \hat{\mathbf{x}}_t\|$, we have

$$\|\mathbf{x}_t - \hat{\mathbf{x}}_t\| \leq \kappa^2 t(\kappa\epsilon_A + \xi)^{t-1} \|\mathbf{x}_0\| \epsilon_A + \kappa^2 \frac{1+t(\kappa\epsilon_A + \xi)^t}{1-\kappa\epsilon_A - \xi} (\bar{B} + \epsilon_B)\bar{u}\epsilon_A + \frac{\kappa}{1-\xi}\bar{u}\epsilon_B, \quad (35)$$

where the following facts are implicitly used: (i) $\kappa \geq 1$ by definition; (ii) $\kappa\epsilon_A + \xi < 1$ according to the premise. Finally, note that we assume perturbation $\epsilon_A \leq \frac{1-\xi}{2\kappa}$ and $\epsilon_B \leq \bar{B}$, we have $\|\mathbf{x}_t - \hat{\mathbf{x}}_t\| \leq t\xi_0^{t-1}\kappa^2 \|\mathbf{x}_0\| \epsilon_A + \frac{2(1+\xi_0^t)\kappa^2\bar{B}\bar{u}}{1-\xi_0}\epsilon_A + \frac{\kappa\bar{u}}{1-\xi}\epsilon_B$, which concludes the proof. \blacksquare

A. PROOF FOR THEOREM 4 (T2)

To ease the proof exposition, we first define a few notations and concepts. For the original system Σ , fixing the initial state \mathbf{x}_0 and input sequence $\mathbf{u}_{0:t-1}$, there can be at most s^t possible \mathbf{x}_t , each of which correspond to one possible mode switching sequence $\omega_{0:t-1} \in [s]^t$. We use $g \in [s^t]$ to index these states

and mode sequences, i.e., mode sequence $\omega_{0:t-1}^{(g)}$ generates state $\mathbf{x}_t^{(g)}$. Then, the reachable set \mathcal{X}_t defined in Section V satisfies $\mathcal{X}_t = \bigcup_{g \in [s^t]} \{\mathbf{x}_t^{(g)}\}$. Define probability measure $q_t(g) := \mathbb{P}(\omega_{0:t-1} = \omega_{0:t-1}^{(g)})$, then we see $p_t(\mathbf{x})$ defined in Section V satisfies $p_t(\mathbf{x}) = \sum_{g: \mathbf{x}_t^{(g)} = \mathbf{x}} q_t(g)$. For the reduced $\hat{\Sigma}$ and for all $\hat{g} \in [r^t]$, we similarly define notations $\hat{\omega}_{0:t-1}^{(\hat{g})}$ for the mode sequence, $\hat{\mathbf{x}}_t^{(\hat{g})}$ for the state, and $\hat{q}_t(\hat{g}) := \mathbb{P}(\hat{\omega}_{0:t-1} = \hat{\omega}_{0:t-1}^{(\hat{g})})$ for the measure. Then, the following holds: $\hat{\mathcal{X}}_t = \bigcup_{\hat{g} \in [r^t]} \{\hat{\mathbf{x}}_t^{(\hat{g})}\}$ and $\hat{p}_t(\hat{\mathbf{x}}) = \sum_{\hat{g}: \hat{\mathbf{x}}_t^{(\hat{g})} = \hat{\mathbf{x}}} \hat{q}_t(\hat{g})$. Next, we introduce the following relation regarding mode sequences between Σ and $\hat{\Sigma}$.

Definition 4 (Mode Sequence Synchrony): For any $g \in [s^t]$, $\hat{g} \in [r^t]$, we say $\omega_{0:t-1}^{(g)}$ is synchronous to $\hat{\omega}_{0:t-1}^{(\hat{g})}$ (denoted by $g \triangleright \hat{g}$) if $\omega_h \in \Omega_h$ for all $h = 0, 1, \dots, t-1$.

Note that the synchrony definition here coincides with the mode synchrony in Setup S2. With this synchrony relation, we first present a preliminary result.

Lemma 13: For any $\hat{g} \in [r]^t$, we have $|\hat{q}_t(\hat{g}) - \sum_{g: g \triangleright \hat{g}} q_t(g)| \leq (t-1)(\bar{\mathcal{T}} + \epsilon_T)^{t-2}\epsilon_T$.

