ROBUST OPTIMAL TRANSPORT MODEL PREDICTIVE CONTROL FOR DISTRIBUTION STEERING

BRENDON G. ANDERSON* AND OWEN FRAUSTO†

Abstract. We consider a novel control scheme, termed ROT-MPC, for robustly steering a decoupled system of constrained agents from one finite distribution to another, subject to agent-wise bounded uncertainty. The approach blends discrete-time robust tube model predictive control (MPC) for the computation of provably safe control actions, together with a reach-constrained modification of optimal transport for the distribution mapping. We prove that the feasible set of the reach-constrained optimal transport problem is a convex polytope in the case that every agent admits a deadbeat state feedback stabilizing controller, thereby reducing the modified transport problem to a standard linear program. Additionally, we prove that the overall system's closed-loop dynamics enjoy recursive feasibility and recursive constraint satisfaction. We also give a finite-time stability proof—namely, in cases where a feasible permutation transport plan is computed by ROT-MPC, it is shown that every agent's associated terminal invariant set is robustly finite-time attractive. A variety of numerical simulations are conducted that highlight the robustness, stability, and performance advantages of the proposed method, even in situations where related control schemes fail.

Key words. Optimal Transport, Model Predictive Control, Robust Control, Networked Systems

MSC codes. 49Q22, 93B45, 93B35

1. Introduction. Controlling large networks of agents, each with inherent dynamics and operational constraints, is an important and challenging problem that has garnered immense interest amongst researchers. Applications include swarm robotics [19, 2, 4], optimal power flow [5], and platooning and formation control in (autonomous) networked vehicle systems [22, 23]. One approach for solving the networked control problem is to use conventional centralized control schemes, viewing each agent's state as part of the overall system's larger state as a whole. However, this state augmentation technique is computationally prohibitive even for moderately sized populations of agents, which has led to the exploration of more efficient approaches, such as modeling the population as a finite sample of a continuous distribution [25], or turning to distributed control methods [7].

In recent years, optimal transport has become a very popular framework for designing and analyzing mappings between distributions, with both theoretical and computational developments [32, 26]. This surge in popularity has been bolstered by the application of the Sinkhorn algorithm for matrix scaling to achieve "lightspeed" computations for entropy-regularized optimal transport problems arising in machine learning [12]. Consequently, a handful of recent works have incorporated dynamical constraints into (conventionally static) transport problems as a means to apply optimal transport tools to density control and distribution steering [10, 30, 8, 9, 13, 14, 15, 20], building off of Benamou and Brenier's classical fluid dynamics reformulation of optimal transport [3]. However, these optimal transport-based control approaches have thus far largely neglected system disturbances and uncertainties, making them unsuited for safety-critical networked control systems in which stringent constraints and robustness guarantees are paramount. In this article, we aim to address this gap in the literature by introducing robust optimal transport model predictive control

---

(ROT-MPC) for distribution steering.

Most closely related to our work is "Sinkhorn MPC," introduced by [16, 18]. This control scheme alternates between updating transport plans using iterations of the Sinkhorn algorithm, and steering each agent towards their updated target using decentralized optimal control problems. The follow-up work [17] lifts the prior works' restrictive assumption that every agent has an invertible control matrix, yet is incapable of handling state or input constraints. In all three of these works, uncertainties in the dynamics are neglected, resulting in a lack of provable robustness (and, in some cases, failure in practice, as we will demonstrate in our numerical simulations).

1.1. Contributions. The primary contributions of this work include:
1. Key reachability constraints to ensure feasible robust control problems under optimal transport-based target assignments are identified, culminating into a novel, semi-infinite reach-constrained optimal transport problem.
2. The feasible set of the reach-constrained optimal transport problem is shown to be a polytope when agents admit deadbeat state feedback stabilizing controllers (Proposition 3.2), reducing the modified transport problem into a finite-dimensional linear program.
3. The reach-constrained optimal transport problem is coupled with agent-wise optimal control problems to define a semi-decentralized robust optimal transport MPC (ROT-MPC) control scheme (Algorithm 3.1).
4. The closed-loop dynamics of ROT-MPC are proven to be recursively feasible (Theorem 4.6), to satisfy recursive constraint satisfaction (Theorem 4.8), and to exhibit robust finite-time attraction to the agents' terminal invariant sets in the case that a permutation transport plan is computed at some point during the dynamics (Theorem 4.10).
5. Numerical simulations are carried out to illustrate the performance and robustness advantages of ROT-MPC over baseline control schemes.

To streamline the exposition, some proofs are deferred to Appendix A.

1.2. Outline. The remainder of the article is organized as follows. In Section 2, we define our notations, formally state the problem under consideration, and summarize the fundamental tools from robust tube MPC and optimal transport used in the sequel. We then introduce the new reach-constrained optimal transport problem in Section 3, prove when it reduces to a finite-dimensional linear program with a polytope feasible set, and introduce robust optimal transport MPC (ROT-MPC) using the modified transport problem. Key properties—recursive feasibility, recursive constraint satisfaction, and stability—of the closed-loop dynamics of ROT-MPC are theoretically analyzed in Section 4. Numerical simulations are conducted in Section 5, illustrating the effectiveness of ROT-MPC and drawing comparisons to baseline methods. Conclusions and directions for future work are discussed in Section 6.

2. Preliminaries.

2.1. Notations. Throughout, we use $t \in \mathbb{N} \cup \{0\}$ to denote time for actual system dynamics, whereas we typically use $k \in \mathbb{N} \cup \{0\}$ to denote time for predicted system dynamics within a given optimal control problem. For two sets $X, Y \subseteq \mathbb{R}^n$, we denote their Minkowski sum by $X \oplus Y := \{x + y : x \in X, \ y \in Y\}$, and we denote their Pontryagin difference by $X \ominus Y = \{x \in \mathbb{R}^n : \{x\} \oplus Y \subseteq X\}$. If $x \in \mathbb{R}^n$, then we write $x \oplus Y$ as shorthand for $\{x\} \oplus Y$. The n-vector of all ones is $\mathbf{1}_n$ and the $m \times n$ matrix of all ones is $\mathbf{1}_{m \times n}$. The $n \times n$ identity matrix is written I_n, and 0 denotes a scalar, vector, or matrix of all zeros, whose size is inferred from context.

The partial orders $\geq$ and $\leq$ on vectors and matrices are to be interpreted elementwise. If OPT denotes an optimization problem, then Feas(OPT) denotes the feasible set of OPT, and Val(OPT) denotes the optimal value of OPT. For $N \in \mathbb{N}$, we define $\mathcal{B}_N \coloneqq \{P \in \mathbb{R}^{N \times N} : P\mathbf{1}_N = \mathbf{1}_N/N, \; P^\top \mathbf{1}_N = \mathbf{1}_N/N, \; P \geq 0\}$. To simplify exposition, we call $\mathcal{B}_N$ the *Birkhoff polytope*, despite it being equal to a $\frac{1}{N}$-scaled version of the classical definition of the Birkhoff polytope (the polytope of doubly stochastic matrices). Similarly, we use the term *permutation matrix* to mean a matrix with one and only one nonzero element, equal to $1/N$, in every row and column; we consider scaling permutation matrices by a factor of $1/N$ according to their classical definition.

2.2. Problem Statement. Consider $N \in \mathbb{N}$ agents that evolve according to the discrete-time dynamics

(2.1)
$$\begin{aligned} x_i(0) &= x_{i,0}, \\ x_i(t+1) &= A_i x_i(t) + B_i u_i(t) + w_i(t), \; t \in \mathbb{N} \cup \{0\}, \end{aligned}$$

for all $i \in \{1, \dots, N\}$, where all $A_i \in \mathbb{R}^{n \times n}$ and $B_i \in \mathbb{R}^{n \times m_i}$ are known, but $w_i(t) \in \mathbb{R}^n$ is unknown for all i and all t. For all $i \in \{1, \dots, N\}$, let $X_i \subseteq \mathbb{R}^n$ and $U_i \subseteq \mathbb{R}^{m_i}$ be state and input constraint sets, respectively. The unknown quantity $w_i(t)$ can be thought of either as an external disturbance, or as model error (e.g., in the case that (A_i, B_i) is a linearization of a truly nonlinear system). In any case, we make the following bounded uncertainty assumption:

ASSUMPTION 2.1 (Standing). *For all $i \in \{1, \dots, N\}$, there exists a compact polytope $W_i \subseteq \mathbb{R}^n$ such that $0 \in W_i$ and $w_i(t) \in W_i$ for all $t \in \mathbb{N} \cup \{0\}$.*

The problem at hand is to robustly steer the system of uncertain agents to a collection of desired target states $x_1^d, \dots, x_N^d \in \mathbb{R}^n$ while strictly satisfying the state and input constraints. We do not care which agent moves to a particular target state, only that all of the targets get "covered." We assume that these target states are viable equilibria of each agent (neglecting disturbance):

DEFINITION 2.2. *Let $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$, $X \subseteq \mathbb{R}^n$, and $U \subseteq \mathbb{R}^m$. A state $x^e \in \mathbb{R}^n$ is an* equilibrium *of (A, B) subject to (X, U) if $x^e \in X$ and there exists $u^e \in U$ such that $x^e = Ax^e + Bu^e$.*

ASSUMPTION 2.3 (Standing). *For all $i, j \in \{1, \dots, N\}$, the state x_j^d is an equilibrium of (A_i, B_i) subject to (X_i, U_i).*

Assumption 2.3 is natural; if we have any hope in stabilizing an agent i near one of the target states x_j^d in the presence of disturbance, we should expect to have access to some permissible control that makes x_j^d an equilibrium of system i in the absence of disturbance. Going forward, we let u_{ij}^d denote an input making x_j^d an equilibrium of (A_i, B_i) subject to (X_i, U_i).

We also make the following standard stabilizability assumption:

ASSUMPTION 2.4 (Standing). *(A_i, B_i) is stabilizable for all $i \in \{1, \dots, N\}$.*

By Assumption 2.4, there exists a matrix $K_i \in \mathbb{R}^{m_i \times n}$ for all $i \in \{1, \dots, N\}$ such that $A_i + B_i K_i$ is Schur stable (all eigenvalues have modulus strictly less than one). These matrices can be computed, e.g., via the classical linear quadratic regulator problem. Throughout the remainder of this paper, we fix such matrices K_i.

Remark 2.5. One may replace Assumption 2.3 and Assumption 2.4 by assuming the stronger condition that every pair (A_i, B_i) is controllable.

2.3. Robust Control Invariant Sets. As is standard in robust MPC [21], we use notions of disturbance invariant sets and robust control invariant sets to manage the effects of the uncertainty:

DEFINITION 2.6. *Let* $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$, $K \in \mathbb{R}^{m \times n}$, *and* $W \subseteq \mathbb{R}^n$. *A set* $\Omega \subseteq \mathbb{R}^n$ *is a* disturbance invariant set *for* (A, B) *under* K *subject to* W *if* $((A + BK)\Omega) \oplus W \subseteq \Omega$.

It is desirable that disturbance invariant sets are as small as possible [21], as they are used to define terminal constraint sets in MPC, ensuring the system's state remains close to the target state once it reaches the disturbance invariant set. For A, B, K, and W as in Definition 2.6, the minimal disturbance invariant set is $\Omega = \bigoplus_{j=0}^{\infty}(A + BK)^j W$, in the sense that it is a subset of every other closed disturbance invariant set [27]. Generally, the minimal disturbance invariant set is difficult to compute exactly. However, it is easily inner-approximated by $\hat{\Omega} := \bigoplus_{j=0}^{K}(A + BK)^j W$ for some finite $K \in \mathbb{N}$, which is clearly a polytope when W is. Disturbance invariant outer approximations of the minimal disturbance invariant set are also easy to compute [27].

We make the following assumption, requiring the equilibrium to be contained in a shrunken version of the state constraint set (and similarly for the equilibrium input), which is commonly used to ensure robust constraint satisfaction (see, e.g., [24]).

ASSUMPTION 2.7 (Standing). *For all* $i \in \{1, \ldots, N\}$, *there exists a nonempty convex set* $\Omega_i \subseteq \mathbb{R}^n$ *that is a disturbance invariant set for* (A_i, B_i) *under* K_i *subject to* W_i *such that* $x_j^d \in X_i \ominus \Omega_i$ *and* $u_{ij}^d \in U_i \ominus (K_i \Omega_i)$ *for all* $j \in \{1, \ldots, N\}$.

Taking each Ω_i to be minimal makes Assumption 2.7 more likely to be satisfied. The assumption can be interpreted as shifting the disturbance invariant set around the equilibrium of interest, and requiring the shifted set to satisfy the state and input constraints (under the feedback controller K_i). This allows us to efficiently compute robust control invariant sets (defined below) around the equilibria, since we may simply shift the sets Ω_i rather than computing N^2 robust control invariant sets from scratch (one for each agent-target pair):

DEFINITION 2.8. *Let* $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$, $X, W \subseteq \mathbb{R}^n$, *and* $U \subseteq \mathbb{R}^m$. *A set* $\Omega \subseteq \mathbb{R}^n$ *is a* robust control invariant set *for* (A, B) *subject to* (X, U, W) *if* $\Omega \subseteq X$ *and for all* $x \in \Omega$, *there exists* $u \in U$ *such that* $(Ax + Bu) \oplus W \subseteq \Omega$.

PROPOSITION 2.9. *Let* $i, j \in \{1, \ldots, N\}$. *It holds that* $\Omega_{ij}^d := x_j^d \oplus \Omega_i$ *is a robust control invariant set for* (A_i, B_i) *subject to* (X_i, U_i, W_i).

2.4. Robust Tube MPC. In order to steer each agent $i \in \{1, \ldots, N\}$ while taking into account the uncertainty and constraints, we will employ robust tube MPC (cf., [11, 21] as standard references). For all i, let $T_i \in \mathbb{N}$ be the optimal control time horizon, and let $L_i : (\mathbb{R}^n)^{T_i + 1} \times (\mathbb{R}^{m_i})^{T_i} \to \mathbb{R}$ be the optimal control loss function, which is taken to be time-invariant for convenience. Robust tube MPC ensures robustness of the uncertain closed-loop system by shrinking the state and input constraints on the linear model (A_i, B_i) according to the propagation of disturbance through its "disturbance tube." These disturbance tubes are formally defined as follows:

$$R_i^0 = \{0\}, \quad R_i^k = \bigoplus_{j=0}^{k-1}(A_i + B_i K_i)^j W_i, \ k \in \mathbb{N}.$$

Every set R_i^k is a compact polytope, as is every set $K_i R_i^k$.

