Optimal Steering of Ensembles with Heterogeneous Objectives

Isabel Haasler, Yongxin Chen, and Johan Karlsson

***Abstract*— We consider the optimal control problem of steering a collection of agents over a network. The group behavior of an ensemble is often modeled by a distribution, and thus the optimal control problem we study can be cast as a distribution steering problem. While most existing works for steering distributions require the agents in the ensemble to be homogeneous, we consider the setting of agents with heterogeneous objectives. This control problem also resembles a minimum cost network flow problem with a massive number of commodities. We propose a novel framework for this problem and derive an efficient algorithm for solving it. Our method is based on optimal transport theory, and extends it to settings with heterogeneous agents. The proposed method is illustrated on a numerical simulation for traffic planning.**

I. INTRODUCTION

Many problems in control and estimation involve a large number of agents, ranging from applications in engineering and biology to social sciences [1]. Such groups of indistinguishable and almost identically behaving agents are also called ensembles. In general, modelling the dynamics and control for individual agents in a large ensemble is challenging. One remedy is to design a control law that steers the ensemble collectively without designing control laws for each individual [2]. Therefore, the aggregate state information of the agents is often described by a distribution, or density function [2], [3]. A well-studied problem based on this representation is to steer a given distribution to a target one [4]–[6]. In [4] the steering problem is solved based on the optimal transport framework. Similar methods have been applied to estimating distributions [7] and tracking of ensembles [8], [9].

An inherent feature of these optimal transport methods as well as typical ensemble applications, is that one models only aggregate distributions and flows, but no individual behaviours. This reflects the underlying asumption that agents in the ensemble are indistinguishable and behave identically. However, in some applications this assumption does not match reality. For instance, agents in the ensemble may have heteregenous objectives, such as agents in a traffic network, which are each equipped with the objective to reach their destination. Similar problems appear in air traffic planning, railroad traffic scheduling, communication and logistics, and are often treated as multi-commodity flow problems over networks [10]–[12].

In this work, we extend the optimal transport framework in order to allow for modelling ensembles with heterogeneous objectives. More precisely, inspired by the example of agents moving in a traffic network, we consider the problem of steering ensembles with heterogeneous objectives over a network. Note that the framework for steering distributions in [4] has been applied to networks in order to find robust transport plans [13]–[15].

In this work, we extend the existing methods for robust transport over networks to the case of ensembles with heterogeneous objectives. In addition, our method can be seen as new framework for addressing multi-commodity flow problems [11], [12], [16], and allows for solving problems with a massive number of commodities. Moreover, we present an efficient algorithm to solve the optimal transport formulation.

The paper is structured as follows. In Section II we review some background material on optimal transport. The main contribution is presented in Section III, where we formulate the optimal steering problem for heterogeneous ensembles, and develop an efficient algorithm for solving it. In Section IV, we numerically illustrate the proposed framework on a traffic planning simulation. Some proofs are deferred to the appendix.

II. BACKGROUND ON OPTIMAL TRANSPORT

The optimal transport problem is to find a transport plan that minimizes the cost of moving the mass from an initial distribution to a target distribution [17]. In this work, we consider optimal transport over discrete space. In the discrete optimal transport setting the two distributions are represented by non-negative vectors $\mu_0, \mu_1 \in \mathbb{R}_+^n$, and the transport plan is a matrix $M \in \mathbb{R}_+^{n\times n}$, where the element M_{ij} denotes the amount of mass transported from point i to point j. Thus, a transport plan between the two distributions μ_0 and μ_1 satisfies $M\mathbf{1} = \mu_0$ and $M^T\mathbf{1} = \mu_1$, where $\mathbf{1} \in \mathbb{R}^{n\times 1}$ denotes a vector of ones. Given a cost matrix $C \in \mathbb{R}_+^{n\times n}$, where the element C_{ij} denotes the cost for moving a unit mass from i to j, the optimal transport problem can then be formulated as

$$\begin{aligned} &\underset{M\in\mathbb{R}_+^{n\times n}}{\text{minimize}} \ \text{trace}(C^T M) \\ &\text{subject to } M\mathbf{1} = \mu_0, \quad M^T\mathbf{1} = \mu_1. \end{aligned} \tag{1}$$

The optimal transport problem has been extended to settings with several marginals $\mu_0, \dots, \mu_\mathcal{T} \in \mathbb{R}_+^n$ [7], [18]–[20]. In this multi-marginal optimal transport setting, the transport plan and cost are represented by $(\mathcal{T}+1)$-mode tensors $\mathbf{M}, \mathbf{C} \in \mathbb{R}_+^{n\times\cdots\times n}$. For a given tuple $(i_0, \dots, i_\mathcal{T})$,

This work was supported by the Swedish Research Council (VR), grant 2014-5870, SJTU-KTH cooperation grant and the NSF under grant 1901599 and 1942523.

