

Application of Wasserstein Attraction Flows for Optimal Transport in Network Systems

Ferran Arqué, César A. Uribe, Carlos Ocampo-Martinez

Abstract—This paper presents a Wasserstein attraction approach for solving dynamic mass transport problems over networks. In the transport problem over networks, we start with a distribution over the set of nodes that needs to be “transported” to a target distribution accounting for the network topology. We exploit the specific structure of the problem, characterized by the computation of implicit gradient steps, and formulate an approach based on discretized flows. As a result, our proposed algorithm relies on the iterative computation of constrained Wasserstein barycenters. We show how the proposed method finds approximate solutions to the network transport problem, taking into account the topology of the network, the capacity of the communication channels, and the capacity of the individual nodes.

I. INTRODUCTION

Optimal transport (OT) theory has experienced increased interest over the last few years due to its wide range of applications in both theoretical and applied fields of mathematics [21]. In particular, the recent efforts to overcome the high computational cost of the associated linear programming problem [6] have made OT an attractive choice to tackle problems involving a large number of distributions or other high dimensional objects and requiring high accuracy.

Our work focuses on the discrete OT problem, where probability distributions are defined over the nodes of a finite graph. In traditional OT approaches, it is assumed that mass (or a fraction of it) at each point in the support of one of the probability measures can be sent to any of the elements in the support of the other probability measure. As a result, the transport plan is executed effectively in one step. However, we seek to explicitly consider the topology of the underlying graph, which naturally imposes some transportation constraints. This means that there may not be a direct link between two points in the support, and additionally, one may need to account for channel and node capacities. Thus, our goal is to find a sequence of transport plans that move the mass from an initial distribution to a final one along the edges of a connected graph so that the cost of transportation is minimal and the capacity constraints are satisfied.

Finding the amount of mass that needs to be sent through each edge to minimize the total cost of transportation is a well-known problem called the *minimum-cost flow problem*

(MCFP). This problem has been widely studied [1], [13], and different algorithms have been proposed to solve it [14]. More importantly, the Wasserstein distance can be rewritten as a MCFP when considering a complete bipartite graph [3] and this can be extended to more general graphs if one considers the shortest path distance as the cost to send a unit of resource from one node to the other. However, classical methods to solve this problem do not have a condition to discern between paths when the optimal flow is not unique, which leads to unpredictability of the output from the solver [11]. To avoid that case, some algorithms introduce an additional term to the objective function so that it becomes strongly convex. These regularized OT methods, like the well-known Sinkhorn algorithm [6], achieve uniqueness and significantly speed up the computation, compared to solving a large linear programming problem. However, it is at the cost of finding an approximate solution to the original problem.

Our approach is based on the resolution of the *Wasserstein attraction* (WA) problem [17], which requires the computation of a *Wasserstein barycenter* (WB) of two distributions at every iteration. Computing the WB yields an intermediate distribution, defined as the Fréchet mean of the two measures, which is the result of minimizing the sum of the (Wasserstein) distances between itself and each of the two distributions [7]. However, the support of this resulting distribution can include any of the graph nodes. We expand the definition of the WB problem by adding constraints that ensure the mean obtained has the appropriate support and each node does not receive more mass than the amount available from its neighbors. This approach resembles what is called *displacement interpolation* [21]. However, displacement interpolation in the discrete-time case may require a small step size of the weight to avoid some of the mass moving over more than one node in a single step, leading to many more iterations than necessary. Furthermore, certain nodes can receive more mass than the total obtainable from their neighboring nodes. In summary, the main differentiating factor between displacement interpolation and our proposal is the addition of the topology and capacity constraints imposed by a graph. In this regard, [12] recently studied this problem in the context of traffic planning, where edge capacity constraints are taken into account, and proposed a framework based on the Lagrangian dual problem to solve it, which resembles the Sinkhorn algorithm.

Moreover, our proposed approach can be reformulated as a discrete gradient flow problem. Several papers work on discrete gradient flows over graphs (or other discrete

F. Arqué and C. Ocampo-Martinez are with the Automatic Control Department at Universitat Politècnica de Catalunya, Barcelona, Spain. C.A. Uribe is with the Department of Electrical and Computer Engineering at Rice University, Houston TX. Corresponding author: Ferran Arqué (ferran.arque@upc.edu). This paper has been partially supported by the L-BEST project (Ref. PID2020-115905RB-C21) from the Spanish Ministry of Science and Innovation.

domains) [5], [10], [15], [19]. However, such papers focus on the theoretical analysis of differential equations rather than the computational aspect with the regularized approximation of the Wasserstein metric (except for [10] which provides a more in-depth discussion on the topic), and no additional constraints are considered on the elements of the graph. The closest works to our setting with constrained WB are [8], [17]. The former presents a framework to approximate gradient flows for Wasserstein metrics by computing discrete entropy-regularized flows, which are computed as *JKO flows* (named after the authors in [18]). It introduces the concept of Wasserstein attraction, which is used in our work. We expand on this concept by observing that our particular problem formulation allows us to write each iteration of the WA problem as the computation of a WB, which unlocks the use of powerful computational tools found in the literature to solve this problem. Additionally, as previously mentioned, we further generalize the definition of this regularized flow by including the supplemental constraints of the topology of a network and the node and edge capacity bounds, which are features not considered in [17]. The latter work, [8], complements [17] while focusing on the dual formulation of Wasserstein variational problems.

