

Steering Large Agent Populations using Mean-Field Schrödinger Bridges with Gaussian Mixture Models

George Rapakoulias, *Student Member, IEEE*, Ali Reza Pedram, and Panagiotis Tsiotras, *Fellow, IEEE*

Abstract—The Mean-Field Schrödinger Bridge (MFSB) problem is an optimization problem aiming to find the minimum effort control policy to drive a McKean-Vlasov stochastic differential equation from one probability measure to another. In the context of multi-agent control, the objective is to control the configuration of a swarm of identical, interacting cooperative agents, as captured by the time-varying probability measure of their state. Available methods for solving this problem for distributions with continuous support rely either on spatial discretizations of the problem's domain or on approximating optimal solutions using neural networks trained through stochastic optimization schemes. For agents following Linear Time Varying dynamics, and for Gaussian Mixture Model boundary distributions, we propose a highly efficient parameterization to approximate the optimal solutions of the corresponding MFSB in closed form, without any learning step. Our proposed approach consists of a mixture of elementary policies, each solving a Gaussian-to-Gaussian Covariance Steering problem from the components of the initial mixture to the components of the terminal mixture. Leveraging the semidefinite formulation of the Covariance Steering problem, the proposed solver can handle probabilistic constraints on the system's state while maintaining numerical tractability. We illustrate our approach on a variety of numerical examples.

Index Terms—Mean-field control, Multi-agent control, Schrödinger Bridges, Stochastic control.

I. INTRODUCTION

Designing optimal control strategies for large multi-agent systems is a difficult optimization problem, mainly due to the well-known curse of dimensionality related to the exponential computational complexity of the problem in the number of agents. To overcome this problem, assuming homogeneity among the agents, one can view them as a different realization of the same dynamical system, and in the limit of an infinite number of agents, formulate the problem as a control problem in the space of probability distributions [1]. This approach, known as a mean-field control problem [2], [3], was originally presented as a general tool for studying large-scale systems, with applications ranging from finance and social sciences [4] to engineering and control [5], among many others. In the context of swarm robotics and multi-agent control, this macroscopic perspective is highly favorable in large-scale problems since its computation does not depend on the number of agents involved, making it a suitable tool

for applications such as coverage and surveillance, mapping, and task allocation problems [6]–[8].

To be more precise, consider N identical, interacting, cooperative agents, where N is large. Each agent $i \in \{1, 2, \dots, N\}$ follows the same stochastic dynamics of the form

$$dx_t^i = -\nabla V_t(x_t^i) dt - \frac{1}{N-1} \sum_{j \neq i} \nabla W_t(x_t^i - x_t^j) dt + B_t u_t^i dt + D_t dw_t, \quad (1)$$

where V_t and W_t are potentials guiding the system's evolution. These terms correspond to the potential energy of a single agent, and the interaction energy between agents, respectively. The random variables x_t and u_t lie in n - and m -dimensional Euclidean spaces $\mathbb{R}^n$ and $\mathbb{R}^m$, respectively, and dw_t is the q -dimensional Brownian motion increment.

In the specific case where the potentials are quadratic, i.e., $W_t(x) = \frac{1}{2} x^T Q_t x$ and $V_t(x) = \frac{1}{2} x^T P_t x$, (1) simplifies to a system of linear Stochastic Differential Equations (SDEs)

$$dx_t^i = -(P_t + Q_t)x_t^i dt + \frac{1}{N-1} Q_t \sum_{j \neq i} x_t^j dt + B_t u_t dt + D_t dw_t.$$

Instead of controlling each agent individually by assigning a target state, we want to control their collective behavior. In the limit of $N \rightarrow \infty$, this behavior is captured by a time-dependent continuous distribution, denoted $\rho_t(x)$, supported on $\mathbb{R}^n$. Furthermore,

$$\lim_{N \rightarrow \infty} \frac{1}{N-1} \sum_{j \neq i} x_t^j = \int x \rho_t(x) dx = \mathbb{E}_{x \sim \rho_t}[x] \triangleq \bar{x}_t. \quad (2)$$

With (2) in mind, and defining the potentially time-varying matrices $A_t = -P_t - Q_t$, $\bar{A}_t = Q_t$, a generic agent with state x_t evolves according to the well-known linear McKean-Vlasov SDE (MV-SDE)

$$dx_t = A_t x_t dt + \bar{A}_t \bar{x}_t dt + B_t u_t dt + D_t dw_t, \quad (3)$$

while the density of the state of (3) evolves according to the Fokker Planck Kolmogorov (FPK) equation [9]

$$\frac{\partial \rho_t}{\partial t} + \nabla \cdot ((A_t x_t + \bar{A}_t \bar{x}_t + B_t u_t) \rho_t) - \frac{1}{2} \text{tr}(D_t D_t^T \nabla^2 \rho_t) = 0. \quad (4)$$

The control of the density of (3) with prescribed initial and final distributions is known as the Mean Field Schrödinger Bridge (MFSB) [9] problem and reads as follows

$$\min_{u \in \mathcal{U}} J \triangleq \mathbb{E} \left[\int_0^1 \|u_t(x_t)\|^2 dt \right], \quad (5a)$$

$$\text{s.t. } dx_t = A_t x_t dt + \bar{A}_t \bar{x}_t dt + B_t u_t dt + D_t dw_t, \quad (5b)$$

$$x_0 \sim \rho_0, \quad x_1 \sim \rho_1, \quad (5c)$$

Support for this work has been provided by ONR award N00014-18-1-2828 and NASA ULI award #80NSSC20M0163. This article solely reflects the opinions and conclusions of its authors and not of any NASA entity. George Rapakoulias acknowledges financial support from the A. Onassis Foundation Scholarship.

The authors are with the Department of Aerospace Engineering of Georgia Institute of Technology. Corresponding author: grap@gatech.edu

where ρ_0 and ρ_1 denote the given initial and final state densities, respectively, and where the time horizon is set to 1 without loss of generality, assuming appropriate rescaling of the system matrices.

Before discussing our solution approach, we present a concise literature review regarding Problem (5) and its generalizations. First, in the absence of the mean-field term A_t , the problem is the well-studied Schrödinger Bridge (SB) problem over a linear dynamical system [10]. In the special case of Gaussian boundary distributions ρ_0, ρ_1 , the problem reduces to the well-known covariance steering problem [11], which has been studied extensively in the control community in a multitude of different scenarios, for example, under distributional uncertainty [12], [13], for nonlinear systems [14] and distributed [15], [16] and hierarchical schemes [17], to name a few.

