Optimal Transport Based Control of Particle Swarms for Orbiting Rainbows Concept

Carlo Sinigaglia *
Polytechnic University of Milan, 20156 Milan, Italy

Saptarshi Bandyopadhyay †
Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA, 91109-8099,USA

Marco Quadrelli ‡
Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA, 91109-8099,USA

Francesco Braghin §
Polytechnic University of Milan, 20156 Milan, Italy

Nomenclature

n_p	=	Number of particles
n_b	=	Number of bins per side
$\boldsymbol{S}_k^p$	=	State of particle p at time instant k, m
Δt	=	Timestep size, s
$\boldsymbol{f}_k^p$	=	Action acting on particle p at time instant k, $m \cdot s^{-1}$
f_x	=	Action component along the x direction, $m \cdot s^{-1}$
$\tilde{a}_{n,m}$	=	m-th actuator of stack n
$a_{n,m,k}$	=	Intensity of actuator $\tilde{a}_{n,m}$ at time instant k, $m^2 \cdot s^{-1}$
l_b	=	Bin length, m
μ_k	=	Density distribution matrix at time instant k
ν	=	Target distribution matrix
$\Delta_{h,i,j}$	=	Fraction of density flowing horizontally
$\Delta_{v,i,j}$	=	Fraction of density flowing vertically
$d_{i,j}$	=	Ground metric from bin i to bin j, m
d	=	Eucledian distance function, m
$\tilde{d}(\tilde{a}_{1,i}, [i,j])$	=	Minimum distance between actuator $\tilde{a}_{1,i}$ and bin $[i,j]$, m
D_{L_1}	=	L_1 Distance between matrices

*PhD Candidate, Department of Mechanical Engineering,Via La Masa, 1, 20156 Milano MI.
†Robotics Technologist,Maritime and Multi-Agent Autonomy, M/S 198-219
‡Group Supervisor, Mobility and Robotic Systems, M/S 198-219, AIAA Associate Fellow
§Full Professor, Department of Mechanical Engineering,Via La Masa, 1, 20156 Milano MI.

D_{EMD} = Earth's Mover Distance between matrices

g = Density propagation function

$\phi_{i,j}$ = Density flux from bin i to bin j

I. Introduction

Large-aperture space telescopes are extremely important to the space exploration community, especially in the fields of astrophysics and cosmology. Increasing the telescope's aperture size proportionately improves the telescope's angular resolution and signal-to-noise ratio, which enables scientists to observe smaller and fainter objects in the universe. On the other hand, the cost of a large-aperture telescope is largely driven by the cost and complexity (size, mass) of the primary aperture. Strategies to increase the aperture size and accuracy has been an active area of research for many centuries. The James Webb Space Telescope (JWST), scheduled for launch in 2021, has already saturated the capacity of a single launch vehicle [1]. Hence, new solutions are needed to increase the telescope's aperture size, while minimizing complexity and cost.

This paper investigates the control problem of a new architecture for space telescopes, that can increase the aperture size while reducing the mass per unit area by several orders of magnitude. This idea was initially developed under the NASA Innovative Advanced Concepts (NIAC) study [2] named "Orbiting Rainbows" [3–5].

The radical innovation of the "Orbiting Rainbows" concept is to substitute a monolithic aperture with a controllable distributed granular medium, composed of a cloud of very-small, passive particles with appropriate reflective and refractive properties [6]. The conceptual architecture is shown in Figure 1.

The granular cloud serves as distributed primary aperture by forming an adaptive surface with useful electromagnetic characteristics in the optical or microwave bands. Bruggemann effective medium theory is exploited to derive properties equivalent to a monolithic medium from a disordered system of grains so that the granular cloud can be used both as a reflective and refractive system. In particular, the focal length is inversely proportional to the fill factor, which is the ratio between the portion of space filled by the grains and the total space. Figure 1 shows the deployment and control sequence of the granular cloud that consists of four main stages. 1) the cloud is first released; 2) it is contained by laser pressure to avoid dissipation and disruption by gravitational tidal forces, 3) it is shaped by optical manipulation into a two-dimensional object (coarse control), and 4) ultimately into a surface with imaging characteristics (fine control). The key control challenge that we tackle in this work is related to coarse control and confinement of the cloud.

Namely, we address the problem of shaping a large swarm (10^4–10^5) of particles into a large reflective dish, with only a few (40) actuators in the boundary of the workspace. Note that the particles do not carry any onboard actuator. Independently of the physical nature of the actuation used, each actuator simultaneously affects the motion of all the particles within the actuator's sphere of influence. Hence the cloud of particles together form an underactuated system,

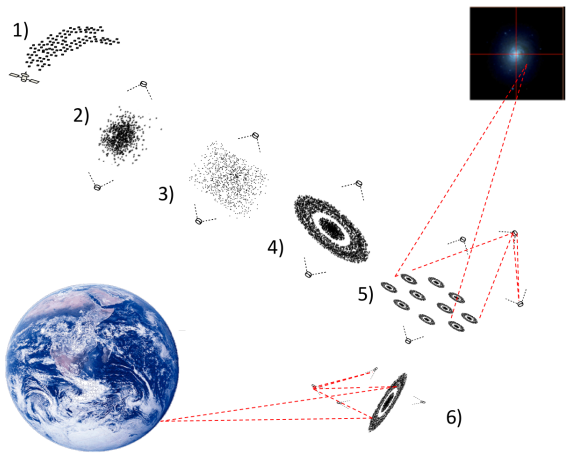

Fig. 1 The "Orbiting Rainbows" conceptual telescope deployment and operations (Reprinted with permission from Ref. [2])

and can only be guided and controlled in bulk.

In recent years, many approaches have been developed for controlling large swarms of active elements with algorithms that are naturally scalable. Most strategies aim at defining a control law for each active agent based on the density (i.e. the spatial distribution), either in a centralized or decentralized manner [7, 8]. The control actions are then applied by each agent independently, i.e., independent actuation is needed for each agent. Usually, control methods for large-scale multi-agent systems rely on strategies to decentralize and divide the computational burden in order to reach an emergent collective behavior, starting from locally stable controllers. Probabilistic density-based approaches [9–12] have emerged as promising methods to scale the problem being independent of the number of agents. Each agent estimates its position through distributed estimation and consensus algorithms, then makes autonomous probabilistic decisions, and executes its actions independently. These swarm control strategies cannot be used in this paper because the particles cannot execute their actions independently.

The most relevant difference in our case is that the cloud of reflective particles is not only actuated remotely but the control actions are global. The number of inputs is orders of magnitude less than the number of particles in the swarm. A possible solution to this problem could be to make use of results from ensemble control theory [13, 14]. This approach aims at controlling simultaneously ensembles of almost identical systems with global control actions. However, this ends up in an open-loop control strategy and it relies on the strong assumption that all elemental systems (i.e. subsystems) are in the same initial state. Although being a powerful approach to control parametric systems, ensemble theory can not be used in our case for this reason.