I. Haasler and J. Karlsson are with the Division of Optimization and Systems Theory, Department of Mathematics, KTH Royal Institute of Technology, Stockholm, Sweden. {haasler, johan.karlsson}@math.kth.se

Y. Chen is with the School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA, USA. yongchen@gatech.edu

the element $\mathbf{M}_{i_0,\dots,i_\mathcal{T}}$ is the associated amount of transported mass, and $\mathbf{C}_{i_0,\dots,i_\mathcal{T}}$ is the associated transportation cost of a unit mass. The multi-marginal optimal transport problem then reads

$$\begin{aligned} &\underset{\mathbf{M}\in\mathbb{R}_+^{n\times\cdots\times n}}{\text{minimize}} \ \langle \mathbf{C}, \mathbf{M}\rangle \\ &\text{subject to } P_t(\mathbf{M}) = \mu_t, \qquad \text{for } t \in \Gamma. \end{aligned} \tag{2}$$

where $P_t(\mathbf{M})$ denotes the projection on the t-th marginal of $\mathbf{M}$ and is defined by

$$(P_t(\mathbf{M}))_{i_t} = \sum_{i_1,\dots,i_{j-1},i_{j+1},i_J} \mathbf{M}_{i_0,\dots,i_{t-1},i_t,i_{t+1},\dots,i_\mathcal{T}}, \tag{3}$$

and $\Gamma \subset \{0,\dots,\mathcal{T}\}$ is an index set that specifies the set of marginals on which constraints are imposed. In the standard multi-marginal optimal transport problem one assumes that all marginals are given, i.e., $\Gamma = \{0,\dots,\mathcal{T}\}$. However, in this work we consider the case where the two sets are not necessarily equal. Such optimal transport problems can be utilized to estimate the distributions on the unknown marginals [7].

Even though the optimal transport problem (2) is a linear program, it suffers from the curse of dimensionality, which makes it computationally infeasible to solve it directly in many applications [7]. A popular method for the classical bi-marginal problem (1) to decrease the computational cost is to regularize it with an entropy term [21]. This approach can also be applied to the multi-marginal problem (2) to partly alleviate the computational burden [7], [18], [22]. In particular, the entropy-regularized multi-marginal optimal transport problem reads

$$\begin{aligned} &\underset{\mathbf{M}\in\mathbb{R}^{n\times\cdots\times n}}{\text{minimize}} \ \langle \mathbf{C}, \mathbf{M}\rangle + \epsilon D(\mathbf{M}) \\ &\text{subject to } P_t(\mathbf{M}) = \mu_t, \qquad \text{for } t \in \Gamma, \end{aligned} \tag{4}$$

where $\epsilon > 0$ is a regularization parameter, and the entropy term is defned as

$$D(\mathbf{M}) = \sum_{i_0,\dots i_\mathcal{T}} \left(\mathbf{M}_{i_0\dots i_\mathcal{T}} \log(\mathbf{M}_{i_0\dots i_\mathcal{T}}) - \mathbf{M}_{i_0\dots i_\mathcal{T}} + 1\right).$$

The entropy regularized optimal transport problem (4) is strictly convex and thus has a unique solution. Moreover, it can be shown that the optimal solution is of the form

$$\mathbf{M} = \mathbf{K} \odot \mathbf{U}, \tag{5}$$

where $\odot$ denotes element-wise multiplication and the tensors are of the form $\mathbf{K} = \exp(-\mathbf{C}/\epsilon)$ and $\mathbf{U} = u_0 \otimes u_1 \otimes \cdots \otimes u_\mathcal{T}$ [7]. The vectors in $\mathbf{U}$ are given by $u_t = \exp(-\lambda_t/\epsilon)$, where $\lambda_t \in \mathbb{R}^n$ is the Lagrange dual variable for the constraint on $P_t(\mathbf{M})$ in problem (4), and $u_t = \mathbf{1}$ if $t \notin \Gamma$ [7]. The vectors u_t for $t \in \Gamma$, can be found as the fixed point of the so called Sinkhorn iterations, which are to iterate

$$u_t \leftarrow u_t \odot \mu_t./P_t(\mathbf{K} \odot \mathbf{U}), \tag{6}$$

for $t \in \Gamma$, where ./ denotes element-wise division.

The Sinkhorn iterations (6) can be seen as iterative Bregman projections [18] or, more generally, a block coordinate ascent in the dual of (4) [7], [22], [23]. The scheme (6) is therefore guaranteed to converge globally [24], [25]. Although the entropy regularized problem (4) can be solved more efficiently than the original formulation (2), the complexity of the computation of the projections $P_t(\mathbf{M})$ in (3) increases exponentially with $\mathcal{T}$, and thus is often too expensive. However, in some applications the cost tensor $\mathbf{C}$ has structures that can be utilized for efficient computations of the projections [18], [22]. Such structures also appear in, e.g., information fusion, interpolation, and tracking applications [7].

III. Steering of heterogeneous ensembles

The optimal transport framework is a useful tool for the estimation and steering of distributions [4], [7]. Based on this, it has also been applied to the tracking of ensembles [8], [9]. A main assumption of the optimal transport framework for ensembles is that agents are indistinguishable and the ensemble can thus be described by aggregate distributions and flows. On a network the optimal transport problem can hence be seen as single-commodity flow problem [14]. However, in many applications groups are better modeled as ensembles with heterogeneous objectives, i.e., agents may have individual objectives. For instance, in traffic control problems agents are usually required to reach specific destinations. Such problems on networks are often modeled as minimum cost multi-commodity flow problems [11], [16]. In this section, we develop an optimal transport based method for ensemble steering with heterogeneous objectives. This extends the network flow interpretation of the optimal transport problem to the multi-commodity case.

