# MULTI-MARGINAL OPTIMAL TRANSPORT WITH A TREE-STRUCTURED COST AND THE SCHRÖDINGER BRIDGE PROBLEM* 

ISABEL HAASLER ${ }^{\dagger}$, AXEL RINGH ${ }^{\ddagger}$, YONGXIN CHEN§, AND JOHAN KARLSSON ${ }^{\dagger}$


#### Abstract

The optimal transport problem has recently developed into a powerful framework for various applications in estimation and control. Many of the recent advances in the theory and application of optimal transport are based on regularizing the problem with an entropy term, which connects it to the Schrödinger bridge problem and thus to stochastic optimal control. Moreover, the entropy regularization makes the otherwise computationally demanding optimal transport problem feasible even for large scale settings. This has led to an accelerated development of optimal transport based methods in a broad range of fields. Many of these applications have an underlying graph structure, for instance information fusion and tracking problems can be described by trees. In this work we consider multi-marginal optimal transport problems with a cost function that decouples according to a tree structure. The entropy regularized multi-marginal optimal transport problem can be viewed as a generalization of the Schrödinger bridge problem with the same tree-structure, and by utilizing these connections we extend the computational methods for the classical optimal transport problem in order to solve structured multi-marginal optimal transport problems in an efficient manner. In particular, the algorithm requires only matrix-vector multiplications of relatively small dimensions. We show that the multi-marginal regularization introduces less diffusion, compared to the commonly used pairwise regularization, and is therefore more suitable for many applications. Numerical examples illustrate this, and we finally apply the proposed framework for tracking of an ensemble of indistinguishable agents.


Key words. Multi-marginal optimal transport, Schrödinger bridge, Hidden Markov chain, graph signal processing, ensemble estimation.

1. Introduction. An optimal transport problem is to find a transport plan that minimizes the cost of moving the mass of one distribution to another distribution [61]. Historically this problem has been important in economics and operations research, but as a result of recent progress in the area it has become a popular tool in a wide range of fields such as control theory $[1,6,15,35,50,63]$, signal processing [30, 42], computer vision $[27,58]$, and machine learning $[2,4,51]$.

An extension to the standard optimal transport framework is multi-marginal optimal transport [52], which seeks a transport plan between not only two, but several distributions. Early works on multi-marginal optimal transport include [33,54,55]. In this work we consider multi-marginal optimal transport problems with cost functions that decouple according to a tree structure. We refer to such a problem as a treestructured multi-marginal optimal transport problem. This should not be confused with optimal transport problems on graphs as in $[16,18,19]$, where the distributions are defined over the nodes of the graphs. Tree-structured cost functions generalize many structures that are commonly used in applications of optimal transport. For example, a path tree naturally appears in tracking and interpolation applications [9, 17, 58]. Similarly, star trees are used in barycenter problems, which occur for instance in

[^0]information fusion applications [23,29].
The optimal transport problem can be formulated as a linear program. However, in many practical applications the problem is too large to be solved directly. For the bi-marginal case, these computational limitations have recently been alleviated by regularizing the problem with an entropy term [22]. The optimal solution to the regularized optimal transport problem can then be expressed in terms of dual variables, which can be efficiently found by an iterative scheme, called Sinkhorn iterations [57]. For the multi-marginal case, although the Sinkhorn iterations can be generalized in a straight-forward fashion [7], the complexity of the scheme increases dramatically with the number of marginals [46]. Thus Sinkhorn iterations alone are not sufficient to address many multi-marginal problems. However, in specific settings, structures in the cost function can be exploited in order to derive computationally feasible methods, e.g., for Euler flows [7] and in tracking and information fusion applications [30].

Unregularized multi-marginal optimal transport problems with transport cost that decouples according to a tree structure can be equivalently formulated as a sum of coupled pairwise optimal transport problems, as for instance, tracking and interpolation problems [9,17, 29], and barycenter problems [ $3,7,23,58]$. Typically these problems are solved using pairwise regularization (see, e.g., $[7$, Sec. 3.2] and $[8,43,45]$ ). However, we have empirically observed in some applications that the multi-marginal formulation yields favourable solutions compared to a corresponding pairwise optimal transport estimate [30]. One main contribution of this work is to develop a framework for solving tree-structured optimal transport problems with a multi-marginal regularization. For these problems, we show that the Sinkhorn algorithm can be performed in an efficient way, requiring only successive matrix-vector multiplications of relatively small size compared to that of the original multi-marginal problem. Thus we extend the computational results from [30] to general trees.

The entropy regularized formulation of optimal transport is connected to another classical topic, the Schrödinger bridge problem [13, 14, 44, 48, 49]. Schrödinger was interested in determining the most likely evolution of a particle cloud observed at two time instances, where the particle dynamics have deviated from the expected Brownian motion [56]. Schrödinger showed that this particle evolution can be characterized as the one out of all the theoretically possible ones, that minimizes the relative entropy to the Wiener measure. This optimal solution may be found by solving a so-called Schrödinger system, which turns out to be tightly connected to the Sinkhorn iterations for the entropy regularized optimal transport problem [13, 48]. This framework has been used in robust and stochastic control problems [15,62]. A version of the Schrödinger bridge problem, that is discrete in both time and space, can be formulated by modeling the evolutions of a number of particles as a Markov chain $[34,53]$. In the infinite particle limit, the maximum likelihood solution of the Markov process can then be approximated by solving a relative entropy problem [16]. An analogous approach has recently been used to develop a framework for modelling ensemble flows on Hidden Markov chains [37]. Other optimal transport based state estimation problems for a continuum of agents and in continuous time have been considered in $[11,12,17]$. For further work on ensemble controllability and observability see, e.g., [10, 66].

Based on [37], we extend the discrete Schrödinger bridge in [53] to trees. In particular, we derive a maximum likelihood estimate for Markov processes defined on the edges of a rooted directed tree. This leads to our second main result: Interestingly, it turns out that the solution to the entropy regularized tree-structured multi-marginal optimal transport problem corresponds to the solution of the generalized Schrödinger
bridge with the same tree-structure. This generalizes the established equivalence of the classical bi-marginal entropy regularized optimal transport problem and the Schrödinger bridge problem [48]. Moreover, this also gives an additional motivation for using multi-marginal optimal transport instead of the often used pairwise optimal transport for tree structured problems.

This paper is organized as follows: Section 2 is an introduction to the problems of optimal transport and Schrödinger bridges. The main results are presented in Section 3 and 4. In Section 3 we define the tree-structured multi-marginal optimal transport problem, and provide an algorithm for efficiently solving its entropy regularized version. In Section 4 we generalize the Schrödinger bridge problem to trees, and show that it is equivalent to an entropy regularized multi-marginal optimal transport problem on the same tree. Section 5 is on the discrepancy between the multi-marginal and pairwise optimal transport formulations with tree structure. A larger numerical simulation is detailed in Section 6, where the proposed framework is used to estimate ensemble flows from aggregate and incomplete measurements. Most of the proofs can be found in the Appendix, however due to space limitations two of the proofs are deferred to the supplementary material.
2. Background. In this section we summarize some background material on optimal transport and Schrödinger bridges. This is also used to set up the notation. To this end, first note that throughout we let $\exp (\cdot), \log (\cdot), \odot$, and ./ denote the elementwise exponential, logarithm, multiplication, and division of vectors, matrices, and tensors, respectively. Moreover, $\otimes$ denotes the outer product. By 1 we denote a column vector of ones, the size of which will be clear from the context.
2.1. Optimal transport. In this work we consider the discrete optimal transport problem. For its continuous counterpart see, e.g., [61]. Let the vectors $\mu_{1}, \mu_{2} \in$ $\mathbb{R}_{+}^{n}$ describe two nonnegative distributions with equal mass. The optimal transport problem is to find a mapping that moves the mass from $\mu_{1}$ to $\mu_{2}$, while minimizing the total transport cost. Here the transport cost is defined in terms of a underlying cost matrix $C \in \mathbb{R}_{+}^{n \times n}$, where $C_{i_{1}, i_{2}}$ denotes the cost of moving a unit mass from location $i_{1}$ to $i_{2}$. Analogously, define a transport plan $M \in \mathbb{R}_{+}^{n \times n}$, where $M_{i_{1}, i_{2}}$ describes the amount of mass that is moved from location $i_{1}$ to $i_{2}$. An optimal transport plan from $\mu_{1}$ to $\mu_{2}$ is then a minimizing solution of

$$
\begin{align*}
T\left(\mu_{1}, \mu_{2}\right):= & \underset{M \in \mathbb{R}_{+}^{n \times n}}{\operatorname{minimize}} \operatorname{trace}\left(C^{T} M\right)  \tag{2.1}\\
& \text { subject to } M \mathbf{1}=\mu_{1}, \quad M^{T} \mathbf{1}=\mu_{2}
\end{align*}
$$

Multi-marginal optimal transport extends the concept of the classical optimal transport problem (2.1) to the setting with a set of marginals $\mu_{1}, \ldots, \mu_{J}$, where $J \geq 2$ [7,30,52]. In this setting, the transport cost and transport plan in (2.1) are described by J-mode tensors $\mathbf{C}, \mathbf{M} \in \mathbb{R}_{+}^{n \times n \cdots \times n}$. For a tuple $\left(i_{1}, \ldots, i_{J}\right)$, the value $\mathbf{C}_{i_{1}, \ldots, i_{J}}$ denotes the transport cost for a unit mass corresponding to the tuple, and similarly $\mathbf{M}_{i_{1}, \ldots, i_{J}}$ represents the amount of transported mass associated with this tuple. The multi-marginal optimal transport problem then reads

$$
\begin{align*}
& \underset{\mathbf{M} \in \mathbb{R}_{+}^{n \times \ldots \times n}}{\operatorname{mimize}}\langle\mathbf{C}, \mathbf{M}\rangle  \tag{2.2}\\
& \text { subject to } P_{j}(\mathbf{M})=\mu_{j}, \text { for } j \in \Gamma
\end{align*}
$$

where $\langle\mathbf{C}, \mathbf{M}\rangle=\sum_{i_{1}, \ldots, i_{J}} \mathbf{C}_{i_{1}, \ldots, i_{J}} \mathbf{M}_{i_{1}, \ldots, i_{J}}$ is the standard inner product, the projec-
tion on the $j$-th marginal of $\mathbf{M}$, denoted by $P_{j}(\mathbf{M})$, is defined as

$$
\begin{equation*}
P_{j}(\mathbf{M})_{i_{j}}=\sum_{i_{1}, \ldots, i_{j-1}, i_{j+1}, \ldots, i_{J}} \mathbf{M}_{i_{1}, \ldots, i_{j-1}, i_{j}, i_{j+1}, \ldots, i_{J}}, \tag{2.3}
\end{equation*}
$$

and $\Gamma$ denotes an index set corresponding to the given marginals. In the original multi-marginal optimal transport formulation, constraints are typically given on all marginals, i.e., for the index set $\Gamma=\{1,2, \ldots, J\}$. However, in this work we consider the case, where constraints typically are only imposed on a subset of marginals, i.e., $\Gamma \subset\{1,2, \ldots, J\}$. Note that the standard bi-marginal optimal transport problem (2.1) is a special case of the multi-marginal optimal transport problem (2.2), where $J=2$ and $\Gamma=\{1,2\}$.
2.2. Entropy regularized optimal transport. Although linear, the number of variables in the multi-marginal optimal transport problem (2.2) is often too large to be solved directly. A popular approach for the bi-marginal setting to bypass the size of the problem has been to add a regularizing entropy term to the objective. In theory the same approach can be used also for the multi-marginal case.

Definition 2.1 ([36, Ch. 4]). Let p and $q$ be two nonnegative vectors, matrices or tensors of the same dimension. The normalized Kullback-Leibler (KL) divergence of $p$ from $q$ is defined as $H(p \mid q):=\sum_{i}\left(p_{i} \log \left(p_{i} / q_{i}\right)-p_{i}+q_{i}\right)$, where $0 \log 0$ is defined to be 0 . Similarly, define $H(p):=H(p \mid \mathbf{1})=\sum_{i}\left(p_{i} \log \left(p_{i}\right)-p_{i}+1\right)$, which is effectively the negative of the entropy of $p$.
Note that $H(p \mid q)$ is jointly convex over $p, q$. For a detailed description of the KL divergence see, e.g., [20,21]. The entropy regularized multi-marginal optimal transport problem is the convex problem

$$
\begin{align*}
& \underset{\mathbf{M} \in \mathbb{R}_{+}^{n \times \ldots \times n} \times{ }_{n}}{\operatorname{minimize}}\langle\mathbf{C}, \mathbf{M}\rangle+\epsilon H(\mathbf{M})  \tag{2.4}\\
& \text { subject to } P_{j}(\mathbf{M})=\mu_{j}, \text { for } j \in \Gamma,
\end{align*}
$$

where $\epsilon>0$ is a regularization parameter. For the bi-marginal case (2.1), where the cost and mass transport tensors are matrices, the entropy regularized problem reads

$$
\begin{align*}
T_{\epsilon}\left(\mu_{1}, \mu_{2}\right): & =\underset{M \in \mathbb{R}_{+}^{n \times n}}{\operatorname{minimize}} \operatorname{trace}\left(C^{T} M\right)+\epsilon H(M)  \tag{2.5}\\
& \text { subject to } M \mathbf{1}=\mu_{1}, \quad M^{T} \mathbf{1}=\mu_{2} .
\end{align*}
$$

It is well established that the entropy regularized bi-marginal optimal transport problem is connected to the Schrödinger bridge problem [13,14,48,49], which is introduced in Section 2.3. More importantly from a computational perspective, the introduction of the entropy term in problem (2.5) allows for expressing the optimal solution $M$ in terms of the Lagrange dual variables, which may be computed by Sinkhorn iterations [22]. This procedure can be generalized to the setting of multi-marginal optimal transport [7].

In particular, for the multi-marginal entropy regularized optimal transport problem (2.4) it can be shown that the optimal solution is of the form [30]

$$
\begin{equation*}
\mathbf{M}=\mathbf{K} \odot \mathbf{U} \tag{2.6}
\end{equation*}
$$

where $\mathbf{K}=\exp (-\mathbf{C} / \epsilon)$ and where $\mathbf{U}$ can be decomposed as

$$
\mathbf{U}=u_{1} \otimes u_{2} \otimes \cdots \otimes u_{J} \quad \text { with } u_{j}= \begin{cases}\exp \left(\lambda_{j} / \epsilon\right), & \text { if } j \in \Gamma  \tag{2.7}\\ 1, & \text { else. }\end{cases}
$$

Here $\lambda_{j} \in \mathbb{R}^{n}$ for $j \in \Gamma$ are optimal dual variables in the dual problem of (2.4):

$$
\begin{equation*}
\underset{\lambda_{j} \in \mathbb{R}^{n}, j \in \Gamma}{\operatorname{maximize}}-\epsilon\langle\mathbf{K}, \mathbf{U}\rangle+\sum_{j \in \Gamma} \lambda_{j}^{T} \mu_{j}, \tag{2.8}
\end{equation*}
$$

where $\mathbf{U}$ depends on the variables $\lambda_{j}$ as specified by (2.7). Note that the dual variable $\lambda_{j}$ corresponds to the constraint on the $j$-th marginal. For details the reader is referred to, e.g., [7,30].