The main contributions of this paper are twofold: first, we propose the mathematical formulation of a Wasserstein attraction-like problem to solve mass transport problems over networks by writing them as the computation of a WB problem with additional constraints. And second, we present a methodology to find an approximation of optimal discrete flows over networks based on Dykstra's projection algorithm and the computation of JKO flow proximal operators for the Kullback-Leibler divergence.

The remainder of this article is structured as follows. In Section II, we provide the necessary background for our work, stating some basic definitions from discrete OT theory and presenting the formal statement of the problem we want to solve. In Section III, we briefly review Dykstra's projection algorithm in the setting of optimization problems involving the Kullback-Leibler divergence and how it can be used to solve the WB problem. Then, we show the additional steps needed on the algorithm to enforce support constraints and capacity bounds on the network's links and nodes. With that, we present our proposed approach. In Section IV, we provide an illustrative example to showcase the performance of the scheme to solve a supply-and-demand problem in the context of flow optimization on drinking water networks. Finally, in Section V, we provide a summary and discuss future investigation directions¹.

Notation

The column vector of all ones is denoted by $\mathbf{1}$ and I is the identity matrix. The adjacency matrix of a graph is denoted by A , and we will write $\bar{A} = A + I$ when considering the connection of one node to itself. \mathbb{R}_+ and \mathbb{R}_{++} refer to non-negative and strictly positive real values respectively.

¹See [2] for an extended version with additional numerical results.

Given $x \in \mathbb{R}^n$, $\|x\|$ stands for its Euclidean norm. Given two matrices $A, B \in \mathbb{R}^{n \times m}$, $\langle A, B \rangle = \sum_{i,j} A_{ij} B_{ij}$. We define the support of a function (or vector) ρ as $\text{SUPP}(\rho) = \{i \mid \rho(i) > 0\}$. We denote $\text{KL}(\pi|\xi)$ as the *Kullback-Leibler divergence* between $\pi \in \mathbb{R}_+^{n \times n}$ and $\xi \in \mathbb{R}_{++}^{n \times n}$, defined as $\text{KL}(\pi|\xi) = \sum_{i,j=1}^n \pi_{ij} \ln \left(\frac{\pi_{ij}}{\xi_{ij}} \right) - \pi_{ij} + \xi_{ij}$, with the convention $0 \ln(0) = 0$. The indicator function of a set \mathcal{C} is defined as $\iota_{\mathcal{C}}(x) = 0$ if $x \in \mathcal{C}$, and $\iota_{\mathcal{C}}(x) = +\infty$ otherwise.

II. PROBLEM STATEMENT: DISCRETE FLOWS AND WASSERSTEIN ATTRACTION ON GRAPHS

A. Discrete Flows on Graphs

Consider a *discrete, finite, fixed and connected* graph $\mathcal{G} = (V, E)$, where V is a set of n nodes $V = (1, \dots, n)$, and $E \subseteq V \times V$ is a set of directed edges so that $(j, i) \in E$ if and only if there is a link between the node $j \in V$ and node $i \in V$. Denote the probability simplex on V as $\text{Prob}(V) = \left\{ \mu \in \mathbb{R}_+^n \mid \sum_{x \in V} \mu(x) = 1 \right\}$. The set of edges E has an associated weight function $c : E \rightarrow \mathbb{R}_+$ where each edge $e \in E$ has a corresponding weight $c_e = c(e)$, i.e., the cost of sending a unit of mass using the edge e . Furthermore, endow the graph \mathcal{G} with its natural metric d which measures the total weight of the shortest path between any two nodes in \mathcal{G} .

We study the discrete flow (i.e., discretization in time) problem of optimally transporting an initial mass distribution $\mu \in \text{Prob}(V)$ to a target mass distribution $\nu \in \text{Prob}(V)$ using the graph \mathcal{G} . The associated weight of each edge allows us to define a cost matrix $C \in \mathbb{R}_+^{n \times n}$, where $[C]_{ji} = d(j, i)$ indicates the cost of transporting a unit mass from node j to node i . Moreover, we endow the space $\text{Prob}(V)$ of probability measures on V with the 1-Wasserstein distance between two probability distributions μ and ν on \mathcal{G} as

$$W_1(\mu, \nu) = \min_{\pi \in \Pi(\mu, \nu)} \sum_{x, y \in V} d(x, y) \pi(x, y),$$

where the minimizer (defined as the *optimal transport plan*) is computed over all couplings on $V \times V$ with marginals μ and ν , i.e., $\Pi(\mu, \nu) = \{ \pi \in \mathbb{R}_+^{n \times n} \mid \pi \mathbf{1} = \mu, \pi^T \mathbf{1} = \nu \}$.