The closed-loop system for agent $i \in \{1, \ldots, N\}$ under robust tube MPC is given by the dynamics (2.1), with the control input $u_i(t)$ taken to be $u_i(t) = \mathbf{u}_i^t(0) \in \mathbb{R}^{m_i}$, where $\mathbf{u}_i^t = (\mathbf{u}_i^t(0), \ldots, \mathbf{u}_i^t(T_i - 1)) \in (\mathbb{R}^{m_i})^{T_i}$ is a T_i-step long optimal control signal associated with steering the agent's current state $x_i(t) \in \mathbb{R}^n$ to a shrunken version of some terminal set $\Gamma_i \subseteq X_i$, with respect to the nominal, disturbance-free model:

$$
\begin{aligned}
(\mathbf{x}_i^t, \mathbf{u}_i^t) & \\
& \in \mathrm{OCP}_i(x_i(t), \Gamma_i) \\
& \coloneqq \underset{\substack{(\overline{x}(0), \ldots, \overline{x}(T_i)) \in (\mathbb{R}^n)^{T_i+1}, \\ (\overline{u}(0), \ldots, \overline{u}(T_i-1)) \in (\mathbb{R}^{m_i})^{T_i}}}{\arg\min} \quad L_i(\overline{x}(0), \ldots, \overline{x}(T_i), \overline{u}(0), \ldots, \overline{u}(T_i - 1)) \\
& \quad \text{subject to} \quad \overline{x}(0) = x_i(t), \\
& \qquad\qquad\qquad \overline{x}(k+1) = A_i \overline{x}(k) + B_i \overline{u}(k), \ k \in \{0, \ldots, T_i - 1\}, \\
& \qquad\qquad\qquad \overline{x}(k) \in X_i \ominus R_i^k, \ k \in \{0, \ldots, T_i - 1\}, \\
& \qquad\qquad\qquad \overline{u}(k) \in U_i \ominus (K_i R_i^k), \ k \in \{0, \ldots, T_i - 1\}, \\
& \qquad\qquad\qquad \overline{x}(T_i) \in \Gamma_i \ominus R_i^{T_i}.
\end{aligned}
$$

(2.2)

This canonical formulation of robust tube MPC uses a fixed terminal set Γ_i. However, in our proposed control scheme, we will allow for the update of the target for agent i to occur at every time step via optimal transport assignments, and therefore the terminal set will be time-varying, as is commonly considered in MPC for reference tracking [29, 34]. More details will be presented in Section 3.

We make the following standard assumptions, which ensure that (2.2) is a convex optimization problem, and that when it is feasible it is attained by a unique minimizer:

ASSUMPTION 2.10 (Standing). *For all $i \in \{1, \ldots, N\}$, the dynamics (2.1) and the optimal control problem (2.2) satisfies all of the following conditions:*
1. *$x_{i,0} \in X_i$,*
2. *X_i, U_i, and Γ_i are closed convex polyhedra,*
3. *U_i and Γ_i are compact, and*
4. *L_i is strictly convex.*

2.5. Optimal Transport. In order to actually steer the agents to the target distribution, each agent must be assigned to some target before solving for their individual control inputs. That is, the optimal control problem, being an agent-wise computation, only generates the local behavior of each agent once a desired global behavior has been set. In order to generate such a global setpoint, we use optimal transport, which "optimally" dictates how to morph one distribution into another. Specifically, the classical Kantorovich formulation of the optimal transport problem from one probability measure μ on a Polish space $\mathcal{X}$ to another probability measure ν on a Polish space $\mathcal{Y}$ is given by

$$
\inf \left\{ \int_{\mathcal{X} \times \mathcal{Y}} c(x, y) d\pi(x, y) : \pi \in \Pi(\mu, \nu) \right\},
$$

where $c \colon \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$ is an integrable transportation cost function, and $\Pi(\mu, \nu)$ denotes the set of probability measures on $\mathcal{X} \times \mathcal{Y}$ with marginal μ on $\mathcal{X}$ and marginal ν on $\mathcal{Y}$ [32]. In our setting with a finite number of agents, optimal transport reduces to a mapping between discrete distributions. Formally, optimal transport between

$\mu = \sum_{i=1}^{m} a_i \delta_{x_i}$ and $\nu = \sum_{j=1}^{n} b_j \delta_{y_j}$ reduces to

$$\inf \left\{ \sum_{i=1}^{m} \sum_{j=1}^{n} c(x_i, y_j) P_{ij} : P\mathbf{1}_n = a, \ P^\top \mathbf{1}_m = b, \ P \geq 0 \right\}.$$

This problem is a linear program, which can be viewed as a relaxation of the classical assignment problem. A feasible matrix P is called a *transport plan*, and the entry P_{ij} represents how much "mass" shall be moved from x_i to y_j.

In our setting, we seek to transport agents at initial states $x_1, \ldots, x_N \in \mathbb{R}^n$ to target states $x_1^d, \ldots, x_N^d \in \mathbb{R}^n$. This corresponds to transporting a discrete distribution $\sum_{i=1}^{N} \rho_i \delta_{x_i}$ to a discrete distribution $\sum_{j=1}^{N} \rho_j^d \delta_{x_j^d}$. We view the targets $x_1^d, \ldots, x_N^d$ as always being fixed, whereas the initial states $x_1, \ldots, x_N$ may change between instances of the optimal transport problem. Therefore, we write the cost associated with moving one unit of "mass" from x_i to x_j^d as $C_{ij}(x_i) := c(x_i, x_j^d)$. Furthermore, since we want one and only one agent at every state, both the initial and terminal probability distributions are uniform; $\rho = \rho^d = \mathbf{1}_N/N$. Under this formulation, the optimal transport problem for directing the population of agents reads

$$(2.3) \qquad \text{OTP}(x_1, \ldots, x_N) := \operatorname*{arg\,min}_{P \in \mathcal{B}_N} \sum_{i,j=1}^{N} C_{ij}(x_i) P_{ij}.$$

The optimal transport problem (2.3) is always feasible; its feasible set equals the Birkhoff polytope $\mathcal{B}_N$. The Birkhoff-von Neumann theorem states that $\mathcal{B}_N$ has $N!$ vertices, each one being equal to some permutation matrix on N elements [6, 33].

Notice that it is not immediately clear whether a transport plan P that is feasible for (2.3) will preserve the fixed "mass" of the individual agents. For example, $P = \frac{1}{4}\mathbf{1}_{2\times 2}$ is feasible for (2.3) in the case that $N = 2$, but this transport plan says to send $1/2$ of agent 1 from x_1 to x_1^d and another $1/2$ of agent 1 to x_2^d, which is not realizable under our considered problem setting of discrete agents. However, a fundamental result in optimal transport theory is that, in this specific case of transport between uniform discrete distributions with finite support, there exists a permutation matrix that solves (2.3), and therefore such an optimal transport plan reduces to a one-to-one assignment between the initial and target states [31, Page 5]. This fact immediately follows from the Birkhoff-von Neumann theorem together with the fact that linear programs over compact feasible sets have at least one vertex solution.

3. Robust Optimal Transport MPC (ROT-MPC). We now describe our proposed blending of optimal transport with robust tube MPC. First, we remark that the optimal transport problem (2.3) makes no considerations for the reachability aspects of the agents. That is, solving (2.3) may yield a transport plan that tells agent i to move to a target location, even though this particular steering may be difficult or impossible for the agent to achieve. Since the agents have limited control authority, this may result in breaking recursive feasibility in the robust tube MPC. The previous works [16, 18] considering optimal transport-based MPC did not need to consider this challenge, since they assumed that the agent's had invertible input matrices B_i, no control constraints (that is, $U_i = \mathbb{R}^{m_i}$), and no uncertainties. However, in our more challenging setting with underactuation, control constraints, and uncertainty, we must take care to ensure that the optimal transport plans do not result in infeasible control problems. Furthermore, in the prior works [16, 18], the optimal transport plan is used to generate a target state for each agent. However, as discussed earlier, reaching an

exact target state may be infeasible in our setting with constraints and disturbance, and therefore we consider using the transport plan to define a target set to which we can feasibly and robustly steer the agent.

In particular, at every time t, we use a transport plan $P(t)$ to generate a temporary target state and temporary terminal set,

$$\chi_{\text{tmp},i}(t) := N \sum_{j=1}^{N} P_{ij}(t)x_j^d, \quad \Omega_{\text{tmp},i}(t) := N \bigoplus_{j=1}^{N} P_{ij}(t)\Omega_{ij}^d.$$

The set $\Omega_{\text{tmp},i}(t)$ will take the place of the terminal set Γ_i in the optimal control problem. This temporary terminal set can be viewed as a generalization of the Barycentric projection from discrete optimal transport theory that returns a deterministic point-to-point assignment from a transport plan (where the transport plan may in general split "mass") [26, Remark 4.11]. In our case where the robust control invariant sets satisfy $\Omega_{ij}^d = x_j^d \oplus \Omega_i$ with Ω_i being convex, it is easy to show that $\Omega_{\text{tmp},i}(t) = N \bigoplus_{j=1}^{N} P_{ij}(t)(x_j^d \oplus \Omega_i) = \left(N \sum_{j=1}^{N} P_{ij}(t)x_j^d\right) \oplus \Omega_i = \chi_{\text{tmp},i}(t) \oplus \Omega_i$. Since $0 \in W_i$ and hence $0 \in \Omega_i$ for all i, it holds that $\chi_{\text{tmp},i}(t) \in \Omega_{\text{tmp},i}(t)$ for all i, and hence the target state $\chi_{\text{tmp},i}(t)$ can be thought of as the "center" of the set $\Omega_{\text{tmp},i}(t)$. In practice, the set $\Omega_{\text{tmp},i}(t)$ can be efficiently computed by shifting the (fixed) set Ω_i by $\chi_{\text{tmp},i}(t)$, rather than performing the associated Minkowski sums.

3.1. Reach-Constrained Optimal Transport. In accordance with the above temporary targets and terminal sets, we modify the optimal transport problem as follows. First, in addition to the current agent states $x_i \in \mathbb{R}^n$, we assume that we have predictions $\hat{x}_i(T_i) \in \mathbb{R}^n$ of where the agents *could have been* robustly steered to from their previous states, within a time horizon T_i. These predictions will come from past optimal control problems. Then, the modified optimal transport problem reads

$$(3.1)\quad \begin{aligned} &(P^\star, \hat{u}_1^\star, \ldots, \hat{u}_N^\star) \\ &\quad \in \text{OTP}_{\text{reach}}(x_1, \ldots, x_N, \hat{x}_1(T_1), \ldots, \hat{x}_N(T_N)) \\ &\quad := \underset{\substack{P \in \mathcal{B}_N, \\ \hat{u}_1 \in \mathbb{R}^{m_1}, \ldots, \hat{u}_N \in \mathbb{R}^{m_N}}}{\arg\min} \quad \sum_{i,j=1}^{N} C_{ij}(x_i)P_{ij} \\ &\quad \text{subject to} \quad \hat{u}_i \in U_i \ominus (K_i R_i^{T_i}), \; i \in \{1, \ldots, N\}, \\ &\qquad\qquad\qquad (A_i\hat{x}_i(T_i) + B_i\hat{u}_i) \oplus ((A_i + B_iK_i)^{T_i}W_i) \\ &\qquad\qquad\qquad\qquad \subseteq \left(N \bigoplus_{j=1}^{N} P_{ij}\Omega_{ij}^d\right) \ominus R_i^{T_i}, \; i \in \{1, \ldots, N\}. \end{aligned}$$

This problem has additional constraints when compared to (2.3), and therefore the optimal value of (3.1) upper-bounds that of (2.3) for all predictions $\hat{x}_1(T_1), \ldots, \hat{x}_N(T_N)$:

PROPOSITION 3.1. *Let* $x_1, \ldots, x_N, \hat{x}_1(T_1), \ldots, \hat{x}_N(T_N) \in \mathbb{R}^n$ *be arbitrary. If* $(P, \hat{u}_1, \ldots, \hat{u}_N)$ *is feasible for* $\text{OTP}_{\text{reach}}(x_1, \ldots, x_N, \hat{x}_1(T_1), \ldots, \hat{x}_N(T_N))$, *then* P *is feasible for* $\text{OTP}(x_1, \ldots, x_N)$. *Consequently,*

$$\text{Val}(\text{OTP}(x_1, \ldots, x_N)) \leq \text{Val}(\text{OTP}_{\text{reach}}(x_1, \ldots, x_N, \hat{x}_1, \ldots, \hat{x}_N)).$$

Intuitively, the added constraints ensure that the temporary terminal set $N \bigoplus_{j=1}^{N} P_{ij}\Omega_{ij}^d$ generated by the transport plan P is one-step robustly reachable from $\hat{x}_i(T_i)$ (hence the subscript "reach" on the optimization problem). As we will show in

Section 4, this in turn ensures (recursive) feasibility of the robust optimal control problems using the updated temporary terminal set at the next time step. In light of the added reachability constraints, we call the above optimization the *reach-constrained optimal transport problem*.