The Sinkhorn scheme for finding $\mathbf{U}$ in (2.7), is to iteratively update $u_{j}$ as

$$
\begin{equation*}
u_{j} \leftarrow u_{j} \odot \mu_{j} . / P_{j}(\mathbf{K} \odot \mathbf{U}) \tag{2.9}
\end{equation*}
$$

for all $j \in \Gamma$. This scheme may for instance be derived as Bregman projections [7] or a block coordinate ascend in the dual (2.8), [30,41,60]. As a result, global convergence of the Sinkhorn scheme (2.9) is guaranteed [5]. For the sake of completeness, we also provide a result on the linear convergence rate in our presentation (see Theorem 3.5). We also note that (2.9) reduces to standard Sinkhorn iterations,

$$
\begin{equation*}
u_{1} \leftarrow \mu_{1} \cdot /\left(K u_{2}\right), \quad u_{2} \leftarrow \mu_{1} \cdot /\left(K^{T} u_{1}\right) \tag{2.10}
\end{equation*}
$$

for the two-marginal case (2.5). The iterations (2.10) converge linearly to an optimal solution $u_{1}, u_{2}$, which is unique up to multiplication/division with a constant [13,32].

The computational bottleneck of the Sinkhorn iterations (2.9) is computing the projections $P_{j}(\mathbf{M})$, for $j \in \Gamma$, which in general scales exponentially in $J$. In fact, even storing the tensor $\mathbf{M}$ is a challenge as it consists of $n^{J}$ elements. However, in many cases of interest, structures in the cost tensors can be exploited to make the computation of the projections feasible. In [30] this is shown on the example of cost functions that decouple sequentially, centrally, or a combination of both. Computing the projections requires then only repeated matrix vector multiplications. In this work, we show that these efficient methods can be generalized to the setting where the cost tensor decouples according to a tree structure, that is, when the marginals of the optimal transport problem are associated with the nodes of a tree, and cost matrices are defined on its edges (see Section 3).
2.3. Schrödinger bridge problem. The Schrödinger bridge problem is to determine the most likely evolution of a particle cloud observed at two time instances [56]. In the case that the second observation cannot be explained as a Brownian motion of the initially observed particle cloud, Schrödinger aimed to find the most likely particle evolution connecting, hence bridging, the two distributions.

A useful mathematical framework to solve this problem, however not yet developed in Schrödingers time, is the theory of large deviations, which studies so called rare events, meaning deviations from the law of large numbers [24, 25]. The probability of these rare events approaches zero, as the number of trials goes to infinity, and large deviation theory analyzes the rate of this decay, which can often be expressed as the exponential of a so called rate function $[25,28]$.

For a large deviation interpretation of the Schrödinger bridge [31, Sec. II.1.3], the particle evolutions are modelled as independent identically distributed random variables on path space, and the Schrödinger bridge is the probability measure $\mathcal{P}$ on path space that is most likely to describe the rare event of observing the two particle distributions. Let $\mathcal{W}$ be the Wiener measure, which corresponds to the probability law of the Brownian motion. Then the Schrödinger bridge $\mathcal{P}$ is found by minimizing the corresponding rate function $\int \log (d \mathcal{P} / d \mathcal{W}) d \mathcal{P}$ over all probability measures that
are absolutely continuous with respect to $\mathcal{W}$ and have the given particle cloud distributions as marginals. A space and time discrete Schrödinger bridge problem for Markov chains is treated in [34,53].

We follow the exposition in [37], where the space and time discrete Schrödinger bridge has been derived as a maximum likelihood approximation for Markov chains. Consider a cloud of $N$ particles and assume that each particle evolves according to a Markov chain. Denote the states of the Markov chain by $X=\left\{X_{1}, X_{2}, \ldots, X_{n}\right\}$ and let the transition probability matrix be $A^{j} \in \mathbb{R}_{+}^{n \times n}$, with elements $A_{k \ell}^{j}=P\left(q_{j+1}=\right.$ $X_{\ell} \mid q_{j}=X_{k}$ ), where $q_{j}$ denotes the state at time $j$. Let the vector $\mu_{j}$ describe the particle distribution over the discrete state space $X$ at time $j$, for $j=1, \ldots, J$. As in the optimal transport framework we define the mass transport matrix $M^{j}$ where element $M_{k \ell}^{j}$ describes the number of particles transitioning from state $k$ at time $j$ to state $\ell$ at time $j+1$.

Given the distribution $\mu_{j}$, and transition probability matrix $A^{j}$, a large deviation argument similar to the continuous setting can be used to approximate the most likely transfer matrix $M^{j}$ by the minimizer of the rate function $H\left(\cdot \mid \operatorname{diag}\left(\mu_{j}\right) A^{j}\right)$. Assuming that the initial and final marginal, $\mu_{1}$ and $\mu_{J}$, are given, these large deviation theoretic considerations motivate the formulation of an optimization problem to find the most likely mass transfer matrices $M^{j}$, for $j=1, \ldots, J-1$, and intermediate distributions $\mu_{j}$, for $j=2, \ldots, J-1$, as

$$
\begin{align*}
\underset{M_{[1: J-1]}, \mu_{[2: J-1]}}{\operatorname{minimize}} & \sum_{j=1}^{J-1} H\left(M^{j} \mid \operatorname{diag}\left(\mu_{j}\right) A^{j}\right)  \tag{2.11}\\
\quad \text { subject to } & M^{j} \mathbf{1}=\mu_{j}, \quad\left(M^{j}\right)^{T} \mathbf{1}=\mu_{j+1}, \quad \text { for } \quad j=1, \ldots, J-1 .
\end{align*}
$$

Indeed, this problem is equivalent to the discrete time and space Schrödinger bridge problem in [53] (for details see [37] or Section 5.1).

An extension to a Hidden Markov Model formulation with aggregate indirect observations of the distributions was described in [37]. In this article the framework is extended to general tree structures, where each vertex is associated with a distribution, and each edge with a Markov process.
3. Tree-structured multi-marginal optimal transport. In this section we introduce tree-structured multi-marginal optimal transport problems. For the entropy regularized version of these problems, the projections (2.3) can be computed by only matrix-vector multiplications. This yields an efficient Sinkhorn algorithm, which we present in this section. Moreover, some properties of tree-structured multi-marginal optimal transport problems are discussed.

Definition 3.1. A graph $\mathcal{T}=(\mathcal{V}, \mathcal{E})$, with vertices $\mathcal{V}$ and edges $\mathcal{E}$, is a tree if it is acyclic and connected [26]. The vertices with degree 1 are called leaves, and we denote the set of leaves by $\mathcal{L}$. For a vertex $j \in \mathcal{V}$, the set of neighbours $\mathcal{N}_{j}$ is defined as the set of vertices, which have a common edge with $j$. The path between two vertices $j_{1}$ and $j_{L}$ is a sequence of edges that connect $j_{1}$ and $j_{L}$, such that all edges are distinct. We also denote a path by the set of intermediate vertices.

Let $\mathcal{T}=(\mathcal{V}, \mathcal{E})$ be a tree. In this section we consider a multi-marginal optimal transport problem where the marginals correspond to the nodes of the tree, i.e., $\mathcal{V}=\{1,2, \ldots, J\}$. Assume that the cost tensor decouples as

$$
\begin{equation*}
\mathbf{C}_{i_{1}, \ldots, i_{J}}=\sum_{\left(j_{1}, j_{2}\right) \in \mathcal{E}} C_{i_{j_{1}}, i_{j_{2}}}^{\left(j_{1}, j_{2}\right)} \tag{3.1}
\end{equation*}
$$

where $C^{\left(j_{1}, j_{2}\right)} \in \mathbb{R}_{+}^{n \times n}$ is a cost matrix characterizing the cost of transportation between marginal $j_{1}$ and $j_{2}$, for all $\left(j_{1}, j_{2}\right) \in \mathcal{E}$. Since we consider an undirected tree, the expression (3.1) should not depend on the ordering of the indices $\left(j_{1}, j_{2}\right)$ in the edges. This is achieved by simply letting $C^{\left(j_{2}, j_{1}\right)}=\left(C^{\left(j_{1}, j_{2}\right)}\right)^{T}$ for $\left(j_{1}, j_{2}\right) \in \mathcal{E}$. We refer to problem (2.2) with a cost of the form (3.1) as a multi-marginal optimal transport problem on the tree $\mathcal{T}$.

In the following, we consider discretized and entropy regularized multi-marginal optimal transport with tree structure. The transport plan can then be expressed as $\mathbf{M}=\mathbf{K} \odot \mathbf{U}$, where $\mathbf{K}=\exp (-\mathbf{C} / \epsilon)(c f .(2.6))$. Due to the structured cost (3.1), this tensor decouples as

$$
\begin{equation*}
\mathbf{K}_{i_{1}, \ldots, i_{J}}=\prod_{\left(j_{1}, j_{2}\right) \in \mathcal{E}} K_{i_{j_{1}}, i_{j_{2}}}^{\left(j_{1}, j_{2}\right)}, \tag{3.2}
\end{equation*}
$$

with the matrices defined as $K_{i_{j_{1}}, i_{j_{2}}}^{\left(j_{1}, j_{2}\right)}=\exp \left(-C_{i_{j_{1}}, i_{j_{2}}}^{\left(j_{1}, j_{2}\right)} / \epsilon\right)$, for $\left(j_{1}, j_{2}\right) \in \mathcal{E}$. Thus, it holds $K^{\left(j_{2}, j_{1}\right)}=\left(K^{\left(j_{1}, j_{2}\right)}\right)^{T}$ for $\left(j_{1}, j_{2}\right) \in \mathcal{E}$. We show that for these problems the projections (2.3) can therefore be computed by successive matrix vector multiplications. As the computation of the projections account for the bottleneck of the Sinkhorn iterations (2.9), this yields an efficient algorithm for solving entropy regularized multi-marginal optimal transport problems with tree structure.

We first illustrate the computation of the projections on a small example.

Example 1. Consider the tree $\mathcal{T}=(\mathcal{V}, \mathcal{E})$ with vertices $\mathcal{V}=\{1,2,3,4\}$ and edges $\mathcal{E}=\{(1,2),(2,3),(1,4)\}$ as depicted in Figure 1. Assume that the cost in problem (2.4) decouples as $\mathbf{C}_{i_{1}, i_{2}, i_{3}, i_{4}}=C_{i_{1}, i_{2}}^{(1,2)}+C_{i_{2}, i_{3}}^{(2,3)}+C_{i_{1}, i_{4}}^{(1,4)}$, where $C^{\left(j_{1}, j_{2}\right)}$, for $\left(j_{1}, j_{2}\right) \in \mathcal{E}$, are cost matrices defined on the respective edges. The transport tensor is then of the form $\mathbf{M}=\mathbf{K} \odot \mathbf{U}$, where $\mathbf{K}=\exp (-\mathbf{C} / \epsilon)$ and $\mathbf{U}=$ $u_{1} \otimes u_{2} \otimes u_{3} \otimes u_{4}$. Thus, denoting $K^{\left(j_{1}, j_{2}\right)}=\exp \left(-C^{\left(j_{1}, j_{2}\right)} / \epsilon\right)$


Figure 1. Illustration of the tree in Example 1. it holds that $\mathbf{K}_{i_{1}, i_{2}, i_{3}, i_{4}}=K_{i_{1}, i_{2}}^{(1,2)} K_{i_{2}, i_{3}}^{(2,3)} K_{i_{1}, i_{4}}^{(1,4)}$. Hence, the projection on the first marginal of $\mathbf{K} \odot \mathbf{U}$ can be computed as
$P_{1}(\mathbf{K} \odot \mathbf{U})_{i_{1}}=\sum_{i_{2}, i_{3}, i_{4}} \mathbf{K}_{i_{1}, i_{2}, i_{3}, i_{4}} \mathbf{U}_{i_{1}, i_{2}, i_{3}, i_{4}}=\left(u_{1}\right)_{i_{1}} \sum_{i_{2}, i_{3}, i_{4}} K_{i_{1}, i_{2}}^{(1,2)} K_{i_{2}, i_{3}}^{(2,3)} K_{i_{1}, i_{4}}^{(1,4)}\left(u_{2}\right)_{i_{2}}\left(u_{3}\right)_{i_{3}}\left(u_{4}\right)_{i_{4}}$.
The sum can be computed as successive matrix-vector multiplications, starting from the leaves of the tree. In particular, denote the products starting in the two leaves
$\sum_{i_{3}} K_{i_{2}, i_{3}}^{(2,3)}\left(u_{3}\right)_{i_{3}}=\left(K^{(2,3)} u_{3}\right)_{i_{2}}=:\left(\alpha_{(2,3)}\right)_{i_{2}}, \quad \sum_{i_{4}} K_{i_{1}, i_{4}}^{(1,4)}\left(u_{4}\right)_{i_{4}}=\left(K^{(1,4)} u_{4}\right)_{i_{1}}=:\left(\alpha_{(1,4)}\right)_{i_{1}}$.
On the lower branch of the tree, we have thereby brought the expression to a vector indexed by $i_{1}$. On the upper branch, another multiplication, corresponding to the edge $(1,2)$, is required. This leads to defining

$$
\sum_{i_{2}} K_{i_{1}, i_{2}}^{(1,2)}\left(u_{2} \odot \alpha_{(2,3)}\right)_{i_{2}}=\left(K^{(1,2)}\left(u_{2} \odot \alpha_{(2,3)}\right)\right)_{i_{1}}=:\left(\alpha_{(1,2)}\right)_{i_{1}}
$$

The full expression for the projection thus reads

$$
P_{1}(\mathbf{K} \odot \mathbf{U})=u_{1} \odot K^{(1,2)}\left(u_{2} \odot K^{(2,3)} u_{3}\right) \odot K^{(1,4)} u_{4}=u_{1} \odot \alpha_{(1,2)} \odot \alpha_{(1,4)} .
$$

This thus allows for decomposing the computation of $P_{1}(\mathbf{K} \odot \mathbf{U})$ into the three parts $u_{1}, \alpha_{(1,2)}, \alpha_{(1,4)}$, which represents the contributions from node 1; from nodes 2 and 3; and from node 4, respectively. Analogously, the projection on the second marginal is computed as $P_{2}(\mathbf{K} \odot \mathbf{U})=u_{2} \odot \alpha_{(2,3)} \odot \alpha_{(2,1)}$, where $\alpha_{(2,1)}:=K^{(2,1)}\left(u_{1} \odot \alpha_{(1,4)}\right)$. In particular, note here that the tuple $(j, k)$ in $\alpha_{(j, k)}$ is ordered. For instance, it holds $\alpha_{(1,2)} \neq \alpha_{(2,1)}$. The other marginals can be computed similarly by performing the corresponding matrix vector products, starting from the leaves of the tree.
Illustratively, the procedure in Example 1 for projecting the information of the full tensor down to one marginal can be understood as passing down the information from all branches of the tree to the desired vertex, where we interpret the vectors $u_{j}$ as local information in each node. Starting from the leaves of the tree, this is done according to the following rules:

1. Information is passed down the edge $\left(j_{1}, j_{2}\right) \in \mathcal{E}$ by multiplication with the matrix $K^{\left(j_{1}, j_{2}\right)}$.
2. Information is collected in the node $j \in \mathcal{V}$ by elementwise multiplication of the vector $u_{j}$, which contains local information, with information passed down from the connected branches.
In this interpretation the vector $\alpha_{(j, k)}$ can be understood as the collected information that is sent from node $k$ to node $j$. The following theorem shows that these rules hold for any tensor of the form $\mathbf{K} \odot \mathbf{U}$, with $\mathbf{U}=u_{1} \otimes \cdots \otimes u_{j}$ and $\mathbf{K}$ decoupling according to an arbitrary tree.