The control of the collective behaviour of identical elements has been addressed in [15, 16] but the nature of the actuation is still local. However, the idea of transporting the physical positions of the elements into a density of mass, then normalized to be a probability distribution, can be exploited to design a control law. Compared to the state of the

art of probabilistic density control of large-scale multi-agent systems, the Wasserstein distance, from Optimal Transport theory [17] in its discrete formulation, has been used to solve the agents' assignment problem in [18] and it has turned out to be more efficient than previous methods. The Wasserstein distance in discrete form is called the Earth's Mover Distance (EMD) and it has been introduced in [19] as a measure of similarity between images. This metric can define a distance not only between amplitudes but also according to a ground metric (i.e. in a spatial sense). This feature has many useful advantages compared to classical Euclidean norms [20]. In particular, it has been used to better tune numerical models based on experimental data [21] or as a cost function for image retrieval since it performs better than histogram-based techniques [20]. In this paper, a tailored version of the Earth's Mover Distance is used as a performance functional to recast the density control of the swarm of particles as an optimal control problem.

The idea of using the optimal transport based control strategy was first introduced in our prior conference paper [22]. This paper builds on this idea developing a more accurate density dynamics and a control strategy that uses the Wasserstein distance between probability distributions as a cost functional for the optimal control problem. This paper's modelling strategy follows from [22] where the configuration layout has been developed together with a first-order dynamic model. The system is controlled by remote actuators that confine the particles in a two-dimensional workspace. In [22], the L_1 norm has been used as a metric between actual and target distributions to derive the control strategy. The main limitation of this metric is that it does not provide any information on the spatial distribution of the density but, being an element-wise operator, it only provides *local* distances between the elements of the predicted and actual density.

The idea of using optimal transport based techniques to shape the particles swarm with laser radiation pressure has also been considered in [23]. Optimal Transport has been used to find the optimal assignment (i.e. target positions) for the particles. The optimal control problem is not solved in the space of probability distributions but the optimization algorithm considers every single grain, hence suffering from the *curse of dimensionality*. Instead, by shifting the problem in the probability space and using Wasserstein distance as cost function, the control actions computation is completely independent of the number of particles.

The main contributions of this paper, compared to exisiting literature and our prior conference paper [22] are:

1) It is shown that the collective control of the swarm of particles can be recast as an optimal control problem in the space of probability distributions. The optimizer seeks the best actuation intensities to minimize the Wasserstein distance between the predicted and target distribution. Despite retaining a single-step prediction horizon, the spatial properties of this metric allow finding minimizing actuation intensities in almost every density configuration. This method has proven to be effective and able to reach fast convergence.
2) The density dynamic model presented in [22] is refined and expanded. The new model is more accurate and consistent with the particle dynamics by satisfying the mass conservation property and by following the local behaviour in a probabilistic averaged sense.

This paper is organized as follows. In Section II, a dynamic model of the particle swarm is described together

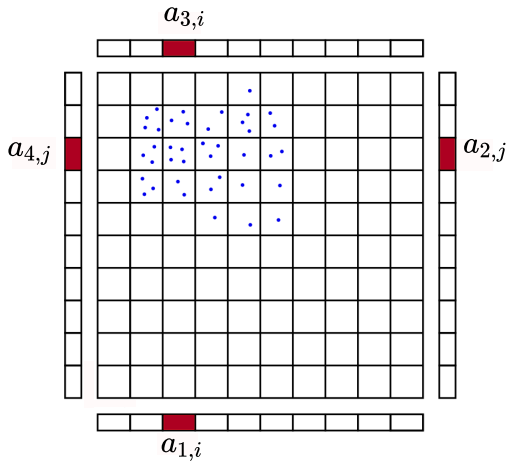

Fig. 2 The particles in bin $[i, j]$ experience the forces from actuators $a_{1,i}$,$a_{2,j}$,$a_{3,i}$,$a_{4,j}$, however, all the neighboring bins are influenced by the same actuation.

with the layout of the control problem. Then, the actual dynamic is abstracted and a probabilistic density model is developed to describe the collective behaviour of the swarm. Furthermore, the previous strategy used in [22] is briefly reviewed and its limitations discussed. In Section III, the Earth Mover's Distance (EMD) is introduced as the discrete counterpart of the Wasserstein distance. Its main features are examined and then it is inserted as a cost functional in the density control problem. Finally, the formal setting of the optimal density control problem is explained. In Section IV, a numerical simulation is carried out proving convergence and demonstrating scalability for a very large system of particles. In Section V, an experimental test of the control algorithm is performed in the Robotarium testbed. This paper is concluded in Section VI.

II. Collective Control of Particles

A. Microscopic dynamic model

The cloud of particles is modeled as a large ensemble of point masses. As shown in Figure 2, a two-dimensional workspace, delimited by actuator stacks on four sides, is considered. Each actuator exerts influence over all the particles located in its row or column. The workspace is discretized according to the number of actuators. Each element of the grid is called bin. Ten actuators per side are used, as a consequence, the workspace is gridded with one-hundred bins. The number of actuators is orders of magnitude smaller than the number of particles. Hence, as previously stated, the system is largely underactuated.

It is assumed that the particle dynamics can be reduced to:

$$\boldsymbol{S}^p_{k+1} = \boldsymbol{S}^p_k + \boldsymbol{f}^p_k \Delta t \qquad \boldsymbol{S}^p_k = \begin{bmatrix} x^p_k \\ y^p_k \end{bmatrix} \tag{1}$$

where $\boldsymbol{S}_k^p \in \mathbb{R}^2$ is the p^{th} particle's state at time instant k (i.e. a vector describing its position), Δt is the time step and $\boldsymbol{f}_k^p$ the force acting on the p^{th} particle at time instant k. Furthermore, it is assumed that the discretized dynamics behaves like a set of decoupled single integrators where the time index is denoted by a right subscript and the local agent index is denoted by a right superscript. Thus, Equation (1) describes the motion of particle p located in bin $[i, j]$ at time instant k. This force depends on the position of the particle at time instant k, i.e. on the distance from the four actuators acting on bin [i,j]. In fact, making reference to laser actuators, this force is always oriented from the actuator to the particle and proportional to the inverse of the relative Eucledian distance. $\tilde{a}_{n,m}$ is defined as the m-th actuator located on side n. Sides are numbered counter-clockwise starting from the bottom, while the actuator's number is left-to-right and bottom-up. Consequently, the Eucledian distance between an actuator and a generic particle is $d(\tilde{a}_{n,m}, p)$ $(n = 1, ..., 4; m = i, j)$. It is now possible to write down the functional dependence between the forces and the actuators as:

$$\begin{aligned} f_x &= \frac{a_{4,j}}{d(\tilde{a}_{4,j}, p)} - \frac{a_{2,j}}{d(\tilde{a}_{2,j}, p)} \\ f_y &= \frac{a_{1,i}}{d(\tilde{a}_{1,i}, p)} - \frac{a_{3,i}}{d(\tilde{a}_{3,i}, p)} \end{aligned} \tag{2}$$

Where $a_{n,m}$ is the intensity associated to actuator $\tilde{a}_{n,m}$.