Our objective is to design a discrete flow $\{\rho_t\}_{t \geq 0}$ on \mathcal{G} , where $\rho_t \in \text{Prob}(V)$, by constructing a sequence of transport plans $\{\pi_t\}_{t \geq 0}$ such that $\rho_0 = \mu$, $\rho_{t+1} = \pi_t \mathbf{1}$, $\rho_t = \pi_t^T \mathbf{1}$ and $\lim_{t \rightarrow \infty} \rho_t = \nu$. Moreover, the transport cost at each iteration should be minimized and the sequence $\{\pi_t\}_{t \geq 0}$ is required to satisfy the following constraints imposed by the network:

- (a) A node can only send mass to its neighbors, i.e., $[\pi_t]_{ij} > 0$ if $[\rho_t]_j > 0$ and $(j, i) \in E$. In other words, the flow should follow the sparsity pattern induced by the graph topology. Intuitively, a flow can only be assigned between two nodes if and only if there is an edge connecting them. Hence, for a transport plan π_t it must hold that $\text{SUPP}(\rho_{t+1}) \subseteq \{ \text{SUPP}(\rho_t) \cup \{j \mid (j, i) \in E\} \}$.
- (b) The mass sent over an edge cannot surpass the associated edge capacity, i.e., $\pi_t \leq \bar{C}$, for a capacity matrix $\bar{C} \in \mathbb{R}_+^{n \times n}$, where $[\bar{C}]_{ij}$ is the capacity of the edge $(j, i) \in E$ (the inequality is understood element-wise).

- (c) The mass at a node i at some time instant $t \geq 0$ must not exceed its local storage capacity, i.e., $\rho_t \leq \rho$, for a vector of storage capacities $\rho \in \mathbb{R}_+^n$ (again, the inequality is understood entry-wise).
- (d) The mass transported from a node j to a node i cannot exceed the mass held at node j , i.e., $[\pi_t]_{ij} \leq [\rho_t]_j$.

B. Wasserstein Attraction Flows

We formulate the dynamic transport problem described in Section II-A as a constrained Wasserstein attraction problem [17]. Our main technical tool will be the JKO flow proximal operators which we introduce next. We first present the JKO flow proximal operator with respect to a functional f . For all $q \in \text{Prob}(V)$,

$$\text{Prox}_{\tau, f}^{W_1}(q) \triangleq \underset{p \in \text{Prob}(V)}{\text{argmin}} \{W_1(p, q) + \tau f(p)\},$$

where τ is a step-size. Thus, starting from an initial distribution $\rho_0 = \mu$, the discrete JKO flow with respect to f is defined as

$$\rho_{t+1} \triangleq \text{Prox}_{\tau, f}^{W_1}(\rho_t). \quad (1)$$

Wasserstein attraction refers to the flow generated by the implicit gradient steps in (1), known as *JKO stepping*, with respect to the potential function defined as $W_1(\rho_t, \nu)$ for some fixed distribution ν . Informally, the potential function drives the flow to minimize its Wasserstein distance to a *target* distribution. Thus, we define the WA discrete flow as

$$\begin{aligned} \rho_{t+1} &= \text{Prox}_{\tau, W_1(\cdot, \nu)}^{W_1}(\rho_t) \\ &= \underset{p \in \text{Prob}(V)}{\text{argmin}} \{W_1(p, \rho_t) + \tau W_1(p, \nu)\}. \end{aligned} \quad (2)$$

The WA defined in (2) has a precise optimization structure. However, the computation of each proximal operation is computationally intense [17]. Moreover, the constraints imposed by the graph are not taken into account. In the next subsection, we describe our proposed approach for the efficient computation of the discrete WA, taking into account the constraints imposed by the network.

C. Approximate Wasserstein Attraction Flow on Graphs

Initially, we present the entropy-regularized discrete JKO flow for the WA problem following the ideas introduced in [17]. The main contribution in [17] is to replace the Wasserstein metrics with their entropy-regularized versions.

Definition 1: Given a cost matrix $C \in \mathbb{R}_+^{n \times n}$, the **discrete entropy-regularized Wasserstein distance** between $\mu, \nu \in \text{Prob}(V)$ is defined as

$$\mathcal{W}_\gamma(\mu, \nu) = \min_{\pi \in \Pi(\mu, \nu)} \langle C, \pi \rangle + \gamma H(\pi), \quad (3)$$

where $H(\pi) = \sum \pi_{ij} (\ln \pi_{ij} - 1) = \langle \pi, \ln \pi - \mathbf{1}\mathbf{1}^\top \rangle$ is the negative entropy and $\gamma \geq 0$ is the regularization parameter.

Now, we can define the approximate entropy-regularized WA flow as

$$\begin{aligned} \rho_{t+1} &= \text{Prox}_{\tau, \mathcal{W}_\gamma(\cdot, \nu)}^{\mathcal{W}_\gamma}(\rho_t) \\ &= \underset{p \in \text{Prob}(V)}{\text{argmin}} \{\mathcal{W}_\gamma(p, \rho_t) + \tau \mathcal{W}_\gamma(p, \nu)\}. \end{aligned} \quad (4)$$

Note $\mathcal{W}_\gamma(\cdot, \cdot)$ is a strictly convex and coercive function, therefore the operator in (4) is uniquely defined.