Proof: Recall ζ_t indexes the active cluster of Σ at time t , i.e., $\zeta_t = k$ if and only if $\omega_t \in \Omega_k$. First observe that $\sum_{g: g \triangleright \hat{g}} q_t(g) = \sum_{g: g \triangleright \hat{g}} \mathbb{P}(\omega_{0:t-1} = \omega_{0:t-1}^{(g)}) = \mathbb{P}(\omega_{t-1} \in \Omega_{\hat{\omega}_{t-1}^{(\hat{g})}}, \dots, \omega_0 \in \Omega_{\hat{\omega}_0^{(\hat{g})}}) = \mathbb{P}(\zeta_{0:t-1} = \hat{\omega}_{0:t-1}^{(\hat{g})})$. Also note that $\hat{q}_t(\hat{g}) = \mathbb{P}(\hat{\omega}_{0:t-1} = \hat{\omega}_{0:t-1}^{(\hat{g})})$. So, to show the claim, it suffices to show for any $\sigma_{0:t} \in [r]^t$,

$$|\mathbb{P}(\hat{\omega}_{0:t} = \sigma_{0:t}) - \mathbb{P}(\zeta_{0:t} = \sigma_{0:t})| \leq t(\bar{\mathcal{T}} + \epsilon_T)^{t-1}\epsilon_T. \quad (36)$$

For the LHS of (36), we have

$$\mathbb{P}(\hat{\omega}_{0:t} = \sigma_{0:t}) = \mathbb{P}(\hat{\omega}_0 = \sigma_0) \cdot \prod_{h=1}^t \hat{\mathbf{T}}(\sigma_{h-1}, \sigma_h) \quad (37)$$

$$\mathbb{P}(\zeta_{0:t} = \sigma_{0:t}) = \mathbb{P}(\omega_0 \in \Omega_{\sigma_0}) \cdot \prod_{h=1}^t \tilde{\mathbf{T}}_h \quad (38)$$

where $\tilde{\mathbf{T}}_h := \mathbb{P}(\omega_h \in \Omega_{\sigma_h} \mid \omega_{h-1} \in \Omega_{\sigma_{h-1}}, \dots, \omega_0 \in \Omega_{\sigma_0})$. Note that $\zeta_{0:t}$ may not be a Markov process when $\epsilon_T \neq 0$, so we cannot drop the past conditional events in (38).

Let $\alpha_i := \mathbb{P}(\omega_{h-1} = i \mid \omega_{h-2} \in \Omega_{\sigma_{h-2}}, \dots, \omega_0 \in \Omega_{\sigma_0})$, then $\tilde{\mathbf{T}}_h = \sum_{i \in \Omega_{\sigma_{h-1}}} [\mathbb{P}(\omega_h \in \Omega_{\sigma_h} \mid \omega_{h-1} = i) \cdot \mathbb{P}(\omega_{h-1} = i \mid \omega_{h-2}, \dots, \omega_0 \in \Omega_{\sigma_0})] = \sum_{i \in \Omega_{\sigma_{h-1}}} [(\sum_{j \in \Omega_{\sigma_h}} \mathbf{T}(i, j))\alpha_i]$.

Let $\beta_i := |\Omega_{\sigma_{h-1}}|^{-1}$. For $\hat{\mathbf{T}}(\sigma_{h-1}, \sigma_h)$, by definition in Algorithm 1 and the assumption $\hat{\Omega}_{1:r} = \Omega_{1:r}$, we know $\hat{\mathbf{T}}(\sigma_{h-1}, \sigma_h) = |\Omega_{\sigma_{h-1}}|^{-1} \cdot \sum_{i \in \Omega_{\sigma_{h-1}}} (\sum_{j \in \Omega_{\sigma_h}} \mathbf{T}(i, j)) = \sum_{i \in \Omega_{\sigma_{h-1}}} [(\sum_{j \in \Omega_{\sigma_h}} \mathbf{T}(i, j))\beta_i]$.