Notice that the reach-constrained optimal transport problem (3.1) is a (semi-infinite) linear program, and as such has a convex feasible set. In general, semi-infinite linear programs need not have polyhedral feasible sets. In the case that every K_i is a deadbeat controller (meaning that $(A_i + B_i K_i)^{k_i} = 0$ for some finite $k_i \in \mathbb{N} \cup \{0\}$), then the set of feasible reach-constrained transport plans does reduce to a polytope:

PROPOSITION 3.2. *If, for every $i \in \{1, \ldots, N\}$, there exists $k_i \in \mathbb{N} \cup \{0\}$ such that $(A_i + B_i K_i)^{k_i} = 0$, then the feasible set of (3.1) is a polytope. In this case, the reach-constrained optimal transport problem (3.1) is equivalent to a standard (finite-dimensional) linear program.*

Proof. Suppose that, for every $i \in \{1, \ldots, N\}$, there exists $k_i \in \mathbb{N} \cup \{0\}$ such that $(A_i + B_i K_i)^{k_i} = 0$. Then $\Omega_i = \bigoplus_{j=0}^{\infty} (A_i + B_i K_i)^j W_i = \bigoplus_{j=0}^{k_i - 1} (A_i + B_i K_i)^j W_i = R_i^{k_i}$, and therefore $\Omega_{ij}^d = x_j^d \oplus \Omega_i = x_j^d \oplus R_i^{k_i}$. Therefore, the subset constraint in (3.1) corresponding to agent $i \in \{1, \ldots, N\}$ reduces to

$$(A_i \hat{x}_i(T_i) + B_i \hat{u}_i) \oplus ((A_i + B_i K_i)^{T_i} W_i) \subseteq \left(N \bigoplus_{j=1}^{N} P_{ij} (x_j^d \oplus R_i^{k_i}) \right) \ominus R_i^{T_i}.$$

This is equivalent to $-N \sum_{j=1}^{N} P_{ij} x_j^d + B_i \hat{u}_i \in (-A_i \hat{x}_i(T_i)) \oplus (R_i^{k_i} \ominus R_i^{T_i+1})$. Since $R_i^{k_i}$ and $R_i^{T_i+1}$ are both polytopes, the above constraint on $(P, \hat{u}_1, \ldots, \hat{u}_N)$ is polyhedral. Since every other constraint in (3.1) is also polyhedral, and since the feasible set is clearly a subset of the polytope $\mathcal{B}_N \times U_1 \times \cdots \times U_N$, we conclude that the feasible set of (3.1) is a polytope. Certainly, this further implies that (3.1) is equivalent to a finite-dimensional linear program. □

Even if the feedback controllers K_i are not deadbeat, there are a variety of ways to reformulate (3.1) into equivalent finite-dimensional problems (both convex and nonconvex) that are amenable to numerical computations. In practice, however, the reformulations of the subset constraints in (3.1) should be done with care, for the choice of reformulation can affect the solution quality of the transport plan. We discuss such computational considerations in Appendix B, and, in particular, we describe a (nonconvex) bilinear reformulation of the subset constraints that we find to yield the best balance between solution quality and overall efficiency of the proposed control scheme in our numerical simulations. The recursive feasibility, recursive constraint satisfaction, and stability of the overall control scheme (given in Section 4) still hold even when using nonconvex reformulations of the reach-constrained optimal transport problem, as their proofs rely solely on (robust) feasibility, rather than on global optimality of the transport plan used by the system. As such, we only require finding a feasible point of (3.1). However, in practice, this feasible point will be found through (locally) minimizing the bilinear reformulation of the reach-constrained problem (3.1).

It is not immediately clear that the feasible set of (3.1) is nonempty, i.e., the added constraints may have made the reach-constrained optimal transport problem infeasible if there are no transport plans in the Birkhoff polytope that yield a robustly one-step reachable temporary terminal set. However, we will show in Section 4 that this problem inherits feasibility from the robust control invariance of the terminal sets Ω_{ij}^d, and furthermore that the optimal control problems inherit feasibility from these added reachability constraints on the transport plan.

Algorithm 3.1 ROT-MPC: Robust optimal transport MPC

---

input: targets x_j^d, dynamics A_i, B_i, X_i, U_i, W_i, $x_{i,0}$, control parameters T_i, L_i, and transport cost functions $C_{ij} \colon x_i \mapsto c(x_i, x_j^d)$

1: **initialize** $x_i(0) = x_{i,0}$ for all i

2: **initialize** $\mathbf{x}_i^{-1}(T_i) := \frac{1}{N} \sum_{j=1}^{N} x_j^d + (A_i + B_i K_i)^{T_i} \left(x_{i,0} - \frac{1}{N} \sum_{j=1}^{N} x_j^d \right)$ for all i

3: **initialize** $P(-1) \in \mathcal{B}_N$ as non-permutation matrix

4: **for** times $t \in \mathbb{N} \cup \{0\}$:

5:     **if** $P(t-1)$ is not a permutation matrix:

6:        **compute** reach-constrained transport plan:

$$(P(t), \hat{u}_1(t), \dots, \hat{u}_N(t))$$
$$\in \mathrm{Feas}\left(\mathrm{OTP}_{\mathrm{reach}}(x_1(t), \dots, x_N(t), \mathbf{x}_1^{t-1}(T_1), \dots, \mathbf{x}_N^{t-1}(T_N))\right)$$

7:        **for** agents $i \in \{1, \dots, N\}$:

8:           **compute** temporary target and terminal set:

$$\chi_{\mathrm{tmp},i}(t) = N \sum_{j=1}^{N} P_{ij}(t) x_j^d, \quad \Omega_{\mathrm{tmp},i}(t) = \chi_{\mathrm{tmp},i}(t) \oplus \Omega_i$$

9:           **compute** optimal control trajectory $(\mathbf{x}_i^t, \mathbf{u}_i^t) \in \mathrm{OCP}_i(x_i(t), \Omega_{\mathrm{tmp},i}(t))$

10:          **set** control input $u_i(t) = \mathbf{u}_i^t(0)$

11:          **if** $P(t)$ is a permutation matrix:

12:             **set** fixed transport plan start time $t^\star = t$

13:     **else**:

14:        **set** fixed transport plan $P(t) = P(t-1)$

15:        **for** agents $i \in \{1, \dots, N\}$:

16:           **compute** target index $j_i \in \{1, \dots, N\}$ such that $P_{ij_i}(t) = 1/N$

17:           **if** $t - t^\star < T_i$:

18:             **set** control input $u_i(t) = \mathbf{u}_i^{t^\star}(t - t^\star) + K_i(x_i(t) - \mathbf{x}_i^{t^\star}(t - t^\star))$

19:           **else**:

20:             **set** control input $u_i(t) = u_{ij_i}^d + K_i(x_i(t) - x_{j_i}^d)$

21:     **update** state $x_i(t+1) = A_i x_i(t) + B_i u_i(t) + w_i(t)$

3.2. Summary of the Control Scheme. Overall, the proposed control scheme, which we call ROT-MPC in short for "robust optimal transport MPC," is summarized in Algorithm 3.1.[1] We remark the step where, once the (robustly feasible) transport plan recovers a permutation, that plan and its subsequent optimal control policies are no longer re-optimized. Then, once an agent i reaches its target set $\Omega_{\mathrm{tmp},i}(t)$ that is generated by the permutation plan, that agent's control law switches to the simple and efficient static feedback control defined by K_i and the equilibrium state-input pair, which robustly maintains agent i within their final target set for all subsequent time (cf., Theorem 4.10). This tri-mode control scheme can

---

[1]The initializations $\mathbf{x}_i^{-1}(T_i)$ are chosen to reflect a T_i-step ahead prediction from the initial condition under a sort of "average" unconstrained closed-loop control given by $u_i(t) = \frac{1}{N} \sum_{j=1}^{d} u_{ij}^d + K_i(x_i(t) - \frac{1}{N} \sum_{j=1}^{N} x_j^d)$ generated by the target equilibria. Other initializations are possible, e.g., letting $\mathbf{x}_i^{-1}(T_i)$ be a T_i-step ahead prediction that comes about from steering the system from $x_{i,0}$ to the convex hull of $\bigcup_{j=1}^{d} \Omega_{ij}^d$ via optimal control with robustified constraints. An open problem for future research is to determine how these initializations can be chosen to ensure that the first optimal transport problem is feasible.

be thought of as having three phases: 1) alternating centralized transport planning and decentralized policy optimization, 2) decentralized fixed-policy control to achieve robust finite-time stability, and 3) decentralized static feedback control to robustly maintain the distribution around the desired equilibrium distribution. Although the second and third phases are not strictly necessary for feasibility or stability in practice, they are key in rigorously proving stability, as those phases guarantee the absence of "switching" behavior where an agent gets redirected towards a different target in the state space after it has already stabilized near another target. Fixing the transport plan at a robustly feasible permutation also aids with computational efficiency, as the centralized reach-constrained optimal transport plan can be skipped entirely at subsequent time steps, while maintaining feasibility.

4. Theoretical Analysis of ROT-MPC. In this section, we mathematically analyze the properties of ROT-MPC. Specifically, we consider recursive feasibility in Section 4.1, recursive constraint satisfaction and uniform boundedness in Section 4.2, and stability in Section 4.3.

4.1. Recursive Feasibility. We begin our theoretical analysis of ROT-MPC by showing that its optimization problems exhibit advantageous feasibility properties.

DEFINITION 4.1. *The system ROT-MPC is called* recursively feasible *if, for all $t \in \mathbb{N} \cup \{0\}$ for which $P(t-1)$ is not a permutation matrix, it holds that* $\mathrm{OTP}_{\mathrm{reach}}(x_1(t+1), \ldots, x_N(t+1), \mathbf{x}_1^t(T_1), \ldots, \mathbf{x}_N^t(T_N))$ *and every* $\mathrm{OCP}_i(x_i(t+1), \Omega_{\mathrm{tmp},i}(t+1))$ *are feasible whenever* $\mathrm{OTP}_{\mathrm{reach}}(x_1(t), \ldots, x_N(t), \mathbf{x}_1^{t-1}(T_1), \ldots, \mathbf{x}_N^{t-1}(T_N))$ *and every* $\mathrm{OCP}_i(x_i(t), \Omega_{\mathrm{tmp},i}(t))$ *are feasible.*

Note that if the optimization problems in ROT-MPC are feasible at the time $t^\star$ where $P(t)$ first becomes a permutation matrix, then clearly the remainder of the algorithm (the "second" stage and the "third" stage, after the optimizations terminate) is well-defined without any further concerns of feasibility. Therefore, one really only needs to be concerned with recursive feasibility for times t leading up to the computation of a permutation transport plan $P(t^\star)$, as in Definition 4.1.

To establish the recursive feasibility of ROT-MPC, we will utilize the following lemma ensuring that the temporary terminal sets always satisfy the state constraints:

LEMMA 4.2. *For all $i \in \{1, \ldots, N\}$ and all $P \in \mathcal{B}_N$, it holds that $\Omega_{\mathrm{tmp},i} := N \bigoplus_{j=1}^N P_{ij} \Omega_{ij}^d$ is a subset of X_i.*

We now show that the reach-constrained optimal transport problem is always feasible whenever the predictions $\hat{x}_i(T_i)$ are known to be contained in robustified "past" temporary terminal sets generated by a valid transport plan.

PROPOSITION 4.3. *Let $T_1, \ldots, T_N \in \mathbb{N}$ and consider arbitrary states $x_1, \ldots, x_N, \hat{x}_1(T_1), \ldots, \hat{x}_N(T_N) \in \mathbb{R}^n$. If there exists $Q \in \mathcal{B}_N$ such that $\hat{x}_i(T_i) \in \left(N \bigoplus_{j=1}^N Q_{ij} \Omega_{ij}^d \right) \ominus R_i^{T_i}$ for all $i \in \{1, \ldots, N\}$, then (3.1) is feasible for some $(P, \hat{u}_1, \ldots, \hat{u}_N)$ with $P = Q$.*

Proof sketch. The Minkowski convex combination $N \bigoplus_{j=1}^N Q_{ij} \Omega_{ij}^d$ is robust control invariant, and therefore, there exists a feasible input that maintains agent i within this set when taking one step from $\hat{x}_i(T_i) + w_i$, with $w_i \in R_i^{T_i}$ being some compounded disturbance. This corresponds to satisfaction of the subset constraint in (3.1). We explicitly identify such an input,

$$\hat{u}_i = N \sum_{j=1}^N Q_{ij} u_{ij}^d + K_i \left(\hat{x}_i(T_i) - N \sum_{j=1}^N Q_{ij} x_j^d \right),$$

and show that it satisfies the tightened control constraint in (3.1). A complete proof is deferred to Appendix A to streamline the flow of the article. □

The following corollary shows that, if $P(t) \in \mathcal{B}_N$ is a feasible transport plan for the reach-constrained optimal transport problem at time t, then it is guaranteed to remain feasible at the next time step when it is used to construct the temporary targets and terminal sets:

COROLLARY 4.4. *Consider the closed-loop dynamics of ROT-MPC at some time* $t \in \mathbb{N} \cup \{0\}$. *If* $(\overline{P}, \hat{u}_1, \ldots, \hat{u}_N)$ *is a feasible point for* $\mathrm{OTP}_{\mathrm{reach}}(x_1(t), \ldots, x_N(t), \mathbf{x}_1^{t-1}(T_1), \ldots, \mathbf{x}_N^{t-1}(T_N))$ *and is used to define* $\chi_{\mathrm{tmp},i}(t) = N \sum_{j=1}^N P_{ij}(t) x_j^d$ *and* $\Omega_{\mathrm{tmp},i}(t) = \chi_{\mathrm{tmp},i}(t) \oplus \Omega_i$ *with* $P(t) = \overline{P}$, *and if every* $\mathrm{OCP}_i(x_i(t), \Omega_{\mathrm{tmp},i}(t))$ *is feasible, then* $(\overline{P}, \hat{u}_1', \ldots, \hat{u}_N')$ *is feasible for* $\mathrm{OTP}_{\mathrm{reach}}(x_1(t+1), \ldots, x_N(t+1), \mathbf{x}_1^t(T_1), \ldots, \mathbf{x}_N^t(T_N))$ *for some* $\hat{u}_1' \in \mathbb{R}^{m_1}, \ldots, \hat{u}_N' \in \mathbb{R}^{m_N}$.

Proof. It holds by the dynamics of ROT-MPC that $\mathbf{x}_i^t(T_i) \in \left(N \bigoplus_{j=1}^N \overline{P}_{ij} \Omega_{ij}^d \right) \ominus R_i^{T_i}$ for all $i \in \{1, \ldots, N\}$, and hence, by Proposition 4.3, $(\overline{P}, \hat{u}_1', \ldots, \hat{u}_N')$ is feasible for $\mathrm{OTP}_{\mathrm{reach}}(x_1(t+1), \ldots, x_N(t+1), \mathbf{x}_1^t(T_1), \ldots, \mathbf{x}_N^t(T_N))$ for some $\hat{u}_1', \ldots, \hat{u}_N'$. □

Remark 4.5. There are two primary implications of Corollary 4.4. First, one may decide to skip solving the reach-constrained optimal transport problem at any time steps in the dynamics of ROT-MPC that they would like, so long as a previously (robustly) feasible transport plan is used in the control problems. This clearly gives rise to computational benefits, as the reach-constrained optimal transport problem need not be solved at every time step. Second, if one would like to continue re-optimizing the control policies at every time step after finding a permutation transport plan, instead of fixing the control policies as is done in ROT-MPC, then the robustly feasible permutation plan will continue to give rise to feasible control problems. This may be desired in the case that minimizing control trajectory losses is more important than the computational burdens of solving the optimal control problems. However, we remark that this approach of re-optimizing the control policies (and hence never entering the "third" stage of ROT-MPC) violates the assumptions of our stability guarantees in Theorem 4.10.