A. Optimal transport over networks

Consider the optimal transport problem over a network where the set of edges is denoted by E, with $|E| = n$, and let μ_t denote a distribution over E at time $t = 0,\dots,\mathcal{T}$. Define a cost tensor $\mathbf{C} \in \mathbb{R}_+^{n\times\cdots\times n}$ as

$$\mathbf{C}_{i_0,\dots,i_\mathcal{T}} = \sum_{t=1}^{\mathcal{T}} C_{i_{t-1}i_t}, \tag{7}$$

where $C \in \mathbb{R}_+^{n\times n}$ with C_{ij} denoting the cost for a unit mass to move from $i \in E$ to $j \in E$. Note that the network topology is encoded in the optimal transport problem through this cost matrix. In particular, $C_{ij} := \infty$ if the edges i and j do not share a node [13]. Multi-marginal optimal transport problems (2) with a cost tensor that decouples as (7) appear naturally in applications with a sequential structure, such as interpolation or tracking problems [7]. In the network setting, the element $\mathbf{M}_{i_0,i_1,\dots,i_\mathcal{T}}$ in (2) can then be understood as the amount of mass that is transported from time $t = 0$ to $t = \mathcal{T}$ over the edges $i_0, i_1, \dots, i_\mathcal{T}$. With a sequentially decoupling cost (7) the projections (3) are of the form

$$P_t(\mathbf{K} \odot \mathbf{U}) = u_t \odot \varphi_t \odot \hat{\varphi}_t, \tag{8}$$

where

$$\begin{aligned} \hat{\varphi}_t &= K^T \text{diag}(u_{t-1}) K^T \dots \text{diag}(u_1) K^T u_0, \\ \varphi_t &= K \text{diag}(u_{t+1}) K \dots \text{diag}(u_{\mathcal{T}-1}) K u_\mathcal{T}, \end{aligned}$$

and $K = \exp(-C/\epsilon)$ [7], [22]. The projections for the Sinkhorn scheme (6) can thus be efficiently computed by sequential matrix-vector multiplications.

Remark 1: It is worth noting that the entropy regularized multi-marginal optimal transport problem (4) can also be viewed as a so called Schrödinger bridge [9], [22]. The Schrödinger bridge problem has been studied as a framework for robust transport over networks in [13]–[15]. Therein it is argued that the entropy term not only makes the problem computationally feasible, but also induces desirable smoothing to the solution, and thus yields robust transport plans over networks.

B. Steering of heterogeneous ensembles over networks

In this section, we introduce a framework for optimally steering an ensemble of agents with heterogeneous objectives over a network. That is, given an initial distribution of agents, we want to find an optimal steering plan to move all agents to their given destination without violating capacity constraints on the networks edges.

Similarly to the optimal transport framework for networks in Section III-A, define a cost tensor $\mathbf{C} \in \mathbb{R}_+^{n\times\cdots\times n}$ that decouples as in (7), where C_{ij} denotes the cost for an agent on $i \in E$ to move to $j \in E$ in one time step. Moreover, let $\mathbf{M} \in \mathbb{R}_+^{n\times\cdots\times n}$ be a transport plan, where $\mathbf{M}_{i_0\dots i_\mathcal{T}}$ denotes the number of agents that take the path $i_0,\dots,i_\mathcal{T}$. Thus, the total cost of transporting the ensemble is given by $\langle \mathbf{C}, \mathbf{M}\rangle$. Furthermore, let $R \in \mathbb{R}_+^{n\times n}$, where R_{ij} denotes the number of agents that are to be steered from edge $i \in E$ at time 0 to node $j \in E$ at time $\mathcal{T}$. Then the transport plan $\mathbf{M}$ is required to fulfill $P_{0,\mathcal{T}}(\mathbf{M}) = R$, where the bi-marginal projection between the two times t_1 and t_2, with $0 \le t_1 < t_2 \le \mathcal{T}$, is defined as

$$P_{t_1,t_2}(\mathbf{M}) = \sum_{i_0,\dots,i_\mathcal{T}\setminus\{i_{t_1},i_{t_2}\}} \mathbf{M}_{i_0\dots i_\mathcal{T}}.$$

Moreover, define $d \in \mathbb{R}^n$, where d_i denotes the maximal capacity of agents on edge $i \in E$. Then, the transport plan has to satisfy the additional constraint $P_t(\mathbf{M}) \le d$, for $t = 1,\dots,\mathcal{T}-1$. The resulting optimal steering problem reads then

$$\begin{aligned} \underset{\mathbf{M}\in\mathbb{R}_+^{n\times\cdots\times n}}{\text{minimize}} \quad & \langle \mathbf{C}, \mathbf{M}\rangle \\ \text{subject to} \quad & P_{0,\mathcal{T}}(\mathbf{M}) = R \\ & P_t(\mathbf{M}) \le d \quad t = 1,\dots,\mathcal{T}-1. \end{aligned} \tag{10}$$

Problem (10) is similar to a multi-commodity network flow problem, where the number of nonzero elements in R denotes the number of commodities. Note that problem (10) is a multi-marginal optimal transport problem (2) with modified constraints. Based on this connection, we develop a numerical scheme for finding an approximate solution to (10) by solving the entropy regularized problem

$$\begin{aligned} \underset{\mathbf{M}\in\mathbb{R}^{n\times\cdots\times n}}{\text{minimize}} \quad & \langle \mathbf{C}, \mathbf{M}\rangle + \epsilon D(\mathbf{M}) \\ \text{subject to} \quad & P_{0,\mathcal{T}}(\mathbf{M}) = R && \text{(11a)} \\ & P_t(\mathbf{M}) \le d \quad t = 1,\dots,\mathcal{T}-1. && \text{(11b)} \end{aligned}$$

Remark 2: Note that problem (11) may also be seen as a constrained Schrödinger bridge. In the light of Remark 1, the entropy regularization is thus not only needed from a computational perspective, but it also yields robust transport plans over networks.

Similarly to the standard multi-marginal optimal transport problem (cf. (5)), the optimal solution to (11) can be expressed in terms of the dual variables.