Then, it follows that the difference $|\tilde{\mathbf{T}}_h - \hat{\mathbf{T}}(\sigma_{h-1}, \sigma_h)| = |\sum_{i, i' \in \Omega_{\sigma_{h-1}}} [(\sum_{j \in \Omega_{\sigma_h}} \mathbf{T}(i, j))\alpha_i \beta_{i'}] - \sum_{i, i' \in \Omega_{\sigma_{h-1}}} [(\sum_{j \in \Omega_{\sigma_h}} \mathbf{T}(i', j))\alpha_i \beta_{i'}]| \leq \sum_{i, i' \in \Omega_{\sigma_{h-1}}} [|\sum_{j \in \Omega_{\sigma_h}} \mathbf{T}(i, j) - \sum_{j \in \Omega_{\sigma_h}} \mathbf{T}(i', j)| \alpha_i \beta_{i'}] \leq \epsilon_T$, where the first inequality follows from triangle inequality on the absolute values; the second inequality holds since the definition of perturbation ϵ_T in either Problem P1 or P2 gives $|\sum_{j \in \Omega_{\sigma_h}} \mathbf{T}(i, j) - \sum_{j \in \Omega_{\sigma_h}} \mathbf{T}(i', j)| \leq \epsilon_T$ for any $i, i' \in \Omega_{\sigma_{h-1}}$.

We have established upper bounds for the differences between each multiplier in (37) and (38), by Corollary 3, we obtain $|\mathbb{P}(\hat{\omega}_{0:t} = \sigma_{0:t}) - \mathbb{P}(\zeta_{0:t} = \sigma_{0:t})| \leq t(\bar{\mathcal{T}} + \epsilon_T)^{t-1}\epsilon_T$ which shows (36) and concludes the proof. ■

Main Proof for Theorem 4 (T2): To lower bound the Wasserstein distance $W_\ell(p_t, \hat{p}_t)$ defined in the mass transportation problem (8), we consider the objective value given by a constrained mass transportation scheme. Recall with measures q_t and \hat{q}_t , we have $p_t(\mathbf{x}) = \sum_{g: \mathbf{x}_t^{(g)} = \mathbf{x}} q_t(g)$ and $\hat{p}_t(\hat{\mathbf{x}}) = \sum_{\hat{g}: \hat{\mathbf{x}}_t^{(\hat{g})} = \hat{\mathbf{x}}} \hat{q}_t(\hat{g})$. With these relations, we consider the following transportation scheme in terms of q_t and \hat{q}_t : for all the mass $q_t(g)$ with mode sequence $\omega_t^{(g)}$ synchronous to mode sequence $\hat{\omega}_t^{(\hat{g})}$, it is prioritized to be moved to location \hat{g} ; if there is surplus, i.e., $\sum_{g: g \triangleright \hat{g}} q_t(g) > \hat{q}_t(\hat{g})$, we move the surplus portion $\sum_{g: g \triangleright \hat{g}} q_t(g) - \hat{q}_t(\hat{g})$ elsewhere.

Under this moving scheme, let $\bar{W}_\ell(q_t, \hat{q}_t)$ denote the optimal objective value of the mass transportation problem (8). Let $\hat{\mathcal{G}}_1 = \{\hat{g} : \hat{g} \in [r'], \sum_{g: g \triangleright \hat{g}} q_t(g) \leq \hat{q}_t(\hat{g})\}$, $\hat{\mathcal{G}}_2 = [r'] \setminus \hat{\mathcal{G}}_1$. Then, $\bar{W}_\ell(q_t, \hat{q}_t)$ can be viewed as the optimal objective of the following problem:

$$\min_{f \geq 0} \left(\sum_{g \in [s'], \hat{g} \in [r']} f(g, \hat{g}) \|\mathbf{x}_t^{(g)} - \hat{\mathbf{x}}_t^{(\hat{g})}\|^\ell \right)^{1/\ell} \quad (39)$$

$$\text{s.t. } \sum_{g \in [s']} f(g, \hat{g}) = \hat{q}_t(\hat{g}), \forall \hat{g}$$

$$\sum_{\hat{g} \in [r']} f(g, \hat{g}) = q_t(g), \forall g$$

$$f(g, \hat{g}) = q_t(g), \quad \forall g \triangleright \hat{g}, \forall \hat{g} \in \hat{\mathcal{G}}_1 \quad (40)$$