Although covariance steering solvers work primarily with Gaussian boundary distributions, recent works have proposed methods to handle the case of Gaussian Mixture Model (GMM) boundary distributions. In [18], the authors proposed a randomized feedback policy that can steer an initial GMM to a terminal GMM subject to deterministic discrete-time linear dynamics, while in [19] the method was extended to solve constrained problems. These methods provide a feasible solution to the problem; however, they sacrifice optimality for computational tractability and work well only for deterministic systems. In [20], the authors proposed a closed-form deterministic non-linear feedback law to solve the SB problem with continuous-time linear prior dynamics and GMM boundary conditions, using a dynamically weighted sum of elementary policies where each one solves a covariance steering problem from each component of the initial to each component of the terminal Gaussian mixture. Also, in [21], the authors propose a neural-network-based scheme to approximate a feasible policy for the SB problem and arrive at a similar policy to [20] for the special case of GMM boundary distributions.

Returning to problems with mean-field terms, for the specific case of unconstrained Linear Time Varying systems with linear mean-field terms, as in equation (5), it has been shown in [5] that the problem can be decomposed into two parts: a standard SB problem for a shifted set of zero-mean boundary distributions, which is still challenging to compute, and an easy-to-compute feedforward term to compensate for the mean-field effect. For more general problems with boundary distributions available only through samples, and for non-linear prior dynamics with mean-field terms, solutions are difficult to obtain since the problem consists of a combination of a general optimal transport problem and a non-linear optimal control problem, neither of which can be solved in closed form, except for very few special cases [22]. Notable works addressing the mean-field SB in this more general setting, under some assumptions on the structure of the dynamical system are [23]–[25], which utilize neural networks to approximate the optimal control policies, [26], which uses a proximal recursion scheme leveraging the Jordan-Kinderlehrer-Otto (JKO) scheme and evaluates the optimal solution at the finite number of samples rather than the entire space, and [1] that discretizes the state space and formulates the problem as a multi-marginal discrete optimal transport problem.

As evident from the above discussion, the available methods in the literature for solving Problem (5) can be categorized into three types: Covariance steering type methods that are computationally efficient but are limited to Gaussian boundary distributions, domain-discretization-based methods, which although accurate, do not scale well to high-dimensional problems, and methods utilizing neural networks, which despite being general, are computationally very demanding and lack performance guarantees. In this context, we propose a new method that maintains the numerical tractability and accuracy of covariance steering algorithms, but handles complicated boundary distributions using mixture models. More specifically, in this paper, we present the following main contributions:

- Building on the decomposition technique introduced in [5], we extend the results of [20] to solve the MFSSB problem. Our method splits (5) into a deterministic mean steering problem and an SB problem, with the latter solved via the mixture-based approach from [20].
- We then introduce a more general problem formulation that allows us to solve SB/MFSSB problems with probabilistic state constraints by leveraging the semidefinite formulation of the covariance steering problem subject to chance constraints.
- We use our proposed algorithm to solve the path planning problem for stochastic multi-agent systems in the limit of infinite agents.

II. PRELIMINARIES

A. The Gaussian Schrödinger Bridge

The SB problem with Gaussian marginals, known as the Optimal Covariance Steering (OCS) problem in the control community [11], has been extensively studied in the literature and can be solved either analytically for simple choices of prior dynamics [27] or as a convex semidefinite optimization problem for general linear dynamical systems both for continuous and discrete time cases [28], [29]. Because we use the OCS solution as a building block to construct a policy that works with general GMM boundary distributions, we briefly review the available methods for its solution here.

To this end, consider the optimization problem

$$\min_{u \in \mathcal{U}} \mathbb{E} \left[\int_0^1 \|u_t(x_t)\|^2 dt \right], \quad (6a)$$

$$\text{s.t. } dx_t = A_t x_t dt + B_t u_t dt + D_t dw_t, \quad (6b)$$

$$x_0 \sim \mathcal{N}(\hat{\mu}_0, \hat{\Sigma}_0), \quad x_1 \sim \mathcal{N}(\hat{\mu}_1, \hat{\Sigma}_1), \quad (6c)$$

where $\hat{\mu}_0, \hat{\Sigma}_0, \hat{\mu}_1, \hat{\Sigma}_1$ are the means and covariances of the initial and final Gaussian boundary distributions, respectively. Since the optimal distribution of the state that solves (6) remains Gaussian [11], [29], i.e., $x_t \sim \mathcal{N}(\mu_t, \Sigma_t)$, Problem (6) reduces to that of the control of the mean and the covariance of the state. It can be shown [10], [29] that the optimal policy solving (6) has the form

$$u_t(x) = K_t(x - \mu_t) + v_t. \quad (7)$$

Here, $K_t \in \mathbb{R}^{m \times n}$ is a time-varying feedback gain matrix, and $v_t \in \mathbb{R}^m$ is a feedforward control term. The vector $\mu_t \in \mathbb{R}^n$ denotes the mean of the state distribution at time t . Substituting

(7) in (6b), performing the change of variables $U_t = K_t \Sigma_t$ and introducing the new variable $Y_t = U_t \Sigma_t U_t^\top$, Problem (6) becomes

$$\min \mathbb{E} \left[\int_0^1 \text{tr}(Y_t) + v_t^\top v_t dt \right], \quad (8a)$$

$$\text{s.t. } Y_t = U_t \Sigma_t U_t^\top, \quad (8b)$$

$$\dot{\Sigma}_t = A_t \Sigma_t + \Sigma_t A_t^\top + B_t U_t + U_t^\top B_t^\top + D_t D_t^\top, \quad (8c)$$

$$\dot{\mu}_t = A_t \mu_t + B_t v_t, \quad (8d)$$

$$\mu_0 = \hat{\mu}_0, \Sigma_0 = \hat{\Sigma}_0, \mu_1 = \hat{\mu}_1, \Sigma_1 = \hat{\Sigma}_1, \quad (8e)$$

where variables are U_t, Σ_t, Y_t, v_t , and μ_t . Relaxing (8b) to a semidefinite inequality and using the Schur complement Lemma yields [28]

$$\begin{bmatrix} \Sigma_k & U_k^\top \\ U_k & Y_k \end{bmatrix} \succeq 0, \quad (9)$$

and turns (8) into a semidefinite program (SDP), which can be discretized in time and solved efficiently with any off-the-shelf SDP optimizer.

The semidefinite formulation (8) offers the flexibility of adding probabilistic constraints on the state of the dynamical system (6b). However, because the problem is formulated in continuous time, imposing constraints on the entire trajectory is difficult. To this end, in a similar setting to the discrete-time chance constraint formulation [30], we may constrain the probability that the state belongs to a safe set at a time t , simply by imposing constraints on the moments μ_t, Σ_t .