Theorem 1: The optimal solution to (11) is given by

$$\mathbf{M} = \mathbf{K} \odot \mathbf{U}, \tag{12}$$

where the two tensors $\mathbf{K}, \mathbf{U} \in \mathbb{R}^{n\times\cdots\times n}$ are given by $\mathbf{K} = \exp(-\mathbf{C}/\epsilon)$, and

$$\mathbf{U}_{i_0,\dots,i_\mathcal{T}} = U_{i_0 i_\mathcal{T}} \cdot (u_1)_{i_1} \cdot \ldots \cdot (u_{\mathcal{T}-1})_{i_{\mathcal{T}-1}}, \tag{13}$$

with $U = \exp(-\Lambda/\epsilon)$ and $u_t = \exp(-\lambda_t/\epsilon)$, for $t = 1,\dots,\mathcal{T}-1$. Here, $\Lambda \in \mathbb{R}^{n\times n}$ and $\lambda_t \in \mathbb{R}_+^n$ are the Lagrange dual variable for constraint (11a) and (11b), respectively. Moreover, the corresponding Lagrangage dual of (11) reads

$$\underset{\substack{\Lambda\in\mathbb{R}^{n\times n},\\ \lambda_1,\dots,\lambda_{\mathcal{T}-1}\in\mathbb{R}_+^n}}{\text{maximize}} \quad -\epsilon\langle \mathbf{K}, \mathbf{U}\rangle - \langle \Lambda, R\rangle - \sum_{t=1}^{\mathcal{T}-1} \langle \lambda_t, d\rangle, \tag{14}$$

where $\mathbf{U}$ depends on the dual variables as in (13).

Proof: See Appendix A. ■

C. Numerical method

We next develop a numerical method for solving problem (11). Since the Sinkhorn iterations (6) can be viewed as a block coordiante ascent in a dual problem of the standard multi-marginal optimal transport problem (4) [23], which resembles the steering problem (11), we follow the same approach. In fact, this leads to an efficient scheme for solving (11) which resembles the Sinkhorn iterations (6).

Proposition 1: Let $\mathbf{K}, \mathbf{U}$ be defined as in Theorem 1, and iteratively perform

$$U \leftarrow U \odot R./P_{0,\mathcal{T}}(\mathbf{K} \odot \mathbf{U}) \tag{15a}$$

$$u_t \leftarrow \min\left(u_t \odot d./P_t(\mathbf{K} \odot \mathbf{U}), \mathbf{1}\right), \tag{15b}$$

for $t = 1,\dots,\mathcal{T}-1$, where the minimization in (15b) is element-wise. Then the iterates (15) converge, and in the limit point the variables $\Lambda = -\epsilon\log(U)$, and $\lambda_t = -\epsilon\log(u_t)$ for $t = 1,\dots,\mathcal{T}-1$, are an optimal solution to (14).

Proof: We derive the scheme as a block coordinate ascent method, which is to maximize the objective with respect to one set of dual variables, while keeping the other dual variables fixed. Applied to (14) this is to iteratively perform the updates

$$\Lambda \leftarrow \underset{\Lambda\in\mathbb{R}^{n\times n}}{\arg\max} \; -\epsilon\langle \mathbf{K}, \mathbf{U}\rangle - \langle \Lambda, R\rangle, \tag{16a}$$

$$\lambda_t \leftarrow \underset{\lambda_t\in\mathbb{R}_+^n}{\arg\max} \; -\epsilon\langle \mathbf{K}, \mathbf{U}\rangle - \langle \lambda_t, d\rangle, \tag{16b}$$

for $t=1,\dots,\mathcal{T}-1$. Since the objective in the unconstrained maximization step (16a) is strictly concave, a necessary and sufficient condition for optimality is that its gradient

$$\exp(-\Lambda/\epsilon)\odot\Bigg(\sum_{i_1,\dots,i_{\mathcal{T}-1}}\mathbf{K}_{i_0\dots i_{\mathcal{T}}}(u_1)_{i_1}\dots(u_{\mathcal{T}-1})_{i_{\mathcal{T}-1}}\Bigg)-R$$

vanishes. This yields the update (15a). Note that the objective in (16b) can be written as

$$\sum_{i_t}\Bigg(-\epsilon e^{-(\lambda_t)_{i_t}/\epsilon}\Big(\sum_{\substack{i_0,\dots,i_{t-1}\\ i_{t+1},\dots,i_{\mathcal{T}}}}\mathbf{K}_{i_0\dots,i_{\mathcal{T}}}U_{i_0 i_{\mathcal{T}}}\prod_{\substack{s=1\\ s\neq t}}^{\mathcal{T}-1}(u_s)_{i_s}\Big)-(\lambda_t)_{i_t}d_{i_t}\Bigg).$$

Thus, the maximization in (16b) can be performed in each element of λ_t individually. If the derivative of the objective in (16b) with respect to $(\lambda_t)_{i_t}$ vanishes for a feasible, i.e., non-negative, point, then this is the global maximizer. Otherwise, the maximizer is $(\lambda_t)_{i_t}=0$. This yields (15b). The convergence of the scheme follows from [24]. ■

Note that given the solution to (14), the optimal solution to (11) can then be recovered as in (12).

Similarly to the standard multi-marginal Sinkhorn iterations, the computational bottleneck of the method in Proposition 1 is to compute the projections $P_{0,\mathcal{T}}(\mathbf{K}\odot\mathbf{U})$ and $P_t(\mathbf{K}\odot\mathbf{U})$ for $t=1,\dots,\mathcal{T}-1$. For a cost tensor that decouples sequentially as in (7), the tensor $\mathbf{K}$ decouples as

$$\mathbf{K}_{i_0\dots i_{\mathcal{T}}}=\prod_{t=1}^{\mathcal{T}}K_{i_{t-1}i_t}, \tag{17}$$

where $K=\exp(-C/\epsilon)$. Given this structure, the bi-marginal projections of $\mathbf{K}\odot\mathbf{U}$ can be computed efficiently as described in the next theorem.