Next, we state one important observation about the entropy-regularized WA flow in (4). Without loss of generality, one can multiply the argument in the optimization problem (4) by a constant $\omega = 1/(1 + \tau)$. Thus, we obtain

$$\rho_{t+1} = \underset{p \in \text{Prob}(V)}{\text{argmin}} \{\omega \mathcal{W}_\gamma(p, \rho_t) + (1 - \omega) \mathcal{W}_\gamma(p, \nu)\}, \quad (5)$$

which is precisely the entropy-regularized Wasserstein barycenter between ρ_t and ν [7]. Recall that for a finite set of probability distributions $\{\mu_i\}_{i=1}^m$ where $\mu_i \in \text{Prob}(V)$, the entropy-regularized Wasserstein barycenter is defined as

$$\mu \triangleq \underset{p \in \text{Prob}(V)}{\text{argmin}} \sum_{i=1}^m \omega_i \mathcal{W}_\gamma(p, \mu_i),$$

where $\omega_i \geq 0$ and $\sum_{i=1}^m \omega_i = 1$.

We interpret the Wasserstein attraction problem as the sequential computation of Wasserstein barycenters. This introduces an additional weight parameter that can be modified to give preference to one measure or the other. Such parameter consequently alters how the mass is transported across the graph.

Approximate solutions to problems of the form (5) can be efficiently computed by reformulating the entropy-regularized OT problem (3) as

$$W_\gamma(\mu, \nu) = \min_{\pi \in \Pi(\mu, \nu)} \text{KL}(\pi | \xi), \quad (6)$$

where $\xi = e^{-C/\gamma}$ (entry-wise exponential) [4]. Note that (6) can be extended for higher dimensional arrays (such as the tuples $\pi = (\pi_1, \dots, \pi_m)$ introduced in the definition of the WB) by summing over the indices (i, j, k, \dots) . Thus, following [4], we can rewrite (5) as

$$\min_{\pi \in \mathcal{C}_1 \cap \mathcal{C}_2} \text{KL}_\omega(\pi | \xi) = \omega \text{KL}(\pi_1 | \xi) + (1 - \omega) \text{KL}(\pi_2 | \xi), \quad (7)$$

where

$$\mathcal{C}_f = \{\pi_1, \pi_2 \mid \pi_1 \mathbf{1} = \rho_t, \pi_2 \mathbf{1} = \nu\}, \quad (8)$$

$$\mathcal{C}_e = \{\pi_1, \pi_2 \mid \pi_1^\top \mathbf{1} = \pi_2^\top \mathbf{1} = p\}. \quad (9)$$

Finally, taking into account in Problem (7) the constraints (a), (b), (c) and (d) presented in Section II-A, we can state our main contribution regarding the design of the entropy-regularized discrete WA flow.

Problem 1: Consider a discrete, finite, fixed and connected graph with n vertices, $\tilde{C} \in \mathbb{R}_+^{n \times n}$ the capacity matrix, and $\mu, \nu \in \text{Prob}(V)$ the initial and final distributions respectively. We design the sequence of probability measures $\{\rho_t\}_{t \geq 0}$ by finding, for each $t \geq 0$, the transport plan that solves the optimization problem

$$\{\pi_t\} = \underset{\substack{\pi \in \mathcal{C}_f \cap \mathcal{C}_e \\ \pi \in \mathcal{C}_1 \cap \mathcal{C}_2 \cap \mathcal{C}_3}}{\text{argmin}} \omega \text{KL}(\pi_1 | \xi) + (1 - \omega) \text{KL}(\pi_2 | \xi), \quad (10a)$$

where

$$\mathcal{C}_f = \{ \boldsymbol{\pi} \in \mathbb{R}_+^{n \times n} \times \mathbb{R}_+^{n \times n} \mid \boldsymbol{\pi}_1 \mathbf{1} = \boldsymbol{\rho}_t, \boldsymbol{\pi}_2 \mathbf{1} = \boldsymbol{\nu} \} \quad (10b)$$

$$\mathcal{C}_e = \{ \boldsymbol{\pi} \in \mathbb{R}_+^{n \times n} \times \mathbb{R}_+^{n \times n} \mid \boldsymbol{\pi}_1^\top \mathbf{1} = \boldsymbol{\pi}_2^\top \mathbf{1} = p \} \quad (10c)$$

$$\mathcal{C}_1 = \{ \boldsymbol{\pi} \in \mathbb{R}_+^{n \times n} \times \mathbb{R}_+^{n \times n} \mid \boldsymbol{\pi}_1 \leq \tilde{\mathcal{C}} \} \quad (10d)$$

$$\mathcal{C}_2 = \{ \boldsymbol{\pi} \in \mathbb{R}_+^{n \times n} \times \mathbb{R}_+^{n \times n} \mid \boldsymbol{\pi}_1^\top \mathbf{1} \leq \boldsymbol{\rho}, \boldsymbol{\pi}_2^\top \mathbf{1} \leq \boldsymbol{\rho} \} \quad (10e)$$

$$\mathcal{C}_3 = \left\{ \boldsymbol{\pi} \in \mathbb{R}_+^{n \times n} \times \mathbb{R}_+^{n \times n} \mid [\boldsymbol{\pi}_1^\top \mathbf{1}]_i \leq \sum_{j:(j,i) \in E} [\boldsymbol{\rho}_t]_j \right\} \quad (10f)$$