$$\sum_{g \triangleright \hat{g}} f(g, \hat{g}) = \hat{q}_t(\hat{g}), \quad \forall \hat{g} \in \hat{\mathcal{G}}_2 \quad (41)$$

where constraints (40) and (41) characterize the moving scheme outlined above. Without them, the problem reduces to (8), thus $W_\ell(p_t, \hat{p}_t) \leq \bar{W}_\ell(q_t, \hat{q}_t)$. To prove the main claim, it suffices to show

$$\begin{aligned} \bar{W}_\ell(q_t, \hat{q}_t) &\leq t \xi_0^{t-1} \kappa^2 \|\mathbf{x}_0\| \epsilon_A \\ &+ 2r^2 t \kappa \|\mathbf{x}_0\| r^t (\kappa \epsilon_A + \xi)^t (\bar{\mathcal{T}} + \epsilon_T)^{\frac{t-2}{\ell}} \epsilon_T^{\frac{1}{\ell}}. \end{aligned} \quad (42)$$

For all $\hat{g} \in [r']$, define its synchrony set $\mathcal{S}(\hat{g}) := \{g : g \in [s'], g \triangleright \hat{g}\}$ and the asynchrony set $\mathcal{S}^c(\hat{g}) := [s'] \setminus \mathcal{S}(\hat{g})$. For the synchrony flow, define total flow $F_s := \sum_{\hat{g} \in [r'], g \in \mathcal{S}(\hat{g})} f(g, \hat{g})$ and maximum travel distance $D_s := \max_{\hat{g} \in [r'], g \in \mathcal{S}(\hat{g})} \|\mathbf{x}_t^{(g)} - \hat{\mathbf{x}}_t^{(\hat{g})}\|$. For the asynchrony flow, similarly define $F_a := \sum_{\hat{g} \in [r'], g \in \mathcal{S}^c(\hat{g})} f(g, \hat{g})$ and $D_a := \max_{\hat{g} \in [r'], g \in \mathcal{S}^c(\hat{g})} \|\mathbf{x}_t^{(g)} - \hat{\mathbf{x}}_t^{(\hat{g})}\|$. Then, we have $\bar{W}_\ell(q_t, \hat{q}_t) \leq (F_s D_s^\ell + F_a D_a^\ell)^{\frac{1}{\ell}} \leq F_s^{\frac{1}{\ell}} D_s + F_a^{\frac{1}{\ell}} D_a$. We next bound F_s, D_s, F_a, D_a separately.

For the synchrony maximum travel distance D_s , since $g \triangleright \hat{g}$, by Theorem 4 (T1), we know $D_s \leq t \xi_0^{t-1} \kappa^2 \|\mathbf{x}_0\| \epsilon_A$. For the synchrony total flow F_s , we simply bound it with $F_s \leq 1$.

Now we consider the asynchrony maximum travel distance D_a . First note that for any g and \hat{g} , we have $\|\mathbf{x}_t^{(g)} - \hat{\mathbf{x}}_t^{(\hat{g})}\| = \|\prod_{h=0}^{t-1} \mathbf{A}_{\omega_h^{(g)}} \mathbf{x}_0 - \prod_{h=0}^{t-1} \hat{\mathbf{A}}_{\hat{\omega}_h^{(\hat{g})}} \hat{\mathbf{x}}_0\| \leq \|\prod_{h=0}^{t-1} \mathbf{A}_{\omega_h^{(g)}}\|$

$\|\mathbf{x}_0\| + \|\prod_{h=0}^{t-1} \hat{\mathbf{A}}_{\hat{\omega}_h^{(\hat{g})}} \hat{\mathbf{x}}_0\| \leq 2\kappa(\kappa \epsilon_A + \xi)^t \|\mathbf{x}_0\|$, where the second inequality follows from Lemma 10. Then, it follows that $D_a \leq 2\kappa(\kappa \epsilon_A + \xi)^t \|\mathbf{x}_0\|$.