B. Macroscopic model

The aggregated particle dynamics can be transformed in a density time evolution with mass conservation properties (i.e. the integral of the density in the control volume is constant). It should be noted that describing the control problem in terms of density evolution, instead of tracking the states of every single particle, allows to reduce the dimensionality (i.e. the variables are in the order of the number of bins instead of the number of particles).

In the deterministic continuous time case, the system is described by the Liouville equation, while if stochastic effects are considered, by the Fokker-Planck equation [24]. On the other hand, in discrete time, the density evolution is described by a matrix-valued function with associated transition probabilities. A suitable control strategy is used to guide the distribution as *close* as possible to the target. The main idea is to find the actuator intensities reformulating optimal transport* techniques for image classification [19] to find the optimal way of moving mass from an initial distribution to a final one with the least amount of cost [25].

This problem can be reformulated in terms of probability distributions after a unitary rescaling [17, 27]. A probability mass function (pmf) [18] is defined as the probability of finding a certain amount of mass in a specific location.

The optimization objective is to find a control law able to steer the probability mass function towards a target

*Originally, the optimal transport problem has been formulated in a static fashion [25], thus concerned in finding an optimal map between an initial and a final distribution (i.e. *what* goes *where*), then it has been used to find optimal plans mostly for resource allocation problems [26].

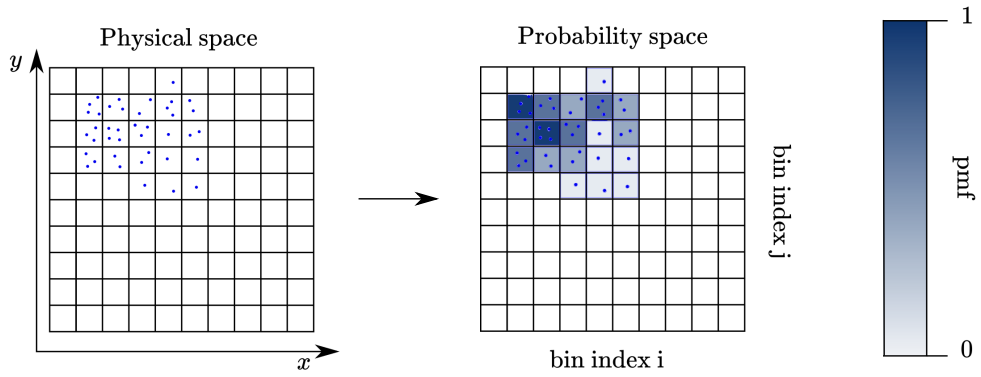

Fig. 3 The cloud of particles is transformed into a probability distribution: the darker the bin's color the higher the mass probability density.

minimizing a certain cost functional [7, 15, 22].

The deterministic system of particles is transported in the distribution space defining the discrete probability density matrix $\mu \in \mathbb{R}^{n_b \times n_b}$, where $\mu[i, j]$ represents the present probability mass in the bin $[i, j]$. For μ to be a probability density function the conservation property $\sum_{i=1}^{n_b} \sum_{i=1}^{n_b} \mu[i, j] = 1$ must hold. This constraint is automatically taken into account in the density evolution derived in the next section. Note that the information about the specific position of the particles inside each bin is lost while the average discrete spatial distribution over the workspace is considered.

If l_b represents the length (width) of each bin, it follows from Equation (1)[†] that the force (action) that actuator $a_{1,i}$ should apply to move all particles from bin $[i, j]$ to bin $[i, j + 1]$ is equal to:

$$a_{1,i,j,max} = \frac{1}{\Delta t} l_b \, \tilde{d}(\tilde{a}_{1,i}, [i, j]) \tag{3}$$

where $\tilde{d}(\tilde{a}_{1,i}, [i, j])$ represents the minimum distance between actuator $\tilde{a}_{1,i}$ and bin $[i, j]$. Let's assume that, if $a_{1,i} \leq a_{1,i,j,max}$, the fraction of density that will be transported from bin $[i, j]$ to bin $[i, j + 1]$ due to the action $a_{1,i}$ is $\frac{a_{1,i}}{a_{1,i,j,max}}$. This concept was first introduced in our previous paper [22].

The probabilistic transition model used in [22] present some limitations in describing the behavior of the particles in the workspace: the effect of the four actuators acting on a bin element is considered simply superimposed. Furthermore, transitions towards diagonally adjacent bins are not allowed, thus inserting non physical constraints in the optimization problem. This has the effect of violating the mass conservation property that must hold for a consistent macroscopic model.

A more accurate density model can be derived considering the simultaneous effect of the four actuators active on a generic bin. Let's consider the set of actuators $\{a_{1,i}, a_{2,j}, a_{3,i}, a_{4,j}\}$ active on bin $[i, j]$ having different effects that depend on the relative distance. The set of maximum actions relative to that bin is $\{a_{1,i,j,max}, a_{2,i,j,max}, a_{3,i,j,max}, a_{4,i,j,max}\}$.

[†]The discrete time evolution of the density function μ is derived from the particle dynamics model in Equation (1).

If only actuator $a_{1,i}$ is active, the probability of transition will be $\frac{a_{1,i}}{a_{1,i,j,max}}$ as described in the original formulation. If also the actuator $a_{3,i}$ (i.e. the opposite one to $a_{1,i}$) is active, it would push the particles in the opposite direction. Therefore, given the same maximum actuation, the probability flow towards vertically neighboring bins is equal to $\mid \frac{a_{1,i}}{a_{1,i,j,max}} - \frac{a_{3,i}}{a_{3,i,j,max}} \mid$. Both the terms inside the absolute value are positive and bounded to one. Hence, the absolute value of their difference is also bounded to one. In the same way, the fraction of density that is flowing out horizontally due to the actuator pair $\{a_{2,j}, a_{4,j}\}$ is given by $\mid \frac{a_{4,j}}{a_{4,i,j,max}} - \frac{a_{2,j}}{a_{3,i,j,max}} \mid$. Following this idea, a new density model is derived considering that the fraction of density distribution that remains in the bin is: $(1- \mid \frac{a_{1,i}}{a_{1,i,j,max}} - \frac{a_{3,i}}{a_{3,i,j,max}} \mid)(1- \mid \frac{a_{4,j}}{a_{4,i,j,max}} - \frac{a_{2,j}}{a_{3,i,j,max}} \mid)$.