Theorem 2: Consider two $(\mathcal{T}+1)$-mode tensors $\mathbf{K},\mathbf{U}\in\mathbb{R}^{n\times\dots\times n}$, that decouple as in (17) and (13), respectively. Then, for $t=1,\dots,\mathcal{T}-1$, the tensor $\mathbf{K}\odot\mathbf{U}$ has the bi-marginal projections

$$P_{0,t}(\mathbf{K}\odot\mathbf{U})=\Big(\hat\Phi_t\odot\big(U\Phi_t^T\big)\Big)\operatorname{diag}(u_t), \tag{18a}$$

$$P_{t,\mathcal{T}}(\mathbf{K}\odot\mathbf{U})=\operatorname{diag}(u_t)\left(\Big(\hat\Phi_t^T U\Big)\odot\Phi_t\right), \tag{18b}$$

where

$$\hat\Phi_t=K\operatorname{diag}(u_1)K\operatorname{diag}(u_2)\dots\operatorname{diag}(u_{t-1})K,$$
$$\Phi_t=K\operatorname{diag}(u_{t+1})K\operatorname{diag}(u_{t+2})\dots\operatorname{diag}(u_{\mathcal{T}-1})K.$$

Moreover,

$$P_{0,\mathcal{T}}(\mathbf{K}\odot\mathbf{U})=\hat\Phi_{\mathcal{T}}\odot U=U\odot\Phi_0. \tag{19}$$

Proof: See Appendix B. ■

Note that expression (19) can be used for computing the projections in (15a). Moreover, the results from Theorem 2 can be used to compute the projections $P_t(\mathbf{K}\odot\mathbf{U})$ for $t=1,\dots,\mathcal{T}-1$, which are required for (15b), as follows.

Corollary 1: Let $\mathbf{K},\mathbf{U}$ be defined as in Theorem 2. For $t=1,\dots,\mathcal{T}-1$, the marginals of the tensor $\mathbf{K}\odot\mathbf{U}$ are then of the form

$$P_t(\mathbf{K}\odot\mathbf{U})=u_t\odot v_t, \tag{20}$$

where $v_t=\Big(\hat\Phi_t^T\odot\big(\Phi_t U^T\big)\Big)\mathbf{1}=\Big(\big(\hat\Phi_t^T U\big)\odot\Phi_t\Big)\mathbf{1}$.

Proof: The claim follows directly from Theorem 2, since $P_t(\mathbf{K}\odot\mathbf{U})=P_{0,t}(\mathbf{K}\odot\mathbf{U})^T\mathbf{1}=P_{t,\mathcal{T}}(\mathbf{K}\odot\mathbf{U})\mathbf{1}$. ■

Algorithm 1 Block coordinate ascent for problem (14).

Initialize $\mathbf{U}$
while Not converged **do**
  Compute $P_{0,\mathcal{T}}(\mathbf{K}\odot\mathbf{U})$ as in (19).
  $U\leftarrow U\odot R./P_{0,\mathcal{T}}(\mathbf{K}\odot\mathbf{U})$.
  Update $\mathbf{U}$ as in (13).
  for $t=1,\dots,\mathcal{T}-1$ **do**
    Compute $P_t(\mathbf{K}\odot\mathbf{U})$ as in (20).
    $u_t\leftarrow\min(u_t\odot d./P_t(\mathbf{K}\odot\mathbf{U}),\mathbf{1})$.
    Update $\mathbf{U}$ as in (13).
  end for
end while
return $\mathbf{U}$

With the expressions for the projections in Theorem 2 and Corollary 1, we can efficiently perform the updates in (15) in order to solve the steering problem (11). The full method is summarized in Algorithm 1. Recall that upon convergence the optimal solution to the primal problem (11) can be recovered as detailed in Theorem 1.

It is worth noting that the method in Algorithm 1 has a similar sturucture as the standard Sinkhorn iterations (6). Compared to the standard multi-marginal optimal transport problem, the number of optimization variables in the dual (14) is increased due to the matrix constraint in (10), and the computation of the projections becomes more expensive, as detailed in the following remark.

Remark 3: Computing the projections in Theorem 2, requires matrix-matrix multiplications and the complexity is thus of a factor n larger as compared to the standard Sinkhorn case (cf. (8)). More precisely, computing one of the projections in Theorem 2 is of complexity $\mathcal{O}(\mathcal{T}n^3)$, whereas the complexity of computing a projection in (8) is only $\mathcal{O}(\mathcal{T}n^2)$. Moreover, the computations of the projections in Algorithm 1 can be arranged so that the outer loop is computed in $\mathcal{O}(\mathcal{T}n^3)$. By extending the graph over time, it resembles a multi-commodity network with $n(\mathcal{T}-1)$ edges and at most n^2 commodities. Standard frameworks for multi-commodity flow problems thus require solving a linear program with $n^3(\mathcal{T}-1)$ variables [12].

IV. Traffic planning simulation

In order to illustrate the method for steering an ensemble with heterogeneous objectives, which was introduced in Section III, we consider a traffic planning example. Figure 1 illustrates a network of streets, inspired by north Atlanta, where the wide edges symbolize highways and the thin edges symbolize local streets. We consider an ensemble of agents starting in two point of departures, e.g., residential areas, and travelling to two points of destinations, e.g., offices and university campus near the city center.