An essential feature of the scheme presented is that, unlike in the computation of the Wasserstein distance (or, for that matter, solving the MCFP), we do not compute the complete flow in a single step, which would also entail having to store the shortest path between each node (or at least the first step of each path). In this regard, our method not only does not need to store this additional information, but it is also *memoryless*, in the sense that, at each step, the algorithm solves a new problem with initial and final distributions (hence it adapts to changes in the measures and parameters during the transport). Here lies the main difference between the flow we compute, which is discrete, and the one found by solving a MCFP, which is continuous.

III. ITERATIVE PROJECTIONS FOR THE COMPUTATION OF TRANSPORT PLANS

Now that we have the necessary background on discrete OT and have introduced the problem we want to solve, we describe the approach that we propose. We will solve the regularized version of the WB problem, with the additional constraints (10d), (10e) and (10f). To do so, we use a well-known algorithm for solving regularized OT problems called *Dijkstra's projection algorithm* [9], which, in our setting, is a generalization of the widely used Iterative Bregman Projections (IBP) algorithm [4]. We use Dijkstra's method because the convergence of IBP cannot be guaranteed in the presence of inequality constraints.

In Section III-A, we give some background on how this algorithm is used to compute the regularized WB. In Section III-B, we show how one can modify it to compute the WB with the added constraints, and finally, in Section III-C, we move on to the description of the proposed algorithm.

A. Computation of the WB Using Dijkstra's Projection Algorithm

Dijkstra's projection algorithm can be used to solve problems of the form

$$\min_{\boldsymbol{\pi} \in \cap_i \mathcal{C}_i} \text{KL}(\boldsymbol{\pi} \mid \boldsymbol{\xi}),$$

much like Problem 1 defined in Section II. It is based on the computation of the proximal operators for the KL divergence. This is done iteratively, cycling through each constraint set \mathcal{C}_i , and since $\mathcal{C} = \cap_i \mathcal{C}_i$ is a finite intersection of L sets, we shall define, for every index i , $\mathcal{C}_{i+L} = \mathcal{C}_i$. Then, for each $k > 0$ we compute

$$\boldsymbol{\pi}^{(k)} = \text{Prox}_{\iota_{\mathcal{C}_k}}^{\text{KL}} \left(\boldsymbol{\pi}^{(k-1)} \cdot q^{(k-L)} \right), \quad q^{(k)} = q^{(k-L)} \frac{\boldsymbol{\pi}^{(k-1)}}{\boldsymbol{\pi}^{(k)}},$$

with initial values $\boldsymbol{\pi}^{(0)} = \boldsymbol{\xi}$ and $q^{(0)} = q^{(-1)} = \dots = q^{(-L+1)} = \mathbf{1}\mathbf{1}^\top$. The product and division of matrices are considered element-wise. We slightly abuse notation by omitting the step-size τ in the definition of the proximal operator, since we are multiplying the argument in the optimization problem (4) by $\omega = 1/(1 + \tau)$, as noted in Section II.

The next propositions state how we can compute in closed form the proximal operator corresponding to each constraint in the WB problem (7).

Proposition 1 (Proposition 1 in [4]): The proximal operator of the indicator function $\iota_{\mathcal{C}_f}$, corresponding to the constraint set \mathcal{C}_f in (8), has the closed form

$$\left[\text{Prox}_{\iota_{\mathcal{C}_f}}^{\text{KL}\omega}(\boldsymbol{\pi}) \right]_l = \text{Prox}_{\iota_{\{\boldsymbol{\pi}_1 \mathbf{1} = P_l\}}}^{\text{KL}}(\boldsymbol{\pi}_l) = \text{diag} \left(\frac{P_l}{\boldsymbol{\pi}_l \mathbf{1}} \right) \boldsymbol{\pi}_l. \quad (11)$$

Proposition 2 (Proposition 2 in [4]): The proximal operator of the indicator function $\iota_{\mathcal{C}_e}$, corresponding to the constraint set \mathcal{C}_e in (9), has the closed form

$$\left[\text{Prox}_{\iota_{\mathcal{C}_e}}^{\text{KL}\omega}(\boldsymbol{\pi}) \right]_l = \boldsymbol{\pi}_l \text{diag} \left(\frac{p}{\mathbf{1}^\top \boldsymbol{\pi}_l} \right), \quad (12)$$

where $p = \prod_{l=1}^m (\mathbf{1}^\top \boldsymbol{\pi}_l)^{\omega_l}$ (the products and exponentiation are considered element-wise), and $m = 2$ in our case.