B. The Schrödinger Bridge Problem with GMM boundary distributions

Before solving the MFSB Problem (5), we discuss the solution of the corresponding SB problem. Although an exact, globally optimal solution to this problem is difficult to obtain when the boundary distributions are GMMs, approximate solutions can be achieved with the technique proposed in [20]. In this section, we review this technique, since it serves as the backbone of our analysis. To this end, consider a generalization of (6) where the boundary distributions are now Gaussian mixture models, that is,

$$\min_{u \in \mathcal{U}} J_{\text{GMM}} \triangleq \mathbb{E} \left[\int_0^1 \|u_t(x_t)\|^2 dt \right], \quad (10a)$$

$$dx = A_t x_t dt + B_t u_t dt + D_t dw, \quad (10b)$$

$$x_0 \sim \sum_{i=1}^{N_0} \alpha_0^i \mathcal{N}(\mu_0^i, \Sigma_0^i), x_1 \sim \sum_{j=1}^{N_1} \alpha_1^j \mathcal{N}(\mu_1^j, \Sigma_1^j). \quad (10c)$$

A set of feasible solutions to (10) can be constructed using the following theorem.

Theorem 1 ([20] Theorem 1). *Consider Problem (10), with N_0 components in the initial mixture and N_1 components in the terminal mixture. Assume that $u_{t|ij}$ is the conditional policy that solves the (i, j) -OCS problem, that is, a policy of the form (7) that steers the system (10b) from the i -th component of the initial mixture to the j -th component of the terminal mixture, and let the resulting probability flow be $\rho_{t|ij} = \mathcal{N}(\mu_{t|ij}, \Sigma_{t|ij})$. Furthermore, let $\lambda_{ij} \geq 0$ such that, for all $j \in \{1, 2, \dots, N_1\}$, $\sum_i \lambda_{ij} = \alpha_1^j$ and for all $i \in \{1, 2, \dots, N_0\}$, $\sum_j \lambda_{ij} = \alpha_0^i$. Then, the policy*

$$u_t(x) = \sum_{i,j} u_{t|ij}(x) \frac{\rho_{t|ij}(x) \lambda_{ij}}{\sum_{i,j} \rho_{t|ij}(x) \lambda_{ij}}, \quad (11)$$

is a feasible policy for Problem (10), and the corresponding probability flow is

$$\rho_t(x) = \sum_{i,j} \rho_{t|ij}(x) \lambda_{ij}. \quad (12)$$

The mixture policy (11) is a weighted average of conditional policies, weighted according to $\lambda_{ij} \rho_{t|ij}(x)$, while the denominator $\sum_{i,j} \rho_{t|ij}(x) \lambda_{ij}$ is just a normalizing constant. Since $\rho_{t|ij}(x)$ is a Gaussian distribution centered at the mean of the (i, j) -OCS flow at time t , this weighting scheme prioritizes the conditional policies whose mean is closer to the value of x at the time t .

Finding the optimal policy within the feasible set defined by Theorem 1 is challenging. To this end, instead of optimizing (10a) directly, we derive an upper bound, which is linear with respect to the transport plan λ_{ij} , and is therefore easier to optimize. More formally, we have the following theorem.

Theorem 2. ([20, Theorem 2]) *Let J_{ij} be the optimal cost of the (i, j) -OCS subproblem with marginal distributions given by the i -th component of the initial mixture and the j -th component of the terminal mixture. Then, the cost function of the linear optimization problem*

$$\min_{\lambda_{ij}} J_{\text{OT}} \triangleq \sum_{i,j} \lambda_{ij} J_{ij}, \quad (13a)$$

such that for all $i = 1, 2, \dots, N_0$, $j = 1, 2, \dots, N_1$

$$\sum_j \lambda_{ij} = \alpha_0^i, \sum_i \lambda_{ij} = \alpha_1^j, \lambda_{ij} \geq 0, \quad (13b)$$

provides an upper bound for (10a), that is, $J_{\text{GMM}} \leq J_{\text{OT}}$, for all values of λ_{ij} satisfying (13b).

Proof. An immediate proof follows from the discrete version of Jensen's inequality, as shown in [20]. Here, we provide an alternative proof using the variance decomposition identity, which allows us to quantify the gap of the upper bound. Specifically, recall that for any random variable $x \in \mathbb{R}^n$, we have $\|\mathbb{E}[x]\|^2 = \mathbb{E}[\|x\|^2] - \mathbb{E}[\|x - \mathbb{E}[x]\|^2]$. Expressing (11) as $u_t(x) = \mathbb{E}[\omega_t(x)]$, where the discrete variable $\omega_t(x)$ is defined by $\omega_t(x) \triangleq \{u_{t|ij}(x) \text{ w.p. } \frac{\rho_{t|ij}(x) \lambda_{ij}}{\rho_t(x)}\}$ we obtain:

$$\begin{aligned} J_{\text{GMM}} &= \int_0^1 \int_{\mathbb{R}^n} \rho_t \|u_t\|^2 dx dt \\ &= J_{\text{OT}} - \int_0^1 \int_{\mathbb{R}^n} \sum_{i,j} \lambda_{ij} \rho_{t|ij} \|u_{t|ij} - u_t\|^2 dx dt. \end{aligned} \quad (14)$$

The second term in (14) is positive by construction, hence dropping it yields the desired result. $\square$

Theorem 2 allows for a very efficient computation of the SB policy since the conditional policies in Theorem 1 can be computed independently and are decoupled from the calculation of the transport plan λ_{ij} , which is computed as the solution of (13). Furthermore, although the upper bound in Theorem 2 introduces some sub-optimality, in practice, we observed that policy (11) results in lower transport costs in SB problems with GMM boundary distributions compared to neural approaches [20], which is attributed to the often imperfect neural network training.

Remark 1. By substituting the policy (11) in (14), it can be shown that when the components of the flow (12) are well-separated, the second term in (14) tends to zero, making the upper bound in Theorem 2 tight. A more rigorous discussion of this case is deferred to Appendix A.