We describe the street network in Figure 1 by a directed graph with edges E, that is, each edge in the street network

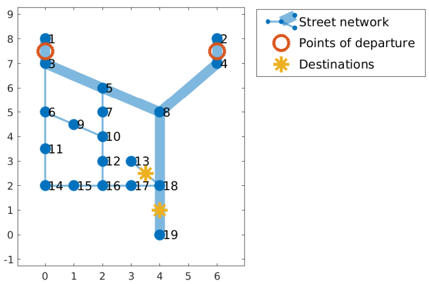

Fig. 1. Street network for the example in Section IV.

is represented by two directed edges, and $n = |E| = 42$. Denote the set of edges on the highway as $E_H \subset E$, the set of departure edges, as $E_{in} \subset E$, and the set of destinations, as $E_{out} \subset E$. Furthermore, let l_i be the Euclidean length of edge $i \in E$, as seen in Figure 1.

In order to apply the framework in Section III, let $\mu_t \in \mathbb{R}^n$ denote the distribution of agents over E at time t. To define the cost matrix $C \in \mathbb{R}_+^{n\times n}$ in (7), we assume that for agents it is desirable to spend as little time as possible travelling, i.e., to reach their destination fast, or otherwise wait with their departure. The cost for an agent to move to another edge, or stay on an edge in the network, is dependent on the length of the edge. Moreover, in one time step, agents should only be allowed to move to neighbouring edges. Thus, we define the cost for an agent to move from $i \in E$ to $j \in E$ as

$$C_{ij} = \begin{cases} (l_i + l_j)/2, & \text{if } j \in N(i) \\ 0.06 & \text{if } j = i \in E_{in} \\ 0, & \text{if } j = i \in E_{out} \\ l_i, & \text{if } j = i \notin E_{in} \cup E_{out} \\ 1000, & \text{oherwise,} \end{cases}$$

where $N(i)$ denotes the set of edges that start in the end node of edge i. Moreover, the capacity on the edges is defined by a vector $d \in \mathbb{R}^n$, where

$$d_i = \begin{cases} 14 l_i, & \text{if } i \in E_H \\ 7 l_i, & \text{if } i \in E \setminus (E_H \cup E_{in} \cup E_{out}) \\ 1000, & \text{if } i \in E_{in} \cup E_{out}. \end{cases}$$

Utilizing Algorithm 1 with regularization parameter $\epsilon = 0.1$, we solve the steering problem (11) for two scenarios. In particular, we consider two different ensemble flows defined by the matrix R in (11a), which defines the number of agents moving between each pair of nodes from time 0 to $\mathcal{T} = 20$.

In a first scenario, an ensemble of 500 agents starts on edge $(1, 3)$, where half of the agents are steered to each of the two edges of destination. The result of solving the steering problem (11) can be seen in Figure 2, where the width of the edges corresponds to the log-scaled number of agents on it. The optimal scheduling plan sends the majority of agents to travel via the highway, only a small group of agents is steered over the local roads. Moreover, one can see that the traffic is spread out over time. Thus, some agents travel earlier, and reach their destination within five time steps, while others stay in the point of departure and are scheduled to travel later.

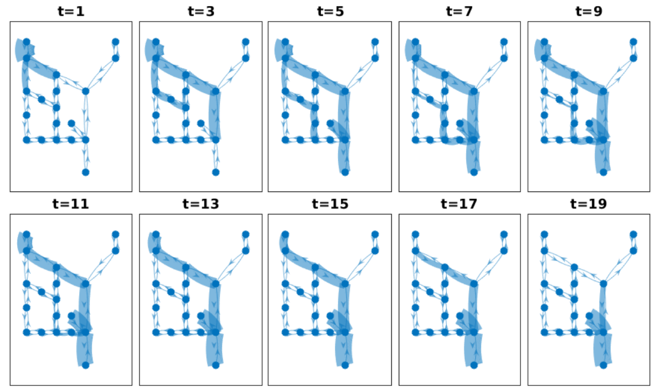

Fig. 2. Optimal traffic flow from edge $(1, 3)$ to the two destinations.

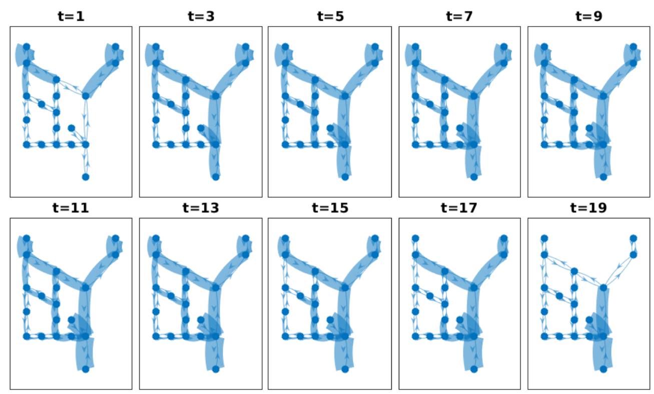

Fig. 3. Optimal traffic flow during rush hour.

A second scenario models the case of heavy traffic, e.g., during rush hour. Here, an additional 500 agents enter the network on edge $(2, 4)$, and half of the additional agents are steered to each of the two edges of destination. The result is shown in Figure 3. In this scenario the highway is used to full capacity, and thus a large group of agents from the previous scenario is instead steered over the local roads. Moreover, more agents are scheduled for departure later as compared to the scenario with lighter traffic.

In order to illustrate the method we have selected only two points of departure and two points of destination, i.e., four commodities. However, note that the computational complexity of each iteration in Algorithm 1 is $\mathcal{O}(\mathcal{T} n^3)$, which is independent of the number of points of departure and points of destination.

V. Conclusion and Future Work

We formulated and studied an optimal steering problem for ensembles with heterogeneous objectives. We also presented an efficient algorithm for solving it, and applied it to a traffic planning example. Seen as a multi-commodity flow problem, note that, independently of the number of commodities, the complexity for each iteration is $\mathcal{O}(\mathcal{T} n^3)$, and allows for n^2 commodities, i.e., one commodity for *every* pair of edges.