For the asynchrony total flow F_a , define $F_{a,\hat{g}} := \sum_{g \in \mathcal{S}^c(\hat{g})} f(g, \hat{g})$, then $F_a = \sum_{\hat{g} \in [r']} F_{a,\hat{g}}$. By constraints (40) and (41), $F_{a,\hat{g}} = \sum_{g \in \mathcal{S}^c(\hat{g})} f(g, \hat{g}) = \hat{q}_t(\hat{g}) - \sum_{g: g \triangleright \hat{g}} f(g, \hat{g}) = \hat{q}_t(\hat{g}) - \sum_{g: g \triangleright \hat{g}} q_t(g)$. Thus, if $\hat{g} \in \hat{\mathcal{G}}_2$, $F_{a,\hat{g}} = 0$; and if $\hat{g} \in \hat{\mathcal{G}}_1$, $F_{a,\hat{g}} = \hat{q}_t(\hat{g}) - \sum_{g: g \triangleright \hat{g}} q_t(g) > 0$. For the latter case, according to Lemma 13, we have $F_{a,\hat{g}} = |\hat{q}_t(\hat{g}) - \sum_{g: g \triangleright \hat{g}} q_t(g)| \leq (t-1)(\bar{\mathcal{T}} + \epsilon_T)^{t-2} \epsilon_T$, which further implies that $F_a = \sum_{\hat{g} \in [r']} F_{a,\hat{g}} \leq r^2 t (\bar{\mathcal{T}} + \epsilon_T)^{t-2} \epsilon_T$.

Finally, (42) can be shown by plugging the upper bounds for F_s, D_s, F_a, D_a into the relation that $\bar{W}_\ell(q_t, \hat{q}_t) \leq F_s^{\frac{1}{\ell}} D_s + F_a^{\frac{1}{\ell}} D_a$, which concludes the proof. ■

REFERENCES

- [1] L. Blackmore, S. Funiak, and B. C. Williams, "Combining stochastic and greedy search in hybrid estimation," in *Proc. AAAI Conf. Artif. Intell.*, 2005, pp. 282–287.
- [2] D. Cajueiro, "Stochastic optimal control of jumping Markov parameter processes with applications to finance," Ph.D. dissertation, Electron. Eng. Comput. Sci. Instituto Tecnológico de Aeronáutica-ITA, Brazil, 2002.
- [3] K. Loparo and F. Abdel-Malek, "A probabilistic approach to dynamic power system security," *IEEE Trans. Circuits Syst.*, vol. 37, no. 6, pp. 787–798, Jun. 1990.
- [4] L. E. Svensson and N. Williams, "Optimal monetary policy under uncertainty: A Markov jump-linear-quadratic approach," *Federal Reserve Bank St Louis Rev.*, vol. 90, no. 4, pp. 275–293, 2008.
- [5] V. Ugrinovskii and H. R. Pota, "Decentralized control of power systems via robust control of uncertain Markov jump parameter systems," *Int. J. Control.*, vol. 78, no. 9, pp. 662–677, 2005.
- [6] B. Sinopoli, L. Schenato, M. Franceschetti, K. Poolla, and S. Sastry, "An LQG optimal linear controller for control systems with packet losses," in *Proc. 44th IEEE Conf. Decis. Control*, 2005, pp. 458–463.
- [7] T. H. Truong, P. Seiler, and L. E. Linderman, "Analysis of networked structural control with packet loss," *IEEE Trans. Control Syst. Technol.*, vol. 30, no. 1, pp. 344–351, Jan. 2022.
- [8] L. Zhang, B. Huang, and J. Lam, "\mathcal{H}_\infty model reduction of Markovian jump linear systems," *Syst. Control Lett.*, vol. 50, no. 2, pp. 103–118, 2003.
- [9] M. Zamani and A. Abate, "Approximately Bisimilar symbolic models for randomly switched stochastic systems," *Syst. Control Lett.*, vol. 69, pp. 38–46, 2014.
- [10] Z. Du, N. Ozay, and L. Balzano, "Clustering-based mode reduction for Markov jump systems," in *Proc. Learn. Dyn. Control Conf.*, 2022, pp. 689–701.
- [11] Y. Sattar, Z. Du, D. A. Tarzanagh, L. Balzano, N. Ozay, and S. Oymak, "Identification and adaptive control of Markov jump systems: Sample complexity and regret bounds," 2021, *arXiv:2111.07018*.
- [12] J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panangaden, "Metrics for labelled Markov processes," *Theor. Comput. Sci.*, vol. 318, no. 3, pp. 323–354, 2004.
- [13] I. Tkachev and A. Abate, "On approximation metrics for linear temporal model-checking of stochastic systems," in *Proc. 17th Int. Conf. Hybrid Syst.: Comput. Control*, 2014, pp. 193–202.
- [14] G. Bian and A. Abate, "On the relationship between Bisimulation and trace equivalence in an approximate probabilistic context," in *Proc. Int. Conf. Found. Softw. Sci. Comput. Struct.*, 2017, pp. 321–337.
- [15] Y. Z. Lun, J. Wheatley, A. D'Innocenzo, and A. Abate, "Approximate abstractions of Markov chains with interval decision processes," *IFAC-PapersOnLine*, vol. 51, no. 16, pp. 91–96, 2018.