III. UNCONSTRAINED MEAN FIELD SCHÖRINGER BRIDGES

In this section, we extend the above approach to solve the corresponding MFSB problem with GMM boundary conditions. We first study the problem without state constraints, that is, we solve

$$\min_{u \in \mathcal{U}} \mathbb{E} \left[\int_0^1 \|u_t(x_t)\|^2 dt \right], \quad (15a)$$

$$dx_t = A_t x_t dt + \bar{A}_t \bar{x}_t dt + B_t u_t dt + D_t dw_t, \quad (15b)$$

$$x_0 \sim \sum_{i=1}^{N_0} \alpha_0^i \mathcal{N}(\mu_0^i, \Sigma_0^i), \quad x_1 \sim \sum_{j=1}^{N_1} \alpha_1^j \mathcal{N}(\mu_1^j, \Sigma_1^j), \quad (15c)$$

where $\bar{x}_t = \mathbb{E}[x_t]$. We follow an approach similar to [5], but tailored to the specific case of the mixture policy defined in Theorem 1. Defining $\bar{u}_t = \mathbb{E}[u_t(x_t)]$ and $\tilde{x}_t = x_t - \bar{x}_t$, $\tilde{u}_t = u_t - \bar{u}_t$, and observing that, for all $t \in [0, 1]$, $\mathbb{E}[\tilde{x}_t] = 0$ and $\mathbb{E}[\tilde{u}_t] = 0$, one can rewrite (15) as

$$\min_{u \in \mathcal{U}} \mathbb{E} \left[\int_0^1 \|\tilde{u}_t(x_t)\|^2 + \|\bar{u}_t\|^2 dt \right], \quad (16a)$$

$$d\tilde{x}_t = A_t \tilde{x}_t dt + B_t \tilde{u}_t dt + D_t dw, \quad (16b)$$

$$d\bar{x}_t = (A_t + \bar{A}_t) \bar{x}_t dt + B_t \bar{u}_t dt, \quad (16c)$$

$$\mathbb{E}[\tilde{x}_t] = 0, \quad \mathbb{E}[\tilde{u}_t] = 0, \quad (16d)$$

$$\tilde{x}_0 \sim \sum_{i=1}^{N_0} \alpha_0^i \mathcal{N}(\mu_0^i - \hat{x}_0, \Sigma_0^i), \quad (16e)$$

$$\tilde{x}_1 \sim \sum_{j=1}^{N_1} \alpha_1^j \mathcal{N}(\mu_1^j - \hat{x}_1, \Sigma_1^j), \quad (16f)$$

$$\bar{x}_0 = \hat{x}_0, \quad \bar{x}_1 = \hat{x}_1, \quad (16g)$$

where $\hat{x}_0 = \sum_i \mu_0^i w^i$, $\hat{x}_1 = \sum_j \mu_1^j w^j$ are fixed quantities deduced by the parameters of the mixture models representing the means of the boundary distributions (15c). In (16), the problem of calculating the mean control $\bar{u}_t$ and the mean state trajectory $\bar{x}_t$

$$\min_{\bar{u}} \int_0^1 \|\bar{u}_t\|^2 dt \quad \text{s.t.} \quad (16c), (16g), \quad (17)$$

is decoupled from the rest of the problem, which reads

$$\min_{\tilde{u}_t \in \mathcal{U}} \mathbb{E} \left[\int_0^1 \|\tilde{u}_t(x_t)\|^2 dt \right] \quad \text{s.t.} \quad (16b), (16d) - (16f). \quad (18)$$

The solution of (17) can be obtained in closed form, while (18) is of the form (10) with the additional constraints in (16d). Next, consider a relaxation of (18) where we drop the constraint (16d) and solve the remaining problem with the technique of Section II-B. Most importantly, we show that for any policy in the feasible set of policies defined in Theorem 1, the conditions $\mathbb{E}[\tilde{x}_t] = 0$, $\mathbb{E}[\tilde{u}_t] = 0$ will hold automatically due to the structure of the optimal solution of the (i, j) -OCS and the fact that the boundary distributions in (16e) and (16f) have zero-mean. We formalize this in the following proposition.

Proposition 1. *Consider Problem (18) without the constraints (16d), and let $\tilde{u}_t, \tilde{\rho}_t$ be a feasible solution pair of the form (11), (12) for some λ_{ij} satisfying the conditions of Theorem 1. Then, for any such feasible pair, it holds that $\mathbb{E}[\tilde{x}_t] = 0$, $\mathbb{E}[\tilde{u}_t] = 0 \forall t \in [0, 1]$.*

Proof. First, taking the expectations of Equations (11), (12), for the boundary distributions (16e), (16f), we get

$$\mathbb{E}[\tilde{x}_t] = \sum_{ij} \lambda_{ij} \mathbb{E}[\tilde{x}_{t|ij}] = \sum_{ij} \lambda_{ij} \tilde{\mu}_{t|ij}, \quad (19a)$$

$$\mathbb{E}[\tilde{u}_t] = \sum_{ij} \lambda_{ij} \mathbb{E}[\tilde{u}_{t|ij}] = \sum_{ij} \lambda_{ij} \tilde{v}_{t|ij}, \quad (19b)$$

where the expectations are taken with respect to the underlying distribution of the random variable in the expectation and $\tilde{\mu}_{t|ij}, \tilde{v}_{t|ij}$ are the mean state and control effort of the (i, j) -OCS problem with boundary distributions of the i -th component of the initial mixture (16e) and the j -th component of the terminal mixture (16f). Calculating $\tilde{\mu}_{t|ij}, \tilde{v}_{t|ij}$ amounts to a minimum-effort control problem for the linear dynamical system defined by the pair (A_t, B_t) with boundary conditions $\tilde{\mu}_{0|ij} = \mu_0^i - \hat{x}_0$, $\tilde{\mu}_{1|ij} = \mu_1^j - \hat{x}_1$. Their optimal values are [10]

$$\begin{aligned} \tilde{\mu}_{t|ij}^* &= \Phi(t, 1) M(1, t) M_{10}^{-1} \Phi_{10} (\mu_0^i - \hat{\mu}_0) \\ &\quad + M(t, 0) \Phi(1, t)^T M_{10}^{-1} (\mu_1^j - \hat{\mu}_1), \end{aligned} \quad (20a)$$

$$\begin{aligned} \tilde{v}_{t|ij}^* &= B_t^T \Phi(1, t) M_{10}^{-1} (\mu_1^j - \hat{\mu}_1) \\ &\quad - B_t^T \Phi(1, t)^T M_{10}^{-1} \Phi_{10} (\mu_0^i - \hat{\mu}_0), \end{aligned} \quad (20b)$$

where $\Phi(t, s)$ is the state transition matrix of (A_t, B_t) calculated from time s to t , $M(t, s)$ is the respective controllability Grammian, and $\Phi_{10} \equiv \Phi(1, 0)$, $M_{10} = M(1, 0)$. It is easy to verify that substituting (20) to (19) indeed yields $\mathbb{E}[\tilde{x}_t] = 0$, $\mathbb{E}[\tilde{u}_t] = 0$ regardless of the choice of λ_{ij} . $\square$

An immediate consequence of Proposition 1 is that given the solution $\tilde{\rho}_t, \tilde{u}_t$ of Problem (18) and the solution $\bar{x}_t, \bar{u}_t$ of Problem (17), a solution for the MFSB Problem (15) can be constructed using the formulas $\rho_t(x) = \tilde{\rho}_t(x - \bar{x}_t)$ and $u_t(x) = \tilde{u}_t(x - \bar{x}_t) + \bar{u}_t$.