The above properties give rise to our main feasibility result, i.e., recursive feasibility of the closed-loop dynamics of ROT-MPC:

THEOREM 4.6. *The system ROT-MPC is recursively feasible.*

Proof sketch. Based on the optimal control solution $(\mathbf{x}_i^t, \mathbf{u}_i^t)$ at time t, and the input $\hat{u}_i(t+1)$ that is feasible for the reach-constrained transport problem at time $t+1$ (per Proposition 4.3), a candidate control trajectory is defined for time $t+1$:

$$
\begin{aligned}
\mathbf{x}_i^{t+1}(0) &:= x_i(t+1), \\
\mathbf{u}_i^{t+1}(k) &:= \mathbf{u}_i^t(k+1) + K_i(\mathbf{x}_i^{t+1}(k) - \mathbf{x}_i^t(k+1)), \ k \in \{0, \ldots, T_i - 2\}, \\
\mathbf{u}_i^{t+1}(T_i - 1) &:= \hat{u}_i(t+1) + K_i(\mathbf{x}_i^{t+1}(T_i - 1) - \mathbf{x}_i^t(T_i)), \\
\mathbf{x}_i^{t+1}(k+1) &:= A_i \mathbf{x}_i^{t+1}(k) + B_i \mathbf{u}_i^{t+1}(k), \ k \in \{0, \ldots, T_i - 1\}.
\end{aligned}
$$
(4.1)

This candidate trajectory is shown to be feasible for the optimal control problem at time $t+1$ by bounding its deviation from $(\mathbf{x}_i^t, \mathbf{u}_i^t)$ using the disturbance tube polytopes. A complete proof is deferred to Appendix A to streamline the flow of the article. □

4.2. Recursive Constraint Satisfaction. We now show that ROT-MPC enjoys recursive constraint satisfaction, meaning that the closed-loop states and inputs

always satisfy the constraints defined by X_i and U_i, even in the presence of the uncertainty. This result, as well as our stability theory in Section 4.3, are rooted in the following key lemma. The lemma amounts to combining ideas from the single-policy disturbance invariant MPC scheme studied in [21] and the "growing" tube introduced in [11]. The result shows that, once an agent begins following a fixed control policy that robustly steers the disturbance-free system to the terminal set, the true dynamics can be bounded within the growing tube surrounding the nominal trajectory:

LEMMA 4.7. *Consider a trajectory of the system ROT-MPC. Suppose that, for some $t^\star \in \mathbb{N} \cup \{0\}$, the computed transport plan $P(t^\star) \in \mathcal{B}_N$ is a permutation matrix, and that $P(t)$ is non-permutation for all $t < t^\star$. Then, for all agents $i \in \{1, \dots, N\}$, the following all hold:*

1. $x_i(t) \in \mathbf{x}_i^{t^\star}(t - t^\star) \oplus R_i^{t-t^\star} \subseteq X_i$ *for all $t \in \{t^\star, \dots, t^\star + T_i - 1\}$,*
2. $u_i(t) \in \mathbf{u}_i^{t^\star}(t - t^\star) \oplus (K_i R_i^{t-t^\star}) \in U_i$ *for all $t \in \{t^\star, \dots, t^\star + T_i - 1\}$, and*
3. $x_i(t^\star + T_i) \in \mathbf{x}_i^{t^\star}(T_i) \oplus R^{T_i} \subseteq \Omega_{ij_i}^d$,

where $j_i \in \{1, \dots, N\}$ is the unique target index for which $P_{ij_i}(t^\star) = 1/N$.

Proof. Let $i \in \{1, \dots, N\}$ be arbitrary. We start by relating the true state $x_i(t)$ to the nominal state $\mathbf{x}_i^{t^\star}(t - t^\star)$ for times $t \in \{t^\star, \dots, t^\star + T_i\}$, via induction on t. It is clear that $x_i(t^\star) = \mathbf{x}_i^{t^\star}(0)$ by the initial condition constraint in the optimal control problem solved at time $t^\star$. At the subsequent time, we see that $x_i(t^\star + 1) = A_i x_i(t^\star) + B_i u_i(t^\star) + w_i(t^\star) = \mathbf{x}_i^{t^\star}(1) + w_i(t^\star)$, since $x_i(t^\star) = \mathbf{x}_i^{t^\star}(0)$ and $u_i(t^\star) = \mathbf{u}_i^{t^\star}(0) + K_i(x_i(t^\star) - \mathbf{x}_i^{t^\star}(0)) = \mathbf{u}_i^{t^\star}(0)$. These relations constitute the base case for induction. Now, suppose that $t \in \{t^\star, \dots, t^\star + T_i - 1\}$ is such that

$$(4.2) \qquad x_i(t) = \mathbf{x}_i^{t^\star}(t - t^\star) + \sum_{j=0}^{t-t^\star-1} (A_i + B_i K_i)^j w_i(t - j - 1),$$

which serves as the inductive hypothesis. Then, it holds that

$$\begin{aligned}
&x_i(t+1) - \mathbf{x}_i^{t^\star}(t + 1 - t^\star) \\
&= (A_i + B_i K_i)(x_i(t) - \mathbf{x}_i^{t^\star}(t - t^\star)) + w_i(t) \\
&= \sum_{j=0}^{t-t^\star-1} (A_i + B_i K_i)^{j+1} w_i(t - j - 1) + w_i(t) \\
&= \sum_{j=0}^{(t+1)-t^\star-1} (A_i + B_i K_i)^j w_i((t+1) - j - 1),
\end{aligned}$$

which shows that (4.2) also holds at time $t + 1$. Thus, by induction, it must be that (4.2) holds for all $t \in \{t^\star, \dots, t^\star + T_i\}$.

We now use the relation (4.2) to prove that the three enumerated claims hold. The first is obvious from (4.2), since, for all $t \in \{t^\star, \dots, t^\star + T_i - 1\}$, it holds that

$$\sum_{j=0}^{t-t^\star-1} (A_i + B_i K_i)^j w_i(t - j - 1) \in \bigoplus_{j=0}^{t-t^\star-1} (A_i + B_i K_i)^j W_i = R_i^{t-t^\star},$$

and since $\mathbf{x}_i^{t^\star}(t - t^\star) \in X_i \ominus R_i^{t-t^\star}$ by the robustified state constraints in the optimal

control problem at time $t^\star$. Next, we see for all such t that

$$u_i(t) = \mathbf{u}_i^{t^\star}(t - t^\star) + K_i \sum_{j=0}^{t-t^\star-1} (A_i + B_i K_i)^j w_i(t - j - 1)$$

$$\in \mathbf{u}_i^{t^\star}(t - t^\star) \oplus (K_i R_i^{t-t^\star})$$

$$\subseteq U_i,$$

where, again, the final subset inclusion follows from the fact that $\mathbf{u}_i^{t^\star}(t - t^\star) \in U_i \ominus (K_i R_i^{t-t^\star})$ from the optimal control constraints. Thus, the second enumerated claim holds. The final enumerated claim holds similarly:

$$x_i(t^\star + T_i) = \mathbf{x}_i^{t^\star}(T_i) + \sum_{j=0}^{T_i-1} (A_i + B_i K_i)^j w_i(t - j - 1) \in \mathbf{x}_i^{t^\star}(T_i) \oplus R_i^{T_i} \subseteq \Omega_{ij_i}^d,$$

where the final subset inclusion holds due to the robust terminal state constraint in the optimal control problem at time $t^\star$, and the fact that $P(t^\star)$ is a permutation matrix and hence $\Omega_{\mathrm{tmp},i}(t^\star) = \Omega_{ij_i}^d$ for some unique index $j_i \in \{1, \ldots, N\}$ satisfying $P_{ij_i} = 1/N$. □

THEOREM 4.8. *Assume that all of the initial optimization problems* $\mathrm{OTP}_{\mathrm{reach}}(x_1(0), \ldots, x_N(0), \mathbf{x}_1^{-1}(T_1), \ldots, \mathbf{x}_N^{-1}(T_N))$ *and* $\mathrm{OCP}_i(x_i(0), \Omega_{\mathrm{tmp},i}(0))$ *are feasible, so that ROT-MPC is feasible at every time* $t \in \mathbb{N} \cup \{0\}$ *per Theorem 4.6. Then, for all* $i \in \{1, \ldots, N\}$, *it holds that* $x_i(t) \in X_i$ *and* $u_i(t) \in U_i$ *for all* $t \in \mathbb{N} \cup \{0\}$ *for the closed-loop system ROT-MPC.*

Proof sketch. At times t with $P(t-1)$ a non-permutation matrix, one can show that $x_i(t) = \mathbf{x}_i^{t-1} + w_i(t-1) \in (X_i \ominus R_i^1) \oplus W_i \subseteq X_i$ and that $u_i(t) = \mathbf{u}_i^t(0) \in U_i \ominus (K_i R_i^0) = U_i$. At times t with $P(t-1)$ a permutation matrix, Lemma 4.7 gives that $x_i(t) \in X_i$ and $u_i(t) \in U_i$ if $t < t^\star + T_i$, and otherwise $x_i(t)$ is in some robust control invariant set $\Omega_{ij_i}^d \subseteq X_i$ and thus the static feedback controller $u_i(t) = u_{ij_i}^d + K_i(x_i(t) - x_{j_i}^d)$ feasibly maintains the agent's state to be within that set. A complete proof is deferred to Appendix A to streamline exposition. □

The recursive constraint satisfaction of Theorem 4.8 immediately implies uniform boundedness of the closed-loop system within bounded state constraints:

COROLLARY 4.9. *Assume that all of the initial optimization problems* $\mathrm{OCP}_i(x_i(0), \Omega_{\mathrm{tmp},i}(0))$ *and* $\mathrm{OTP}_{\mathrm{reach}}(x_1(0), \ldots, x_N(0), \mathbf{x}_1^{-1}(T_1), \ldots, \mathbf{x}_N^{-1}(T_N))$ *are feasible. If every* X_i *is compact, then the closed-loop trajectory* $t \mapsto (x_1(t), \ldots, x_N(t))$ *of ROT-MPC is uniformly bounded in* $X_1 \times \cdots X_N$, *meaning that* $\|x_i(t)\|_2 \leq M_i := \max_{x \in X_i} \|x\|_2 < \infty$ *for all* $i \in \{1, \ldots, N\}$ *and all* $t \in \mathbb{N} \cup \{0\}$.

4.3. Stability. We now turn to our main stability result: if ROT-MPC starts off feasible and, at some time, a permutation transport plan is computed, then every agent reaches, and stays within, its final target set in finite time.

THEOREM 4.10. *Assume that all of the initial optimization problems* $\mathrm{OTP}_{\mathrm{reach}}(x_1(0), \ldots, x_N(0), \mathbf{x}_1^{-1}(T_1), \ldots, \mathbf{x}_N^{-1}(T_N))$ *and* $\mathrm{OCP}_i(x_i(0), \Omega_{\mathrm{tmp},i}(0))$ *are feasible, so that ROT-MPC is feasible at every time* $t \in \mathbb{N} \cup \{0\}$ *per Theorem 4.6. Furthermore, suppose that for some* $t^\star \in \mathbb{N} \cup \{0\}$, *the computed transport plan* $P(t^\star) \in \mathcal{B}_N$ *is a permutation matrix, and that* $P(t)$ *is non-permutation for all* $t < t^\star$.

Then, for all agents $i \in \{1, \ldots, N\}$, it holds that $\Omega_{ij_i}^d$ is robustly finite-time attractive, meaning that

$$x_i(t) \in X_i \text{ for all } t \in \mathbb{N} \cup \{0\},$$
$$u_i(t) \in U_i \text{ for all } t \in \mathbb{N} \cup \{0\}, \text{ and}$$
$$x_i(t) \in \Omega_{ij_i}^d = x_{j_i}^d \oplus \Omega_i \text{ for all } t \in \{t^\star + T_i, t^\star + T_i + 1, \ldots\},$$

where $j_i \in \{1, \ldots, N\}$ is the unique target index for which $P_{ij_i}(t^\star) = 1/N$.

Proof. First, the constraint satisfactions $x_i(t) \in X_i$ and $u_i(t) \in U_i$ for all $t \in \mathbb{N} \cup \{0\}$ follow immediately from Theorem 4.8. Next, Lemma 4.7 gives that $x_i(t^\star + T_i) \in \Omega_{ij_i}^d$. At this time and subsequent times, the control input becomes the static feedback control $u_i(t) = u_{ij_i}^d + K_i \left(x_i(t) - x_{j_i}^d \right)$. If $t \in \{t^\star + T_i, t^\star + T_i + 1, \ldots\}$ and $x_i(t) \in \Omega_{ij_i}^d$, then $x_i(t) = x_{j_i}^d + \omega_i(t)$ for some $\omega_i(t) \in \Omega_i$ and hence

$$x_i(t+1) = A_i(x_{j_i}^d + \omega_i(t)) + B_i(u_{ij_i}^d + K_i\omega_i(t)) + w_i(t)$$
$$= x_{j_i}^d + (A_i + B_iK_i)\omega_i(t) + w_i(t),$$

which shows that $x_i(t+1) \in x_{j_i}^d \oplus (((A_i + B_iK_i)\Omega_i) \oplus W_i) \subseteq x_{j_i}^d \oplus \Omega_i = \Omega_{ij_i}^d$. Thus, by induction, it holds that $x_i(t) \in \Omega_{ij_i}^d$ for all $t \in \{t^\star + T_i, t^\star + T_i + 1, \ldots\}$. $\square$

Remark 4.11. We find in our simulations (Section 5) that the reach-constrained transport problem typically returns a permutation transport plan at either the first or second time step of the control scheme. Therefore, the assumption in Theorem 4.10 of computing a permutation transport plan appears mild, in practice. Indeed, sufficiently long optimal control time horizons and sufficiently lax input constraint sets U_i both increase the likelihood of a permutation transport plan becoming feasible.