Theorem 3.2. Let $\mathbf{K}=\exp (-\mathbf{C} / \epsilon)$ with $\mathbf{C}$ as in (3.1) for the tree $\mathcal{T}=(\mathcal{V}, \mathcal{E})$, and let $\mathbf{U}=u_{1} \otimes u_{2} \otimes \cdots \otimes u_{J}$. Define $K^{\left(j_{1}, j_{2}\right)}=\exp \left(-C^{\left(j_{1}, j_{2}\right)} / \epsilon\right)$, for $\left(j_{1}, j_{2}\right) \in \mathcal{E}$. Then the projection on the $j$-th marginal of $\mathbf{K} \odot \mathbf{U}$ is of the form

$$
P_{j}(\mathbf{K} \odot \mathbf{U})=u_{j} \odot \bigodot_{k \in \mathcal{N}_{j}} \alpha_{(j, k)}
$$

The vectors $\alpha_{(j, k)}$, for all ordered tuples $(j, k) \in \mathcal{E}$, can be computed recursively starting in the leaves of the tree according to

$$
\begin{array}{ll}
\alpha_{(j, k)}=K^{(j, k)} u_{k}, & \text { for } k \in \mathcal{L}, \\
\alpha_{(j, k)}=K^{(j, k)}\left(u_{k} \odot \bigodot_{\ell \in \mathcal{N}_{k} \backslash\{j\}} \alpha_{(k, \ell)}\right), & \text { for } k \notin \mathcal{L} . \tag{3.3}
\end{array}
$$

Proof. See the Appendix.
Analogously to the marginal projections (2.3), let $P_{j_{1}, j_{2}}(\mathbf{M})$ denote the bi-marginal projections on marginals $j_{1}$ and $j_{2}$ as $P_{j_{1}, j_{2}}(\mathbf{M})_{i_{j_{1}}, i_{j_{2}}}=\sum_{i_{1}, \ldots, i_{J} \backslash\left\{i_{j_{1}}, i_{j_{2}}\right\}} \mathbf{M}_{i_{1}, \ldots, i_{J}}$. These pairwise projections can be expressed similarly to the marginal projections in Theorem 3.2.

Proposition 3.3. Let the assumptions in Theorem 3.2 hold, let $j_{1}, j_{L} \in \mathcal{V}$, and let $j_{1}, j_{2}, \ldots, j_{L}$ denote the path between $j_{1}$ and $j_{L}$. Then the pairwise projection of
$\mathbf{K} \odot \mathbf{U}$ on the marginals $j_{1}$ and $j_{L}$, denoted by $P_{j_{1}, j_{L}}(\mathbf{K} \odot \mathbf{U})$, is

$$
\begin{equation*}
\left(\prod_{\ell=1}^{L-1} \operatorname{diag}\left(u_{j_{\ell}} \odot \bigodot_{\substack{k \in \mathcal{N}_{j_{\ell}} \\ k \notin\left\{j_{1}, \ldots, j_{L}\right\}}} \alpha_{\left(j_{\ell}, k\right)}\right) K^{\left(j_{\ell}, j_{\ell}+1\right)}\right) \operatorname{diag}\left(u_{j_{L}} \odot \bigodot_{k \in \mathcal{N}_{j_{L}} \backslash\left\{j_{L-1}\right\}} \alpha_{\left(j_{L}, k\right)}\right), \tag{3.4}
\end{equation*}
$$

where $\alpha_{\left(j_{1}, j_{2}\right)}$, for $\left(j_{1}, j_{2}\right) \in \mathcal{E}$, are defined as in Theorem 3.2
Proof. See the supplementary material.
Multi-marginal optimal transport problems on some specific trees have previously been introduced in [30]. It is worth noting that the expressions for the projections of the mass tensor $\mathbf{M}$ for the examples in [30] are special cases of the results in Theorem 3.2 and Proposition 3.3. For instance, in tracking problems each marginal in the optimal transport problem is associated with a time instance, and the corresponding graph is a path graph, as in [30, Section 3.1 and 5.1]. Sensor fusion applications can be cast as barycenter problems, which are described by star-shaped graphs, as in [30, Section 3.2 and 5.2]. Moreover, a combination of these two applications is the tracking of barycenters over time, which is described by a graph as in Figure 6(a), see [30, Section 3.3 and 5.3].

The expressions for the projections in Theorem 3.2 can be used to solve a multimarginal optimal transport problem, with cost structure according to the tree $\mathcal{T}=$ $(\mathcal{V}, \mathcal{E})$, by a Sinkhorn method as in (2.9). To this end, without loss of generality, we consider tree-structured multi-marginal optimal transport problems (2.2), where the constraints are given on the set of leaves, i.e., $\Gamma=\mathcal{L}$. The following proposition shows that if a marginal is given for a non-leaf node, then the problem can be decomposed into smaller problems on subtrees, where the marginals are known exactly on the set of leaves of the subtrees.

Proposition 3.4. Consider problem (2.4) with constraint set $\Gamma$ and cost structured according to the tree $\mathcal{T}$ with leaves $\mathcal{L}$. The following holds:

1. If there is a marginal $k \in \Gamma$, but $k \notin \mathcal{L}$, then the solution to (2.4) can be found by solving problems of form (2.4) on the subtrees of $\mathcal{T}$, which are obtained by cutting $\mathcal{T}$ in node $k$.
2. If there is a marginal $\ell \in \mathcal{L}$, but $\ell \notin \Gamma$, then the solution to (2.4) can be found by solving the problem on the subtree of $\mathcal{T}$ that is obtained by removing the node $\ell$ and its adjoining edge from $\mathcal{T}$.
Proof. See the Appendix.
Note that since $\Gamma=\mathcal{L}$, the factors $u_{j}$, for $j \in \mathcal{V} \backslash \mathcal{L}$, can be neglected in Theorem 3.2 and Proposition 3.3. Moreover, when applying the iterative scheme (2.9), in each iteration step some of the factors $\alpha_{\left(j_{1}, j_{2}\right)}$ do not change. Specifically, between the update of $u_{j_{1}}$ and $u_{j_{2}}$ only the factors on all edges that lie on the path between nodes $j_{1}$ and $j_{2}$ need to be updated. The full method for solving tree-structured multimarginal optimal transport problems is summarized in Algorithm 3.1. Moreover, the algorithm has in fact linear convergence speed.

Theorem 3.5. Let $\left\{u_{j}^{k}\right\}_{j \in \mathcal{L}}$ be the set of vectors after the $k$ th iteration in the while-loop in Algorithm 3.1. Then the sequence $\left\{u_{j}^{k}\right\}_{j \in \mathcal{L}}$ converges at least linearly to an optimal solution of (2.8), as $k \rightarrow \infty$.

Proof. Algorithm 3.1 implements the Sinkhorn iterations (2.9), which are a block coordinate ascent in the dual multi-marginal optimal transport problem (2.8) [30]. Thus, the result follows from Application 5.3 and Theorem 2.1 in [47].

```
Algorithm 3.1 Sinkhorn method for the tree-structured multi-marginal optimal
transport problem.
    Given: Tree \(\mathcal{T}=(\mathcal{V}, \mathcal{E})\) with leaves \(\mathcal{L}=\{1,2, \ldots,|\mathcal{L}|\}\);
    Initial guess \(u_{j}\), for \(j \in \mathcal{L}\)
    for all ordered tuples \(\left(j_{1}, j_{2}\right) \in \mathcal{E}\) do
        Initialize \(\alpha_{\left(j_{1}, j_{2}\right)}\) according to (3.3)
    end for
    Initialize \(j \in \mathcal{L}\)
    while Sinkhorn not converged do
        \(u_{j} \leftarrow \mu_{j} . / \bigodot_{k \in \mathcal{N}_{j}} \alpha_{(j, k)}\)
        for \(\left(j_{1}, j_{2}\right) \in \mathcal{E}\) on the path from \(j\) to \((j+1 \bmod |\mathcal{L}|)\) do
            Update \(\alpha_{\left(j_{1}, j_{2}\right)}\) according to (3.3)
        end for
        \(j \leftarrow j+1 \bmod |\mathcal{L}|\)
    end while
    return \(u_{j}\) for \(j \in \mathcal{L}\)
```

It should be noted that computing the marginals $P_{j}(\mathbf{K} \odot \mathbf{U})$ by summing over the indices as defined in (2.3) scales exponentially in the number of marginals $J$ of the tensor $\mathbf{K} \odot \mathbf{U}$. Thus, the complexity of one Sinkhorn update in (2.9) is in general of the order $\mathcal{O}\left(n^{J}\right)$. Algorithm 3.1 utilizes the result in Theorem 3.2 in order to exploit the graph-structure in the tensor $\mathbf{K} \odot \mathbf{U}$ to compute the projections. This decreases the computational complexity of the Sinkhorn iterations substantially. In particular, if the direct path between node $j-1$ and node $j$ is of length $p$, then the update of the vector $u_{j}$ requires $p$ matrix-vector multiplications. The complexity of one update is thus of the order $\mathcal{O}\left(p n^{2}\right)$. As a rule of thumb the leaves in the underlying tree $\mathcal{T}$ should thus be labeled such that the paths between any two nodes that are updated successively are short. In case there are additional structures in the cost matrices $C^{\left(j_{1}, j_{2}\right)}$, for $\left(j_{1}, j_{2}\right) \in \mathcal{E}$, e.g., they describe the squared Euclidean distance, the cost of the matrix-vector multiplications can be decreased further, see, e.g., [30, Rem. 4], [41, Rem. 3.11].

Note that Proposition 3.4.1 implies that if the tensor $\mathbf{K}$ is normalized to be a discrete probability distribution, it defines a Markov random field (see also [38]). Every Markov process is also a reciprocal process [39], which were first introduced in the context of studying Schrödinger bridges [40]. Recall that the maximum likelihood estimation problem for a Markov chain in Section 2.3 is equivalent to a discrete Schrödinger bridge. This suggests a connection between the multi-marginal optimal transport problem and the maximum likelihood estimation problem in Section 2.3. In fact, under certain conditions a generalization of the Markov chain problem (2.11) to trees, which is introduced in Section 4.1, yields the same solution as the regularized multi-marginal optimal transport problem (2.4) on the same tree, as we will see in Section 4.2.
4. Multi-marginal optimal transport vs. Schrödinger bridge with treestructure. In this section, we extend the Schrödinger bridge problem which is defined on a path tree to a formulation valid on an arbitrary tree graph. We then show that under certain conditions this problem yields a solution which is equivalent to the solution of the entropy regularized multi-marginal optimal transport problem (2.4) on the same tree, thus extending the well-known equivalence from the optimal transport
problem and the Schrödinger bridge problem on a path tree [13, 14, 48, 49].
4.1. Tree-structured Schrödinger bridges. The notion of time of the Markov chain introduces a direction to the Schrödinger bridge problem, which can be seen directly in the optimization problem (2.11). This directionality needs to be taken into account when extending the problem to a general tree. To this end, we introduce the following notation: Consider a tree $\mathcal{T}_{r}=\left(\mathcal{V}, \mathcal{E}_{r}\right)$ that is rooted in a vertex $r \in \mathcal{L}$, i.e., one of the leaves is defined to be the root of the tree. This defines a partial ordering on the tree, and we write $k<j$ if node $k$ lies on the path between $r$ and node $j$. To formulate the following results, we assume that all edges are directed according to this partial ordering, i.e., $j_{1}<j_{2}$ for all $\left(j_{1}, j_{2}\right) \in \mathcal{E}_{r}$. For a vertex $j \in \mathcal{V} \backslash r$, its parent is then defined as the (unique) vertex $p(j)$ such that $(p(j), j) \in \mathcal{E}_{r}$. In the following, without loss of generality we denote the root vertex by $r=1$.

Let $A^{\left(j_{1}, j_{2}\right)}$ be the probability transition matrix on edge $\left(j_{1}, j_{2}\right) \in \mathcal{E}_{r}$. Then, the Schrödinger bridge problem in (2.11) may be naturally extended to the tree structure as

$$
\begin{align*}
& \underset{\substack{M^{\left(j_{1}, j_{2}\right)},\left(j_{1}, j_{2}\right) \in \mathcal{E}_{r} \\
\mu_{j}, j \in \mathcal{V} \backslash \Gamma}}{\operatorname{minimize}} \sum_{\substack{\left.j_{1}, j_{2}\right) \in \mathcal{E}_{r}}} H\left(M^{\left(j_{1}, j_{2}\right)} \mid \operatorname{diag}\left(\mu_{j_{1}}\right) A^{\left(j_{1}, j_{2}\right)}\right)  \tag{4.1}\\
& \quad \text { subject to } \quad M^{\left(j_{1}, j_{2}\right)} \mathbf{1}=\mu_{j_{1}}, \quad\left(M^{\left(j_{1}, j_{2}\right)}\right)^{T} \mathbf{1}=\mu_{j_{2}}, \quad \text { for } \quad\left(j_{1}, j_{2}\right) \in \mathcal{E}_{r} .
\end{align*}
$$

Moreover, we can assume without loss of generality that $\Gamma=\mathcal{L}$, by similar reasoning as in Section 3. The extension of (2.11) to (4.1) builds on the large deviation principle in [37, Prop. 1], which requires that an initial distribution is known. This is why we restrict our analysis to the case when the directed tree $\mathcal{T}_{r}\left(\mathcal{V}, \mathcal{E}_{r}\right)$ is rooted in a leaf.

Similarly to the multi-marginal optimal transport problem, the optimal solution to (4.1) can be expressed in terms of its dual variables. In fact this is a natural generalization of the forward and backward propagation factors in the Schrödinger system $[34,53]$ to tree graphs.

Theorem 4.1. Let $A^{\left(j_{1}, j_{2}\right)}$ be probability transition matrices, for $\left(j_{1}, j_{2}\right) \in \mathcal{E}_{r}$, with strictly positive elements and assume that $\mathbf{1}^{T} \mu_{j_{1}}=\mathbf{1}^{T} \mu_{j_{2}}$ for all leaves $j_{1}, j_{2} \in \mathcal{L}$. Then an optimal solution to (4.1) can be written as

$$
\begin{aligned}
\mu_{j} & =\varphi_{j} \odot \hat{\varphi}_{j} \\
M^{\left(j_{1}, j_{2}\right)} & =\operatorname{diag}\left(\hat{\varphi}_{j_{1}} \odot \varphi_{j_{1} \backslash j_{2}}\right) A^{\left(j_{1}, j_{2}\right)} \operatorname{diag}\left(\varphi_{j_{2}}\right)
\end{aligned}
$$

for a set of vectors $\varphi_{j}, \hat{\varphi}_{j}$, for $j \in \mathcal{V}$, and $\varphi_{j_{1} \backslash j_{2}}$ for $\left(j_{1}, j_{2}\right) \in \mathcal{E}_{r}$.
Moreover, there exists a set of vectors $v_{j}$, for $j \in \Gamma$, such that the vectors $\varphi_{j}, \hat{\varphi}_{j}$ and $\varphi_{j_{1} \backslash j_{2}}$ can be written as

$$
\varphi_{j}= \begin{cases}v_{j}, & \text { if } j \in \mathcal{L} \backslash\{1\}  \tag{4.2}\\ \bigodot_{k:(j, k) \in \mathcal{E}_{r}} A^{(j, k)} \varphi_{k}, & \text { else },\end{cases}
$$

$$
\hat{\varphi}_{j}= \begin{cases}1 . / v_{1}, & \text { if } j=1  \tag{4.3}\\ \left(A^{(p(j), j)}\right)^{T}\left(\hat{\varphi}_{p(j)} \odot \varphi_{p(j) \backslash j}\right), & \text { else },\end{cases}
$$

and $\varphi_{j_{1} \backslash j_{2}}=\varphi_{j_{1}} \cdot /\left(A^{\left(j_{1}, j_{2}\right)} \varphi_{j_{2}}\right)$.