[16] A. Zhang and M. Wang, "Spectral state compression of Markov processes," *IEEE Trans. Inf. Theory*, vol. 66, no. 5, pp. 3202–3231, May 2020.

[17] Z. Du, N. Ozay, and L. Balzano, "Mode clustering for Markov jump systems," in *Proc. IEEE 8th Int. Workshop Comput. Adv. Multi-Sensor Adaptive Process.*, 2019, pp. 126–130.

[18] A. Bittracher and C. Schütte, "A probabilistic algorithm for aggregating vastly undersampled large Markov chains," *Physica D: Nonlinear Phenomena*, vol. 416, 2021, Art. no. 132799.

[19] A. Abate, A. D'Innocenzo, and M. D. Di Benedetto, "Approximate abstractions of stochastic hybrid systems," *IEEE Trans. Autom. Control*, vol. 56, no. 11, pp. 2688–2694, Nov. 2011.

[20] S. E. Z. Soudjani and A. Abate, "Adaptive gridding for abstraction and verification of stochastic hybrid systems," in *Proc. 8th Int. Conf. Quantitative Eval. Syst.*, 2011, pp. 59–68.

[21] A. A. Julius, A. Girard, and G. J. Pappas, "Approximate bisimulation for a class of stochastic hybrid systems," in *Proc. Amer. Control Conf.*, 2006, pp. 4724–4729.

[22] A. A. Julius and G. J. Pappas, "Approximations of stochastic hybrid systems," *IEEE Trans. Autom. Control*, vol. 54, no. 6, pp. 1193–1203, Jun. 2009.

[23] M. Zamani, M. Rungger, and P. M. Esfahani, "Approximations of stochastic hybrid systems: A compositional approach," *IEEE Trans. Autom. Control*, vol. 62, no. 6, pp. 2838–2853, Jun. 2017.

[24] Y. Shen, Z.-G. Wu, P. Shi, and C. K. Ahn, "Model reduction of Markovian jump systems with uncertain probabilities," *IEEE Trans. Autom. Control*, vol. 65, no. 1, pp. 382–388, Jan. 2020.

[25] G. Kotsalis and A. Rantzer, "Balanced truncation for discrete time Markov jump linear systems," *IEEE Trans. Autom. Control*, vol. 55, no. 11, pp. 2606–2611, Nov. 2010.

[26] M. Sun and J. Lam, "Model reduction of discrete Markovian jump systems with time-weighted H_2 performance," *Int. J. Robust Nonlinear Control*, vol. 26, no. 3, pp. 401–425, 2016.

[27] K. G. Larsen and A. Skou, "Bisimulation through probabilistic testing," *Inf. Comput.*, vol. 94, no. 1, pp. 1–28, 1991.

[28] J. Desharnais, A. Edalat, and P. Panangaden, "Bisimulation for labelled Markov processes," *Inf. Comput.*, vol. 179, no. 2, pp. 163–193, 2002.

[29] A. Abate, "Approximation metrics based on probabilistic bisimulations for general state-space Markov processes: A survey," *Electron. Notes Theor. Comput. Sci.*, vol. 297, pp. 3–25, 2013.