IV. CONSTRAINED MEAN FIELD SCHRÖDINGER BRIDGES

In this section, we are interested in further extending the computation of the mixture policy (11) to handle constrained problems of the form

$$\min_{u \in \mathcal{U}} \mathbb{E} \left[\int_0^1 \|u_t(x_t)\|^2 dt \right], \quad (21a)$$

$$(15b), (15c), \mathbb{P}(x_t \in \mathcal{X}) \geq 1 - \delta_t. \quad (21b)$$

where $\mathcal{X} = \bigcup_n \mathcal{X}_n \triangleq \bigcup_n \{x \in \mathbb{R}^n : \alpha_n^T x \leq \beta_n\}$. Contrary to the technique in Section III, we do not split the state and control action into mean and deviation components, as in the decomposition of Problem (15) to (17) and (18) because the two problems are coupled through the chance constraint (21b). In the same way we combine conditional OCS policies to solve the SB problem for GMM boundary conditions in Theorem 1, our approach will now be to combine conditional policies, each solving a constrained mean-field OCS problem and assemble them together to solve the MFSB problem with GMM boundary conditions. Similarly to the case without the mean-field term, the policy will have the form (11), and the resulting probability flow will have the form (12).

Before presenting the details for the calculation of the conditional mean-field OCS problem, we reformulate the chance constraint (21b) in terms of the parameters of the conditional i - j policies. To do that, we assume that, similarly to the standard SB problem, the adopted policy will yield a probability flow of the form (12). We will later prove that this is indeed the case. To transform the constraint (21b) to constraints on the moments of the conditional probabilities $\rho_{t|ij} = \mathcal{N}(\mu_{t|ij}, \Sigma_{t|ij})$, we allocate the total constraint violation budget δ_t to the probability of the i - j component to exit the half-spaces $\mathcal{X}_n$, denoted $\delta_{t|ijn}$. Using the union bound [19], (21b) can be conservatively replaced by

$$\mathbb{P}(x_{t|ij} \in \mathcal{X}_n) \geq 1 - \delta_{t|ijn}, \quad \text{and} \quad \sum_{i,j,n} \lambda_{ij} \delta_{t|ijn} \leq \delta_t. \quad (22)$$

The simplest option for implementing (22) in practice is to allocate δ_t equally on all components and half-spaces, selecting $\delta_{t|ijn} = \delta_t/N_c$. If this is too conservative for a given application, the results could be refined with an iterative allocation scheme such as the one proposed in [19], [31]. Furthermore, since $x_{t|ij} \sim \mathcal{N}(\mu_{t|ij}, \Sigma_{t|ij})$, the constraint $\mathbb{P}(x_{t|ij} \in \mathcal{X}_n) \geq 1 - \delta_{t|ijn}$ is equivalent to

$$\Phi^{-1}(1 - \delta_{t|ijn})\sqrt{a_n^\top \Sigma_{t|ij} a_n} + a_n^\top \mu_{t|ij} \leq \beta_n, \quad (23)$$

where $\Phi^{-1}(\cdot)$ is the inverse cumulative distribution function for the normal distribution. Following the technique in [30], the constraint (23) is linearized around some reference values of the (i, j) problem, yielding a linear constraint of the form

$$\ell_{ijn}^\top \Sigma_{t|ij} \ell_{ijn} + \mu_{t|ij}^\top a_{ijn} + b_{ijn} \leq 0, \quad (24)$$

where $b_{ijn} = -\beta_n - \Phi^{-1}(1 - \delta_{t|ijn})\frac{1}{2}\sqrt{a_n^\top \Sigma_r a_n}$ and $\ell_{ijn} = \sqrt{\Phi^{-1}(1 - \delta_{t|ijn})\frac{1}{2\sqrt{a_n^\top \Sigma_r a_n}}}a_n$, and Σ_r being a reference value used for the linearization.

We now present the (i, j) constrained mean-field OCS problem, which reads

$$\min_{u_{t|ij} \in \mathcal{U}} \mathbb{E} \left[\int_0^1 \|u_{t|ij}(x_{t|ij})\|^2 dt \right], \quad (25a)$$

$$dx_{t|ij} = A_t x_{t|ij} dt + \bar{A}_t \bar{x}_t dt + B_t u_{t|ij} dt + D_t dw, \quad (25b)$$

$$x_0 \sim \mathcal{N}(\mu_0^i, \Sigma^i), x_1 \sim \mathcal{N}(\mu_1^j, \Sigma_1^j), \quad (25c)$$

$$\mathbb{P}(x_{t|ij} \in \mathcal{X}) \geq 1 - \delta_{t|ij}. \quad (25d)$$

Although solving (25) requires the unknown mean-field term $\bar{x}_t$, we provide the following result, which allows us to solve all the mean-field OCS policies jointly.

Theorem 3. *Given the optimal solution $u_{t|ij}$ and $\rho_{t|ij} = \mathcal{N}(\mu_{t|ij}, \Sigma_{t|ij})$ to (25) and a transport plan λ_{ij} satisfying the conditions of Theorem 1, the mixture policy (11) is a feasible policy for (25), and the resulting probability flow is (12).*

Proof. The proof follows from the fact that when the conditional policies solve (25), the pair of (11) and (12) satisfies the mean-field FPK Equation (4). Due to its similarity with the proof of Theorem 1, the details are omitted. $\square$

With (3) in mind, taking the expectation of (12), the mean-field term $\bar{x}_t$ can be calculated explicitly, yielding $\bar{x}_t = \sum_{ij} \lambda_{ij} \mu_{t|ij}$, which suggests a coupling between the Gaussian mean-field OCS problems. Using Theorems 3 and 2, we can assemble all the conditional mean-field OCS problems into a single coupled optimization problem using the moment version of (25), which yields

$$\min \sum_{ij} \lambda_{ij} \int_0^1 \text{tr}(Y_{t|ij}) + \|v_{t|ij}\|^2 dt, \quad (26a)$$

such that for all $i = 1, 2, \dots, N_0, j = 1, 2, \dots, N_1$

$$\begin{bmatrix} \Sigma_{t|ij} & U_{t|ij}^\top \\ U_{t|ij} & Y_{t|ij} \end{bmatrix} \succeq 0, \quad (26b)$$

$$\dot{\Sigma}_{t|ij} = A_t \Sigma_{t|ij} + \Sigma_{t|ij} A_t^\top + B_t U_{t|ij} + U_{t|ij}^\top B_t^\top + D_t D_t^\top, \quad (26c)$$

$$\dot{\mu}_{t|ij} = A_t \mu_{t|ij} + \bar{A} \left(\sum_{ij} \lambda_{ij} \mu_{t|ij} \right) + B_t v_{t|ij}, \quad (26d)$$

$$(13b), \mu_{0|ij} = \mu_0^i, \Sigma_{0|ij} = \Sigma_0^i, \mu_{1|ij} = \mu_1^j, \Sigma_{1|ij} = \Sigma_1^j, \quad (26e)$$

$$\ell_{ijn}^\top \Sigma_{t|ij} \ell_{ijn} + \mu_{t|ij}^\top a_{ijn} + b_{ijn} \leq 0, \quad (26f)$$

TABLE I: Results for Problem 2.