Proof. See the Appendix.
Note that the vectors $v_{j}$ for $j \in \mathcal{L}$ in Theorem 4.1 satisfy the nonlinear equations $\varphi_{1} . / v_{1}=\mu_{1}$, and $v_{j} \odot \hat{\varphi}_{j}=\mu_{j}$, for $j \in \mathcal{L} \backslash\{1\}$. When considering the Schrödinger bridge, which is the special case with two leaves, these equations correspond to the boundary conditions in the Schrödinger system [53, Thm. 4.1]. The factors $\varphi_{j}$ and $\hat{\varphi}_{j}$ can be seen as propagating information from the leaves to the other nodes in a similar manner as $\alpha_{(j, k)}$ in Theorem 3.2.

Next, we show that the tree-structured Schrödinger bridge problem (4.1) is independent of the choice of root $r \in \mathcal{L}$. This follows as a corollary to the next proposition.

Proposition 4.2. Let $\mathcal{T}_{r}=\left(\mathcal{V}, \mathcal{E}_{r}\right)$ be a rooted directed tree with root $r \in \mathcal{L}$. Then problem (4.1) is equivalent to the problem

$$
\begin{align*}
\underset{\substack{M^{\left(j_{1}, j_{2}\right)},\left(j_{1}, j_{2}\right) \in \mathcal{E}_{r}, \mu_{j}, j \in \mathcal{V} \backslash \Gamma}}{\operatorname{minimize}} & \sum_{\left(j_{1}, j_{2}\right) \in \mathcal{E}_{r}} H\left(M^{\left(j_{1}, j_{2}\right)} \mid A^{\left(j_{1}, j_{2}\right)}\right)-\sum_{j \in \mathcal{V} \backslash \mathcal{L}}(\operatorname{deg}(j)-1) H\left(\mu_{j}\right)  \tag{4.4}\\
\text { subject to } & M^{\left(j_{1}, j_{2}\right)} \mathbf{1}=\mu_{j_{1}}, \quad\left(M^{\left(j_{1}, j_{2}\right)}\right)^{T} \mathbf{1}=\mu_{j_{2}}, \quad \text { for } \quad\left(j_{1}, j_{2}\right) \in \mathcal{E}_{r} .
\end{align*}
$$

Proof. See the Appendix.
Corollary 4.3. Let $\mathcal{T}_{r}=\left(\mathcal{V}, \mathcal{E}_{r}\right)$ be a directed tree with root $r \in \mathcal{L}$, and $\mathcal{T}_{\hat{r}}=$ $\left(\mathcal{V}, \mathcal{E}_{\hat{r}}\right)$ be a directed tree with root $\hat{r} \in \mathcal{L}$ with the same strucure as $\mathcal{T}_{r}$ and the edges on the path between $r$ and $\hat{r}$ reversed. Then, the solution to (4.1) on $\mathcal{T}_{r}$ and on $\mathcal{T}_{\hat{r}}$ are equivalent in the sense that the marginal distributions $\mu_{j}$, for $j \in \mathcal{V}$, and transport plans $M^{\left(j_{1}, j_{2}\right)}$, for $\left(j_{1}, j_{2}\right) \in \mathcal{E}_{r} \cap \mathcal{E}_{\hat{r}}$, are the same, and on the reversed edges it holds $M^{\left(j_{2}, j_{1}\right)}=\left(M^{\left(j_{1}, j_{2}\right)}\right)^{T}$.

Proof. See the supplementary material.
4.2. Equivalence of problems (2.4) and (4.1) for tree-structures. We will now verify that the generalized Schrödinger bridge (4.1) on a rooted tree $\mathcal{T}_{r}$ is equivalent to an entropy regularized multi-marginal optimal transport problem (2.4) on an undirected tree $\mathcal{T}$ with the same structure as $\mathcal{T}_{r}$. In particular, given positive transition probability matrices $A^{\left(j_{1}, j_{2}\right)}$ for $\left(j_{1}, j_{2}\right) \in \mathcal{E}_{r}$ and a regularization parameter $\epsilon$, there is a natural choice of cost matrices $C^{\left(j_{1}, j_{2}\right)}$ for $\left(j_{1}, j_{2}\right) \in \mathcal{E}$, so that the minimizers to both problems represent the same solution. For the original bimarginal optimal transport problem (2.1) the equivalence between entropy regularized optimal transport and the Schrödinger bridge problem has been studied extensively [13, 14, 48, 49].

Remark 1 (cf. [37, Rem. 1]). In the discrete setting, it is easy to see that with a cost matrix defined as $C=-\epsilon \log (A)$, the entropy regularized optimal transport problem (2.5) is equivalent to the bi-marginal Schrödinger bridge (2.11), by noting that the objective can then be written as

$$
\operatorname{trace}\left(C^{T} M\right)+\epsilon H(M)=\sum_{i, j=1}^{n} \epsilon\left(M_{i j} \log \left(\frac{M_{i j}}{A_{i j}}\right)-M_{i j}+1\right)=\epsilon H(M \mid A) .
$$

Both bi-marginal problems thus yield the same optimal solution $M$.
We will show that this equivalence holds true even for the multi-marginal case with tree-structured cost.

THEOREM 4.4. Let $\mathcal{T}_{r}=\left(\mathcal{V}, \mathcal{E}_{r}\right)$ be a rooted directed tree with root $r=1 \in \mathcal{L}$, let $\mathcal{T}=(\mathcal{V}, \mathcal{E})$ be its undirected counterpart, and let $\mu_{j}$ for $j \in \mathcal{L}$ be given marginals. Let
$\epsilon>0$ and $C^{\left(j_{1}, j_{2}\right)}$ be such that

$$
K^{\left(j_{1}, j_{2}\right)}=A^{\left(j_{1}, j_{2}\right)}, \quad \text { for }\left(j_{1}, j_{2}\right) \in \mathcal{E}_{r}
$$

Let $\mathbf{M}=\mathbf{K} \odot \mathbf{U}$ be the solution to the entropy regularized multi-marginal optimal transport problem (2.4) on $\mathcal{T}$, where $\mathbf{K}$ and $\mathbf{U}$ are defined as in Theorem 3.2. Moreover, let $\mu_{j}$ for $j \in \mathcal{V}$, and $M^{\left(j_{1}, j_{2}\right)}$ for $\left(j_{1}, j_{2}\right) \in \mathcal{E}_{r}$, be the solution to the generalized Schrödinger bridge (4.1) on $\mathcal{T}_{r}$. Then,

$$
\begin{aligned}
P_{j}(\mathbf{M}) & =\mu_{j}, & & \text { for } j \in \mathcal{V}, \\
P_{\left(j_{1}, j_{2}\right)}(\mathbf{M}) & =M^{\left(j_{1}, j_{2}\right)}, & & \text { for }\left(j_{1}, j_{2}\right) \in \mathcal{E}_{r}
\end{aligned}
$$

Furthermore, if $v_{j}$ for $j \in \mathcal{L}$ are the corresponding vectors in Theorem 4.1, then it holds that

$$
u_{j}= \begin{cases}1 . / v_{1}, & \text { if } j=1 \\ v_{j}, & \text { otherwise } .\end{cases}
$$

Proof. See the Appendix.
In particular, given a maximum likelihood problem (4.1) structured according to the tree $\mathcal{T}_{r}=\left(\mathcal{V}, \mathcal{E}_{r}\right)$, a corresponding multi-marginal optimal transport problem can be formulated by definining the cost matrices as $C^{\left(j_{1}, j_{2}\right)}=-\epsilon \log \left(A^{\left(j_{1}, j_{2}\right)}\right)$, for all $\left(j_{1}, j_{2}\right) \in \mathcal{E}_{r}$. Vice versa, given an entropy regularized multi-marginal optimal transport problem (2.4) with cost structured according to the tree $\mathcal{T}=(\mathcal{V}, \mathcal{E})$, a corresponding maximum likelihood problem can be formulated by transforming the undirected tree into a rooted tree $\mathcal{T}_{r}=\left(\mathcal{V}, \mathcal{E}_{r}\right)$, where $r \in \mathcal{L}$, and defining the matrices $A^{\left(j_{1}, j_{2}\right)}=\exp \left(-C^{\left(j_{1}, j_{2}\right)} / \epsilon\right)$, for all $\left(j_{1}, j_{2}\right) \in \mathcal{E}_{r}$. However, if these matrices are not row-stochastic, they cannot be interpreted as transition probability matrices of a Markov process.

Remark 2. For a nonnegative matrix $A$ with at least one positive element in each row, define the vector $b=1 . / A 1$. Then the matrix $\hat{A}=\operatorname{diag}(b) A$ is row stochastic. Note that we can write $H(M \mid \operatorname{diag}(\mu) A)=H(M \mid \hat{A})+H(\mu \mid b)$. For a given rooted tree $\mathcal{T}_{r}=\left(\mathcal{V}, \mathcal{E}_{r}\right)$ assume that there is a set of matrices $A^{\left(j_{1}, j_{2}\right)}$, for $\left(j_{1}, j_{2}\right) \in \mathcal{E}_{r}$, which are not necessarily row-stochastic, but are such that for every $j_{1} \in \mathcal{E}_{r}$ there is a vector $b_{j_{1}}$ such that it holds that $b_{j_{1}}=\mathbf{1} . / A^{\left(j_{1}, j_{2}\right)} \mathbf{1}$ for all $j_{2}$ such that $\left(j_{1}, j_{2}\right) \in \mathcal{E}_{r}$. Then, according to Proposition 4.2, the objective function of the maximum likelihood estimation problem of a Markov process on $\mathcal{T}_{r}$ in (4.1) may be written as

$$
\begin{equation*}
\sum_{\left(j_{1}, j_{2}\right) \in \mathcal{E}_{r}} H\left(M^{\left(j_{1}, j_{2}\right)} \mid \hat{A}^{\left(j_{1}, j_{2}\right)}\right)-\sum_{j \in \mathcal{V} \backslash \mathcal{L}}(\operatorname{deg}(j)-1) H\left(\mu_{j} \mid b_{j}\right) \tag{4.5}
\end{equation*}
$$

The Sinkhorn scheme (2.9) can equivalently be formulated in terms of the notation from the maximum likelihood estimation problem in Theorem 4.1. Given an initial set of positive vectors $v_{j}$ for $j \in \mathcal{L}$, the updates are then expressed as

$$
\begin{align*}
& v_{1}=\varphi_{1} \cdot / \mu_{1} \\
& v_{j}=\mu_{j} \cdot / \hat{\varphi}_{j}, \text { for } j \in \mathcal{L} \backslash 1, \tag{4.6}
\end{align*}
$$

where $\varphi_{1}$ is computed as in (4.2) before each update of $v_{1}$, and $\hat{\varphi}_{j}$ is computed as in (4.3) before updating $v_{j}$ for each $j \in \mathcal{L} \backslash 1$. This can be done efficiently, and by
storing and reusing intermediate results appropriately, this algorithm is equivalent to Algorithm 3.1. It was shown that the classical Sinkhorn iterations are a block coordinate ascent in a Lagrange dual problem [41,60]. In this sense, the scheme (4.6) can be understood as Sinkhorn iterations even for the problem (4.1).

Proposition 4.5. Let $v_{j}$, for $j \in \mathcal{L}$, be an initial set of positive vectors. Iteratively performing the updates (4.6) corresponds to a block coordinate ascent in a Lagrange dual problem to (4.1).

In the light of Proposition 4.5 it is not surprising, that the algorithm presented in [37], for the special case where the tree represents a hidden Markov chain, is of the form (4.6), where the vectors $\varphi_{j}$ and $\hat{\varphi}_{j}$ are constructed according to the special tree structure.

REmARK 3. The tree-structured optimal transport problem, i.e., problem (2.2) where the cost is of the form (3.1) for a tree-graph $\mathcal{G}=(\mathcal{V}, \mathcal{E})$, can be formulated as a sum of pairwise optimal transport costs

$$
\begin{equation*}
\underset{\mu_{j}, j \in \mathcal{V} \backslash \Gamma}{\operatorname{minimize}} \sum_{\left(j_{1}, j_{2}\right) \in \mathcal{E}} T^{\left(j_{1}, j_{2}\right)}\left(\mu_{j_{1}}, \mu_{j_{2}}\right), \tag{4.7}
\end{equation*}
$$

where $T^{\left(j_{1}, j_{2}\right)}(\cdot, \cdot)$ is the optimal transport problem (2.1) with cost matrix $C^{\left(j_{1}, j_{2}\right)}$ for $\left(j_{1}, j_{2}\right) \in \mathcal{E}$. In particular, both problems have common optimal solutions which are identical in the sense that the projections of an optimal tensor $\mathbf{M}$ in (2.2) coincide with a set of optimal transport plans in (4.7) in the sense that $P_{\left(j_{1}, j_{2}\right)}(\mathbf{M})=M^{\left(j_{1}, j_{2}\right)}$ for $\left(j_{1}, j_{2}\right) \in \mathcal{E}$.

This, however, can not be generalized to graphs $\mathcal{G}=(\mathcal{V}, \mathcal{E})$ that contain cycles. A counterexample can be found for a complete graph with three nodes (Figure 2). Let the cost matrices be

$$
C^{(1,2)}=C^{(2,3)}=J=\left[\begin{array}{ll}
0 & 1 \\
1 & 0
\end{array}\right], \quad C^{(1,3)}=I=\left[\begin{array}{ll}
1 & 0 \\
0 & 1
\end{array}\right]
$$



Figure 2. Illustration of the cyclic graph in Remark 3.
where $\mu_{1}=[1,1]^{T}$, and $\Gamma=\{1\}$. The unique optimal
solution to (4.7) with these parameters are $\mu_{2}=\mu_{3}=\left[\begin{array}{ll}1 & 1\end{array}\right]^{T}, \quad M^{(1,2)}=M^{(2,3)}=$ I, $M^{(1,3)}=J$, and the objective value is 0 . However, note that the corresponding multimarginal cost tensor given by $\mathbf{C}_{i_{1}, i_{2}, i_{3}}=C_{i_{1}, i_{2}}^{(1,2)}+C_{i_{1}, i_{3}}^{(1,3)}+C_{i_{2}, i_{3}}^{(2,3)}$ is elementwise strictly positive, and thus the multimarginal problem cannot attain the value 0 . In fact, there is no tensor $\mathbf{M} \in \mathbb{R}_{+}^{2 \times 2 \times 2}$ that is consistent with these projections $P_{1,2}(\mathbf{M})=$ $P_{2,3}(\mathbf{M})=I$ and $P_{1,3}(\mathbf{M})=J$.