To avoid formal complexity, it is useful to define:

$$\Delta_{v,i,j} = \frac{a_{1,i}}{a_{1,i,j,max}} - \frac{a_{3,i}}{a_{3,i,j,max}} \tag{4}$$

$$\Delta_{h,i,j} = \frac{a_{4,j}}{a_{4,i,j,max}} - \frac{a_{2,j}}{a_{2,i,j,max}} \tag{5}$$

Equations (4) and (5) represent the fraction of density (with sign) that is flowing out vertically or horizontally from bin $[i, j]$. The direction is given by the sign (the positive convention is taken left to right and down to up). The density evolution can now be written as:

$$\begin{aligned}
\mu_{k+1}[i,j] = \; & \mu_k[i,j](1- \mid \Delta_{v,i,j} \mid)(1- \mid \Delta_{h,i,j} \mid) \\
& + \mu_k[i-1,j-1] \left|\Delta_{v,i-1,j-1}\right| \left|\Delta_{h,i-1,j-1}\right| \\
& + \mu_k[i,j-1] \mid \Delta_{v,i,j-1} \mid (1- \mid \Delta_{h,i,j-1} \mid) \\
& + \mu_k[i+1,j-1] \mid \Delta_{v,i+1,j-1} \mid\mid \Delta_{h,i+1,j-1} \mid \\
& + \mu_k[i+1,j] \mid \Delta_{h,i+1,j} \mid (1- \mid \Delta_{v,i+1,j} \mid) \\
& + \mu_k[i+1,j+1] \mid \Delta_{v,i+1,j+1} \mid\mid \Delta_{h,i+1,j+1} \mid \\
& + \mu_k[i,j+1] \mid \Delta_{v,i,j+1} \mid (1- \mid \Delta_{h,i,j+1} \mid) \\
& + \mu_k[i-1,j+1] \mid \Delta_{v,i-1,j+1} \mid\mid \Delta_{h,i-1,j+1} \mid \\
& + \mu_k[i-1,j] \mid \Delta_{h,i-1,j} \mid (1- \mid \Delta_{v,i-1,j} \mid) \\
& = g(\mu_k, a_{1,i}, a_{2,j}, a_{3,i}, a_{4,j})
\end{aligned} \tag{6}$$

The first term of Equation (6) represents the probability mass flowing out of bin $[i, j]$ while the other eight terms represent the *mass* flowing into bin $[i, j]$ from adjacent bins. The presence of these *input* terms is conditioned by the sign of the respective Δ that gives the direction of propagation. The information on the direction of the flow is needed for

the conservation property to hold and for the overall propagation to be consistent with the propagation of the particles. Furthermore, note that this layout allows to take into account diagonal transitions. It is important to highlight that the density evolution model is based on the particle dynamics and it is used inside the controller only. The dynamics of the particles is propagated according to Equation (1) and it is not affected by the logic adopted inside the controller.

C. Optimal Transport based Control Problem

The objective of the control strategy is to steer the density distribution of the particles $\mu \in \mathbb{R}^{n_b \times n_b}$ towards a desired target distribution $\nu \in \mathbb{R}^{n_b \times n_b}$. The problem can be formulated as an optimization problem whose objective is to find the actions that minimize a certain distance between the actual and the desired distributions at each step.

The choice of a consistent metric for the cost function of the optimal control problem in the space of distributions is essential. In [22], the L_1 matrix norm is used. This is an element-wise convex operator between two matrix-valued functions. The structure of this metric and the horizon considered present some limitations: the optimization problem is solved considering only a single step, the density can only flow between adjacent bins in one time step due to the dynamic constraint and the maximum distance being bounded. In fact, if the two distributions are sufficiently "distant' to be disjointed (i.e. there is no overlap between them, see Figure 4) then $D_{L_1,max}(\mu, \nu) = 2$. This means that the cost function is not affected by the flow of the distribution at the next step because both $D_{L_1,max}(\mu_k, \nu) = D_{L_1,max}(\mu_{k+1}, \nu) = 2$ for every combination of the actuator intensities. Furthermore, being D_{L_1} an element-wise norm, it is a limited indicator of the spatial distance between the distributions on the two-dimensional grid. These problems are shown in Figures 4, 5.

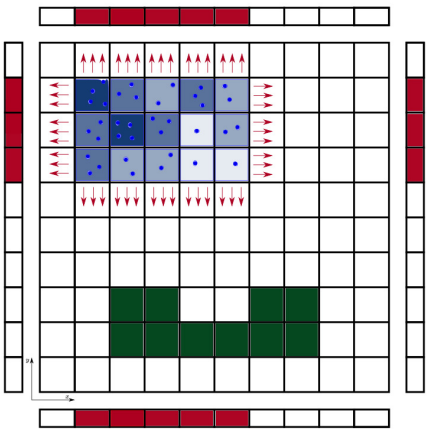

Fig. 4 The actual mass probability function μ_k (in blue) can only be transported between adjacent bins (red arrows), the L_1 distance with respect to the desired distribution ν (in green) is unchanged for every combination of the active actuators (in red).

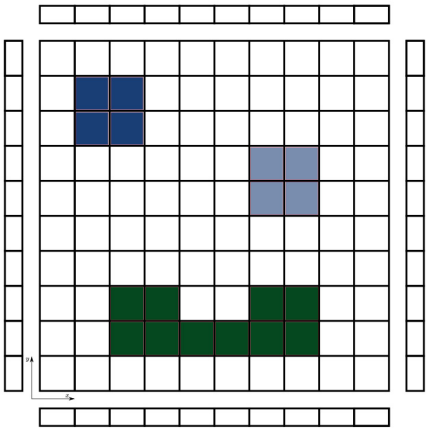

Fig. 5 The L_1 distances between the density function in blue and in light-blue with respect to the desired one ν (in green) have the same value while it is clear that the light-blue density is *closer* to the desired one.

III. Reformulation of the Optimal Transport problem

Optimal Transport theory provides a way to define a metric that simultaneously takes into account the distance over the grid and the *amount* of mass to move. This metric is the Wasserstein distance [17] that represents the optimal, minimum cost solution of Monge-Kantorovich's transport problem [20, 25, 26]. It can be seen as the geodesic between the two distributions and thus the ideal objective that the optimization problem should minimize to make the distribution at a certain time step as close as possible to the desired one.

A. Earth's Mover Distance

In our discrete two dimensional case, the probability distributions are matrix-valued functions. Each element of the matrix is associated to a bin and its value represents the probability mass in that bin. The discrete Monge's optimal transport problem seeks for a transport plan that determines how much and where the *mass* in each bin should be moved from the initial distribution to match the desired one, while minimizing a transportation cost.[‡]

The choice of the cost determines the structure of the overall problem [17]. The cost of transporting a certain amount of distribution from one bin to another can be associated with the ground distance between the bins times the amount of mass to be transported. Let's translate these sentences into equations. The general layout [19, 28] can be reformulated for our particular case as follows. The matrix-valued functions μ and ν representing the initial and final distributions defined above are vectorized per rows. With a slight abuse of notation, the vectorized functions are defined as $\mu \in \mathbb{R}^{n_b^2}$ and $\nu \in \mathbb{R}^{n_b^2}$. The conservation property is now expressed as $\sum_{i=1}^{n_b^2} \mu[i] = 1$ and it obviously holds for the final distribution ν as well $\sum_{i=1}^{n_b^2} \nu[i] = 1$. The clusters are associated with the location of each bin and its weight is given by the mass probability function evaluated at that bin. The signature describing the mass probability function μ is

[‡]Intuitively, given two distributions, one can be seen as a mass of earth properly spread in space, the other as a collection of holes in that same space. Then, the EMD measures the least amount of work needed to fill the holes with earth. Here, a unit of work corresponds to transporting a unit of earth by a unit of ground distance [19].