[30] A. Girard and G. J. Pappas, "Approximation metrics for discrete and continuous systems," *IEEE Trans. Autom. Control*, vol. 52, no. 5, pp. 782–798, May 2007.

[31] R. Alur, T. A. Henzinger, G. Lafferriere, and G. J. Pappas, "Discrete abstractions of hybrid systems," *Proc. IEEE*, vol. 88, no. 7, pp. 971–984, Jul. 2000.

[32] E. M. Clarke Jr., O. Grumberg, D. Kroening, D. Peled, and H. Veith, *Model Checking*. Cambridge, MA, USA: MIT Press, 2018.

[33] R. P. Kurshan, *Computer-Aided Verification of Coordinating Processes: The Automata-Theoretic Approach*, vol. 302. Princeton, NJ, USA: Princeton Univ. Press, 2014.

[34] O. Maler, A. Pnueli, and J. Sifakis, "On the synthesis of discrete controllers for timed systems," in *Proc. Annu. Symp. Theor. Aspects Comput. Sci.*, 1995, pp. 229–242.

[35] A. Abate, J.-P. Katoen, J. Lygeros, and M. Prandini, "Approximate model checking of stochastic hybrid systems," *Eur. J. Control*, vol. 16, no. 6, pp. 624–641, 2010.

[36] S. Gugercin and A. C. Antoulas, "A survey of model reduction by balanced truncation and some new results," *Int. J. Control*, vol. 77, no. 8, pp. 748–766, 2004.

[37] P. Buchholz, "Exact and ordinary Lumpability in finite Markov chains," *J. Appl. Probability*, vol. 31, no. 1, pp. 59–75, 1994.

[38] K. H. Hoffmann and P. Salamon, "Bounding the lumping error in markov chain dynamics," *Appl. Math. Lett.*, vol. 22, no. 9, pp. 1471–1475, 2009.

[39] L. Schulman and B. Gaveau, "Coarse grains: The emergence of space and order," *Found. Phys.*, vol. 31, no. 4, pp. 713–731, 2001.

[40] B. Gaveau and L. Schulman, "Dynamical distance: Coarse grains, pattern recognition, and network analysis," *Bull. des Sci. Mathématiques*, vol. 129, no. 8, pp. 631–642, 2005.

[41] M. Meilă and J. Shi, "A random walks view of spectral segmentation," in *Proc. Int. Workshop Artif. Intell. Statist.*, 2001, pp. 203–208.

[42] L. Bottou and Y. Bengio, "Convergence properties of the k-means algorithms," in *Proc. Adv. Neural Inf. Process. Syst.*, 1994, vol. 7, pp. 585–592.

[43] T. F. Gonzalez, "Clustering to minimize the maximum intercluster distance," *Theor. Comput. Sci.*, vol. 38, pp. 293–306, 1985.

[44] A. Kumar, Y. Sabharwal, and S. Sen, "A simple linear time $(1 + \epsilon)$ -approximation algorithm for k-means clustering in any dimensions," in *Proc. 45th Annu. IEEE Symp. Found. Comput. Sci.*, 2004, pp. 454–462.

[45] M. Song and S. Rajasekaran, "Fast algorithms for constant approximation k-means clustering," *Trans. Mach. Learn. Data Min.*, vol. 3, no. 2, pp. 67–79, 2010.

[46] O. L. V. Costa, M. D. Fragoso, and R. P. Marques, *Discrete-Time Markov Jump Linear Systems*. Berlin, Germany: Springer, 2006.

[47] D. Kuhn, P. M. Esfahani, V. A. Nguyen, and S. Shafeezadeh-Abadeh, "Wasserstein distributionally robust optimization: Theory and applications in machine learning," in *Operations Research & Management Science in the Age of Analytics*, Seattle, WA, USA: Informs, 2019, pp. 130–166.

[48] Y. Chen, T. T. Georgiou, and M. Pavon, "Optimal steering of a linear stochastic system to a final probability distribution, Part I," *IEEE Trans. Autom. Control*, vol. 61, no. 5, pp. 1158–1169, May 2016.