Here we have illustrated a fundamental difficulty that must be handled when extending tree-graphs to graphs with cycles. Thus, naively extending the Schrödinger bridge from trees to general graphs by defining transition probability matrices $A^{\left(j_{1}, j_{2}\right)}=$ $\exp \left(-C^{\left(j_{1}, j_{2}\right)} / \epsilon\right)$, for all $\left(j_{1}, j_{2}\right) \in \mathcal{E}$, and solving (4.1) does not yield an equivalence result with a multi-marginal optimal transport problem as in Theorem 4.4.
5. Multi-marginal vs. pairwise regularization. Another natural way to define an optimal transport problem on the tree $\mathcal{T}$ is to minimize the sum of all bimarginal transport costs on the edges of $\mathcal{T}$, as in (4.7). In fact, the unregularized multi-marginal optimal transport problem (2.2) structured according to $\mathcal{T}=(\mathcal{V}, \mathcal{E})$ is equivalent to this pairwise problem. An alternative computational approach to the one taken in this paper is thus to regularize each term in the sum of (4.7) by an


Figure 3. Illustration of the linear path tree in Section 5.1.
entropy term. The entropy regularized pairwise optimal transport problem on $\mathcal{T}$ is then

$$
\begin{equation*}
\underset{\mu_{j}, j \in \mathcal{V} \backslash \Gamma}{\operatorname{minimize}} \sum_{\left(j_{1}, j_{2}\right) \in \mathcal{E}} T_{\epsilon}^{\left(j_{1}, j_{2}\right)}\left(\mu_{j_{1}}, \mu_{j_{2}}\right), \tag{5.1}
\end{equation*}
$$

where $T_{\epsilon}^{\left(j_{1}, j_{2}\right)}(\cdot, \cdot)$ is defined as the regularized bi-marginal optimal transport problem (2.5) with cost matrix $C^{\left(j_{1}, j_{2}\right)}$. In particular, it is common to formulate barycenter problems in this pairwise regularized manner $[7,8,23,43,45,58]$. Recently, we have empirically observed that in some applications the barycenter problem with multimarginal regularization as in (2.4) gives better results as compared to the pairwise regularization in (5.1), see [30, Sec. 6.3]. Therein solutions using the multi-marginal regularization are less smoothed out as compared to the pairwise regularization. In the following we investigate the difference between the two problems. In particular, we show that the latter is not equivalent to the generalized Schrödinger bridge (4.1).

Consider a rooted version of the tree, denoted $\mathcal{T}_{r}=\left(\mathcal{V}, \mathcal{E}_{r}\right)$, where $r \in \mathcal{L}$. In the case that $A^{\left(j_{1}, j_{2}\right)}=K^{\left(j_{1}, j_{2}\right)}$, for all $\left(j_{1}, j_{2}\right) \in \mathcal{E}_{r}$, the multi-marginal optimal transport problem has the same solution as problem (4.4). Now note that using Proposition 4.2 this problem can be written as

$$
\begin{equation*}
\underset{\mu_{j}, j \in \mathcal{V} \backslash \Gamma}{\operatorname{minimize}} \sum_{\left(j_{1}, j_{2}\right) \in \mathcal{E}} T_{\epsilon}^{\left(j_{1}, j_{2}\right)}\left(\mu_{j_{1}}, \mu_{j_{2}}\right)-\sum_{j \in \mathcal{V} \backslash \mathcal{L}}(\operatorname{deg}(j)-1) H\left(\mu_{j}\right) . \tag{5.2}
\end{equation*}
$$

Thus, the generalized Schrödinger bridge problem (4.1), or equivalently (5.2), is not equivalent to the pairwise regularized optimal transport problem (5.1), unless for the trivial case of a tree with only 2 vertices. Moreover, we can see from (5.2) that the multi-marginal optimal transport problem penalizes not only the transport cost between the marginals, but in addition favors marginal distributions with low entropy. ${ }^{1}$ One can thus expect less smoothed out distributions when solving the multi-marginal optimal transport problem, which is desirable in many applications, such as localization problems [30] and computer vision applications [58]. Solving tree-structured problems by using multi-marginal regularization thus has some advantages as compared to pairwise regularization. Firstly, it yields less smoothed out solutions, and secondly it preserves the connections to the Schrödinger bridge problem. Finally, empirical study suggests that the multi-marginal problem is better conditioned compared to the pairwise problem, which allows for smaller values of the regularization parameter $\epsilon$, while still yielding a numerically stable algorithm. We give empirical evidence for this behaviour in Section 5.2.

The solution of the pairwise optimal transport problem (5.1) is compared to the solution of the two equivalent problems (2.4) and (4.1) on the example of a path graph in Section 5.1, and on a more complex tree in Section 5.2.
5.1. The discrete time Schrödinger bridge problem. Consider a path tree $\mathcal{T}_{r}=\left(\mathcal{V}, \mathcal{E}_{r}\right)$, where $\mathcal{V}=\{1,2, \ldots, J\}$ and $\mathcal{E}_{r}=\{(j, j+1) \mid j=1, \ldots, J-1\}$, as sketched

[^1]

Figure 4. Solutions to the multi-marginal and pairwise entropy regularized optimal transport problem on a path graph for varying regularization parameter $\epsilon$.
in Figure 3. Let $A^{j}$ denote the probability transition matrix on $(j, j+1) \in \mathcal{E}_{r}$. This model corresponds to a Markov chain of length $J$. Assume that the distributions on the leaves $j=1$ and $j=J$ are known. The most likely particle evolutions between them are then found by solving (2.11). Following [37], we see that this problem is equivalent to the discrete time and discrete space Schrödinger bridge in [53]. Assume that the distributions $\mu_{j}$, for $j=1, \ldots, J$, are strictly positive, and define the row stochastic matrices $\bar{A}^{j}=\operatorname{diag}\left(\mu_{j}\right)^{-1} M^{j}$, for $j=1, \ldots, J$. In terms of these matrices problem (2.11) reads

$$
\begin{align*}
\underset{\left.\bar{A}_{[1: J-1]}, \mu_{[2, J-1]}\right]}{\operatorname{minimize}} & \sum_{j=1}^{J-1} \sum_{i=1}^{n}\left(\mu_{j}\right)_{i} H\left(\bar{A}_{i \cdot}^{j} \cdot \mid A_{i \cdot}^{j}\right)  \tag{5.3}\\
\text { subject to } & \bar{A}^{j} \mathbf{1}=\mathbf{1}, \quad \mu_{j+1}=\left(\bar{A}^{j}\right)^{T} \mu_{j}, \text { for } j=1, \ldots, J-1 .
\end{align*}
$$

Here $A_{i}$. denotes the $i$-th row of $A$. Problem (5.3) is exactly the formulation of a Schrödinger bridge over a Markov chain from [53, eq. (24)]. In [53] it is shown that a unique solution to a corresponding Schrödinger system exists if $\mu_{J}$ is a strictly positive distribution and the matrix $\prod_{j=1}^{J-1} A^{j}$ has only positive elements. The solution to the Schrödinger system may be obtained as a fixed point iteration [34], which is linked to the Sinkhorn iterations for entropy regularized optimal transport problems. We recall from [37] that the optimization problem (5.3) is non-convex, whereas the equivalent formulation (2.11) is convex.

We note that with cost matrices defined as $C^{j}=-\epsilon \log \left(A^{j}\right)$, for $j=1, \ldots, J-1$, the discrete Schrödinger bridge problem may be written as

$$
\underset{\mu_{2}, \ldots, \mu_{J-1}}{\operatorname{minimize}} \sum_{j=1}^{J-1} T_{\epsilon}^{(j, j+1)}\left(\mu_{j}, \mu_{j+1}\right)-\sum_{j=2}^{J-1} H\left(\mu_{j}\right) .
$$

In comparison to the pairwise optimal transport problem on the path graph $\mathcal{T}$, the Schrödinger bridge thus favors intermediate distributions with lower entropy, resulting in less smoothed out solutions. We illustrate this behaviour for a path tree $\mathcal{T}_{r}=\left(\mathcal{V}, \mathcal{E}_{r}\right)$, with vertices $\mathcal{V}=\{1,2, \ldots, 6\}$, and where the initial and final distribution are given as $\mu_{1}(x)=\exp \left(-\left(\frac{x-0.2}{10}\right)^{2}\right)$, and $\mu_{J}(x)=\exp \left(-\left(\frac{x-0.8}{10}\right)^{2}\right)$, and the cost function is defined as the Euclidean distance. The results to problem (5.3) and problem (5.1) are compared for different values of the regularization parameter $\epsilon$ in Figure 4. Both optimal transport solutions describe a smooth way of shifting the mass from distribution $\mu_{1}$ to $\mu_{J}$. The larger the regularization parameter $\epsilon$ is chosen, the more smoothed out are the intermediate solutions to both problems. However,


Figure 5. Estimated marginals of the pairwise (b) and multi-marginal (c,d) optimal transport solutions on the tree in (a).
for each value of $\epsilon$ the multi-marginal optimal transport solution infers substantially less smoothing than the pairwise optimal transport solution.
5.2. Optimal transport with tree-structure. In this section we compare the entropy regularized multi-marginal and pairwise optimal transport solutions on a more general tree. Consider the tree $\mathcal{T}=(\mathcal{V}, \mathcal{E})$ illustrated in Figure 5(a) with 15 nodes, each representing a $50 \times 50$ pixel image. The marginal images on the 8 leaves, coloured in gray, are known. Each edge on the tree $\mathcal{T}$ is associated with a cost function defined by the Euclidean distance between any two pixels. Using this choice of cost function in the corresponding optimal transport problems yields smooth translations in power for the intermediate marginals.

We solve the entropy regularized pairwise optimal transport problem (5.1) with regularization parameter $\epsilon=2 \cdot 10^{-3}$ on $\mathcal{T}$. The solution can be seen in Figure 5(b). Compared to the pairwise optimal transport estimate, the solution to the entropy regularized multi-marginal optimal transport problem (2.4) on the same tree $\mathcal{T}$ and with the same regularization parameter $\epsilon$ is significantly sharper and less smoothed out, see Figure 5(c). For the pairwise optimal transport problem, the method diverges with a smaller regularization parameter, e.g., $\epsilon=10^{-3}$. In contrast, for the multi-marginal formulation the regularization parameter can be decreased further, still yielding a numerically stable algorithm. We have found that the method is still stable for a regularization parameter of $\epsilon=4 \cdot 10^{-4}$, which results in very clear estimates on the intermediate nodes.
6. Estimating ensemble flows on a hidden Markov chain. We consider the problem of tracking an ensemble of agents on a network based on aggregate measurements from sensors distributed around the network. This is similar to [37], where ensemble flows of indistinguishable agents have been estimated as the maximum likelihood solution on a hidden Markov chain. The present work generalizes this method and provides an algorithm for solving the problem introduced therein. In particular, the framework in [37] is a special case of the method in Theorem 4.1, and can therefore be solved with Algorithm 3.1. Herein, we study different observation models and robustness of the estimates with respect to the number of agents.

Consider a graph $\mathcal{G}=\left(\mathcal{V}_{\mathcal{G}}, \mathcal{E}_{\mathcal{G}}\right)$ with vertices $\mathcal{V}_{\mathcal{G}}$ and edges $\mathcal{E}_{\mathcal{G}}$. Let the set of nodes $\mathcal{V}_{\mathcal{G}}$ be the set of states of a Markov model with transition probability matrix $A^{t} \in \mathbb{R}^{n \times n}$, where $n=\left|\mathcal{V}_{\mathcal{G}}\right|$, for $t=1, \ldots, \tau-1$. We simulate a finite number of agents to evolve according to this Markov model for a number of time steps $t=1, \ldots, \tau$. Let $\mu_{t} \in \mathbb{R}^{n}$ denote the distribution of agents over the set of states for the times $t=$


Figure 6. Ensemble flow estimation example in Section 6.
$1, \ldots, \tau$. An observation model is represented by the detection probability matrices $B^{s} \in \mathbb{R}^{n \times m}$ where $m$ is the size of the observation space, and $s=1, \ldots, S$ where $S$ denotes the number of uncoupled observations at each time point. Let $\Phi_{t, s} \in \mathbb{R}^{m}$ denote the aggregate measurements at time $t$ for observation $s$. We then estimate the ensemble evolution as the solution of

$$
\begin{array}{cl}
\underset{M_{[1: \tau-1]}, D_{[1: \tau],[1: S]}, \mu_{[1: \tau]}}{\operatorname{minimize}} & \sum_{t=1}^{\tau-1} H\left(M^{t} \mid \operatorname{diag}\left(\mu_{t}\right) A^{t}\right)+\sum_{t=1}^{\tau} \sum_{s=1}^{S} H\left(D^{s, t} \mid \operatorname{diag}\left(\mu_{t}\right) B^{s}\right) \\
\text { subject to } & M^{t} \mathbf{1}=\mu_{t-1}, \quad\left(M^{t}\right)^{T} \mathbf{1}=\mu_{t}, D^{t, s} \mathbf{1}=\mu_{t}, \quad\left(D^{t, s}\right)^{T} \mathbf{1}=\Phi_{t, s}, \\
& \text { for } t=1, \ldots, \tau, \text { and } s=1, \ldots, S .
\end{array}
$$

The tree corresponding to the Markov model (6.1) is sketched in Figure 6(a) This optimization problem is similar to the one in [37], but therein the initial distribution of agents, i.e., the marginal $\mu_{1}$, is assumed to be known.

In the following we consider a number of $N$ agents evolving on the network displayed in Figure 6(b) with $n=100$ nodes and 180 edges. Figure 6(b) also shows the location of the $N_{S}=15$ sensors. In a first step, we compute the discrete Schrödinger bridge (2.11) between the distribution $\mu_{1}$, where all agents are in node 1 , and the distribution $\mu_{\tau}$, where all agents are in node 100 , with the random walk on $\mathcal{G}$ as a prior. The resulting particle evolutions $M^{t}$, for $t=1, \ldots, \tau-1$, in (2.11) define the transition probability matrices $\bar{A}^{t}=\operatorname{diag}\left(\mu_{t}\right)^{-1} M^{t}$, for $t=1, \ldots, \tau-1$, as in (5.3), used to simulate a number of $N$ agents with initial distribution $\mu_{1}$. By construction, the final distribution is then $\mu_{\tau}$.

We consider two observation models, one where the sensors make uncoupled measurements, similar to the example in [37], and one where the sensors are coupled and form a joint measurement. In the uncoupled setting, we consider an observation space of 2 states for each sensor, where one state denotes that an agent is detected, and the other one that the agent is undetected by the sensor. Hence, we define an observation probability matrix $B^{s} \in \mathbb{R}^{n \times 2}$ for each sensor $s=1, \ldots, S=N_{S}$, where the probability for an agent on node $i$ to be detected by the sensor $s$ is defined as $B_{i 1}^{s}=\min \left(0.99,2 e^{-d_{s, i}}\right)$, where $d_{s, i}$ denotes the Euclidean distance between the location of sensor $s$ and the node $i$. Consequently the probability of not being detected is $B_{i 2}^{s}=1-B_{i 1}^{s}$. For the coupled observation model, the observation space consists of all possible sets of sensors that detect a given agent, i.e., all subsets of the set $\left\{1, \ldots, N_{S}\right\}$. The size of the observation space is thus $2^{N_{S}}$, and there is only one


Figure 7. True ensemble flow and estimates with the two observation models for a varying number of agents.
observation at each time instance, i.e., $S=1$. Hence, we define an observation probability matrix $B^{\text {joint }} \in \mathbb{R}^{n \times 2^{S}}$, where the probability for an agent in node $i$ to be detected by exactly the set $\mathfrak{S}$ of sensors is given by $\left(\prod_{s \in \mathfrak{S}} B_{i 1}^{s}\right)\left(\prod_{s \notin \mathfrak{S}} B_{i 2}^{s}\right)$.