$M = \{\text{bin}_i, \mu[i]\}$ where the bin_i and its associated weight $\mu[i]$ are properly ordered. In the same way, the signature describing the target distribution ν is $N = \{\text{bin}_i, \nu[i]\}$. A ground metric $d_{i,j}$ is defined between the bins i and j as the Eucledian distance between the centroids of these bins in a normalized space:

$$d_{i,j} = D_{L_2}(\text{bin}_i, \text{bin}_j) = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2} \tag{7}$$

Note that bin_i is a two dimensional object and its elements are the location of the centroid. The distance between vertically or horizontally adjacent bins is defined to be unitary, irrespective of the physical bin length (i.e. the discretization of the workspace). The flow $\phi_{i,j}$ is defined as the amount of *mass* to be transported from bin i with total *mass* $\mu[i]$ to bin j with total *mass* $\nu[j]$. The cost (work) of transporting a certain amount of *mass* $\phi_{i,j}$ from bin_i to bin_j is defined as $W(\phi_{i,j}) = d_{i,j}\phi_{i,j}$, i.e. the cost of transporting is the amount of *mass* times the distance measured in the ground metric. The optimal objective is thus to achieve the transport plan with the minimum amount of total work. This problem is recast as a linear program (LP):

$$\min_{\phi_{i,j}} \quad \sum_{i=1}^{n_b^2} \sum_{j=1}^{n_b^2} d_{i,j}\phi_{i,j} \tag{8a}$$

$$\text{subject to} \quad \sum_{i=1}^{n_b^2} \phi_{i,j} = \nu[j] \qquad \forall j \in \{1, .., n_b^2\} \tag{8b}$$

$$\sum_{j=1}^{n_b^2} \phi_{i,j} = \mu[i] \qquad \forall i \in \{1, .., n_b^2\} \tag{8c}$$

$$\phi_{i,j} \geq 0 \qquad \forall i, j \in \{1, .., n_b^2\} \tag{8d}$$

The cost function in Equation (8a) represents the total work. The constraint Equation (8b) expresses the fact that the total flow from each bin, associated with the initial distribution μ, that is directed to bin j must equal the amount of *mass* of the desired distribution at that bin (i.e. $\nu[j]$). The constraint in Equation (8c), instead, is the dual of the previous one: the total flow that exits from bin i must equal the initial distribution present in that bin (i.e. $\mu[i]$). Finally, the constraint in Equation (8d) imposes a forward flow from distribution μ to distribution ν and not vice versa. The solution to Problem (8) gives the optimal map $\phi^*_{i,j}$ (i.e. *what* goes *where*) and the optimal cost. The later, normalized by the total flow, is the so-called Earth's Mover Distance (EMD):

$$D_{EMD}(\mu, \nu) = \frac{\sum_{i=1}^{n_b^2} \sum_{j=1}^{n_b^2} d_{i,j}\phi^*_{i,j}}{\sum_{i=1}^{n_b^2} \sum_{j=1}^{n_b^2} \phi^*_{i,j}} = \sum_{i=1}^{n_b^2} \sum_{j=1}^{n_b^2} d_{i,j}\phi^*_{i,j} \tag{9}$$

since the denominator $\sum_{i=1}^{n_b^2}\sum_{j=1}^{n_b^2}\phi_{i,j}^* = 1$.

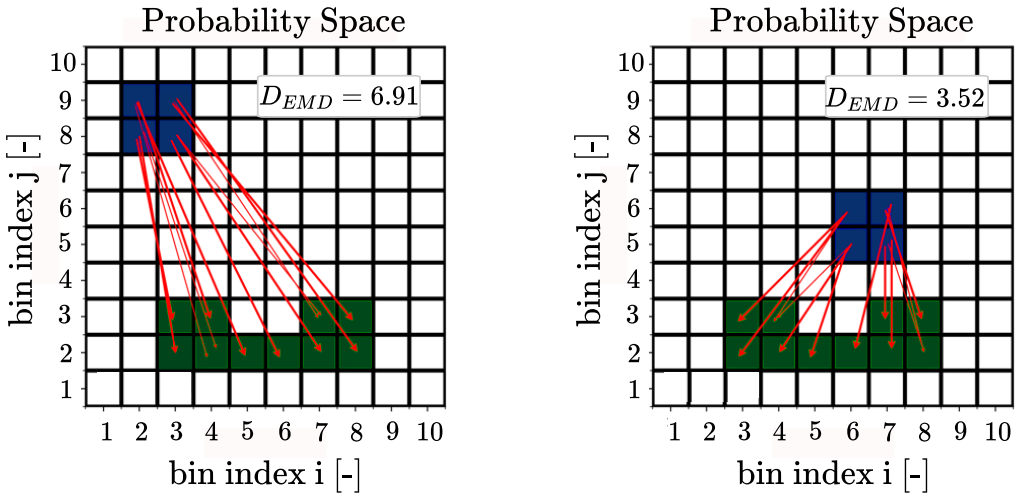

Fig. 6 Optimal transport maps between the initial distribution μ (in blue) and the desired distribution ν (in green) for two cases. The optimal transport flow $\phi_{i,j}^*$ is represented by red arrows. The thicker the arrow the more the amount of transported mass.

The EMD is thus the minimum amount of normalized *work* to be done to transport the initial distribution μ into the desired distribution ν. In this case, the total amount of mass is normalized to one, being μ and ν mass probability functions. However, the classical problem formulation is unconstrained (the mass can be moved to every bin in one single step). Thus the optimal flow $\phi_{i,j}^*$ and its associated EMD can be interpreted as the ideal set of actions to be performed in the unconstrained case. The optimal transport problem is solved for the two cases represented in Figure 5 and the results are shown in Figure 6. It is shown that the EMD is able to capture the difference between the two situations. Therefore, it is a more accurate metric for the density control problem.