B. Capacity and Support Constrained WB

In the context of networks, it is reasonable to restrict how much mass can be sent from one node to another, i.e. to add a capacity to the edges connecting the nodes. This constraint is imposed on each transport plan by defining a capacity matrix $\tilde{\mathcal{C}} \in \mathbb{R}^{n \times n}$ such that $[\tilde{\mathcal{C}}]_{ij}$ is the maximum mass that can be sent from node i to node j . The following proposition concerns the computation of the proximal operator for the set \mathcal{C}_1 in (10d).

Proposition 3 (Section 5.2 in [4]): The proximal map for the function $\iota_{\{\boldsymbol{\pi}_1 \leq \tilde{\mathcal{C}}\}}$ is defined as

$$\text{Prox}_{\iota_{\{\boldsymbol{\pi}_1 \leq \tilde{\mathcal{C}}\}}}^{\text{KL}}(\boldsymbol{\pi}_1) = \min \left(\boldsymbol{\pi}_1, \tilde{\mathcal{C}} \right), \quad (13)$$

with the minimum computed element-wise.

We can also have capacity limits on some of the nodes, meaning that even though the optimal solution might send a certain amount of mass to one of these nodes, it may not be possible to hold that much quantity. This corresponds to the constraint set \mathcal{C}_2 in (10e). We can adapt the results for partial transport problems in [4] for the computation of the projection on this set in closed form, as follows.

Proposition 4 (Proposition 5 in [4]): For the the indicator function $\iota_{\mathcal{C}_2}$, corresponding to the constraint set \mathcal{C}_2 in (10e), one has

$$\begin{aligned} \left[\text{Prox}_{\iota_{\mathcal{C}_2}}^{\text{KL}\omega}(\boldsymbol{\pi}) \right]_l &= \text{Prox}_{\iota_{\{\boldsymbol{\pi}_1^\top \mathbf{1} \leq \boldsymbol{\rho}\}}}^{\text{KL}}(\boldsymbol{\pi}_l) \\ &= \boldsymbol{\pi}_l \text{diag} \left(\min \left(\frac{\boldsymbol{\rho}}{\boldsymbol{\pi}_l^\top \mathbf{1}}, \mathbf{1} \right) \right), \end{aligned} \quad (14)$$

where the minimum and division of vectors are considered element-wise.

Regarding the support constraint (10f), since nodes in $\text{SUPP}(\boldsymbol{\rho}_t)$ can still send mass to non-neighboring nodes, our

Algorithm 1 Conceptual procedure of the proposed approach

Input: Initial and final distributions ρ_0 and ν , adjacency matrix A , cost matrix C , vector of storage capacities ρ , regularization parameter $\gamma(t)$ and weight $\omega(t)$ depending on t and such that $\gamma(t), \omega(t) \rightarrow 0$ as $t \rightarrow +\infty$, accuracy parameter $\varepsilon > 0$

- 1: $t = 0$
- 2: **while** $\frac{1}{2} \|\nu - \rho_t\|_1 > \varepsilon$ **do**
- 3: Define the capacity matrix \tilde{C} as seen in (15)
- 4: Compute the WB ρ_{t+1} with weights $\omega_1 = \omega(t)$ and $\omega_2 = 1 - \omega(t)$ and the additional support and capacity constraints by using Dykstra's projection algorithm with initial conditions $\pi_1^{(0)} = \pi_2^{(0)} = e^{-\frac{C}{\gamma(t)}}$ and the proximal operators defined on (12), (11) and (14) (with ρ) for both transport plans, and (13) only for transport plan π_1 to enforce the capacity constraint (10d) with capacity matrix \tilde{C}
- 5: $t \leftarrow t + 1$
- 6: **end while**

Output: $\{\rho_t\}_t$

fix is to take advantage of constraint (10d) and adapt it to circumvent this issue. We redefine the capacity matrix \tilde{C} for π_1 from ρ_t to ρ_{t+1} , such that for the nodes in the support of ρ_t , if there is no connection between one of them and another node, the “link” between them has zero capacity, i.e.,

$$[\tilde{C}]_{ij} = \begin{cases} 0 & \text{if } j \in \text{SUPP}(\rho_t) \text{ and } \bar{A}_{ij} = 0, \\ [\tilde{C}]_{ij} & \text{otherwise.} \end{cases} \quad (15)$$

C. Description of the Proposed Approach

Now, we can present the proposed algorithm to solve Problem 1. We use Dykstra's projection algorithm, and together with the support and capacity constraints, we can impose the additional restrictions introduced in the problem statement (Section II). Algorithm 1 summarizes the proposed method. It is important to remark that our entropy-regularized approach does not allow the scheme to converge exactly to the target distribution ν . Since the additional entropy term in the definition of the Wasserstein distance (3) forces every node to send a small amount of mass to the rest, the solution obtained can be more or less diffused depending on the regularization strength γ . Moreover, we cannot guarantee the convergence of Algorithm 1 for a fixed weight ω , and to our knowledge, there is no proof for it as of yet. However, if instead of taking fixed values for both γ and ω we consider, at each step t , $\gamma(t), \omega(t)$ such that $\gamma(t), \omega(t) \rightarrow 0$ as $t \rightarrow +\infty$, we can ensure its convergence [4] [17]. Despite that, in the simulations carried out in Section IV, we consider the weight ω to be both fixed and tending to zero since we have observed how, for a constant $\omega < 1/2$, the mass reaches the target distribution as well. We state the following lemma regarding the convergence of the computation of each intermediate distribution in the discrete flow.