5. Numerical Simulations. In this section, we test ROT-MPC on three distribution steering tasks. We consider a homogeneous population of agents, all having the discrete-time dynamics considered in [18]:

$$A_i = \begin{bmatrix} 1.04 & 0.026 \\ -0.01 & 1.02 \end{bmatrix}, \quad B_i = 0.02I_2, \quad i \in \{1, \ldots, N\}.$$

The open-loop agent dynamics are naturally unstable, as A_i has two complex eigenvalues outside of the unit circle. We subject every agent to the state constraints defined by $X_i = [-2, 2]^2, \quad i \in \{1, \ldots, N\}$, and we consider various input constraints and uncertainty bounds in each of the three tasks. In all tasks, the uncertainty at each time step is a random vertex of the uncertainty polytope W_i. To ensure a fair comparison between the three methods tested (the baselines are described below), the actual sequences of realized uncertainties, $w_i(t)$, are made the same for all three control schemes. All three methods are simulated for $T_{\text{sim}} = 40$ time steps on every task. Additional details regarding the simulation setup, as well as the implementation of ROT-MPC and the baselines, are provided in Appendix B.

Baselines. On all three tasks, we compare ROT-MPC against the baseline "Sinkhorn MPC," introduced in [16, 18]. Sinkhorn MPC is the closest related technique to ROT-MPC, as it also performs distribution steering by alternating between optimal transport updates for target assignment, and agent-wise optimal control steps for determining how to steer each agent to the updated target. However, Sinkhorn MPC is not designed to explicitly handle uncertainties, and as such has no guarantees

of feasibility, constraint satisfaction, or stability, in such uncertain cases. Sinkhorn MPC sets the optimal transport cost C_{ij} equal to the optimal control cost of steering agent i to target j (with non-robust state and input constraints). As such, it requires solving N^2 optimal control problems at every time step in order to update the targets.

The first task (described below) is small-scale, allowing us to also compare ROT-MPC against a centralized robust MPC control scheme. The centralized MPC scheme simultaneously optimizes the transport plan and *all* of the agent-wise controls at every time step by solving the following problem:

$$\min_{P \text{ permutation}} \quad \min_{\substack{(\overline{x}_1(0),\ldots,\overline{x}_1(T_1)), \\ (\overline{u}_1(0),\ldots,\overline{u}_1(T_1-1)),\ldots \\ (\overline{x}_N(0),\ldots,\overline{x}_N(T_N)), \\ (\overline{u}_N(0),\ldots,\overline{u}_N(T_N-1))}} \quad \sum_{i=1}^{N} L_i(\overline{x}_i(0),\ldots,\overline{x}_i(T_i),\overline{u}_i(0),\ldots,\overline{u}_i(T_i-1))$$

$$\text{subject to} \quad \overline{x}_i(0) = x_i(t) \text{ for all } i,$$
$$\overline{x}_i(k+1) = A_i\overline{x}(k) + B_i\overline{u}(k) \text{ for all } i,k,$$
$$\overline{x}_i(k) \in X_i \ominus R_i^k \text{ for all } i,k,$$
$$\overline{u}_i(k) \in U_i \ominus (K_iR_i^k) \text{ for all } i,k,$$
$$\overline{x}_i(T_i) \in \left(N\bigoplus_{j=1}^{N} P_{ij}\Omega_{ij}^d\right) \ominus R_i^{T_i} \text{ for all } i.$$

Here, the optimization variables are $P \in \mathcal{B}_N$, $\overline{x}_i(k) \in \mathbb{R}^n$, and $\overline{u}_i(k) \in \mathbb{R}^{m_i}$, where the variables and constraint indices range over $k \in \{0,\ldots,T_i-1\}$ and $i \in \{1,\ldots,N\}$. At every time t, every agent i applies its associated optimal input $\overline{u}_i(0)$ from the solution to the above centralized optimization problem. Since the number of permutation matrices $P \in \mathcal{B}_N$ is $N!$, this centralized MPC scheme quickly becomes intractable at moderate population sizes N.

Performance Metrics. Since the three methods tested utilize different loss functions in their optimization problems, we choose to measure performance using the following method-agnostic metrics:

$$\ell_{\text{state}} := \frac{1}{T_{\text{sim}}+1} \sum_{t=0}^{T_{\text{sim}}} \frac{1}{N} \sum_{i=1}^{N} \min_{j\in\{1,\ldots,N\}} \|x_i(t) - x_j^d\|_2,$$

$$\ell_{\text{input}} := \frac{1}{T_{\text{sim}}} \sum_{t=0}^{T_{\text{sim}}-1} \frac{1}{N} \sum_{i=1}^{N} \|u_i(t)\|_2^2.$$

The value ℓ_{state} quantifies the average deviation of the agents' states from their nearest targets, over the course of the simulation. Low values of ℓ_{state} indicate that the method works well at steering the population towards the target distribution quickly and maintaining it there. The value ℓ_{input} quantifies the average control energy expenditure throughout the simulation.

We also compare the speed of convergence to the target distribution by computing

$$T_{\text{target}} := \inf\left\{t \in \mathbb{N} \cup \{0\} : x_i(t) \in \Omega_{ij_i}^d \text{ for all } i \in \{1,\ldots,N\}\right\},$$

which is the first time at which every agent is within their terminal invariant set $\Omega_{ij_i}^d$. The index j_i denotes the index of the target state to which agent i is assigned by the end of the simulation. If at least one agent never enters its associated terminal set $\Omega_{ij_i}^d$, then $T_{\text{target}} = \infty$, indicating that convergence to the target distribution has not been achieved. We still compute this metric for Sinkhorn MPC using the invariant sets Ω_{ij}^d constructed in Section 2.3, even though it is *not* guaranteed that the agents will remain in such sets under the Sinkhorn MPC control scheme.

Task 1: Small-Scale Comparison to Centralized Robust MPC. In this task, we consider $N = 3$ agents, initialized uniformly at random within the region $[-2, -1]^2$. The target states are $x_1^d = (1, 0)$, $x_2^d = (0, 1)$, and $x_3^d = (1, 1)$. For all $i \in \{1, \ldots, N\}$, the inputs are constrained by $U_i = [-20, 20]^2$ and the disturbance polytope is $W_i = [-0.1, 0.1]^2$. All three methods utilize control time horizons of $T_i = 10$ time steps for every agent.

The agent trajectories for ROT-MPC, Sinkhorn MPC, and centralized MPC are shown in Figure 1. All three methods are seen to eventually stabilize the population around the target distribution. However, Sinkhorn MPC exhibits "switching" behavior, where some agents are redirected to a different target states after they have already reached the vicinity of another state. Fixing the transport plan, once a feasible permutation has been found, avoids this issue, as illustrated in the trajectories of ROT-MPC. This also aids with computation speeds; even though Sinkhorn MPC enjoys much faster optimal control solve times (0.005 seconds versus 0.274 seconds) and transport plan update times (0.066 seconds versus 0.452 seconds) than ROT-MPC, the continued re-optimization at every time step eventually causes the entire task to take longer for Sinkhorn MPC, as tabulated in Table 1. The performance of ROT-MPC is qualitatively near-identical to centralized MPC, which can be considered as the optimal baseline due to its simultaneous optimization over transport plans and all of the control sequences. As shown in Table 1, the quantitative performance is very similar as well, with nearly equal state and input performance metric values, both of which outperform Sinkhorn MPC.

TABLE 1
Performance of the various methods on Tasks 1, 2, and 3. Compute time is the total time taken to complete the simulation. Best values on a given task are bold.

Simulation	ℓ_{state}	ℓ_{input}	T_{target}	Compute time (seconds)
Task 1: ROT-MPC	0.389	**113.953**	**8**	**2.131**
Task 1: Sinkhorn MPC	0.456	124.257	24	3.329
Task 1: Centralized MPC	**0.388**	114.299	**8**	388.083
Task 2: ROT-MPC	**0.133**	**31.058**	9	81.395
Task 2: Sinkhorn MPC	0.155	82.508	26	**38.655**
Task 3: ROT-MPC	**0.261**	**48.691**	9	**19.805**
Task 3: Sinkhorn MPC	1.013	174.018	∞	170.530

Task 2: Scalability to Larger Populations. In this task, we test the performance of ROT-MPC and Sinkhorn MPC with an increased population size of $N = 12$. The initial states are chosen uniformly at random from $[-0.25, 0.25]^2$, shifted by either $(1, 1)$, $(-1, 1)$, $(-1, -1)$, or $(1, -1)$. The target states are arranged in a uniform grid on the line segment $\{(x, y) \in \mathbb{R}^2 : x = 0, \ y \in [-1.75, 1.75]\}$. For all $i \in \{1, \ldots, N\}$, we use the same input constraints as in Task 1, namely, $U_i = [-20, 20]^2$, but we shrink the disturbance polytopes to $W_i = [-0.05, 0.05]^2$. The time horizon is again chosen as $T_i = 10$ time steps for every agent.

The agent trajectories for ROT-MPC and Sinkhorn MPC are shown in Figure 2. Both methods eventually solve the steering task. However, ROT-MPC is seen to yield much more "practical" trajectories than Sinkhorn MPC, which appears to send agents to targets much further away than is necessary. This suboptimality in the transport planning is likely due to ROT-MPC solving the (reach-constrained) transport problem
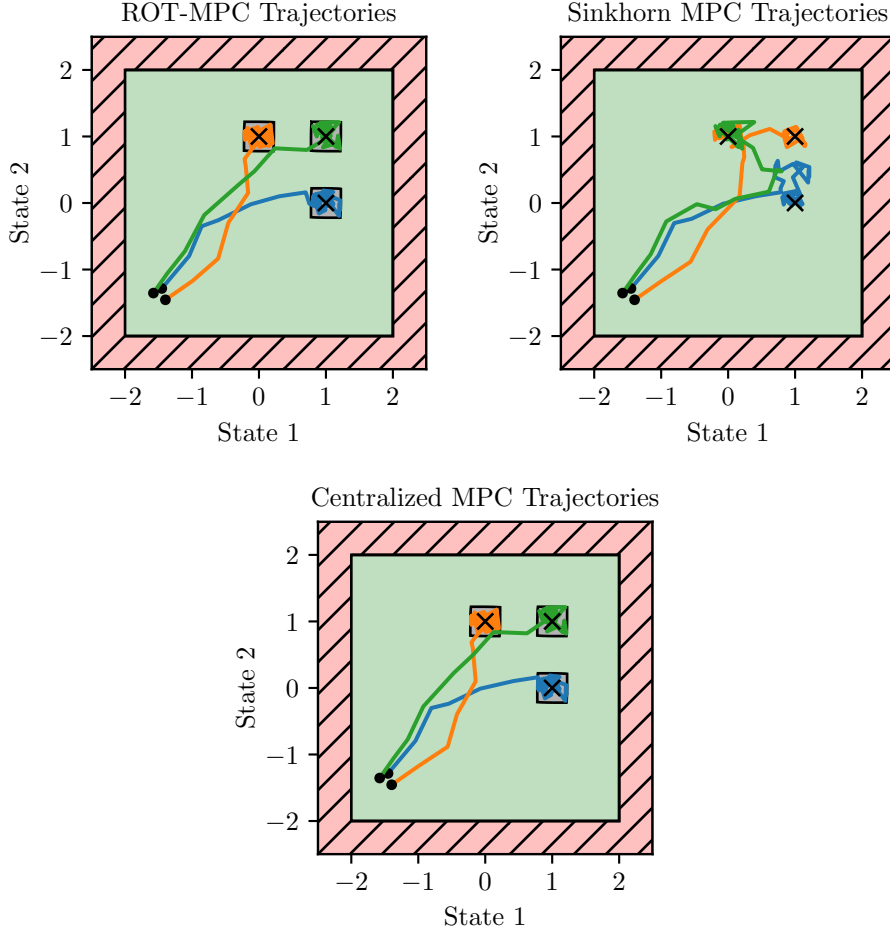
FIG. 1. *Agent trajectories for Task 1. Black circles indicate initial states and crosses indicate target states. Grey regions are the terminal invariant sets $\Omega^d_{ij_i}$ for the two methods with robustness guarantees (ROT-MPC and centralized MPC). The red hatched region is the "unsafe" set $\mathbb{R}^n \setminus X_i$.*

to optimality, whereas Sinkhorn MPC's transport plans may try to "hold on" to the initialization used by the Sinkhorn algorithm. Interesting to note is that, since the targets are quite close together and Sinkhorn MPC makes no guarantees that the agents will remain within terminal regions surrounding their associated target states, the agents appear to move somewhat freely between the targets, along the line segment $\{(x, y) \in \mathbb{R}^2 : x = 0, \ y \in [-1.75, 1.75]\}$, due to the disturbances faced by Sinkhorn MPC. For instance, the pink agent is seen to reach the third-from-top target, but then meander down to the fourth-from-bottom target. Contrarily, the agents perfectly remain within their associated terminal invariant sets under ROT-MPC. As seen in Table 1, ROT-MPC is seen to quantitatively outperform Sinkhorn MPC in all metrics on this task, except the overall compute time.

Task 3: Resilience to Tighter Constraints and Shorter Horizon. In this task, we test how ROT-MPC and Sinkhorn MPC adapt when the task becomes more difficult, with more stringent constraints. We consider $N = 10$ agents, with initial states positioned on a circle of radius 0.25, and target states positioned on a circle
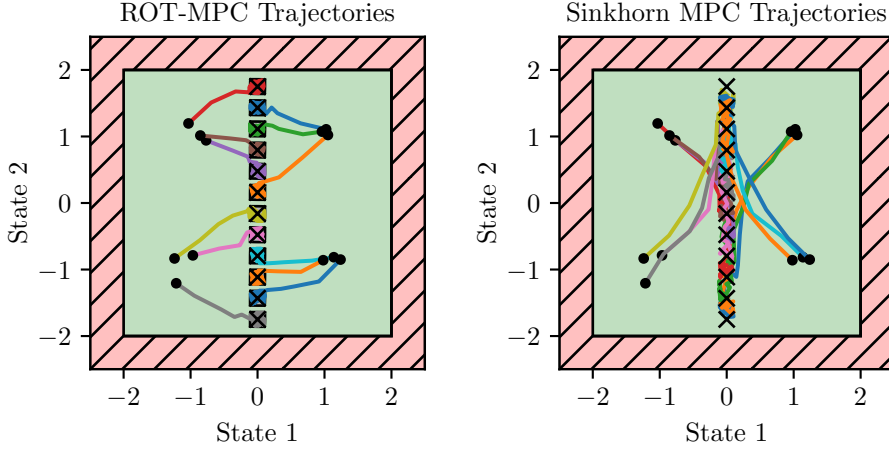
FIG. 2. *Agent trajectories for Task 2. Black circles indicate initial states and crosses indicate target states. Grey regions are the terminal invariant sets $\Omega_{ij_i}^d$ for ROT-MPC. The red hatched region is the "unsafe" set $\mathbb{R}^n \setminus X_i$.*

of radius 1.75. For all $i \in \{1, \ldots, N\}$, the input constraints are halved, to $U_i = [-10, 10]^2$, and the uncertainties are set back to the larger disturbances from Task 1: $W_i = [-0.1, 0.1]^2$. In addition to the tightened input constraints, the optimal control time horizon is reduced to $T_i = 4$ time steps for every agent.