Appendix

A. *Proof of Theorem 1*

With the Lagrange multipliers $\Lambda \in \mathbb{R}^{n\times n}$ and $\lambda_t \in \mathbb{R}_+^n$, for $t = 1, \ldots, \mathcal{T} - 1$, as in the theorem, define the Lagrangian as

$$
L(\mathbf{M}, \Lambda, \lambda_1, \dots \lambda_{\mathcal{T}-1}) = \langle \mathbf{C}, \mathbf{M} \rangle + \epsilon D(\mathbf{M}) \\
+ \langle \Lambda, P_{0,\mathcal{T}}(\mathbf{M}) - R \rangle + \sum_{t=1}^{\mathcal{T}-1} \langle \lambda_t, P_t(\mathbf{M}) - d \rangle. \quad (21)
$$

The minimum of (21) with respect to $\mathbf{M}_{i_0 \dots i_\mathcal{T}}$ is achieved when its derivative vanishes, i.e., when

$$
\mathbf{C}_{i_0 \dots i_\mathcal{T}} + \epsilon \log\left(\mathbf{M}_{i_0 \dots i_\mathcal{T}}\right) + \Lambda_{i_0 i_\mathcal{T}} + \sum_{t=1}^{\mathcal{T}-1} \lambda_t = 0.
$$

Thus, the optimal transport tensor is of the form $\mathbf{M} = \mathbf{K} \odot \mathbf{U}$ with $\mathbf{K}$ and $\mathbf{U}$ as defined in the theorem. Note that the entropy term $D(\mathbf{K} \odot \mathbf{U})$ reads

$$
\sum_{i_0,\dots,i_\mathcal{T}} \Bigg(\mathbf{K}_{i_0 \dots i_\mathcal{T}} \mathbf{U}_{i_0 \dots i_\mathcal{T}} \frac{1}{\epsilon} \Big(- \mathbf{C}_{i_0 \dots i_\mathcal{T}} - \Lambda_{i_0 i_\mathcal{T}} - \sum_{t=1}^{\mathcal{T}-1} \lambda_t \Big) \\
- \mathbf{K}_{i_0 \dots i_\mathcal{T}} \mathbf{U}_{i_0 \dots i_\mathcal{T}} + 1 \Bigg) \\
= - \frac{1}{\epsilon} \langle \mathbf{K} \odot \mathbf{U}, \mathbf{C} \rangle - \frac{1}{\epsilon} \langle \Lambda, P_{0,\mathcal{T}}(\mathbf{K} \odot \mathbf{U}) \rangle \\
- \frac{1}{\epsilon} \sum_{t=1}^{\mathcal{T}-1} \langle \lambda_t, P_t(\mathbf{K} \odot \mathbf{U}) \rangle - \langle \mathbf{K}, \mathbf{U} \rangle + n^{\mathcal{T}+1}.
$$

Thus, plugging $\mathbf{M} = \mathbf{K} \odot \mathbf{U}$ into the Lagrangian (21) yields

$$
-\epsilon \langle \mathbf{K}, \mathbf{U} \rangle - \langle \Lambda, R \rangle - \sum_{t=1}^{\mathcal{T}-1} \langle \lambda_t, d \rangle. \quad (22)
$$

The dual to (10) is to maximize (22) with respect to $\Lambda \in \mathbb{R}^{n \times n}$ and $\lambda_t \in \mathbb{R}_+^n$, for $t = 1, \dots, \mathcal{T} - 1$.

B. Proof of Theorem 2

Note that, for $0 < t_1 < t_2 < \mathcal{T}$ it holds

$$
\sum_{i_{t_1},\dots,i_{t_2}} \Big(\prod_{s=t_1}^{t_2} K_{i_{s-1} i_s} \Big) \Big(\prod_{s=t_1}^{t_2-1} (u_s)_{i_s} \Big) \\
= \left(K \mathrm{diag}(u_{t_1}) K \dots K \mathrm{diag}(u_{t_2}) K \right)_{i_{t_1-1} i_{t_2+1}}.
$$

Thus, (18a) follows as

$$
P_{0,t}\left(\mathbf{K} \odot \mathbf{U}\right) = \sum_{\substack{i_1,\dots,i_{t-1} \\ i_{t+1},\dots i_\mathcal{T}}} \Big(\prod_{s=1}^{\mathcal{T}} K_{i_{s-1} i_s} \Big) \Big(\prod_{s=1}^{\mathcal{T}-1} (u_s)_{i_s} \Big) U_{i_0 i_\mathcal{T}} \\
= (u_t)_{i_t} (\hat{\Phi}_t)_{i_0 i_t} \sum_{i_\mathcal{T}} (\Phi_t)_{i_t i_\mathcal{T}} U_{i_0 i_\mathcal{T}} \\
= (u_t)_{i_t} (\hat{\Phi}_t)_{i_0 i_t} \left(U \Phi_t^T \right)_{i_0 i_t}.
$$

Similarly, (18b) is given as

$$
P_{t,\mathcal{T}}\left(\mathbf{K} \odot \mathbf{U}\right) = (u_t)_{i_t} \Big(\sum_{i_0} (\hat{\Phi}_t)_{i_0 i_t} U_{i_0 i_\mathcal{T}} \Big) (\Phi_t)_{i_t i_\mathcal{T}} \\
= (u_t)_{i_t} (\hat{\Phi}_t^T U)_{i_t i_\mathcal{T}} (\Phi_t)_{i_t i_\mathcal{T}}.
$$

Finally, the expression for the projection (19) follows as

$$
P_{0,\mathcal{T}}\left(\mathbf{K} \odot \mathbf{U}\right) = \sum_{i_1,\dots,i_{\mathcal{T}-1}} \Big(\prod_{s=1}^{\mathcal{T}} K_{i_{s-1} i_s} \Big) \Big(\prod_{s=1}^{\mathcal{T}-1} (u_s)_{i_s} \Big) U_{i_0 i_\mathcal{T}} \\
= U_{i_0 i_\mathcal{T}} \left(K \mathrm{diag}(u_1) K \dots K \mathrm{diag}(u_{\mathcal{T}-1}) K \right)_{i_0 i_\mathcal{T}}.
$$

REFERENCES

[1] D. J. Watts and S. H. Strogatz, "Collective dynamics of 'small-world' networks," *nature*, vol. 393, no. 6684, p. 440, 1998.