B. Density Control using the Earth's Mover Distance

In the density control problem, the dynamic constraint of Equation (6) limits the flow of the density to neighboring bins only. As a consequence, the optimal transport plan represented by the EMD cannot be achieved in one step. However, the control problem can be recast as finding the best possible actuator intensities that minimize the EMD at the successive step. In this way, the algorithm searches for the direction of propagation that will have minimum *distance* from the desired distribution. It should be noted that, even if the propagation and the control horizon are limited, the nature of this new metric is able to *see* the objective in every feasible configuration in the probability space. Hence, there will always be a feasible optimal direction that minimizes the cost until a steady-state distribution is reached (i.e. condition of minimum cost). The drawback of this implementation is the complexity and the computational cost. The optimal control problem in the space of density distributions is formulated with the EMD as cost function subject to the density dynamical model in Equation (6) :

$$\min_{a_{1,i},a_{2,j},a_{3,i},a_{4,j}} \quad D_{EMD}(\mu_{k+1}, \nu) \tag{10a}$$

$$\text{subject to} \quad \mu_{k+1}[i,j] = g(\mu_k, a_{1,i}, a_{2,j}, a_{3,i}, a_{4,j}) \tag{10b}$$

$$a_{1,i} \leq a_{max} \quad \forall i \in \{1,..,n_\text{b}\} \tag{10c}$$

$$a_{2,j} \leq a_{max} \quad \forall j \in \{1,..,n_\text{b}\} \tag{10d}$$

$$a_{3,i} \leq a_{max} \quad \forall i \in \{1,..,n_\text{b}\} \tag{10e}$$

$$a_{4,j} \leq a_{max} \quad \forall j \in \{1,..,n_\text{b}\} \tag{10f}$$

At each step, the algorithm determines the actions that minimize the EMD. For each trial solution, the cost is calculated as the solution of a nested linear program of the same form as Problem (8) but obtained replacing μ with the propagated density μ_{k+1}. This is still a single step algorithm, in the sense that, when the approximate optimal solution is obtained, the particle dynamics is propagated and the cycle is repeated.

Since the EMD is a metric between probability distributions and it is obtained from the solution of Problem (8), it is possible to show [29] that if the discretized single-step density state dynamics is linear then the optimization Problem (10) is convex and thus the minimum is global. However, the dynamic constraint (6) is non-smooth and non-convex due to the cross product between optimization variables. Hence, the optimization Problem (10) is non-smooth and non-convex too. Note however that the objective function is still convex in μ_{k+1} and that the optimization variables are positive and box constrained. Without relying on specific algorithms for non-smooth, non-convex problems, we use the quasi-newton method "BFGS" from the SciPy [30] Python library that is able to obtain satisfactory results in most cases [31]. We remark the fact that the minimum found by the method is local in general. However, in our case, we have noticed that the minimum achieved numerically is the same irrespective of the initial guess provided to the solver thus suggesting that the minimum achieved is global.

IV. Numerical Simulations

The previously defined control logic is applied to a high number (N=100000) of particles randomly distributed over the workspace.

It is interesting to notice that the optimization problem does not scale with the number of the particles but with the number of grid elements. The workspace is discretized according to the number of actuators: $n_a = 10$ per side. Hence, there are $n_\text{b} = n_\text{a}$ bins per side, while the total number of bins is $n_\text{T} = n_\text{b}^2$. At each iteration (i.e. discrete time step), the density distribution is computed and it is used as initial condition to solve the optimization Problem (10). The optimal actions are then used as inputs in the physical space for the system of particles. The initial distribution

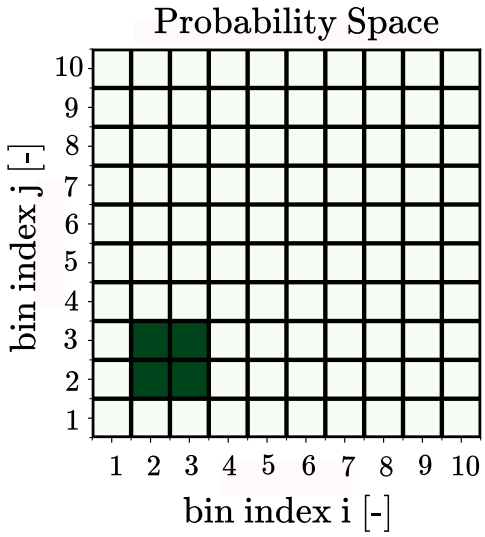

Fig. 7 Target probability distribution ν

shown in Figure 9 both in the physical and probability space is pushed towards the target distribution shown in Figure 7. Successive convergence steps are shown by Figures 10 and 11.

Three simulations are carried out starting from the same random distributions. The cost time history for each case is compared in Figure 8. The continuous line labeled (1) represents the cost progression for the previously described control logic, that is with the EMD as cost function and using the new density model as dynamic constraint. The dashed line labeled as (2) is the optimal control problem with the EMD as cost function and the density model used in [22]. It is straightforward to notice that the convergence is slower and the steady-state value reached higher. This is due to the partial inconsistency between the actual particle dynamics and the density evolution. The control logic is no more able to reduce the cost even if it is possible to reach a lower cost. On the other hand, the dash-dotted line labeled (3) is the control logic that uses the L_1 distance as cost function and the old density model. In order to compare results, it is plotted together with the other simulations with the EMD as ordinate. As pointed out in the previous section, the L_1 lacks a "spatial" indicator of the distance between the actual density and the target. Therefore, it is not able to pick up the correct direction of transport of the density reaching a steady-state value almost without moving the initial random distribution.

On the contrary, the control logic described in this paper (labeled (1)) shows an exponential convergence to a steady-state value. After roughly 75 iterations the algorithm is not able to further reduce the distance. This is a consequence of the layout of the actuators and of the dynamic model assumed for the density function. In fact, each actuator has influence over all its rows/columns and determines a flow for each element in this row/column. Hence, not every target distribution is achievable with zero convergence error, even if in this case the steady-state value is close to zero.

The trade-off between grid refinement and computational cost is now analysed. Retaining the same initial and target

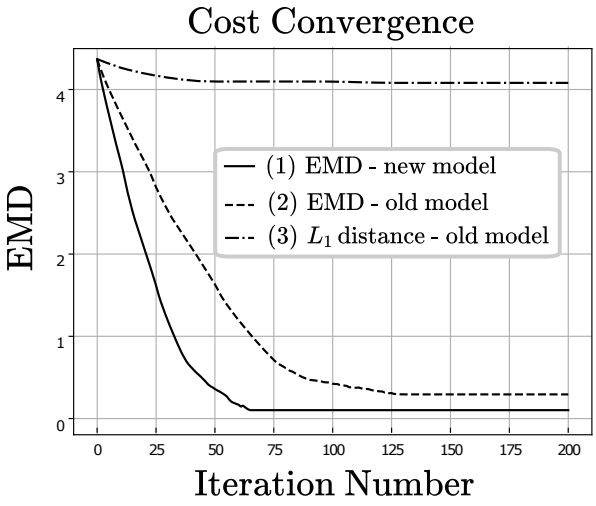

Fig. 8 Convergence of the swarm to the configuration that minimizes the EMD for different macroscopic models and cost functions.