The agent trajectories are shown in Figure 3. Note that the geometry of the targets makes this task have little room for error, as some of the targets are very close to the boundary of the constraint set X_i. It is clear that ROT-MPC successfully solves the steering task, whereas Sinkhorn MPC fails entirely, highlighting the capacity that ROT-MPC has to handle difficult, highly constrained tasks in the presence of uncertainties. Quantitative performance metrics are reported in Table 1.
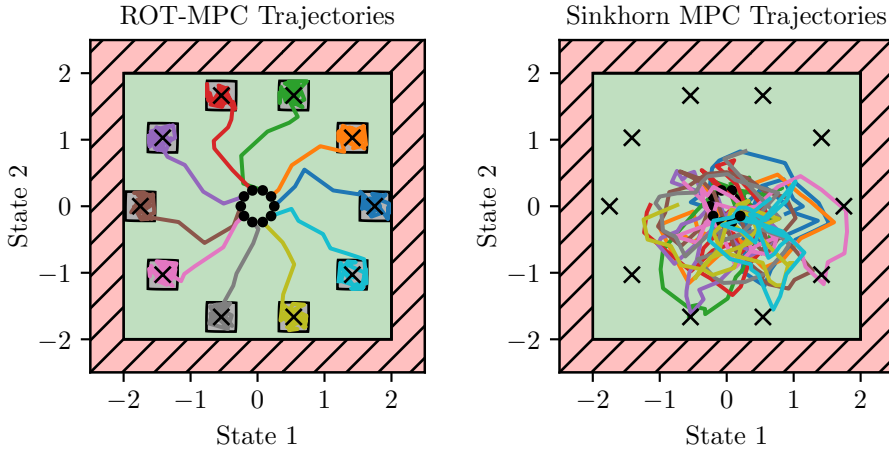

FIG. 3. *Agent trajectories for Task 3. Black circles indicate initial states and crosses indicate target states. Grey regions are the terminal invariant sets $\Omega_{ij_i}^d$ for ROT-MPC. The red hatched region is the "unsafe" set $\mathbb{R}^n \setminus X_i$.*

6. Conclusions. In this article, we introduce ROT-MPC, a method for robustly steering a distribution of decoupled discrete-time dynamical agents in the presence of bounded uncertainty. We prove that a reach-constrained modification to the optimal transport problem results in the closed-loop system inheriting recursive feasibility, recursive constraint satisfaction, and robust finite-time attraction of terminal invariant sets in the practical situation in which the transport problem returns a permutation matrix. Our numerical simulations demonstrate that ROT-MPC generally outperforms baseline optimal transport-based steering techniques, yielding more interpretable agent trajectories, and stabilizing the system to the target distribution even in cases where the baselines fail.

A handful of directions for future research remain. First, the reach-constrained optimal transport problem may pose a practical hurdle in cases where the centralized computation capacity is significantly limited. As such, there is a need to develop and analyze efficient algorithms that are specifically tailored to solving, or even finding a feasible point of, the reach-constrained transport problem. A specific approach of interest is to modify the classical Sinkhorn algorithm in order to provide feasible reach-constrained transport plans. A second interesting direction for future research is to consider continuous-time and continuous-distribution generalizations of the task and methods studied in this article. Finally, incorporating agent-wise collision constraints into the problem setting poses an important challenge to solve in order to ensure safe deployment in practice.

Appendix A. Deferred Proofs.

PROPOSITION 2.9. *Let $i, j \in \{1, \ldots, N\}$. It holds that $\Omega_{ij}^d := x_j^d \oplus \Omega_i$ is a robust control invariant set for (A_i, B_i) subject to (X_i, U_i, W_i).*

Proof of Proposition 2.9. Let $x_{ij} \in \Omega_{ij}^d$ and define $u_{ij} = u_{ij}^d + K_i(x_{ij} - x_j^d)$. Then $x_{ij} = x_j^d + \omega_i$ and $u_{ij} = u_{ij}^d + K_i\omega_i$ for some $\omega_i \in \Omega_i$. Since $x_j^d \in X_i \ominus \Omega_i$ by Assumption 2.7, it holds that $x_{ij} = x_j^d + \omega_i \in x_j^d \oplus \Omega_i \subseteq X_i$, and since $u_{ij}^d \in U_i \ominus (K_i\Omega_i)$ by Assumption 2.7, it holds that $u_{ij} = u_{ij}^d + K_i\omega_i \in u_{ij}^d \oplus (K_i\Omega_i) \subseteq U_i$. By the disturbance invariance of Ω_i, we have that $A_i x_{ij} + B_i u_{ij} + w_i = (A_i + B_i K_i)\omega_i + w_i + x_j^d \in x_j^d \oplus ((A_i + B_i K_i)\Omega_i \oplus W_i) \subseteq x_j^d \oplus \Omega_i$ for all $w_i \in W_i$. □

PROPOSITION 3.1. *Let $x_1, \ldots, x_N, \hat{x}_1(T_1), \ldots, \hat{x}_N(T_N) \in \mathbb{R}^n$ be arbitrary. If $(P, \hat{u}_1, \ldots, \hat{u}_N)$ is feasible for $\mathrm{OTP}_{\mathrm{reach}}(x_1, \ldots, x_N, \hat{x}_1(T_1), \ldots, \hat{x}_N(T_N))$, then P is feasible for $\mathrm{OTP}(x_1, \ldots, x_N)$. Consequently,*

$$\mathrm{Val}(\mathrm{OTP}(x_1, \ldots, x_N)) \leq \mathrm{Val}(\mathrm{OTP}_{\mathrm{reach}}(x_1, \ldots, x_N, \hat{x}_1, \ldots, \hat{x}_N)).$$

Proof of Proposition 3.1. The proof is obvious from the fact that the reach-constrained optimal transport problem includes the constraints of the standard optimal transport problem. □

LEMMA 4.2. *For all $i \in \{1, \ldots, N\}$ and all $P \in \mathcal{B}_N$, it holds that $\Omega_{\mathrm{tmp},i} := N \bigoplus_{j=1}^N P_{ij}\Omega_{ij}^d$ is a subset of X_i.*

Proof of Lemma 4.2. Let $i \in \{1, \ldots, N\}$ and $P \in \mathcal{B}_N$. Suppose that $x_i \in \Omega_{\mathrm{tmp},i}$, so that $x_i = \omega_i + N \sum_{j=1}^N P_{ij}x_j^d$ for some $\omega_i \in \Omega_i$. Since $N \sum_{j=1}^N P_{ij} = 1$, we have that $x_i = \sum_{j=1}^N \frac{1}{N}P_{ij}(x_j^d + \omega_i)$. By Assumption 2.7, it holds that $x_j^d + \omega_i \in X_i$, and therefore since X_i is convex, we conclude that $x_i \in X_i$, which completes the proof. □

PROPOSITION 4.3. *Let $T_1, \ldots, T_N \in \mathbb{N}$ and consider arbitrary states $x_1, \ldots, x_N, \hat{x}_1(T_1), \ldots, \hat{x}_N(T_N) \in \mathbb{R}^n$. If there exists $Q \in \mathcal{B}_N$ such that $\hat{x}_i(T_i) \in$*

$\left(N \bigoplus_{j=1}^{N} Q_{ij}\Omega_{ij}^d\right) \ominus R_i^{T_i}$ *for all* $i \in \{1, \ldots, N\}$, *then* (3.1) *is feasible for some* $(P, \hat{u}_1, \ldots, \hat{u}_N)$ *with* $P = Q$.

Proof of Proposition 4.3. Suppose that there exists $Q \in \mathcal{B}_N$ such that $\hat{x}_i(T_i) \in \left(N \bigoplus_{j=1}^{N} Q_{ij}\Omega_{ij}^d\right) \ominus R_i^{T_i}$ for all $i \in \{1, \ldots, N\}$. It suffices to show that $(P, \hat{u}_1, \ldots, \hat{u}_N)$ with $P = Q$ and

$$\hat{u}_i = N \sum_{j=1}^{N} Q_{ij}u_{ij}^d + K_i \left(\hat{x}_i(T_i) - N \sum_{j=1}^{N} Q_{ij}x_j^d\right), \ i \in \{1, \ldots, N\},$$

is feasible for (3.1). First, it it obvious that $P \in \mathcal{B}_N$ since $Q \in \mathcal{B}_N$.

Next, let $i \in \{1, \ldots, N\}$ be arbitrary. Since $\hat{x}_i(T_i) \in \left(N \bigoplus_{j=1}^{N} Q_{ij}\Omega_{ij}^d\right) \ominus R_i^{T_i} = \left(\left(N \sum_{j=1}^{N} Q_{ij}x_j^d\right) \oplus \Omega_i\right) \ominus R_i^{T_i}$, it holds that $\hat{x}_i(T_i) + \overline{w}_i \in \left(N \sum_{j=1}^{N} Q_{ij}x_j^d\right) \oplus \Omega_i$ for all $\overline{w}_i \in R_i^{T_i}$, and therefore

(A.1) $$\hat{x}_i(T_i) + \overline{w}_i - N \sum_{j=1}^{N} Q_{ij}x_j^d \in \Omega_i \text{ for all } \overline{w}_i \in R_i^{T_i}.$$

Since, for all $j \in \{1, \ldots, N\}$, it holds that $u_{ij}^d \oplus (K_i\Omega_i) \subseteq U_i$ by Assumption 2.7, this gives that $u_{ij}^d + K_i \left(\hat{x}_i(T_i) + \overline{w}_i - N \sum_{j=1}^{N} Q_{ij}x_j^d\right) \in U_i$ for all $\overline{w}_i \in R_i^{T_i}$ and all $j \in \{1, \ldots, N\}$. By convexity of U_i and the fact that $Q \in \mathcal{B}_N$, this yields that

$$\hat{u}_i + K_i\overline{w}_i = N \sum_{j=1}^{N} Q_{ij} \left(u_{ij}^d + K_i \left(\hat{x}_i(T_i) + \overline{w}_i - N \sum_{l=1}^{N} Q_{il}x_l^d\right)\right) \in U_i$$

for all $\overline{w}_i \in R_i^{T_i}$. Thus, we find that the first added constraint in (3.1) is satisfied, namely, that $\hat{u}_i \in U_i \ominus (K_i R_i^{T_i})$.

Finally, let $i \in \{1, \ldots, N\}$ and $w_0, \ldots, w_{T_i} \in W_i$ be arbitrary. Then, it holds that

$$A_i\hat{x}_i(T_i) + B_i\hat{u}_i$$
$$+ (A_i + B_iK_i)^{T_i}w_{T_i} + (A_i + B_iK_i)^{T_i-1}w_{T_i-1} + \cdots + w_0 - N \sum_{j=1}^{N} Q_{ij}x_j^d$$

(A.2) $$= (A_i + B_iK_i)\Big(\hat{x}_i(T_i) - N \sum_{j=1}^{N} Q_{ij}x_j^d$$

$$+ (A_i + B_iK_i)^{T_i-1}w_{T_i} + (A_i + B_iK_i)^{T_i-2}w_{T_i-1} + \cdots + w_1\Big)$$

$$+ w_0 + N \sum_{j=1}^{N} Q_{ij}(A_ix_j^d + B_iu_{ij}^d - x_j^d).$$

Now, we have that $A_ix_j^d + B_iu_{ij}^d - x_j^d = 0$ for all $j \in \{1, \ldots, N\}$, and furthermore, since $(A_i + B_iK_i)^{T_i-1}w_{T_i} + (A_i + B_iK_i)^{T_i-2}w_{T_i-1} + \cdots + w_1 \in R_i^{T_i}$, (A.1) gives that

$$\hat{x}_i(T_i) - N \sum_{j=1}^{N} Q_{ij}x_j^d + (A_i + B_iK_i)^{T_i-1}w_{T_i} + (A_i + B_iK_i)^{T_i-2}w_{T_i-1} + \cdots + w_1 \in \Omega_i.$$

Since Ω_i is a disturbance invariant set for (A_i, B_i) under K_i subject to W_i and $w_0 \in W_i$, we conclude from (A.2) that

$$A_i\hat{x}_i(T_i) + B_i\hat{u}_i$$

$$+ (A_i + B_iK_i)^{T_i}w_{T_i} + (A_i + B_iK_i)^{T_i-1}w_{T_i-1} + \cdots + w_0 - N \sum_{j=1}^{N} Q_{ij}x_j^d \in \Omega_i.$$

Since $w_0, \ldots, w_{T_i} \in W_i$ are arbitrary and hence $(A_i + B_i K_i)^{T_i-1} w_{T_i-1} + \cdots + w_0 \in \bigoplus_{j=0}^{T_i-1} (A_i + B_i K_i)^j W_i = R_i^{T_i}$ is arbitrary, this proves that $(A_i \hat{x}_i(T_i) + B_i \hat{u}_i) \oplus ((A_i + B_i K_i)^{T_i} W_i) \subseteq \left(\left(N \sum_{j=1}^{N} Q_{ij} x_j^d \right) \oplus \Omega_i \right) \ominus R_i^{T_i} = \left(N \bigoplus_{j=1}^{N} Q_{ij} \Omega_{ij}^d \right) \ominus R_i^{T_i}$, and hence the final added constraint in (3.1) is satisfied. This concludes the proof. $\square$

THEOREM 4.6. *The system ROT-MPC is recursively feasible.*

Proof of Theorem 4.6. Let $t \in \mathbb{N} \cup \{0\}$ and assume that $P(t-1)$ is not a permutation matrix. Suppose that

$$\text{OTP}_{\text{reach}}(x_1(t), \ldots, x_N(t), \mathbf{x}_1^{t-1}(T_1), \ldots, \mathbf{x}_N^{t-1}(T_N))$$

and every $\text{OCP}_i(x_i(t), \Omega_{\text{tmp},i}(t))$ are feasible. Let $(P(t), \hat{u}_1(t), \ldots, \hat{u}_N(t))$ denote the feasible point of the reach-constrained optimal transport problem that is used to define $\chi_{\text{tmp},i}(t)$ and $\Omega_{\text{tmp},i}(t)$ for every agent i, as in ROT-MPC. Recall that the optimal control problem (2.2) is attained whenever it is feasible. Thus, let $(\mathbf{x}_i^t, \mathbf{u}_i^t)$ be a minimizer of $\text{OCP}_i(x_i(t), \Omega_{\text{tmp},i}(t))$ for all $i \in \{1, \ldots, N\}$, and let $u_i(t) = \mathbf{u}_i^t(0)$ be the corresponding control for agent i as in ROT-MPC.