[2] X. Chen, "Structure theory for ensemble controllability, observability, and duality," *Mathematics of Control, Signals, and Systems*, vol. 31, no. 2, p. 7, 2019.

[3] S. Zeng, S. Waldherr, C. Ebenbauer, and F. Allgöwer, "Ensemble observability of linear systems." *IEEE Transactions on Automatic Control*, vol. 61, no. 6, pp. 1452–1465, 2016.

[4] Y. Chen, T. Georgiou, and M. Pavon, "Optimal Steering of a Linear Stochastic System to a Final Probability Distribution, Part I," *IEEE Transactions on Automatic Control*, vol. 61, no. 5, pp. 1158–1169, 2016.

[5] R. W. Brockett, "Notes on the control of the Liouville equation," in *Control of partial differential equations*. Springer, 2012, pp. 101–129.

[6] G. Dirr, U. Helmke, and M. Schönlein, "Controlling mean and variance in ensembles of linear systems," *IFAC-PapersOnLine*, vol. 49, no. 18, pp. 1018–1023, 2016.

[7] F. Elvander, I. Haasler, A. Jakobsson, and J. Karlsson, "Multi-Marginal Optimal Mass Transport using Partial Information with Applications in Robust Localization and Sensor Fusion," *Signal Processing*, 2020.

[8] Y. Chen and J. Karlsson, "State tracking of linear ensembles via optimal mass transport," *IEEE Control Systems Letters*, vol. 2, no. 2, pp. 260–265, 2018.

[9] I. Haasler, A. Ringh, Y. Chen, and J. Karlsson, "Estimating ensemble flows on a hidden Markov chain," in *2019 IEEE 58th Conference on Decision and Control (CDC)*. IEEE, 2019, pp. 1331–1338.

[10] D. Bertsimas and S. S. Patterson, "The traffic flow management rerouting problem in air traffic control: A dynamic network flow approach," *Transportation Science*, vol. 34, no. 3, pp. 239–255, 2000.

[11] J. L. Kennington, "A survey of linear cost multicommodity network flows," *Operations Research*, vol. 26, no. 2, pp. 209–236, 1978.

[12] I.-L. Wang, "Multicommodity network flows:A survey, Part I: Applications and Formulations," *International Journal of Operations Research*, vol. 15, no. 4, pp. 145–153, 2018.

[13] Y. Chen, T. Georgiou, M. Pavon, and A. Tannenbaum, "Robust transport over networks," *IEEE Transactions on Automatic Control*, vol. 62, no. 9, pp. 4675–4682, 2016.

[14] Y. Chen, T. T. Georgiou, M. Pavon, and A. Tannenbaum, "Efficient robust routing for single commodity network flows," *IEEE Transactions on Automatic Control*, vol. 63, no. 7, pp. 2287–2294, 2017.

[15] Y. Chen, T. Georgiou, M. Pavon, and A. Tannenbaum, "Relaxed Schrödinger bridges and robust network routing," *IEEE Transactions on Control of Network Systems*, 2019.

[16] J. A. Tomlin, "Minimum-cost multicommodity network flows," *Operations Research*, vol. 14, no. 1, pp. 45–51, 1966.

[17] C. Villani, *Optimal transport: Old and new*. Berlin Heidelberg: Springer, 2008.

[18] J.-D. Benamou, G. Carlier, M. Cuturi, L. Nenna, and G. Peyré, "Iterative Bregman Projections for Regularized Transportation Problems," *SIAM Journal on Scientific Computing*, vol. 37, no. 2, p. A1111–A1138, 2015.

[19] B. Pass, "Multi-marginal optimal transport: theory and applications," *ESAIM: Mathematical Modelling and Numerical Analysis*, vol. 49, no. 6, pp. 1771–1790, 2015.

[20] L. Nenna, "Numerical methods for multi-marginal optimal transportation," Ph.D. dissertation, 2016.

[21] M. Cuturi, "Sinkhorn distances: Lightspeed computation of optimal transport," in *Advances in Neural Information Processing Systems (NIPS)*. Neural Information Processing Systems Foundation, 2013, pp. 2292–2300.

[22] I. Haasler, A. Ringh, Y. Chen, and J. Karlsson, "Multi-marginal optimal transport and Schrödinger bridges on Trees," *Under Review*, 2020.

[23] J. Karlsson and A. Ringh, "Generalized Sinkhorn iterations for regularizing inverse problems using optimal mass transport," *SIAM Journal on Imaging Sciences 10(4), 1935–1962*, 2017.

[24] P. Tseng, "Dual ascent methods for problems with strictly convex costs and linear constraints: A unified approach," *SIAM Journal on Control and Optimization*, vol. 28, no. 1, pp. 214–242, 1990.

[25] H. Bauschke and A. Lewis, "Dykstras algorithm with Bregman projections: A convergence proof," *Optimization*, vol. 48, no. 4, pp. 409–427, 2000.