We solve the multi-marginal optimal transport problem corresponding to (6.1) with assumed probability transition matrix $A$ describing a random walk on $\mathcal{G}$, and observation probability matrices $B^{s}$ for the two described observation models. The results for both observation models are compared for $N \in\{10,100,1000\}$ agents in Figure 7, where the estimated number of agents in each node is plotted as a circle with size corresponding to the log-scaled weigthed number of agents in that node. In Figure $7(\mathrm{a})$, one can see that for $N=10$ agents the coupled estimate localizes the position of the agents slightly better than the uncoupled estimate. However, this effect decreases with an increasing number of agents, as the evolution of a single agent has less influence on the empirical distribution of agents. Already from a number of $N=100$ agents the uncoupled estimate is competitive with the coupled estimate, albeit relying on significantly less information (cf. Figure 7(b)). For an ensemble of $N=1000$ agents one can hardly see any difference between the two estimates, as shown in Figure 7(c).

Note that the number of observations at each time instance for the uncoupled estimate is only of the size $2 N_{S}=30$, whereas the number of coupled observations is $2^{N_{S}}=32768$. The improved observation model thus comes at an increased computational cost, which is exponential in the number of sensors, and becomes infeasible for larger numbers of sensors.
7. Conclusion. In this work we consider multi-marginal optimal transport problems with cost functions that decouple according to a tree, and we show that the entropy regularized formulation of this problem is equivalent to a generalization of a time and space discrete Schrödinger bridge defined on the same tree. Moreover, we derive an efficient algorithm for solving this problem. We also compare the multi-marginally regularized optimal transport problem to a commonly used pairwise regularized optimal transport problem and illustrate the benefits in theory and practice. Finally, we describe how to apply the framework to the problem of tracking an ensemble of indistinguishable agents.

Interestingly, the construction of the marginals in Theorem 3.2 is of the same form as the belief propagation algorithm [59,64] for inference in graphical models. In this interpretation, the vectors $u_{j}$ correspond to the local evidence in node $j \in \mathcal{V}$ and $\alpha_{\left(j_{1}, j_{2}\right)}$ is the message passed from node $j_{2}$ to node $j_{1}$. Moreover, the objective function in (4.5) can be interpreted as the Bethe free energy [65], which is connected to belief propagation [64]. These similarities of our framework to belief propagation algorithms are investigated in [38], and could be a stepping stone to extending our framework to graphs with cycles. Another direction of interest is the extension to continuous state models.

Appendix A. Proofs. For an index set $\Gamma$, we denote the double sum $\sum_{j \in \Gamma} \sum_{i_{j}}$ by $\sum_{i_{j}: j \in \Gamma}$. For simplicity of notation, we let $\mathcal{M}$ denote the set of matrices $M^{\left(j_{1}, j_{2}\right)}$, for $\left(j_{1}, j_{2}\right) \in \mathcal{E}$, and similarly write $\boldsymbol{\mu}$ and $\boldsymbol{\lambda}$ for the respective sets of optimization variables. The proof of Theorem 3.2 is based on the following lemma.

Lemma A. 1 ([30, Lemma 1 and 2]). Let $\mathbf{U}=u_{1} \otimes u_{2} \otimes \cdots \otimes u_{J}$, for a set of vectors $u_{1}, u_{2}, \ldots, u_{J}$, and $\mathbf{K}$ be a tensor of the same size. If $\langle\mathbf{K}, \mathbf{U}\rangle=w^{T} u_{j}$, for a vector $w$ that does not depend on $u_{j}$, then it holds that $P_{j}(\mathbf{K} \odot \mathbf{U})=w \odot u_{j}$.

Similarly, if $\langle\mathbf{K}, \mathbf{U}\rangle=w^{T} \operatorname{diag}\left(u_{j_{1}}\right) W \operatorname{diag}\left(u_{j_{2}}\right) \hat{w}$, for two vectors $w, \hat{w}$, and a matrix $W$ that does not depend on $u_{j_{1}}$ and $u_{j_{2}}$, then it holds that $P_{j_{1}, j_{2}}(\mathbf{K} \odot \mathbf{U})=$ $\operatorname{diag}\left(w \odot u_{j_{1}}\right) W \operatorname{diag}\left(u_{j_{2}} \odot \hat{w}\right)$.

Proof of Theorem 3.2. Due to the decoupling of the cost tensor $\mathbf{C}$ as in (3.1), the tensor $\mathbf{K}=\exp (-\mathbf{C} / \epsilon)$ decouples according to (3.2) with the matrices $K^{\left(j_{1}, j_{2}\right)}=$ $\exp \left(C^{\left(j_{1}, j_{2}\right)} / \epsilon\right)$, for $\left(j_{1}, j_{2}\right) \in \mathcal{E}$. We can therefore write

$$
\langle\mathbf{K}, \mathbf{U}\rangle=\sum_{i_{j}: j \in \mathcal{V}}\left(\prod_{\left(j_{1}, j_{2}\right) \in \mathcal{E}} K_{i_{j_{1}}, i_{j_{2}}}^{\left(j_{1}, j_{2}\right)}\right)\left(\prod_{k \in \mathcal{V}}\left(u_{k}\right)_{i_{k}}\right)=\sum_{i_{j}}\left(u_{j}\right)_{i_{j}}\left(w_{j}\right)_{i_{j}}
$$

where the vector $w_{j}$ is of the form

$$
\left(w_{j}\right)_{i_{\ell}}=\sum_{i_{k}: k \in \mathcal{V} \backslash j}\left(\prod_{\left(j_{1}, j_{2}\right) \in \mathcal{E}} K_{i_{j_{1}}, i_{j_{2}}}^{\left(j_{1}, j_{2}\right)}\right)\left(\prod_{j \in \mathcal{V}, j \neq \ell}\left(u_{j}\right)_{i_{j}}\right)
$$

Consider the underlying tree to be rooted in node $j$. Then this can be written as $\left(w_{j}\right)_{i_{\ell}}=\prod_{k: k \in \mathcal{N}_{j}}\left(\alpha_{(j, k)}\right)_{i_{j}}$, where $\alpha_{(p(k), k)} \in \mathbb{R}^{n}$ is defined by

$$
\left(\alpha_{(p(k), k)}\right)_{i_{p(k)}}=\sum_{i_{\ell}: \ell \geq k} \prod_{m>k} K_{i_{p(m)}, i_{m}}^{(p(m), m)}\left(u_{m}\right)_{i_{m}}
$$

If $k \in \mathcal{L}$, then

$$
\left(\alpha_{(p(k), k)}\right)_{i_{p(k)}}=\sum_{i_{k}} K_{i_{p(k)}, i_{k}}^{(p(k), k)}\left(u_{k}\right)_{i_{k}}=\left(K^{(p(k), k)} u_{k}\right)_{i_{p(k)}}
$$

Otherwise, it holds

$$
\begin{aligned}
\left(\alpha_{(p(k), k)}\right)_{i_{p(k)}} & =\sum_{i_{k}}\left(K_{i_{p(k)}, i_{k}}^{(p(k), k}\left(u_{k}\right)_{i_{k}} \prod_{\substack{\ell \in \mathcal{N}_{k} \\
\ell \neq p(k)}}\left(\sum_{i_{m}: m \geq \ell} \prod_{j>\ell} K_{i_{p(j)}, i_{j}}^{(p(j), j}\left(u_{j}\right)_{i_{j}}\right)\right) \\
& =\sum_{i_{k}}\left(K_{i_{p(k)}, i_{k}}^{(p(k), k}\left(u_{k}\right)_{i_{k}} \prod_{\substack{\ell \in \mathcal{N}_{k} \\
\ell \neq p(k)}}\left(\alpha_{(p(\ell, \ell)}\right)_{i_{p(\ell)}}\right)
\end{aligned}
$$

This inductively defines the vectors $\alpha_{(j, k)}$ as in (3.3). The expression for the projection follows from Lemma A. 1 with $w_{j}=\bigodot_{k \in \mathcal{N}_{j}} \alpha_{(j, k)}$.

Proof of Proposition 3.4. Note that the optimal solution to (2.4) with cost structured according to the tree $\mathcal{T}=(\mathcal{V}, \mathcal{E})$ is of the form $\mathbf{M}=\mathbf{K} \odot \mathbf{U}$, where $\mathbf{K}$ is defined by (3.2) and $\mathbf{U}$ is defined by (2.7). To prove the first claim, let $k \in \mathcal{V}$ lie on the direct path between $j \in \mathcal{V}$ and $\ell \in \mathcal{V}$. Then straightforward computation implies that

$$
\left(P_{j k \ell}(\mathbf{K} \odot \mathbf{U})\right)_{i_{j}, i_{k}, i_{\ell}}\left(P_{k}(\mathbf{K} \odot \mathbf{U})\right)_{i_{k}}=\left(P_{j k}(\mathbf{K} \odot \mathbf{U})\right)_{i_{j} i_{k}}\left(P_{k \ell}(\mathbf{K} \odot \mathbf{U})\right)_{i_{k}, i_{\ell}}
$$

In particular, for any fixed $i_{k}$, the matrix $\left(P_{j k \ell}(\mathbf{K} \odot \mathbf{U})\right)_{\cdot, i_{k}}$, is of rank 1. Hence, given $P_{k}(\mathbf{K} \odot \mathbf{U})$, the marginals $P_{j}(\mathbf{K} \odot \mathbf{U})$ and $P_{\ell}(\mathbf{K} \odot \mathbf{U})$ have no influence on each other. Thus, the optimal transport problem (2.4) on $\mathcal{T}$ can be decoupled into smaller problems, by cutting $\mathcal{T}$ in $k \in \mathcal{V}$, and solving an optimal transport problem (2.4) on each of the resulting subtrees, where $k \in \Gamma$ and $k \in \mathcal{L}$ for each of the subtrees.

To prove the second claim, recall from the definition of the tensor $\mathbf{U}$ in (2.7) that $u_{j}=\mathbf{1}$ for all $j \in \mathcal{V} \backslash \Gamma$. Thus, for $k$ such that $(k, \ell) \in \mathcal{E}$, the corresponding vector (3.3) is $\alpha_{(k, \ell)}=K^{(k, \ell)} \mathbf{1}$, which is constant and does not need to be updated when recomputing the projections (3.4). It thus suffices to solve (2.4) on the subtree of $\mathcal{T}$ that is obtained by removing $\ell \in \mathcal{V}$ and $(k, \ell) \in \mathcal{E}$ from $\mathcal{T}$.

Proof of Theorem 4.1. Assume that the only edge with node $j=1$ is denoted $(1,2)$. We add the trivial constraints $M^{\left(j_{1}, j_{2}\right)} \mathbf{1}=\left(M^{\left(p\left(j_{1}\right), j_{1}\right)}\right)^{T} \mathbf{1}$ for $\left(j_{1}, j_{2}\right) \in \mathcal{E}_{r} \backslash$ $(1,2)$. We relax this constraint, together with the constraint on $\mu_{1}$, and let $\lambda_{\left(j_{1}, j_{2}\right)}$, for $\left(j_{1}, j_{2}\right) \in \mathcal{E}_{r}$, denote the corresponding dual variables. Furthermore, we relax the constraints $\left(M^{\left(j_{p}, j\right)}\right)^{T} \mathbf{1}=\mu_{j}$ with dual variables $\lambda_{j}$, for the leaves $j \in \mathcal{L} \backslash 1$. A Lagrangian for (4.1) is then

$$
\begin{aligned}
& L(\mathcal{M}, \boldsymbol{\mu}, \boldsymbol{\lambda})=\sum_{\left(j_{1}, j_{2}\right) \in \mathcal{E}} H\left(M^{\left(j_{1}, j_{2}\right)} \mid \operatorname{diag}\left(\mu_{j_{1}}\right) A^{\left(j_{1}, j_{2}\right)}\right)+\lambda_{(1,2)}^{T}\left(M^{(1,2)} \mathbf{1}-\mu_{1}\right) \\
& \quad+\sum_{j_{2} \in \mathcal{L}} \lambda_{j_{2}}^{T}\left(\mu_{j_{2}}-\left(M^{\left(j_{1}, j_{2}\right)}\right)^{T} \mathbf{1}\right)+\sum_{\substack{\left(j_{1}, j_{2}\right) \in \mathcal{E} \\
j_{1} \neq 1}} \lambda_{\left(j_{1}, j_{2}\right)}^{T}\left(M^{\left(j_{1}, j_{2}\right)} \mathbf{1}-\left(M^{\left(j_{1}, j_{1}\right)}\right)^{T} \mathbf{1}\right)
\end{aligned}
$$

When $j_{2}$ is an inner node on the tree, i.e., $j_{2} \notin \mathcal{L}$, the derivative with respect to the entries $M_{i_{1} i_{2}}^{\left(j_{1}, j_{2}\right)}$, for all $i_{1}, i_{2}=1, \ldots, n$, and $\left(j_{1}, j_{2}\right) \in \mathcal{E}_{r}$ is

$$
\begin{equation*}
\log \left(\frac{M_{i_{1} i_{2}}^{\left(j_{1} j_{2}\right)}}{\left(\mu_{j_{1}}\right)_{i_{1}} A_{i_{1} i_{2}}^{\left(j_{1}, j_{2}\right)}}\right)+\left(\lambda_{\left(j_{1}, j_{2}\right)}\right)_{i_{2}}-\sum_{k:\left(j_{2}, k\right) \in \mathcal{E}}\left(\lambda_{\left(j_{2}, k\right)}\right)_{i_{2}} \tag{A.1}
\end{equation*}
$$

Since (4.1) is convex, a mass transport plan is optimal if this gradient vanishes, yielding the expression in terms of the other variables,

$$
\begin{equation*}
M^{\left(j_{1}, j_{2}\right)}=\operatorname{diag}\left(\mu_{j_{1}} \cdot / v_{\left(j_{1}, j_{2}\right)}\right) A^{\left(j_{1}, j_{2}\right)} \operatorname{diag}\left(\bigodot_{k:\left(j_{2}, k\right) \in \mathcal{E}} v_{\left(j_{2}, k\right)}\right) \tag{A.2}
\end{equation*}
$$

where $v_{\left(j_{1}, j_{2}\right)}=\exp \left(\lambda_{\left(j_{1}, j_{2}\right)}\right)$ for all $\left(j_{1}, j_{2}\right) \in \mathcal{E}_{r}$. In case $j_{2} \in \mathcal{L}$, the last sum in the derivative (A.1) is replaced by $\left(\lambda_{j_{2}}\right)_{l}$. Defining $v_{j_{2}}=\exp \left(\lambda_{j_{2}}\right)$, the optimal mass transport plan is thus of the form

$$
\begin{equation*}
M^{\left(j_{1}, j_{2}\right)}=\operatorname{diag}\left(\mu_{j_{1}} \cdot / v_{\left(j_{1}, j_{2}\right)}\right) A^{\left(j_{1}, j_{2}\right)} \operatorname{diag}\left(v_{j_{2}}\right) \tag{A.3}
\end{equation*}
$$