The feasibility of $\text{OTP}_{\text{reach}}(x_1(t+1), \ldots, x_N(t+1), \mathbf{x}_1^t(T_1), \ldots, \mathbf{x}_N^t(T_N))$ immediately follows from Proposition 4.3 using the fact that $\mathbf{x}_i^t(T_i) \in \Omega_{\text{tmp},i}(t) \ominus R_i^{T_i} = \left(N \bigoplus_{j=1}^{N} P_{ij}(t) \Omega_j \right) \ominus R_i^{T_i}$ for all $i \in \{1, \ldots, N\}$, in light of the optimal control constraints at time t. Let $(P(t+1), \hat{u}_1(t+1), \ldots, \hat{u}_N(t+1))$ be such a feasible point, and let $\chi_{\text{tmp},i}(t+1)$ and $\Omega_{\text{tmp},i}(t+1)$ respectively be the corresponding temporary target state and terminal set for agent i as in ROT-MPC.

We must now show that $\text{OCP}_i(x_i(t+1), \Omega_{\text{tmp},i}(t+1))$ is feasible for all $i \in \{1, \ldots, N\}$. To this end, let $i \in \{1, \ldots, N\}$. Define $\mathbf{x}_i^{t+1} = (\mathbf{x}_i^{t+1}(0), \ldots, \mathbf{x}_i^{t+1}(T_i)) \in (\mathbb{R}^n)^{T_i+1}$ and $\mathbf{u}_i^{t+1} = (\mathbf{u}_i^{t+1}(0), \ldots, \mathbf{u}_i^{t+1}(T_i-1)) \in (\mathbb{R}^{m_i})^{T_i}$ by (4.1). We claim that $(\mathbf{x}_i^{t+1}, \mathbf{u}_i^{t+1})$ is feasible for $\text{OCP}_i(x_i(t+1), \Omega_{\text{tmp},i}(t+1))$, which we now prove.

Clearly, the dynamics constraint $\mathbf{x}_i^{t+1}(k+1) = A_i \mathbf{x}_i^{t+1}(k) + B_i \mathbf{u}_i^{t+1}(k)$ holds for all $k \in \{0, \ldots, T_i-1\}$ by construction. Similarly, the initial condition constraint $\mathbf{x}_i^{t+1}(0) = x_i(t+1)$ holds by construction. Thus, all that remains to show is that the robust state, input, and terminal set constraints hold.

State constraints. It holds that $\mathbf{x}_i^{t+1}(0) = x_i(t+1) = A_i x_i(t) + B_i u_i(t) + w_i(t) = A_i \mathbf{x}_i^t(0) + B_i \mathbf{u}_i^t(0) + w_i(t) = \mathbf{x}_i^t(1) + w_i(t)$, and

$$\begin{aligned} \mathbf{x}_i^{t+1}(k+1) &= A_i \mathbf{x}_i^{t+1}(k) + B_i \mathbf{u}_i^t(k+1) + B_i K_i(\mathbf{x}_i^{t+1}(k) - \mathbf{x}_i^t(k+1)) \\ &= (A_i + B_i K_i)(\mathbf{x}_i^{t+1}(k) - \mathbf{x}_i^t(k+1)) + \mathbf{x}_i^t(k+2), \end{aligned}$$

for all $k \in \{0, \ldots, T_i-2\}$, and hence

(A.3)
$$\begin{aligned} \mathbf{x}_i^{t+1}(k) &= \mathbf{x}_i^t(k+1) + (A_i + B_i K_i)^k (\mathbf{x}_i^{t+1}(0) - \mathbf{x}_i^t(1)) \\ &= \mathbf{x}_i^t(k+1) + (A_i + B_i K_i)^k w_i(t) \end{aligned}$$

for all $k \in \{0, \ldots, T_i-1\}$. Since $\mathbf{x}_i^t(k+1) \in X_i \ominus R_i^{k+1}$ for all $k \in \{0, \ldots, T_i-2\}$ and $\mathbf{x}_i^t(T_i) \in \Omega_{\text{tmp},i}(t) \ominus R_i^{T_i} \subseteq X_i \ominus R_i^{T_i}$ (as $\Omega_{\text{tmp},i}(t) \subseteq X_i$ by Lemma 4.2), we find from (A.3) that

$$\begin{aligned} \mathbf{x}_i^{t+1}(k) &\in (X_i \ominus R_i^{k+1}) \oplus ((A_i + B_i K_i)^k W_i) \\ &= (X_i \ominus (R_i^k \oplus ((A_i + B_i K_i)^k W_i))) \oplus ((A_i + B_i K_i)^k W_i) \\ &\subseteq X_i \ominus R_i^k \end{aligned}$$

for all $k \in \{0, \ldots, T_i-1\}$. Hence, the robust state constraints are satisfied.

Input constraints. From (A.3), it holds that $\mathbf{u}_i^{t+1}(k) = \mathbf{u}_i^t(k+1) + K_i(\mathbf{x}_i^{t+1}(k) - \mathbf{x}_i^t(k+1)) = \mathbf{u}_i^t(k+1) + K_i(A_i + B_iK_i)^k w_i(t)$ for all $k \in \{0, \ldots, T_i - 2\}$. Since $\mathbf{u}_i^t(k+1) \in U_i \ominus (K_iR_i^{k+1})$ for all $k \in \{0, \ldots, T_i - 2\}$, we find that

$$
\begin{aligned}
\mathbf{u}_i^{t+1}(k) &\in (U_i \ominus (K_iR_i^{k+1})) \oplus (K_i(A_i + B_iK_i)^kW_i) \\
&= (U_i \ominus (K_i(R_i^k \oplus ((A_i + B_iK_i)^kW_i)))) \oplus (K_i(A_i + B_iK_i)^kW_i) \\
&= (U_i \ominus ((K_iR_i^k) \oplus (K_i(A_i + B_iK_i)^kW_i))) \oplus (K_i(A_i + B_iK_i)^kW_i) \\
&\subseteq U_i \ominus (K_iR_i^k)
\end{aligned}
$$

for all $k \in \{0, \ldots, T_i - 2\}$. Since $(P(t+1), \hat{u}_1(t+1), \ldots, \hat{u}_N(t+1))$ is feasible for $\mathrm{OTP}_{\mathrm{reach}}(x_1(t+1), \ldots, x_N(t+1), \mathbf{x}_1^t(T_1), \ldots, \mathbf{x}_N^t(T_N))$, we have that $\hat{u}_i(t+1) \in U_i \ominus (K_iR_i^{T_i})$, and therefore using the same line of analysis as above together with the fact that $\mathbf{u}_i^{t+1}(T_i-1) = \hat{u}_i(t+1) + K_i(\mathbf{x}_i^{t+1}(T_i-1) - \mathbf{x}_i^t(T_i))$ yields that $\mathbf{u}_i^{t+1}(T_i-1) \in U_i \ominus (K_iR_i^{T_i-1})$. Hence, the robust input constraints are satisfied.

Terminal set constraint. Employing (A.3) once again, we find that

$$
\begin{aligned}
\mathbf{x}_i^{t+1}(T_i) &= A_i\mathbf{x}_i^{t+1}(T_i - 1) + B_i\mathbf{u}_i^{t+1}(T_i - 1) \\
&= A_i\mathbf{x}_i^{t+1}(T_i - 1) + B_i\hat{u}_i(t+1) + B_iK_i(\mathbf{x}_i^{t+1}(T_i - 1) - \mathbf{x}_i^t(T_i)) \\
&= (A_i + B_iK_i)(\mathbf{x}_i^{t+1}(T_i - 1) - \mathbf{x}_i^t(T_i)) + A_i\mathbf{x}_i^t(T_i) + B_i\hat{u}_i(t+1) \\
&= (A_i + B_iK_i)^{T_i}w_i(t) + A_i\mathbf{x}_i^t(T_i) + B_i\hat{u}_i(t+1) \\
&\in (A_i\mathbf{x}_i^t(T_i) + B_i\hat{u}_i(t+1)) \oplus ((A_i + B_iK_i)^{T_i}W_i) \\
&\subseteq \left(N \bigoplus_{j=1}^N P_{ij}(t+1)\Omega_{ij}^d\right) \ominus R_i^{T_i} \\
&= \Omega_{\mathrm{tmp},i}(t+1) \ominus R_i^{T_i},
\end{aligned}
$$

where the final inclusion and final equality come from the fact that $(P(t+1), \hat{u}_1(t+1), \ldots, \hat{u}_N(t+1))$ is feasible for $\mathrm{OTP}_{\mathrm{reach}}(x_1(t+1), \ldots, x_N(t+1), \mathbf{x}_1^t(T_1), \ldots, \mathbf{x}_N^t(T_N))$. Thus, the robust terminal set constraint is satisfied. $\square$

THEOREM 4.8. *Assume that all of the initial optimization problems* $\mathrm{OTP}_{\mathrm{reach}}(x_1(0), \ldots, x_N(0), \mathbf{x}_1^{-1}(T_1), \ldots, \mathbf{x}_N^{-1}(T_N))$ *and* $\mathrm{OCP}_i(x_i(0), \Omega_{\mathrm{tmp},i}(0))$ *are feasible, so that ROT-MPC is feasible at every time* $t \in \mathbb{N} \cup \{0\}$ *per Theorem 4.6. Then, for all* $i \in \{1, \ldots, N\}$, *it holds that* $x_i(t) \in X_i$ *and* $u_i(t) \in U_i$ *for all* $t \in \mathbb{N} \cup \{0\}$ *for the closed-loop system ROT-MPC.*

Proof of Theorem 4.8. Let $i \in \{1, \ldots, N\}$ and let $t \in \mathbb{N} \cup \{0\}$. If $t = 0$, then trivially we have that $x_i(t) = x_i(0) = x_{i,0} \in X_i$, and by feasibility and hence attainment of $\mathrm{OCP}_i(x_i(0), \Omega_{\mathrm{tmp},i}(0))$, there exists $(\mathbf{x}_i^0, \mathbf{u}_i^0)$ solving $\mathrm{OCP}_i(x_i(0), \Omega_{\mathrm{tmp},i}(0))$. By the dynamics of ROT-MPC and the input constraints, we find that $u_i(t) = \mathbf{u}_i^0(0) \in U_i \ominus (K_iR_i^0) = U_i \ominus \{0\} = U_i$.

On the other hand, suppose that $t > 0$. As a first case, suppose that $P(t-1)$ is not a permutation matrix. By Theorem 4.6, we have that both control problems $\mathrm{OCP}_i(x_i(t-1), \Omega_{\mathrm{tmp},i}(t-1))$ and $\mathrm{OCP}_i(x_i(t), \Omega_{\mathrm{tmp},i}(t))$ are feasible, and hence attained. Let $(\mathbf{x}_i^{t-1}, \mathbf{u}_i^{t-1})$ and $(\mathbf{x}_i^t, \mathbf{u}_i^t)$ be solutions to these control problems, respec-

tively. By the dynamics of ROT-MPC and the state constraints, we find that

$$
\begin{aligned}
x_i(t) &= A_i x_i(t-1) + B_i u_i(t-1) + w_i(t-1) \\
&= A_i \mathbf{x}_i^{t-1}(0) + B_i \mathbf{u}_i^{t-1}(0) + w_i(t-1) \\
&= \mathbf{x}_i^{t-1}(1) + w_i(t-1) \\
&\in (X_i \ominus R_i^1) \oplus W_i \\
&= (X_i \ominus W_i) \oplus W_i \\
&\subseteq X_i.
\end{aligned}
$$

Furthermore, by the dynamics of ROT-MPC and the input constraints, we find that $u_i(t) = \mathbf{u}_i^t(0) \in U_i \ominus (K_i R_i^0) = U_i \ominus \{0\} = U_i$.