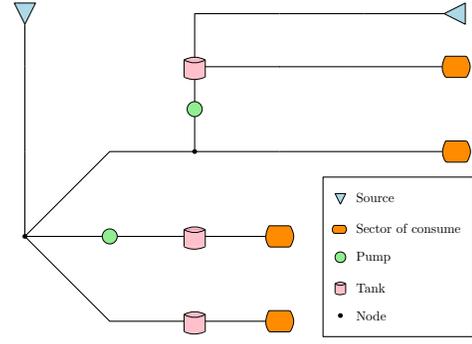


Fig. 1. Topology of the small DWN case study.

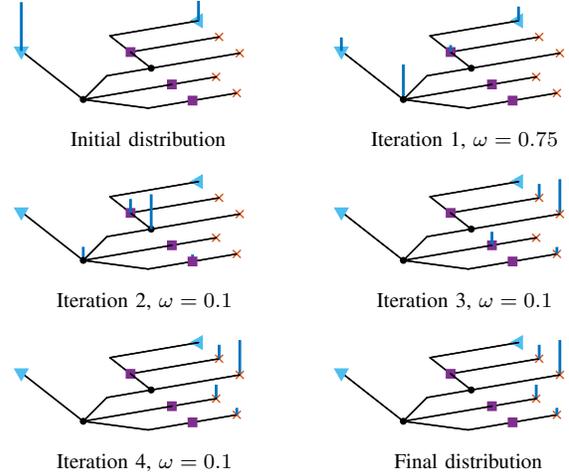


Fig. 2. Steps obtained for the small DWN case study.

Lemma 1: For each step t , let \tilde{C} be the capacity matrix defined in (15) such that it verifies $\tilde{C}^T \mathbf{1} > \rho_t$, and let ρ be the node capacity vector in the constraint set \mathcal{C}_2 such that $\rho_t < \rho$ (both inequalities are considered element-wise). Then, the iterative computation of the proximal steps in Propositions 1, 2, 3 and 4 converges to the solution of (10a).

IV. CASE STUDY

To show the effectiveness of the proposed approach, a simple case study of a drinking water network (DWN) is considered. Fig. 1 depicts a basic topology of a drinking water transport network. Water is moved from the sources towards the network using manipulated actuators to fill retention tanks and supply water to demand sectors. The reader is referred to [16] for further details about this system.

Our objective is to find the (discrete) flow that moves the mass from an initial distribution (water provided by the treatment plants and reservoirs) to a target distribution (expected water in the reservoirs to cover the consumers' water demand) such that it follows the sparsity pattern and constraints induced by the network, and each step is the most cost-efficient (depending on the weight parameter ω).

Fig. 2 shows a simulation on the small network in Fig. 1, ignoring the pumps (which simply add an additional cost) and with no additional capacities on the edges and nodes, for simplicity. Here, we take advantage of the parameter ω

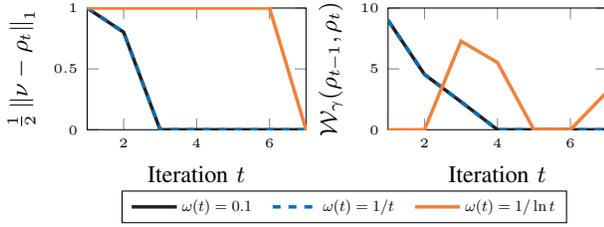


Fig. 3. (Left) Total variation distance between the distribution obtained at iteration t (ρ_t) and the final distribution (ν), and (right) cost of transportation for each iteration, for the example depicted in Fig. 2 (for different weights).

to regulate how the water is transported. In particular, in the first step, we use a fairly high weight $\omega = 0.75$ in favor of the initial distribution so that the transportation is done more gradually. In the following steps, the weight is reduced to $\omega = 0.1$ to cover the demand much faster. Fig. 3 (left), shows the total variation distance between the intermediate distribution ρ_t and the target measure ν , for $\omega(t)$ tending to zero at different rates and also fixed at $\omega(t) = 0.1$. In any case, we see how we eventually converge to the final distribution. For $\omega = 0.1$, since it gives more weight to minimizing the distance to ν rather than the previous distribution, the mass advances faster until it covers the target in few iterations. Similarly, for $\omega(t) = 1/t$, the weight decreases fast enough, and ν is covered in the same number of steps. For $\omega(t) = 1/\ln t$, the decrease rate is slower, and it takes more iterations to cover the target, finally doing so in a single step. On the right plot, we have the Wasserstein distance of each step (the cost of transportation). We can reach the same conclusions we had with the total variation distance, but in particular, we notice how for $\omega(t) = 1/\ln t$ the mass stays still until ω is small enough at the third iteration.