Gap width	Fig.1c	Fig.1d	Fig.1e
Total Cost J	153.89	179.56	504.83
Cost Upper Bound (26a)	306.08	357.16	999.55
$\max_t \mathbb{P}(x_t \notin \mathcal{X})$	0.3%	0.4%	0.5%

Problem (26) is a bilinear semidefinite program due to the products of the mixing weights λ_{ij} and the policy parameters of the (i, j) subproblem. To solve this problem, we alternate between fixing the transport plan λ_{ij} and solving for the conditional OCS parameters, and keeping the conditional policy parameters fixed and optimizing for λ_{ij} by solving (13) until the iteration converges. Terminating the iteration with an update in the (i, j) policy parameters, i.e., solving (26) for a fixed λ_{ij} ensures that the solution will be feasible, although the transport plan λ_{ij} might be sub-optimal. In practice, we observed that a few iterations of this scheme are enough for convergence in all problems we solved.

V. NUMERICAL EXAMPLES

Problem 1: We first apply the algorithm in Section III to solve an unconstrained MFSB problem. We use the LTI dynamical system of the form (5b) with $A = \bar{A} = B = D = I_2$, and for boundary distributions the GMMs of Fig. 1a-1b. We use blue for the initial and light green for the final distributions. We discretize the continuous time problem in 101 time steps and solve the resulting SDP problem with MOSEK [32]. The resulting agent trajectories are illustrated in Fig. 1b. We note that for all our simulations, we use the empirical mean-field term computed from 10,000 agents.

Problem 2: To test the algorithm on a more realistic problem, we considered a swarm of agents following double integrator dynamics, interacting through a negative quadratic potential that incentivizes them to spread apart. The distribution of the initial and final configurations of the agents is depicted in Fig. 1c-1e. The task for this problem is to guide the swarm through an obstacle by either going through a narrow passage or by going around it. The agents follow the dynamics (5b) with parameters

$$A = \begin{bmatrix} 0_2 & I_2 \\ I_2 & 0_2 \end{bmatrix}, B = \begin{bmatrix} 0_2 \\ I_2 \end{bmatrix}, \bar{A} = \begin{bmatrix} 0_2 & 0_2 \\ -I_2 & 0_2 \end{bmatrix}, D = I_4 \quad (27)$$

The violation probability for each wall in the obstacles is set to $\delta_{t|ijn} = 0.3\%$, giving a total violation of up to 0.9%. For each conditional policy connecting a component of the initial mixture to a component of the terminal mixture, we solve the problem through both passages. The transport plan in this case becomes a tensor with three indices; however, the corresponding Problem (13) remains an easy-to-solve, low-dimensional linear program. We solve the problem for three different widths of the passage shown in Fig. 1c-1e. As expected, for a wider passage, the solver routes some of the modes through it, while for smaller passages, the go-around route is preferred. In all problems, the transport plan converges within a few iterations. For each run, we report the total resulting true cost of policy (11), the cost upper bound, which is the value of the optimal cost function (26a), as well as the maximum constraint violation in Table I. As in our first example, we note that during simulation, we use the empirical mean-field term computed by simulating the interactions of 10,000 agents.




Fig. 1: (a): ρ_0 , $\rho_{0.5}$, and ρ_1 for Problem 1. (b): 100 agent trajectories for Problem 1. (c)-(e): ρ_0 to ρ_1 for Problem 2 overlaid with 100 agent trajectories. The shaded regions represent state constraints that the trajectories must avoid.

VI. CONCLUSIONS

In this paper, we present a method for approximating the solutions of the Mean-Field Schrödinger Bridge problem for systems with Linear Time-Varying dynamics and Gaussian mixture model boundary distributions. Our approach is based on a mixture policy consisting of linear feedback laws, each solving a conditional problem from one component of the initial mixture to another component of the terminal mixture. The elementary policies are weighted with an optimal component-level transport plan. In the case of unconstrained problems, the calculation of the transport plan is decoupled from the rest of the problem, and amounts to solving a linear program. For the case where the problems contain state constraints, the calculation of the conditional policies and the transport plan are coupled through the chance constraints and the mean-field term. An iterative optimization method is proposed to solve the resulting optimization problem. Although we do not have a proof of convergence, we observed in practice that the method converges in only a few iterations in all numerical examples.