Next, note that the marginal of the optimal transport plan satisfies

$$
\begin{align*}
\mu_{j_{2}} & =M_{\left(j_{1}, j_{2}\right)}^{T} \mathbf{1}=\hat{\varphi}_{j_{2}} \odot \varphi_{j_{2}}  \tag{A.4}\\
\text { with } \quad \varphi_{j_{2}} & =\bigodot_{k:\left(j_{2}, k\right) \in \mathcal{E}} v_{\left(j_{2}, k\right)}, \quad \text { and } \quad \hat{\varphi}_{j_{2}}=\left(A^{\left(j_{1}, j_{2}\right)}\right)^{T}\left(\mu_{j_{1}} \cdot / v_{\left(j_{1}, j_{2}\right)}\right) . \tag{A.5}
\end{align*}
$$

Since for all $k$ such that $\left(j_{2}, k\right) \in \mathcal{E}_{r}$ it holds that $M_{\left(j_{2}, k\right)} \mathbf{1}=\mu_{j_{2}}$, we get

$$
v_{\left(j_{2}, k\right)}=A^{\left(j_{2}, k\right)}\left(\bigodot_{\ell:(k, \ell) \in \mathcal{E}} v_{(k, \ell)}\right)=A^{\left(j_{2}, k\right)} \varphi_{k}
$$

which completes the definition of the vectors $\left(\varphi_{j}\right)_{j \in \mathcal{V}}$. Similarly to (A.4) it holds that

$$
\begin{equation*}
\mu_{j_{1}}=\left(\bigodot_{k:\left(j_{1}, k\right) \in \mathcal{E}} v_{\left(j_{1}, k\right)}\right) \odot A\left(\mu_{j_{p(1)}} \cdot / v_{\left(j_{p(1)}, j_{1}\right)}\right) . \tag{A.6}
\end{equation*}
$$

Plugging (A.6) into $\hat{\varphi}_{j_{2}}$ in (A.5), yields the recursive definition of the vectors $\left(\hat{\varphi}_{j}\right)_{j \in \mathcal{V}}$. The expression for the mass transport plans $M^{\left(j_{1}, j_{2}\right)}$, for $\left(j_{1}, j_{2}\right) \in \mathcal{E}$, in terms of the vectors $\hat{\varphi}_{j_{1}}, \varphi_{j_{2}}$ and $\varphi_{j_{1} \backslash j_{2}}$ follows by identifying them in expressions (A.2) and (A.3).प

Proof of Proposition 4.2. For any matrices $M, A \in \mathbb{R}_{+}^{n \times n}$ and $\mu \in \mathbb{R}_{+}^{n}$ one can write the term $H(M \mid \operatorname{diag}(\mu) A)$ as

$$
\begin{equation*}
\sum_{i, j=1}^{n}\left(M_{i j} \log \left(\frac{M_{i j}}{A_{i j}}\right)-M_{i j}\right)+\sum_{i, j=1}^{n}\left(\mu_{i} A_{i j}-M_{i j} \log \left(\mu_{i}\right)\right) \tag{A.7}
\end{equation*}
$$

Since $M \mathbf{1}=\mu$ and $A \mathbf{1}=\mathbf{1}$, the second sum can be simplified to $\sum_{i=1}^{n}\left(\mu_{i}-\mu_{i} \log \left(\mu_{i}\right)\right)$. Furthermore, adding the term $\sum_{i, j=1}^{n} A_{i j}-\sum_{i=1}^{n} 1=0$ to (A.7), the expression can be written as $H(M \mid A)-H(\mu)$. Due to the underlying tree structure of problem (4.1), the number of outgoing edges from the root node $j_{r}$ is $\operatorname{deg}\left(j_{r}\right)$, and for all other vertices $j \in \mathcal{V} \backslash\left\{j_{r}\right\}$ the number of outgoing edges is $\operatorname{deg}(j)-1$. In the case that $j_{r} \in \mathcal{L}$, since the marginal $\mu_{j_{r}}$ is known, the term $H\left(\mu_{j_{r}}\right)$ is constant and can be removed from the objective without changing the optimal solution.

Proof of Theorem 4.4. Note that for a rooted directed tree, in each node $j \in \mathcal{V} \backslash \mathcal{L}$, there is one incoming edge and the rest of the connected edges are outgoing. Its neighbouring nodes are therefore given by the set $\mathcal{N}_{j}=p(j) \cup\left\{k:(j, k) \in \mathcal{E}_{r}\right\}$. Thus,

$$
\begin{equation*}
\bigodot_{k \in \mathcal{N}_{j}} \alpha_{(j, k)}=\left(K^{(p(j), j)}\right)^{T} \alpha_{(j, p(j))} \odot \bigodot_{k:(j, k) \in \mathcal{E}_{r}} K^{(j, k)} \alpha_{(j, k)} \tag{A.8}
\end{equation*}
$$

For any edge $(j, k) \in \mathcal{E}$ and for the reverse edge $\left(j, j_{p}\right)$ it holds that

$$
\begin{aligned}
& \alpha_{(j, k)}=\bigodot_{\ell \in \mathcal{N}_{k} \backslash\{j\}} K^{(k, \ell)} \alpha_{(k, \ell)}=\bigodot_{\ell:(k, \ell) \in \mathcal{E}} K^{(k, \ell)} \alpha_{(k, \ell)} \\
& \alpha_{(j, p(j))}=\bigodot_{\ell \in \mathcal{N}_{p(j) \backslash\{j\}}} K^{(p(j), \ell)} \alpha_{(j, \ell)}=\left(K^{(p(p(j)), p(j))}\right)^{T} \alpha_{(p(j), p(p(j))} \bigodot_{\substack{\ell:(p(j), \ell) \in \mathcal{E}, \ell \neq j}} K^{(p(j), \ell)} \alpha_{(p(j), \ell)} .
\end{aligned}
$$

By associating the first term in (A.8) with $\hat{\varphi}_{j}$ and the second term with $\varphi_{j}$, we see that the tensor structure of $\mathbf{K}$ gives rise to the construction of $\varphi_{j}$ and $\hat{\varphi}_{j}$ as in Theorem 4.1.

Hence, we can define a tensor in analogy to $\mathbf{U}$ of the form $\mathbf{V}=\left(\mathbf{1} . / v_{1}\right) \otimes v_{2} \otimes \cdots \otimes v_{J}$, where $v_{j}$ are the vectors from Theorem 4.1 if $j \in \mathcal{L}$, and $v_{j}=\mathbf{1}$ if $j \in \mathcal{V} \backslash \mathcal{L}$. Then, the tensor $\mathbf{K} \odot \mathbf{V}$ has the same marginals as the tensor $\mathbf{K} \odot \mathbf{U}$. Due to an extension of Sinkhorn's theorem to tensors [32], it follows $\mathbf{U}=\mathbf{V}$ (up to scaling with a factor and its inverse in the vectors $\left.u_{1}, \ldots, u_{J}\right)$, and $P_{j}(\mathbf{M})=\mu_{j}$ for all $j \in \mathcal{V}$.

Proof of Proposition 4.5. Based on the proof to Theorem 4.1, a Lagrange dual to problem (4.1) can be formulated as to maximize

$$
\begin{align*}
& -\sum_{\left(j_{1}, j_{2}\right) \in \mathcal{E}}\left(\mu_{j_{1}} \odot \exp \left(-\lambda_{\left(j_{1}, j_{2}\right)}\right)\right) A^{\left(j_{1}, j_{2}\right)}\left(\bigodot_{k:\left(j_{2}, k\right) \in \mathcal{E}} \exp \left(\lambda_{\left(j_{2}, k\right)}\right)\right)  \tag{A.9}\\
& -\sum_{j \in \mathcal{L}}\left(\mu_{j_{p}} \odot \exp \left(-\lambda_{\left(j_{p}, j\right)}\right)\right) A^{\left(j_{1}, j_{2}\right)}\left(\exp \left(\lambda_{j}\right)\right)-\lambda_{(1,2)}^{T} \mu_{1}+\sum_{j \in \mathcal{L}} \lambda_{j}^{T} \mu_{j}
\end{align*}
$$

with respect to $\mu_{j}$, for all inner nodes $j$, and the dual variables $\lambda_{\left(j_{1}, j_{2}\right)}$, for $\left(j_{1}, j_{2}\right) \in \mathcal{E}$, and $\lambda_{j}$, for $j \in \mathcal{L}$. A block coordinate ascent in the dual is then to iteratively maximize (A.9) with respect to one of the dual variable vectors, while keeping the others fixed. Denote $v_{\left(j_{1}, j_{2}\right)}=\exp \left(\lambda_{\left(j_{1}, j_{2}\right)}\right)$, for $\left(j_{1}, j_{2}\right) \in \mathcal{E}$, and $v_{j}=\exp \left(\lambda_{j}\right)$, for $j \in \mathcal{L}$. The gradient of (A.9) with respect to $\lambda_{(1,2)}$ vanishes if it holds that

$$
v_{(1,2)}=A^{\left(j_{1}, j_{2}\right)}\left(\bigodot_{k:(2, k) \in \mathcal{E}} v_{(2, k)}\right)
$$

and the gradient with respect to $\lambda_{j}$ vanishes if $v_{j}=\mu_{j} \cdot /\left(\left(A^{\left(j_{p}, j\right)}\right)^{T}\left(\mu_{j_{p}} \cdot / v_{\left(j_{p}, j\right)}\right)\right)$, for $j \in \mathcal{L}$, and where $j_{p} \in \mathcal{V}$ is the parent of node $j$. Finally, the gradient of (A.9) with respect to $\lambda_{\left(j_{1}, j_{2}\right)}$, where $j_{1} \neq 1$, vanishes if $M_{\left(j_{1}, j_{2}\right)} \mathbf{1}=M_{\left(j_{p}, j_{1}\right)}^{T} \mathbf{1}$, where $j_{p}$ is the parent of node $j_{1}$, which leads to the recursive definition of $\varphi_{j}$ and $\hat{\varphi}_{j}$, for $j \in \mathcal{V}$, as in Theorem 4.1. Hence, given an initial set of positive vectors $v_{(1,2)}$ and $v_{j}$, for $j \in \mathcal{L}$, the scheme (4.6) is a block coordinate ascend in a Lagrange dual of problem (4.1).