Note how the memorylessness property of the proposed scheme is a valuable feature to have in the context of DWN logistics since it allows the algorithm to adapt to different changes as it advances (sudden peaks in demand, changes in the graph topology, or other occurrences).

From the point of view related to the management of a DWN, the proposed approach opens new ways of improving existent management criteria in the sense of scalability and modularity of the control approaches [20], apart from adding robustness capabilities to the system. In any case, a straightforward comparison with existing methods for management and control of DWNs is nowadays not fair since our approach is presented as a proof of concept for the proposed objectives related to the case study, and then some additional design criteria should be considered.

V. CONCLUSIONS AND FUTURE WORK

In this paper, we have presented a mathematical formulation to resolve discrete optimal flows over networks, based on the computation of constrained Wasserstein Barycenters. Using the entropically regularized approximation of the Wasserstein metric allows us to make use of Dykstra's projection algorithm, which is easy to implement and is

competitive in terms of performance speed since it only requires elementary operations such as matrix and vector products. Moreover, with this methodology, the solution obtained is unique. The scheme presented can be extended to consider more than two distributions and is able to adapt to different changes, thus, a new line of investigation could be to use the proposed approach to tackle problems involving decentralized or distributed models, where not all the information is available for every agent.

REFERENCES

- [1] R.K. Ahuja, T.L. Magnanti, and J.B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, 1993.
- [2] F. Arqué, C. Uribe, and C. Ocampo-Martinez. Application of Wasserstein attraction flows for optimal transport in network systems. *ArXiv*, abs/2109.09182, 2021.
- [3] F. Bassetti, S. Gualandi, and M. Veneroni. On the computation of Kantorovich–Wasserstein distances between two-dimensional histograms by uncapacitated minimum cost flows. *SIAM Journal on Optimization*, 30(3):2441–2469, 2020.
- [4] 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(2):A1111–A1138, 2015.
- [5] S.-N. Chow, W. Li, and H.-M. Zhou. A discrete Schrödinger equation via optimal transport on graphs. *arXiv: Dynamical Systems*, 2017.
- [6] M. Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in neural information processing systems*, pages 2292–2300, 2013.
- [7] M. Cuturi and A. Doucet. Fast computation of Wasserstein barycenters. 2014.
- [8] M. Cuturi and G. Peyré. A smoothed dual approach for variational Wasserstein problems. *SIAM J. Imaging Sci.*, 9:320–343, 2016.
- [9] R. L. Dykstra. An algorithm for restricted least squares regression. *Journal of the American Statistical Association*, 78(384):837–842, 1983.
- [10] M. Erbar, M. Rumpf, B. Schmitzer, and S. Simon. Computation of optimal transport on discrete metric measure spaces. *Numerische Mathematik*, 144:157–200, 2020.
- [11] M. Essid and J. Solomon. Quadratically regularized optimal transport on graphs. *SIAM Journal on Scientific Computing*, 40(4):A1961–A1986, 2018.
- [12] Isabel Haasler, Yongxin Chen, and Johan Karlsson. Optimal steering of ensembles with origin-destination constraints. *IEEE Control Systems Letters*, 5(3):881–886, 2021.
- [13] Y. Hu, X. Zhao, J. Liu, B. Liang, and C. Ma. An efficient algorithm for solving minimum cost flow problem with complementarity slack conditions. *Mathematical Problems in Engineering*, 2020:1–5, 2020.
- [14] P. Kovács. Minimum-cost flow algorithms: an experimental evaluation. *Optimization Methods and Software*, 30(1):94–127, 2015.
- [15] A. Mielke. Geodesic convexity of the relative entropy in reversible Markov chains. *Calculus of Variations and Partial Differential Equations*, 48:1–31, 2013.
- [16] C. Ocampo-Martinez, V. Puig, G. Cembrano, and J. Quevedo. Application of predictive control strategies to the management of complex networks in the urban water cycle. *IEEE Control Systems Magazine*, 33(1):15–41, 2013.
- [17] G. Peyré. Entropic approximation of Wasserstein gradient flows. *SIAM Journal on Imaging Sciences*, 8(4):2323–2351, 2015.
- [18] J. Richard, K. David, and O. Felix. The variational formulation of the Fokker-Planck equation. *SIAM Journal on Mathematical Analysis*, 29(1):1–17, 1998.
- [19] P. H. Richmond and B. Maginnis. On Wasserstein reinforcement learning and the Fokker-Planck equation. *ArXiv*, abs/1712.07185, 2017.
- [20] F. Tedesco, C. Ocampo-Martinez, A. Cassavola, and V. Puig. Centralised and distributed command governor approaches for the operational control of drinking water networks. *IEEE Transactions on Systems, Man & Cybernetics: Systems*, 48(4):586–595, 2018.
- [21] C. Villani. *Optimal Transport: Old and New*. Grundlehren der mathematischen Wissenschaften. Springer Berlin Heidelberg, 2008.