REFERENCES

- [1] Y. Chen, "Density control of interacting agent systems," *Transactions on Automatic Control*, vol. 69, no. 1, pp. 246–260, 2024.
- [2] J.-M. Lasry and P.-L. Lions, "Mean field games," *Japanese journal of mathematics*, vol. 2, no. 1, pp. 229–260, 2007.
- [3] M. Huang, R. P. Malhamé, and P. E. Caines, "Large population stochastic dynamic games: closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle," *Communications in Information & Systems*, vol. 6, no. 3, pp. 221 – 252, 2006.
- [4] R. Carmona, F. Delarue, et al., *Probabilistic theory of mean field games with applications I-II*. Springer, 2018.
- [5] Y. Chen, T. T. Georgiou, and M. Pavon, "Steering the distribution of agents in mean-field games system," *Journal of Optimization Theory and Applications*, vol. 179, pp. 332–357, 2018.
- [6] K. Elamvazhuthi and S. Berman, "Mean-field models in swarm robotics: a survey," *Bioinspiration & Biomimetics*, vol. 15, p. 015001, Nov. 2019.
- [7] K. Elamvazhuthi and S. Berman, "Optimal control of stochastic coverage strategies for robotic swarms," in *IEEE International Conference on Robotics and Automation*, pp. 1822–1829, IEEE, 2015.
- [8] K. Elamvazhuthi, Z. Kakish, A. Shirsat, and S. Berman, "Controllability and Stabilization for Herding a Robotic Swarm Using a Leader: A Mean-Field Approach," *IEEE Transactions on Robotics*, vol. 37, pp. 418–432, Apr. 2021.
- [9] J. Backhoff, G. Conforti, I. Gentil, and C. Léonard, "The mean field Schrödinger problem: ergodic behavior, entropy estimates and functional inequalities," *Probability Theory and Related Fields*, vol. 178, pp. 475–530, 2020.
- [10] Y. Chen, T. T. Georgiou, and M. Pavon, "Optimal transport over a linear dynamical system," *Transactions on Automatic Control*, vol. 62, no. 5, pp. 2137–2152, 2016.
- [11] Y. Chen, T. T. Georgiou, and M. Pavon, "Optimal steering of a linear stochastic system to a final probability distribution, part I," *Transactions on Automatic Control*, vol. 61, no. 5, pp. 1158–1169, 2015.
- [12] J. Pilipovsky and P. Tsiotras, "DUST: A framework for data-driven density steering," *arXiv preprint arXiv:2408.02777*, 2024.
- [13] V. Renganathan, J. Pilipovsky, and P. Tsiotras, "Distributionally robust covariance steering with optimal risk allocation," in *American Control Conference*, (Philadelphia, PA), pp. 2607–2614, May 2023.
- [14] N. Kumagai and K. Oguri, "Sequential chance-constrained covariance steering for robust cislunar trajectory design under uncertainties," in *AAS/AIAA Astrodynamics Specialist Conference*, 2024.
- [15] A. D. Saravanos, A. Tzolovikos, E. Bakolas, and E. Theodorou, "Distributed covariance steering with consensus ADMM for stochastic multi-agent systems," *Robotics: Systems and Science*, July 2021.
- [16] A. D. Saravanos, I. M. Balci, E. Bakolas, and E. A. Theodorou, "Distributed model predictive covariance steering," in *International Conference on Intelligent Robots and Systems*, pp. 5740–5747, 2024.
- [17] A. D. Saravanos, Y. Li, and E. A. Theodorou, "Distributed hierarchical distribution control for very-large-scale clustered multi-agent systems," in *Robotics: Science and Systems XIX*, (Daegu, Republic of Korea), July 2023.
- [18] I. M. Balci and E. Bakolas, "Density steering of Gaussian mixture models for discrete-time linear systems," in *2024 American Control Conference (ACC)*, pp. 3935–3940, July 2024.
- [19] N. Kumagai and K. Oguri, "Chance-constrained Gaussian mixture steering to a terminal Gaussian distribution," in *63rd Conference on Decision and Control*, (Milan), pp. 2207–2212, December 2024.
- [20] G. Rapakoulias, A. R. Pedram, and P. Tsiotras, "Go with the flow: Fast diffusion for Gaussian mixture models," *arXiv preprint arXiv:2412.09059*, 2024.
- [21] Y. Mei, M. Al-Jarrah, A. Taghvaei, and Y. Chen, "Flow matching for stochastic linear control systems," in *Proceedings of the 7th Annual Learning for Dynamics & Control Conference*, vol. 283, pp. 484–496, PMLR, Jun 2025.
- [22] A. Terpin, N. Lanzetti, and F. Dörfler, "Dynamic programming in probability spaces via optimal transport," *SIAM Journal on Control and Optimization*, vol. 62, no. 2, pp. 1183–1206, 2024.
- [23] G.-H. Liu, T. Chen, O. So, and E. Theodorou, "Deep generalized Schrödinger bridge," in *Advances in Neural Information Processing Systems*, vol. 35, (Louisiana, LA), pp. 9374–9388, Curran Associates, Inc., Dec. 2022.
- [24] L. Ruthotto, S. J. Osher, W. Li, L. Nurbekyan, and S. W. Fung, "A machine learning framework for solving high-dimensional mean field game and mean field control problems," *Proceedings of the National Academy of Sciences*, vol. 117, no. 17, pp. 9183–9193, 2020.
- [25] A. T. Lin, S. W. Fung, W. Li, L. Nurbekyan, and S. J. Osher, "Alternating the population and control neural networks to solve high-dimensional stochastic mean-field games," *Proceedings of the National Academy of Sciences*, vol. 118, p. e2024713118, Aug. 2021.
- [26] K. F. Caluya and A. Halder, "Wasserstein proximal algorithms for the Schrödinger bridge problem: Density control with nonlinear drift," *Transactions on Automatic Control*, vol. 67, no. 3, pp. 1163–1178, 2021.
- [27] C. Bunne, Y.-P. Hsieh, M. Cuturi, and A. Krause, "The Schrödinger bridge between Gaussian measures has a closed form," in *International Conference on Artificial Intelligence and Statistics*, pp. 5802–5833, PMLR, April 2023.
- [28] Y. Chen, T. T. Georgiou, and M. Pavon, "Optimal steering of a linear stochastic system to a final probability distribution, part II," *Transactions on Automatic Control*, vol. 61, no. 5, pp. 1170–1180, 2015.
- [29] F. Liu, G. Rapakoulias, and P. Tsiotras, "Optimal covariance steering for discrete-time linear stochastic systems," *Transactions on Automatic Control*, pp. 1–16, 2024.
- [30] G. Rapakoulias and P. Tsiotras, "Discrete-time optimal covariance steering via semidefinite programming," in *62nd Conference on Decision and Control*, (Singapore), pp. 1802–1807, Dec. 2023.
- [31] J. Pilipovsky and P. Tsiotras, "Covariance steering with optimal risk allocation," *Transactions on Aerospace and Electronic Systems*, vol. 57, no. 6, pp. 3719–3733, 2021.
- [32] MOSEK ApS, "MOSEK modeling cookbook," 2020.

A. Optimality of the mixture policy

In this section we rigorously study the gap in the upper bound of Theorem 2 and show that it asymptotically converges to zero as the components become well-separated. To this end let $\rho_{t|ij}, u_{t|ij}$ be the solution of the (i, j) Covariance steering problem of the form (6), with boundary distributions the i -th component of the initial and the j -th component of the terminal mixture, and let $\rho_t = \sum_{ij} \lambda_{ij} \rho_{t|ij}$ and

$$u_t(x) = \sum_{i,j} u_{t|ij}(x) \frac{\rho_{t|ij}(x) \lambda_{ij}}{\rho_t(x)} = \mathbb{E}[\omega_t(x)], \quad (28)$$

where $\omega_t(x)$ follows a discrete distribution defined by $\{\omega_t(x) = u_{t|ij}(x) \text{ w.p. } \frac{\rho_{t|ij}(x) \lambda_{ij}}{\rho_t(x)}\}$.