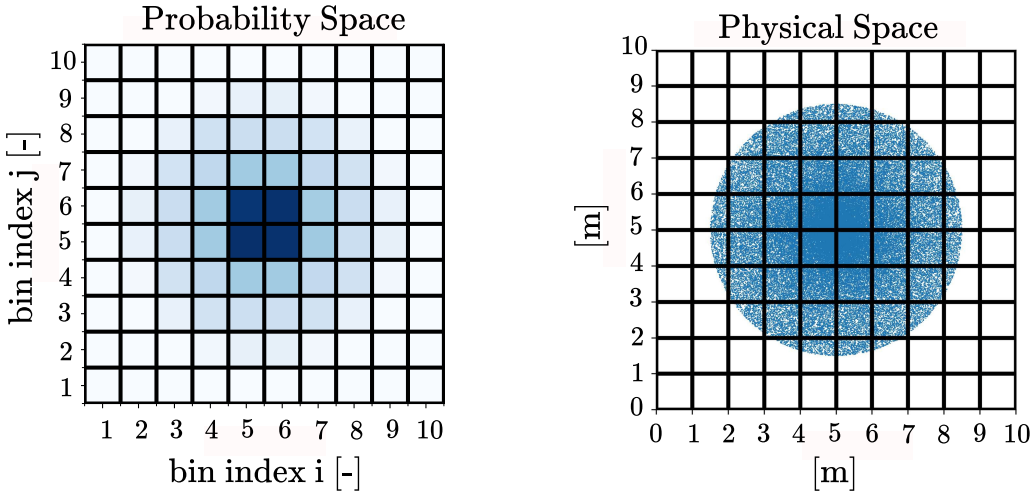

Fig. 9 Iteration number $k = 0$

distribution in the particle space the number of bins per side is varied. The total computational cost of Problem (10) scales as n_b^3 as it is shown in Figure 13. On the other hand, the final EMD achieved by varying the grid size is shown in Figure 12. Note that the points obtained in Figure 12 are computed using the finest grid, that is $n_b = 18$. It is interesting to note that grids finer than $n_b = 10$ do not obtain better results. This is due to the way the target is specified. In the case considered, in order to accurately track the target it is sufficient to use $n_b = 10$ bins per side. In Figure 14, the number of iterations to reach the target (i.e. minimum EMD) is shown. The linear proportionality with the grid size is a result of the maximum actuation strength obtained from Equation (3) that limits the particles to move only across neighboring bins so that the number of time steps needed to reach the minimum is directly proportional to the grid size.

Before closing this section we draw some considerations on the optical performances on the system. First of all, we remark that we have considered the cloud coarse confinement problem and we have shown that the cloud of particles

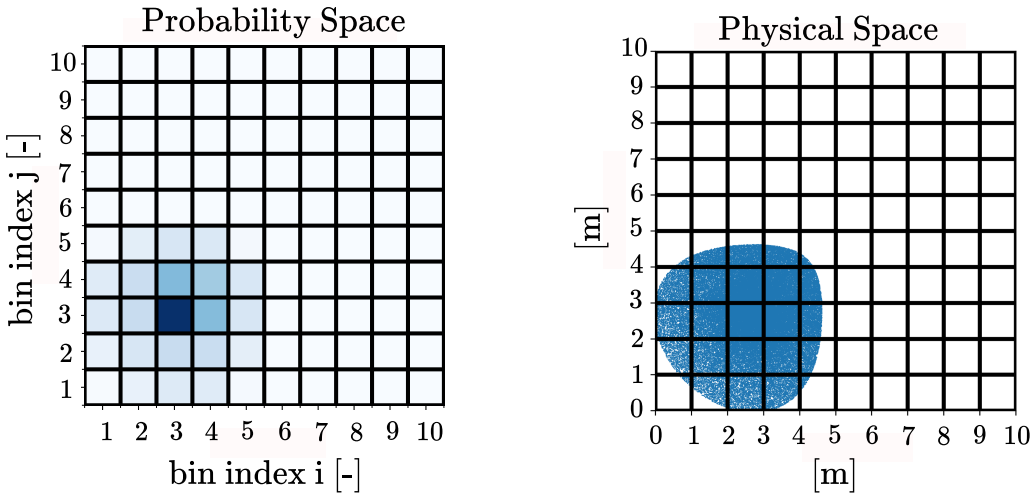

Fig. 10 Iteration number $k = 30$

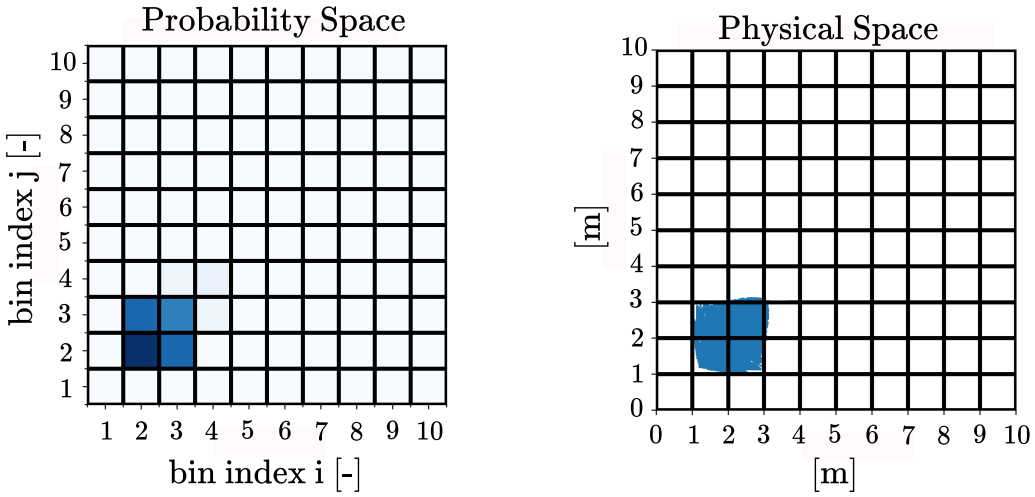

Fig. 11 Iteration number $k = 75$

can achieve the target density with a level of detail that depends on the number of actuators. As mentioned in the introduction, the optical performances of the granular cloud rely on Bruggemann's theory so that the focusing length of the aperture is inversely proportional to the fill factor. Furthermore, in [6] it is shown that optical properties do not deteriorate with a fill factor as small as 0.3 thus justifying the complexity of the control architecture. As it is shown in Figure 11, the control algorithm is able to reach an evenly distributed aperture shape as target density so that the density is 25000 particles per bin, assuming a bin area of 1 $[m^2]$, approximately spherical grains and the minimum working fill factor, particles of radius of roughly 2 $[mm]$ should be considered. Finally, we note that the iterative control algorithm does not scale with the number of particles and thus the same density results can be achieved with micron-sized grains as well. However, simulating such a system would require the computational power to forward propagate 10^{10} particles.

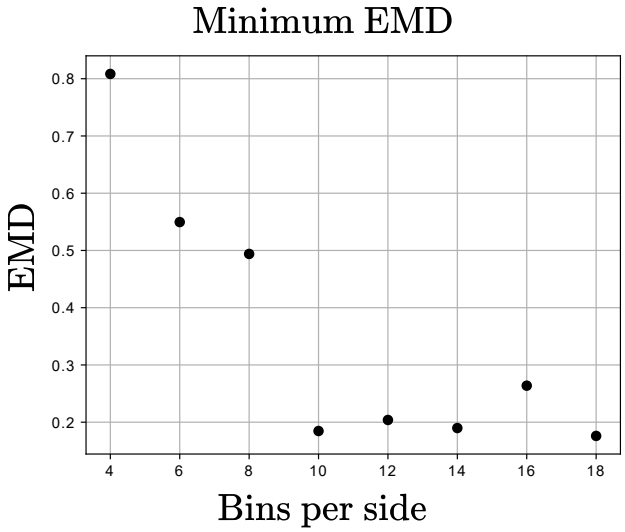

Fig. 12 Minimum EMD achieved varying the grid size. the EMD values are computed with respect to the finest grid.