## REFERENCES

[1] B. Acciaio, J. Backhoff-Veraguas, and R. Carmona, Extended mean field control problems: stochastic maximum principle and transport perspective, SIAM Journal on Control and Optimization, 57 (2019), pp. 3666-3693.
[2] J. Adler, A. Ringh, O. Öktem, and J. Karlsson, Learning to solve inverse problems using Wasserstein loss, arXiv preprint arXiv:1710.10898, (2017).
[3] M. Agueh and G. Carlier, Barycenters in the Wasserstein space, SIAM Journal on Mathematical Analysis, 43 (2011), pp. 904-924.
[4] M. Arjovsky, S. Chintala, and L. Bottou, Wasserstein Generative Adversarial Networks, in Proceedings of the 34th International Conference on Machine Learning, 2017, pp. 214223.
[5] H. Bauschke and A. Lewis, Dykstras algorithm with Bregman projections: A convergence proof, Optimization, 48 (2000), pp. 409-427.
[6] E. Bayraktar, A. M. Cox, and Y. Stoev, Martingale optimal transport with stopping, SIAM Journal on Control and Optimization, 56 (2018), pp. 417-433.
[7] J.-D. Benamou, G. Carlier, M. Cuturi, L. Nenna, and G. Peyré, Iterative Bregman Projections for Regularized Transportation Problems, SIAM Journal on Scientific Computing, 37 (2015), p. A1111-A1138.
[8] N. Bonneel, G. Peyré, and M. Cuturi, Wasserstein barycentric coordinates: histogram regression using optimal transport., ACM Trans. Graph., 35 (2016), pp. 71-1.
[9] N. Bonneel, M. van de Panne, S. Paris, and W. Heidrich, Displacement interpolation using lagrangian mass transport, ACM Transactions on Graphics (TOG), 30 (2011), pp. 158:1158:12.
[10] X. ChEn, Structure theory for ensemble controllability, observability, and duality, Mathematics of Control, Signals, and Systems, 31 (2019), p. 7.
[11] Y. Chen, G. Conforti, and T. Georgiou, Measure-valued spline curves: An optimal transport viewpoint, SIAM Journal on Mathematical Analysis, 50 (2018), pp. 5947-5968.
[12] Y. Chen, G. Conforti, T. Georgiou, and L. Ripani, Multi-marginal Schrodinger bridges, arXiv preprint arXiv:1902.08319, (2019).
[13] Y. Chen, T. Georgiou, and M. Pavon, Entropic and Displacement Interpolation: A Computational Approach Using the Hilbert Metric, SIAM Journal on Applied Mathematics, 76 (2016), pp. 2375-2396.
[14] Y. Chen, T. Georgiou, and M. Pavon, On the relation between optimal transport and Schrödinger bridges: A stochastic control viewpoint, Journal of Optimization Theory and Applications, 169 (2016), pp. 671-691.
[15] 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, 61 (2016), pp. 1158-1169.
[16] Y. Chen, T. Georgiou, M. Pavon, and A. Tannenbaum, Robust transport over networks, IEEE transactions on automatic control, 62 (2016), pp. 4675-4682.
[17] Y. Chen and J. Karlsson, State tracking of linear ensembles via optimal mass transport, IEEE Control Systems Letters, 2 (2018), pp. 260-265.
[18] S.-N. Chow, W. Huang, Y. Li, and H. Zhou, Fokker-planck equations for a free energy functional or Markov process on a graph, Archive for Rational Mechanics and Analysis, 203 (2012), pp. 969-1008.
[19] G. Conforti and C. Léonard, Reciprocal classes of random walks on graphs, Stochastic Processes and their Applications, 127 (2017), pp. 1870-1896.
[20] T. Cover and J. Thomas, Elements of information theory, John Wiley \& Sons, Hoboken, NJ, 2006.
[21] I. Csiszar, Why least squares and maximum entropy? An axiomatic approach to inference for linear inverse problems, The Annals of Statistics, 19 (1991), pp. 2032-2066.
[22] 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.
[23] M. Cuturi and A. Doucet, Fast computation of Wasserstein barycenters, in International Conference on Machine Learning, 2014, pp. 685-693.
[24] D. Dawson, L. Gorostiza, and A. Wakolbinger, Schrödinger processes and large deviations, Journal of mathematical physics, 31 (1990), pp. 2385-2388.
[25] A. Dembo and O. Zeitouni, Large deviations techniques and applications, Springer-Verlag Berlin Heidelberg, 2009.
[26] R. Diestel, Graph Theory, Springer, Berlin, Heidelberg, 2010.
[27] A. Dominitz and A. Tannenbaum, Texture Mapping via Optimal Mass Transport, IEEE Transactions on Visualization and Computer Graphics, 16 (2010), pp. 419-433.
[28] R. Ellis, Entropy, Large Deviations, and Statistical Mechanics, Springer-Verlag Berlin Heidelberg, 2006.
[29] F. Elvander, I. Haasler, A. Jakobsson, and J. Karlsson, Tracking and Sensor Fusion in Direction of Arrival Estimation Using Optimal Mass Transport, in 2018 26th European Signal Processing Conference (EUSIPCO), IEEE, 2018, pp. 1617-1621.
[30] 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).
[31] H. Föllmer, Random fields and diffusion processes, in École d'Été de Probabilités de SaintFlour XV-XVII, 1985-87, P.-L. Hennequin, ed., vol. 1362 of Lecture Notes in Mathematics, Springer, Berlin, Heidelberg, 1988, pp. 101-203.
[32] J. Franklin and J. Lorenz, On the scaling of multidimensional matrices, Linear Algebra and its Applications, 114-115 (1989), pp. $717-735$. Special Issue Dedicated to Alan J. Hoffman.
[33] W. Gangbo and A. ŚWikech, Optimal maps for the multidimensional monge-kantorovich problem, Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences, 51 (1998), pp. 23-45.
[34] T. Georgiou and M. Pavon, Positive contraction mappings for classical and quantum Schrödinger systems, Journal of Mathematical Physics, 56 (2015), p. 033301.
[35] N. Ghoussoub, Y.-H. Kim, and A. Z. Palmer, Optimal transport with controlled dynamics and free end times, SIAM Journal on Control and Optimization, 56 (2018), pp. 3239-3259.
[36] H. Gzyl, The Method of Maximum Entropy, World Scientific, 1995.
[37] 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.
[38] I. HaAsler, R. Singh, Q. Zhang, J. Karlsson, and Y. Chen, Multi-marginal optimal transport and probabilistic graphical models, arXiv preprint arXiv:2006.14113, (2020).
[39] B. Jamison, Reciprocal processes, Probability Theory and Related Fields, 30 (1974), pp. 65-86.
[40] B. Jamison, The markov processes of schrödinger, Probability Theory and Related Fields, 32 (1975), pp. 323-331.
[41] 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).
[42] S. Kolouri, S. R. Park, M. Thorpe, D. Slepcev, and G. K. Rohde, Optimal Mass Transport: Signal processing and machine-learning applications, IEEE Signal Process. Mag., 34 (2017), pp. 43-59.
[43] A. Kroshnin, N. Tupitsa, D. Dvinskikh, P. Dvurechensky, A. Gasnikov, and C. Uribe, On the complexity of approximating Wasserstein barycenters, in International conference on machine learning, 2019, pp. 3530-3540.
[44] C. LÉonard, From the Schrödinger problem to the Monge-Kantorovich problem, Journal of Functional Analysis, 262 (2012), pp. 1879-1920.
[45] T. Lin, N. Ho, X. Chen, M. Cuturi, and M. Jordan, Revisiting fixed support Wasserstein barycenter: Computational hardness and efficient algorithms, arXiv preprint arXiv:2002.04783, (2020).
[46] T. Lin, N. Ho, M. Cuturi, and M. Jordan, On the complexity of approximating multimarginal optimal transport, arXiv preprint arXiv:1910.00152, (2019).
[47] Z.-Q. Luo and P. Tseng, On the convergence of the coordinate descent method for convex differentiable minimization, Journal of Optimization Theory and Applications, 72 (1992), pp. 7-35.
[48] C. LÉOnard, A survey of the Schrödinger problem and some of its connections with optimal transport, Discrete \& Continuous Dynamical Systems - A, 34 (2014), pp. 1533-1574.
[49] T. Mikami, Monge's problem with a quadratic cost by the zero-noise limit of h-path processes, Probability Theory and Related Fields, 129 (2004), pp. 245-260.
[50] T. Mikami and M. Thieullen, Optimal transportation problem by stochastic optimal control, SIAM Journal on Control and Optimization, 47 (2008), pp. 1127-1139.
[51] G. Montavon, K.-R. Müller, and M. Cuturi, Wasserstein training of restricted boltzmann machines, in Advances in Neural Information Processing Systems, 2016, pp. 3718-3726.
[52] B. PASs, Multi-marginal optimal transport: theory and applications, ESAIM: Mathematical Modelling and Numerical Analysis, 49 (2015), pp. 1771-1790.
[53] M. Pavon and F. Ticozzi, Discrete-time classical and quantum Markovian evolutions: Maximum entropy problems on path space, Journal of Mathematical Physics, 51 (2010), p. 042104.
[54] L. RÜSCHENDORF, Optimal solutions of multivariate coupling problems, Applicationes Mathematicae, 23 (1995), pp. 325-338.
[55] L. RüSchendorf and L. Uckelmann, On the n-coupling problem, Journal of multivariate analysis, 81 (2002), pp. 242-258.
[56] E. Schrödinger, Über die Umkehrung der Naturgesetze, Sitzungsberichte der Preussischen Akademie der Wissenschaften, Physikalisch-mathematische Klasse, (1931), pp. 144-153.
[57] R. Sinkhorn, Diagonal Equivalence to Matrices with Prescribed Row and Column Sums, The American Mathematical Monthly, 74 (1967), pp. 402-405.
[58] J. Solomon, F. de Goes, G. Peyré, M. Cuturi, A. Butscher, A. Nguyen, T. Du, and L. Guibas, Convolutional Wasserstein distances: efficient optimal transportation on geometric domains, ACM Transactions on Graphics (TOG), 34 (2015), pp. 66:1-66:11.
[59] Y. Teh and M. Welling, The unified propagation and scaling algorithm, Advances in neural information processing systems, (2002), pp. 953-960.
[60] P. Tseng, Dual ascent methods for problems with strictly convex costs and linear constraints: A unified approach, SIAM Journal on Control and Optimization, 28 (1990), pp. 214-242.
[61] C. Villani, Optimal transport: Old and new, Springer, Berlin Heidelberg, 2008.
[62] I. G. Vladimirov and P. I. R., State distributions and minimum relative entropy noise sequences in uncertain stochastic systems: The discrete-time case, SIAM Journal on Control and Optimization, 53(3) (2015), p. 1107-1153.
[63] I. Yang, A Convex Optimization Approach to Distributionally Robust Markov Decision Processes With Wasserstein Distance, IEEE Control Systems Letters, 1 (2017), pp. 2475-1456.
[64] J. S. Yedidia, W. T. Freeman, and Y. Weiss, Understanding belief propagation and its generalizations, in Exploring Artificial Intelligence in the New Millennium, Morgan Kaufmann

Publishers Inc., San Francisco, CA, USA, 2003, pp. 239-269.
[65] J. S. Yedidia, W. T. Freeman, and Y. Weiss, Constructing free-energy approximations and generalized belief propagation algorithms, IEEE Transactions on Information Theory, 51 (2005), pp. 2282-2312.
[66] S. ZENG, Sample-based population observers, Automatica, 101 (2019), pp. 166-174.

## Appendix B. Supplementary material.

Proof of Proposition 3.3. Let $r \in \mathcal{L}$ and assume that $j_{1}$ lies on the path from $r$ to $j_{L}$ (note that $r=j_{1}$ if $j_{1} \in \mathcal{L}$ ). For the pairwise marginals we would like to write the inner product as in Lemma A.1:

$$
\begin{aligned}
\langle\mathbf{K}, \mathbf{U}\rangle & =\sum_{i_{j}: j \in \mathcal{V}}\left(\prod_{\left(k_{1}, k_{2}\right) \in \mathcal{E}_{r}} K_{i_{k_{1}}, i_{k_{2}}}^{\left(k_{1}, k_{2}\right)}\right)\left(\prod_{j \in \mathcal{V}}\left(u_{j}\right)_{i_{j}}\right) \\
& =\sum_{i_{j_{1}}, i_{j_{L}}}\left(w_{j_{1}}\right)_{i_{j_{1}}}\left(u_{j_{1}}\right)_{i_{j_{1}}} W_{i_{j_{1}} i_{j_{L}}}\left(u_{j_{L}}\right)_{i_{j_{L}}}\left(\hat{w}_{j_{L}}\right)_{i_{j_{L}}} .
\end{aligned}
$$

Note that the nodes $\mathcal{V} \backslash\left\{j_{1}, j_{L}\right\}$ can be partitioned into three sets, which are separated by the nodes $j_{1}$ and $j_{L}:\left\{j \in \mathcal{V}: j \nsupseteq j_{2}, j \neq j_{1}\right\},\left\{j \in \mathcal{V}: j>j_{1}, j \nsupseteq j_{L}\right\}$ and $\left\{j \in \mathcal{V}: j>j_{L}\right\}$. Then $w_{j_{1}}$ corresponds to the contribution from the first set

$$
\left(w_{j_{1}}\right)_{i_{j_{1}}}=\sum_{i_{k}: k \nsupseteq j_{2}, k \neq j_{1}}\left(\prod_{\left(k_{1}, k_{2}\right) \in \mathcal{E}_{r} k_{2} \ngtr j_{2}} K_{i_{k_{1}}, i_{k_{2}}}^{\left(k_{1}, k_{2}\right)}\right)\left(\prod_{\substack{k \in \mathcal{V} \\ k \nsupseteq j_{2}, k \neq j_{1}}}\left(u_{k}\right)_{i_{k}}\right)=\prod_{k \in \mathcal{N}_{j_{1}} \backslash j_{2}}\left(\alpha_{\left(j_{1}, k\right)}\right)_{i_{j_{1}}},
$$

and similarly $\hat{w}_{j_{L}}$ is the contribution from the third set

$$
\left(\hat{w}_{j_{L}}\right)_{i_{j_{L}}}=\sum_{i_{k}: k>j_{L}}\left(\prod_{\substack{\left(k_{1}, k_{2}\right) \in \mathcal{E} \\ k_{2}>j_{L}}} K_{i_{k_{1}}, i_{k_{2}}}^{\left(k_{1}, k_{2}\right)}\right)\left(\prod_{\substack{k \in \mathcal{V} \\ k>j_{L}}}\left(u_{k}\right)_{i_{k}}\right)=\prod_{k \in \mathcal{N}_{j_{L}} \backslash j_{L-1}}\left(\alpha_{\left(j_{L}, k\right)}\right)_{i_{j_{L}}}
$$

The matrix $W$ is then obtained by summing over all the indices corresponding to the second index set $\left\{i_{k}: k>j_{1}, k \nsupseteq j_{L}\right\}$

$$
\begin{aligned}
W_{i_{j_{1}} i_{j_{L}}} & =\sum_{i_{k}: k>j_{1}, k \nexists_{j_{L}}} \prod_{\substack{\left(k_{1}, k_{2}\right) \in \mathcal{E} \\
k_{2}>j_{1}, k_{1} \ngtr j_{L}}} K_{i_{k_{1}}, i_{k_{2}}}^{\left(k_{1}, k_{2}\right)} \prod_{\substack{k \in \mathcal{V} \\
k>j_{1}, k \neq j_{L}}}\left(u_{k}\right)_{i_{k}} \\
& =\sum_{i_{j_{2}}, \ldots, i_{j_{L-1}}}\left(\prod_{\ell=2}^{L-1}\left(K_{i_{j_{\ell-1}}, i_{j_{\ell}}}^{\left(j_{\ell-1}, j_{\ell}\right)}\left(u_{j_{\ell}}\right)_{i_{j_{\ell}}} \prod_{\substack{k \in \mathcal{N}_{j_{\ell}} \\
k \neq j_{\ell-1}, j_{\ell+1}}}\left(\alpha_{\left(j_{\ell}, k\right)}\right)_{i_{j_{\ell}}}\right) K_{i_{j_{L-1},}, i_{j_{L}}}^{\left(j_{L-1}, j_{L}\right)}\right) \\
& =\left(\left(\prod_{\ell=2}^{L-1} K^{\left(j_{\ell-1}, j_{\ell}\right)} \operatorname{diag}\left(u_{j_{\ell}} \odot \prod_{\substack{k \in \mathcal{N}_{j_{\ell}} \\
k \neq j_{\ell-1}, j_{\ell+1}}}\left(\alpha_{\left(j_{\ell}, k\right)}\right)\right)\right) K^{\left(j_{L-1}, j_{L}\right)}\right)_{i_{j_{1}, i_{j_{L}}}} .
\end{aligned}
$$

Applying Lemma A.1, we get the projection on the marginals $j_{1}$ and $j_{L}$.
Proof of Corollary 4.3. Let the set of matrices $M^{\left(j_{1}, j_{2}\right)}$, for $\left(j_{1}, j_{2}\right) \in \mathcal{E}_{r}$, and vectors $\mu_{j}$, for $j \in \mathcal{V} \backslash \mathcal{L}$, be the solution to (4.1) on the rooted directed tree $\mathcal{T}_{r}=$ $\left(\mathcal{V}, \mathcal{E}_{r}\right)$. It is thus also the solution to (4.4) on the same tree. Changing the root to another vertex $\hat{r} \in \mathcal{L}$, requires switching the direction of all edges on the direct path between $r$ and $\hat{r}$. This is done by replacing the respective transition probability matrix $A^{\left(j_{1}, j_{2}\right)}$ by the reverse transition probability matrix $A^{\left(j_{2}, j_{1}\right)}$, which is given
by $A^{\left(j_{2}, j_{1}\right)}=\operatorname{diag}\left(\mathbf{1} \cdot / a_{j_{2}}\right)\left(A^{\left(j_{1}, j_{2}\right)}\right)^{T} \operatorname{diag}\left(a_{j_{1}}\right)$, for two given vectors $a_{j_{1}}$ and $a_{j_{2}}$ [53]. Note that

$$
\begin{aligned}
& H\left(M^{\left(j_{2}, j_{1}\right)} \mid A^{\left(j_{2}, j_{1}\right)}\right)=H\left(M^{\left(j_{2}, j_{1}\right)} \mid \operatorname{diag}\left(\mathbf{1} \cdot / a_{j_{1}}\right)\left(A^{\left(j_{1}, j_{2}\right)}\right)^{T} \operatorname{diag}\left(a_{j_{2}}\right)\right) \\
&=H\left(\left(M^{\left(j_{1}, j_{2}\right)}\right)^{T} \mid A^{\left(j_{1}, j_{2}\right)}\right)+H\left(\mu_{j_{1}} \mid \operatorname{diag}\left(\mathbf{1} \cdot / a_{j_{1}}\right)\right)+H\left(\mu_{j_{2}} \mid \operatorname{diag}\left(a_{j_{2}}\right)\right.
\end{aligned}
$$

Summing over these terms for all edges on the path between $r$ and $\hat{r}$, the last two terms cancel for all inner nodes $j_{1}, j_{2} \in \mathcal{V} \backslash \Gamma$, and are constants for $j_{1}, j_{2} \in\{r, \hat{r}\}$. From problem (4.4) we thus see that the optimal distributions $\mu_{j}$, for $j \in \mathcal{V}$, are unchanged, and for the transport plans on the reversed edges it holds $M^{\left(j_{2}, j_{1}\right)}=\left(M^{\left(j_{1}, j_{2}\right)}\right)^{T}$.


[^0]:    *This work was supported by the Swedish Research Council (VR), grant 2014-5870, KTH Digital Futures, SJTU-KTH cooperation grant, the NSF under grant 1901599 and 1942523, and Knut and Alice Wallenberg foundation under grant KAW 2018.0349.
    ${ }^{\dagger}$ Division of Optimization and Systems Theory, Department of Mathematics, KTH Royal Institute of Technology, Stockholm, Sweden. haasler@kth.se, johan.karlsson@math.kth.se
    $\ddagger$ Department of Electronic and Computer Engineering, The Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong, China. eeringh@ust.hk
    ${ }^{\S}$ School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA, USA. yongchen@gatech.edu

[^1]:    ${ }^{1}$ Recall that the definition of $H\left(\mu_{j}\right)$ essentially corresponds to the negative of the entropy of $\mu_{j}$.