Note that for any random variable $x \in \mathbb{R}^n$, the variance decomposition yields

$$\|\mathbb{E}[x]\|^2 = \mathbb{E}[\|x\|^2] - \mathbb{E}[\|x - \mathbb{E}[x]\|^2]. \quad (29)$$

Using (29) and the expression (28) in (10a) we obtain

$$J_{\text{GMM}} = \int_0^1 \int_{\mathbb{R}^n} \rho_t(x) \|u_t(x)\|^2 dx dt \quad (30)$$

$$= J_{\text{OT}} - \int_0^1 \int_{\mathbb{R}^n} \sum_{i,j} \lambda_{ij} \rho_{t|ij} \|u_{t|ij} - u_t\|^2 dx dt. \quad (31)$$

The fact that the second term in (31) is positive justifies the upper bound in Theorem 2. Now, we show that when the conditional densities are well separated, the second term in (31) goes to zero, and the bound becomes tight.

First, we note two useful identities that we use to calculate the density of the product and quotient of Gaussian densities. To this end, consider two Gaussian densities, $\rho_1 = \mathcal{N}(\mu_1, \Sigma_1)$, $\rho_2 = \mathcal{N}(\mu_2, \Sigma_2)$. Then

$$\rho_1 \rho_2 = c_p \mathcal{N}(\mu_p, \Sigma_p), \quad (32)$$

$$\frac{\rho_1}{\rho_2} = c_q \mathcal{N}(\mu_q, \Sigma_q), \quad (33)$$

where

$$\mu_p = (\Sigma_1^{-1} + \Sigma_2^{-1})^{-1} (\Sigma_1^{-1} \mu_1 + \Sigma_2^{-1} \mu_2), \quad (34)$$

$$\Sigma_p = (\Sigma_1^{-1} + \Sigma_2^{-1})^{-1}, \quad (35)$$

$$c_p = \mathcal{N}(\mu_1; \mu_2, \Sigma_1 + \Sigma_2), \quad (36)$$

$$\mu_q = (\Sigma_1^{-1} - \Sigma_2^{-1})^{-1} (\Sigma_1^{-1} \mu_1 - \Sigma_2^{-1} \mu_2), \quad (37)$$

$$\Sigma_q = (\Sigma_1^{-1} - \Sigma_2^{-1})^{-1}, \quad (38)$$

$$c_q = \frac{|\Sigma_2|}{|\Sigma_2 - \Sigma_1|} \frac{1}{\mathcal{N}(\mu_1; \mu_2, \Sigma_2 - \Sigma_1)}. \quad (39)$$

Using (32), it is easy to show that when the components $\rho_{t|ij}$ are well separated, i.e., $\|\mu_{t|ij} - \mu_{t|i'j'}\| > c \forall, t \in [0, 1]$ for some sufficiently large constant c , whenever $(i, j) \neq (i', j')$:

$$\rho_{t|ij}(x) \rho_{t|i'j'}(x) \leq c_1 \mathcal{N}(\mu, \Sigma) \quad \forall x \in \mathbb{R}^n \quad (40)$$

for some μ, Σ defined through (34), (35), and $c_1 \rightarrow 0$ exponentially as $c \rightarrow \infty$.

With the above tools, expanding the term inside the norm in Equation (14), and dropping the dependence on x for notational convenience, we obtain

$$\|u_{t|i,j} - u_t\|^2 \quad (41)$$

$$= \left\| u_{t|i,j} - \sum_{i',j'} u_{t|i'j'} \frac{\lambda_{i'j'} \rho_{t|i'j'}}{\rho_t} \right\|^2 \quad (42)$$

$$= \left\| \frac{u_{t|i,j} \rho_t - \sum_{i',j'} u_{t|i'j'} \lambda_{i'j'} \rho_{t|i'j'}}{\rho_t} \right\|^2 \quad (43)$$

$$= \left\| \frac{\sum_{i' \neq i, j' \neq j} (u_{t|i,j} - u_{t|i'j'}) \lambda_{i'j'} \rho_{t|i'j'}}{\rho_t} \right\|^2 \quad (44)$$

$$\leq c_2 \sum_{\substack{i' \neq i \\ j' \neq j}} \|u_{t|i,j} - u_{t|i'j'}\|^2 \left(\frac{\lambda_{i'j'} \rho_{t|i'j'}}{\rho_t} \right)^2 \quad (45)$$

$$\leq c_2 \sum_{\substack{i' \neq i \\ j' \neq j}} \|u_{t|i,j} - u_{t|i'j'}\|^2 \frac{\lambda_{i'j'} \rho_{t|i'j'}}{\rho_t} \quad (46)$$

$$\leq c_2 \sum_{\substack{i' \neq i \\ j' \neq j}} \|u_{t|i,j} - u_{t|i'j'}\|^2 \frac{\lambda_{i'j'} \rho_{t|i'j'}}{\max_{i,j} \{\lambda_{ij} \rho_{t|ij}\}} \quad (47)$$

where the max in (47) selects the component with the largest covariance and $c_2 = (N_0 - 1)(N_1 - 1)$, i.e., the cardinality of the sum of policies. Rearranging terms in Equation (31) and substituting the bound (47), we obtain

$$\begin{aligned} 0 &\leq J_{\text{OT}} - J_{\text{GMM}} \\ &\leq \int_0^1 \int_{\mathbb{R}^n} \sum_{i,j} \lambda_{ij} \rho_{t|ij} \|u_{t|ij}(x) - u_t(x)\|^2 dx dt \\ &\leq c_2 \int_0^1 \int_{\mathbb{R}^n} \sum_{i,j} \sum_{\substack{i' \neq i \\ j' \neq j}} \|u_{t|ij} - u_{t|i'j'}\|^2 \frac{\lambda_{i'j'} \lambda_{ij} \rho_{t|i'j'} \rho_{t|ij}}{\max_{i,j} \{\lambda_{ij} \rho_{t|ij}\}} dx dt \\ &= c_2 c_1 \int_0^1 \int_{\mathbb{R}^n} \sum_{i,j} \sum_{\substack{i' \neq i \\ j' \neq j}} \|u_{t|ij} - u_{t|i'j'}\|^2 \frac{\lambda_{i'j'} \lambda_{ij} \mathcal{N}(\mu, \Sigma)}{\max_{i,j} \{\lambda_{ij} \rho_{t|ij}\}} dx dt \end{aligned} \quad (48)$$

$$= c_2 c_1 C \rightarrow 0 \quad \text{as } c \rightarrow \infty, \quad (49)$$

where C is the value of the integral in (48), which is finite, since the quotient density in (48) is well defined, due to (33).

Finally, we note that although the condition that the mixture components are well separated is quite strict, and might not be satisfied in practice for problems with many components, it can be relaxed to only require the components with non-zero weights in the optimal transport plan λ_{ij} to be well-separated. The last condition is hard to check before solving the problem; however, it is satisfied for many problems in practice since the optimal transport plan λ_{ij} is usually sparse.