[49] M. Goldstein and P. Tsotras, "Finite-horizon covariance control of linear time-varying systems," in *Proc. IEEE 56th Annu. Conf. Decis. Control*, 2017, pp. 3606–3611.

[50] K. Okamoto, M. Goldstein, and P. Tsotras, "Optimal covariance control for stochastic systems under chance constraints," *IEEE Contr. Syst. Lett.*, vol. 2, no. 2, pp. 266–271, Apr. 2018.

[51] J.-S. Li and J. Qi, "Ensemble control of time-invariant linear systems with linear parameter variation," *IEEE Trans. Autom. Control*, vol. 61, no. 10, pp. 2808–2820, Oct. 2015.

[52] R. Jungers, *The Joint Spectral Radius: Theory and Applications*, vol. 385. Berlin, Germany: Springer, 2009.

[53] P. A. Parrilo and A. Jadbabaie, "Approximation of the joint spectral radius using sum of squares," *Linear Algebra Appl.*, vol. 428, no. 10, pp. 2385–2402, 2008.

[54] Y. Sattar, Z. Du, D. A. Tarzanagh, S. Oymak, L. Balzano, and N. Ozay, "Certainty equivalent quadratic control for Markov jump systems," in *Proc. Amer. Control Conf.*, 2022, pp. 2871–2878.

[55] J. Lei and A. Rinaldo, "Consistency of spectral clustering in stochastic block models," *Ann. Statist.*, vol. 43, no. 1, pp. 215–237, 2015.

[56] G. E. Cho and C. D. Meyer, "Comparison of perturbation bounds for the stationary distribution of a Markov chain," *Linear Algebra Appl.*, vol. 335, no. 1–3, pp. 137–150, 2001.



ZHE DU (Student Member, IEEE) received the B.Eng. degree in electrical engineering and automation from the Huazhong University of Science and Technology, Wuhan, China. He is currently working toward the Ph.D. degree in electrical and computer engineering with the University of Michigan, Ann Arbor, MI, USA. His research interests mainly include topics that involve both control and learning, including system identification, adaptive control, system reduction, and clustering.



LAURA BALZANO (Senior Member, IEEE) received the B.S. degree in electrical and computer engineering from Rice University, Houston, TX, USA, the M.S. degree in electrical engineering from UCLA, Los Angeles, CA, USA, and the Ph.D. degree in electrical and computer engineering from the University of Wisconsin—Madison, Madison, WI, USA. She is currently an Associate Professor of electrical engineering and computer science with the University of Michigan, Ann Arbor, MI, USA. She is currently an Associate Editor for the IEEE OPEN JOURNAL OF SIGNAL PROCESSING and the *SIAM Journal of the Mathematics of Data Science*. She was the recipient of the NSF Career Award, ARO Young Investigator Award, AFOSR Young Investigator Award, and Faculty Fellowships from Intel and 3M and the Vulcans Education Excellence Award at the University of Michigan. Her main research interests include modeling and optimization with big, messy data — highly incomplete or corrupted data, uncalibrated data, and heterogeneous data — and its applications in a wide range of scientific problems.



NECMIYE OZAY (Senior Member, IEEE) received the B.S. degree in electrical engineering from Bogazici University, Istanbul, Turkey, in 2004, the M.S. degree in electrical engineering from Pennsylvania State University, University Park, PA, USA, in 2006 and the Ph.D. degree in electrical engineering from Northeastern University, Boston, MA, USA, in 2010. She was a Postdoctoral Scholar with the California Institute of Technology, Pasadena, CA, USA, between 2010 and 2013. She joined the University of Michigan, Ann Arbor, MI, USA, in 2013, where she is currently an Associate Professor of electrical engineering and computer science. She is also a Member of the Michigan Robotics Institute. Her research interests include hybrid dynamical systems, control, optimization and formal methods with applications in cyber-physical systems, system identification, verification and validation, autonomy, and dynamic data analysis. She was the recipient of the 1938E Award and a Henry Russel Award from the University of Michigan for her contributions to teaching and research, and five young investigator awards, including NSF CAREER, and the 2021 Antonio Ruberti Young Researcher Prize from the IEEE Control Systems Society for her fundamental contributions to the control and identification of hybrid and cyber-physical systems. Her papers have received several awards.