As a second case, suppose that $t > 0$ is such that $P(t-1)$ is a permutation matrix. Then, there exists $t^\star \in \mathbb{N} \cup \{0\}$ such that $P(t^\star)$ is a permutation matrix, and $P(t')$ is non-permutation for all $t' < t^\star$. It is necessarily the case for the time under consideration that $t \geq t^\star$. Consider the solution $(\mathbf{x}_i^{t^\star}, \mathbf{u}_i^{t^\star})$ to the control problem $\mathrm{OCP}_i(x_i(t^\star), \Omega_{\mathrm{tmp},i}(t^\star))$. If $t < t^\star + T_i$, then $x_i(t) \in X_i$ and $u_i(t) \in U_i$ follows directly from Lemma 4.7. To show that the constraints hold when $t \geq t^\star + T_i$ (and therefore complete the proof), note that we have by Lemma 4.7 that $x_i(t^\star + T_i) \in \Omega_{i j_i}^d$, implying that $x_i(t^\star + T_i) \in X_i$. Suppose that $t \geq t^\star + T_i$ and that $x_i(t) \in \Omega_{i j_i}^d$ (as is the case for $t = t^\star + T_i$). Then, it holds that $x_i(t) = x_{j_i}^d + \omega_i(t)$ for some $\omega_i(t) \in \Omega_i$. Therefore, the dynamics of ROT-MPC give that

$$
u_i(t) = u_{i j_i}^d + K_i \omega_i(t) \in u_{i j_i}^d \oplus (K_i \Omega_i) \subseteq U_i
$$

per Assumption 2.7, and furthermore that

$$
\begin{aligned}
x_i(t+1) &= A_i x_i(t) + B_i(u_{i j_i}^d + K_i(x_i(t) - x_{j_i}^d)) + w_i(t) \\
&= (A_i + B_i K_i) x_i(t) + w_i(t) + B_i u_{i j_i}^d - B_i K_i x_{j_i}^d \\
&= (A_i + B_i K_i) x_i(t) + w_i(t) + x_{j_i}^d - A_i x_{j_i}^d - B_i K_i x_{j_i}^d \\
&= x_{j_i}^d + (A_i + B_i K_i)(x_i(t) - x_{j_i}^d) + w_i(t),
\end{aligned}
$$

implying that $x_i(t+1) \in x_{j_i}^d \oplus (((A_i + B_i K_i) \Omega_i) \oplus W_i) \subseteq x_{j_i}^d \oplus \Omega_i = \Omega_{i j_i}^d$, further implying that $x_i(t+1) \in X_i$. This proves that the constraint satisfactions hold for all cases of t, i.e., $x_i(t) \in X_i$ and $u_i(t) \in U_i$ for all $t \in \mathbb{N} \cup \{0\}$. $\square$

COROLLARY 4.9. *Assume that all of the initial optimization problems* $\mathrm{OCP}_i(x_i(0), \Omega_{\mathrm{tmp},i}(0))$ *and* $\mathrm{OTP}_{\mathrm{reach}}(x_1(0), \ldots, x_N(0), \mathbf{x}_1^{-1}(T_1), \ldots, \mathbf{x}_N^{-1}(T_N))$ *are feasible. If every X_i is compact, then the closed-loop trajectory $t \mapsto (x_1(t), \ldots, x_N(t))$ of ROT-MPC is uniformly bounded in $X_1 \times \cdots X_N$, meaning that $\|x_i(t)\|_2 \leq M_i :=$ $\max_{x \in X_i} \|x\|_2 < \infty$ for all $i \in \{1, \ldots, N\}$ and all $t \in \mathbb{N} \cup \{0\}$.*

Proof of Corollary 4.9. This follows immediately from the fact that $x_i(t)$ is an element of the compact set X_i for all $i \in \{1, \ldots, N\}$ and all $t \in \mathbb{N} \cup \{0\}$, due to Theorem 4.8. $\square$

Appendix B. Additional Implementation Details.

Simulations are conducted with a 2020 Macbook Air M1, Python 3.9.19, NumPy 1.23.2, and SciPy 1.9.1. Polytope computations are carried out by the `pytope` package, and the optimization problems are solved using `scipy.optimize.minimize`.

All optimal control loss functions are defined by

$$L_i(\overline{x}(0), \ldots, \overline{x}(T_i), \overline{u}(0), \ldots, \overline{u}(T_i - 1))$$
$$= \sum_{k=0}^{T_i} (\overline{x}(k) - x_{\text{ref}}(k))^\top Q(\overline{x}(k) - x_{\text{ref}}(k)) + \sum_{k=0}^{T_i-1} (\overline{u}(k) - u_{\text{ref}}(k))^\top R(\overline{u}(k) - u_{\text{ref}}(k)),$$

with $Q = 10I_n$ and $R = 0.01I_{m_i}$, and where $u_{\text{ref}}(0), \ldots, u_{\text{ref}}(T_i - 1) \in \mathbb{R}^{m_i}$ and $x_{\text{ref}}(0), \ldots, x_{\text{ref}}(T_i) \in \mathbb{R}^n$ denote reference input and state trajectories constructed based on the temporary target state at that time, i.e., $x_{\text{ref}}(k) = \chi_{\text{tmp},i}(t)$ for all $k \in \{1, \ldots, T_i\}$, and $u_{\text{ref}}(0) = \cdots = u_{\text{ref}}(T_i - 1) =: \mu_{\text{tmp},i}(t)$ is taken as a solution to the linear system of equations $B_i\mu_{\text{tmp},i}(t) = (I_n - A_i)\chi_{\text{tmp},i}(t)$.

In our implementation of Sinkhorn MPC, we use one iteration of the Sinkhorn algorithm per time step, as proposed in [16]. Although more Sinkhorn iterations are possible, it was shown that one iteration yields comparable performance in [18], so we employ this choice for simplicity. In Task 1, Sinkhorn MPC uses entropy regularization parameter $\epsilon = 1$. In Tasks 2 and 3, this is increased to $\epsilon = 2$, as used in [18].

Our implementation of ROT-MPC uses transport cost functions defined by $C_{ij}(x_i) = \|x_i - x_j^d\|_2$. Our implementation of ROT-MPC also reformulates the modified transport problem's reachability subset constraint,

$$(B.1) \qquad (A_i\hat{x}_i(T_i) + B_i\hat{u}_i) \oplus ((A_i + B_iK_i)^{T_i}W_i) \subseteq \left(N \bigoplus_{j=1}^N P_{ij}\Omega_{ij}^d\right) \ominus R_i^{T_i},$$

into an equivalent finite-dimensional form. Specifically, we apply the reformulation techniques considered in [1, 28] by framing the constraint as an equivalent optimization problem, and then dualize it. The reformulation is summarized as follows. For every i, compute (offline) halfspace representations of $R_i^{T_i+1}$ and Ω_i:

$$R_i^{T_i+1} = \{w \in \mathbb{R}^n : H_iw \le h_i\}, \quad \Omega_i = \{x \in \mathbb{R}^n : G_ix \le g_i\},$$

with $H_i \in \mathbb{R}^{p_i \times n}$, $h_i \in \mathbb{R}^{p_i}$, $G_i \in \mathbb{R}^{q_i \times n}$, and $g_i \in \mathbb{R}^{q_i}$. For every i, compute (offline)

$$\overline{G}_i = \underbrace{\text{blkdiag}(G_i, \ldots, G_i)}_{N\text{-fold block diagonalization}}, \quad \overline{g}_i = \underbrace{(g_i, \ldots, g_i)}_{N\text{-fold vertical concatenation}}.$$

Replace the subset constraint (B.1) by the equivalent system of constraints

$$\overline{g}_i + \overline{G}_i\eta_i \ge 0, \quad \gamma_i \ge 0, \quad \lambda_i \ge 0,$$
$$\overline{G}_i^\top \lambda_i = Y_i(P)^\top H_i^\top \gamma_i,$$
$$H_i(-A_i\hat{x}_i(T_i) - B_i\hat{u}_i + y_i(P) - Y_i(P)\eta_i) \le h_i,$$
$$(\overline{g}_i + \overline{G}_i\eta_i)^\top \lambda_i = 0,$$
$$\gamma_i^\top (H_i(-A_i\hat{x}_i(T_i) - B_i\hat{u}_i + y_i(P) - Y_i(P)\eta_i) - h_i) = 0,$$

in both the original variables $P, \hat{u}_i$ as well as new variables $\lambda_i \in \mathbb{R}^{q_iN}$, $\eta_i \in \mathbb{R}^{nN}$, and $\gamma_i \in \mathbb{R}^{p_i}$, where $y_i(P) = N\sum_{j=1}^N P_{ij}x_j^d$ and $Y_i(P) = N\begin{bmatrix} P_{i1}I_n \cdots P_{iN}I_n \end{bmatrix}$.

These constraints are nonconvex due to the bilinear terms. This nonconvexity does *not* break the recursive feasibility or recursive constraint satisfaction of ROT-MPC, since all that is needed at every timestep of the control scheme is a feasible point for the optimal transport problem, and not necessarily a global optimizer; any point

satisfying the above constraints ensures reachability of the associated transport plan. In practice, we find that off-the-shelf nonlinear programming solvers quickly converge to (local) solutions with permutation transport plans when using this reformulation. In fact, our testing found that this nonconvex formulation results in much faster overall runtime for ROT-MPC than equivalent finite-dimensional convex formulations, as the latter tend to return solutions with non-permutation transport plans.

References.

[1] B. G. ANDERSON, *Bilinear positivity certificates as convex-in-polyhedron set containments*, Preprint, (2025). URL: https://brendon-anderson.github.io/files/publications/anderson2025bilinear.pdf.

[2] B. G. ANDERSON, E. LOESER, M. GEE, F. REN, S. BISWAS, O. TURANOVA, M. HABERLAND, AND A. L. BERTOZZI, *Quantitative assessment of robotic swarm coverage*, in 15th International Conference on Informatics in Control, Automation and Robotics, 2018.

[3] J.-D. BENAMOU AND Y. BRENIER, *A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem*, Numerische Mathematik, 84 (2000), pp. 375–393.

[4] S. BERMAN, A. HALÁSZ, M. A. HSIEH, AND V. KUMAR, *Optimized stochastic policies for task allocation in swarms of robots*, IEEE Transactions on Robotics, 25 (2009), pp. 927–937.

[5] D. BIENSTOCK, M. CHERTKOV, AND S. HARNETT, *Chance-constrained optimal power flow: Risk-aware network control under uncertainty*, SIAM Review, 56 (2014), pp. 461–495.

[6] G. BIRKHOFF, *Tres observaciones sobre el algebra lineal*, Univ. Nac. Tucuman, Ser. A, 5 (1946), pp. 147–154.

[7] F. BULLO, J. CORTÉS, AND S. MARTÍNEZ, *Distributed Control of Robotic Networks: A Mathematical Approach to Motion Coordination Algorithms*, Princeton University Press, 2009.

[8] Y. CHEN, T. T. GEORGIOU, AND M. PAVON, *Optimal steering of a linear stochastic system to a final probability distribution, Part II*, IEEE Transactions on Automatic Control, 61 (2015), pp. 1170–1180.

[9] Y. CHEN, T. T. GEORGIOU, AND M. PAVON, *Optimal transport over a linear dynamical system*, IEEE Transactions on Automatic Control, 62 (2016), pp. 2137–2152.

[10] Y. CHEN, T. T. GEORGIOU, AND M. PAVON, *Optimal transport in systems and control*, Annual Review of Control, Robotics, and Autonomous Systems, 4 (2021), pp. 89–113.

[11] L. CHISCI, J. A. ROSSITER, AND G. ZAPPA, *Systems with persistent disturbances: Predictive control with restricted constraints*, Automatica, 37 (2001), pp. 1019–1028.

[12] M. CUTURI, *Sinkhorn distances: Lightspeed computation of optimal transport*, Advances in Neural Information Processing Systems, 26 (2013).

[13] K. ELAMVAZHUTHI, P. GROVER, AND S. BERMAN, *Optimal transport over deterministic discrete-time nonlinear systems using stochastic feedback laws*, IEEE Control Systems Letters, 3 (2018), pp. 168–173.

[14] K. ELAMVAZHUTHI, S. LIU, W. LI, AND S. OSHER, *Dynamical optimal transport of nonlinear control-affine systems*, Journal of Computational Dynamics, 10

(2023), pp. 425–449.

[15] K. HOSHINO, *Finite-horizon control of nonlinear discrete-time systems with terminal cost of Wasserstein distance*, in 59th IEEE Conference on Decision and Control (CDC), 2020.

[16] K. ITO AND K. KASHIMA, *Sinkhorn MPC: Model predictive optimal transport over dynamical systems*, in American Control Conference (ACC), IEEE, 2022, pp. 2057–2062.

[17] K. ITO AND K. KASHIMA, *Entropic model predictive optimal transport for underactuated linear systems*, IEEE Control Systems Letters, 7 (2023), pp. 2761–2766.

[18] K. ITO AND K. KASHIMA, *Entropic model predictive optimal transport over dynamical systems*, Automatica, 152 (2023).

[19] V. KRISHNAN AND S. MARTÍNEZ, *Distributed control for spatial self-organization of multi-agent swarms*, SIAM Journal on Control and Optimization, 56 (2018), pp. 3642–3667.

[20] V. KRISHNAN AND S. MARTÍNEZ, *Distributed optimal transport for the deployment of swarms*, in 57th IEEE Conference on Decision and Control (CDC), 2018.

[21] W. LANGSON, I. CHRYSSOCHOOS, S. RAKOVIĆ, AND D. Q. MAYNE, *Robust model predictive control using tubes*, Automatica, 40 (2004), pp. 125–133.

[22] S. E. LI, Y. ZHENG, K. LI, L.-Y. WANG, AND H. ZHANG, *Platoon control of connected vehicles from a networked control perspective: Literature review, component modeling, and controller synthesis*, IEEE Transactions on Vehicular Technology, (2017).

[23] Y. LIU AND R. BUCKNALL, *A survey of formation control and motion planning of multiple unmanned vehicles*, Robotica, 36 (2018), pp. 1019–1047.

[24] D. Q. MAYNE, M. M. SERON, AND S. RAKOVIĆ, *Robust model predictive control of constrained linear systems with bounded disturbances*, Automatica, (2005).

[25] D. MILUTINOVIĆ AND P. LIMA, *Modeling and optimal centralized control of a large-size robotic population*, IEEE Transactions on Robotics, 22 (2006), pp. 1280–1285.

[26] G. PEYRÉ AND M. CUTURI, *Computational optimal transport*, Foundations and Trends in Machine Learning, 11 (2019), pp. 355–607.

[27] S. V. RAKOVIC, E. C. KERRIGAN, K. I. KOURAMAS, AND D. Q. MAYNE, *Invariant approximations of the minimal robust positively invariant set*, IEEE Transactions on Automatic Control, 50 (2005), pp. 406–410.

[28] S. SADRADDINI AND R. TEDRAKE, *Linear encodings for polytope containment problems*, in 58th IEEE Conference on Decision and Control (CDC), 2019.

[29] D. SIMON, J. LÖFBERG, AND T. GLAD, *Reference tracking MPC using dynamic terminal set transformation*, IEEE Transactions on Automatic Control, 59 (2014), pp. 2790–2795.

[30] A. TERPIN, N. LANZETTI, AND F. DÖRFLER, *Dynamic programming in probability spaces via optimal transport*, SIAM Journal on Control and Optimization, 62 (2024), pp. 1183–1206.

[31] C. VILLANI, *Topics in Optimal Transportation*, vol. 58, American Mathematical Society, 2003.

[32] C. VILLANI, *Optimal Transport: Old and New*, vol. 338, Springer, 2008.

[33] J. VON NEUMANN, *A certain zero-sum two-person game equivalent to the optimal assignment problem*, Contributions to the Theory of Games, (1950), p. 5.

[34] Z. WAN AND M. V. KOTHARE, *Efficient robust constrained model predictive control with a time varying terminal constraint set*, Systems & Control letters, 48 (2003), pp. 375–383.