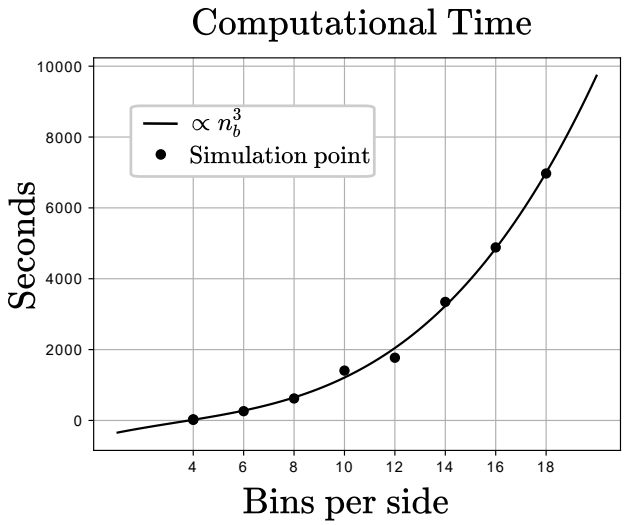

Fig. 13 Computational time for the simulation closely follows a cubic trend. Note that the simulated number of particles is the same for each point.

V. Experimental Validation

In this section, the Robotarium [32] platform is used to perform an experimental validation of the control algorithm. The Robotarium environment consists of a swarm of unicycle robots [33]. The state of each agent is tracked with a motion capture system while the local control signals can be broadcasted to each robot independently. In this framework, each underactuated particle is modelled as a unicycle robot. The control algorithm is run based on the state of each agent and the actuators' input intensities are computed. The simulation algorithm propagates the dynamics of the particles and computes each individual target based on the actuators' values.

The algorithm is implemented as an iterative optimal single-step controller so that the control horizon is equal to

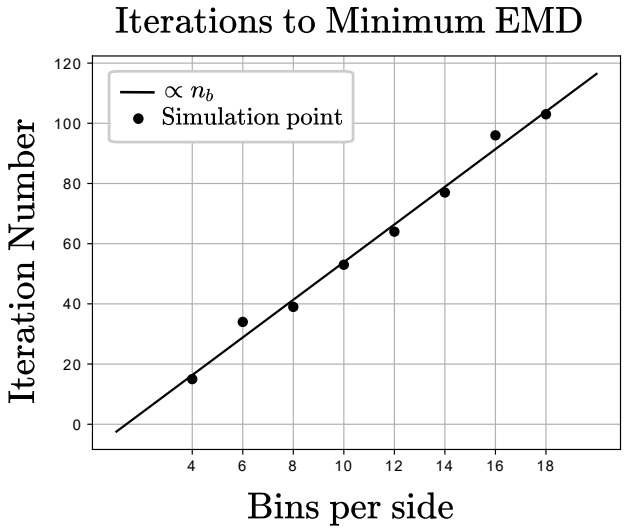

Fig. 14 Number of iterations (i.e. time steps) to reach the minimum EMD as a function of the grid size.

one. The parameters given as input to the controller are consistent with the dimension of the Robotarium's workspace. Hence, the density control algorithm is implemented considering a $2[m] \times 2[m]$ grid with 3 bins for each side. The actuators are assumed to be at a distance of $0.5[m]$ and the forward particle dynamics in Equation (1) is propagated with a time step of $0.5[s]$. The target position of each robot at each step is obtained from Equation (1) using the actuation intensities obtained by the density control algorithm.

Then, low-level controllers are used to guiding each robot to the target.

The robot control architecture relies on barrier control functions to avoid collisions and a Proportional-Derivative (PD) position controller [33]. Since the dimension of the bin is much greater than the single robot dimension and the local density is coarse, the collision avoidance algorithm is only active once the robots are in the same bin. Thus, the macroscopic density dynamics is not affected by the low-level microscopic dynamics. These considerations are confirmed by the experimental results as shown in Figure 15 where the macroscopic EMD convergence is not influenced by the weak activation of the control barrier algorithm. On the other hand, the PD controller is saturated at the maximum robot input speed (i.e. $0.15\left[\frac{m}{s}\right]$) so that the robot does not achieve the target position obtained from Equation (1) in a single step as it can be noticed by the stair-case shape of the EMD convergence in Figure 15.

The experimental test is initialized with a random distribution of a swarm of 8 agents as shown in Figure 16. The values of the density matrix μ are plotted in the background during the entire execution of the experiment. The target distribution is the bottom-right bin. Figure 16 shows a sequence of the iteration of the control algorithm where the agents are pushed towards the target distribution. The control logic is able to reach convergence with almost zero EMD from the target as shown in Figure 15. This situation is depicted in Figure 17. Although being in a controlled environment and with unicycle robots simulating the particles, the experiment proves the performance of the density control algorithm developed in this paper.

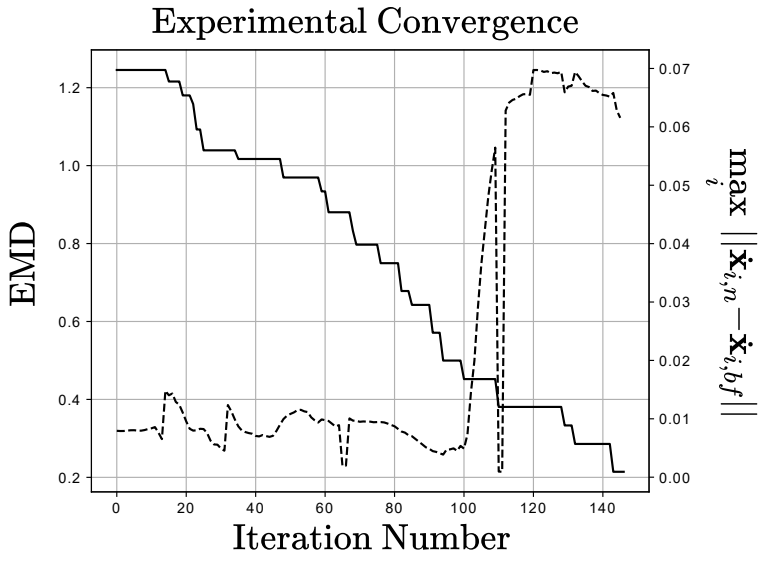

Fig. 15 Comparison of the microscopic and macroscopic dynamics obtained from the experiment. The macroscopic density dynamics (on the left axis) is not influenced by the local collision avoidance logic (on the right axis). The level of activation of the control barrier function is obtained by comparing the nominal velocity $\dot{\mathbf{x}}_{i,n}$ and the input one $\dot{\mathbf{x}}_{i,bf}$ obtained from the barrier function algorithm for each robot in the swarm.

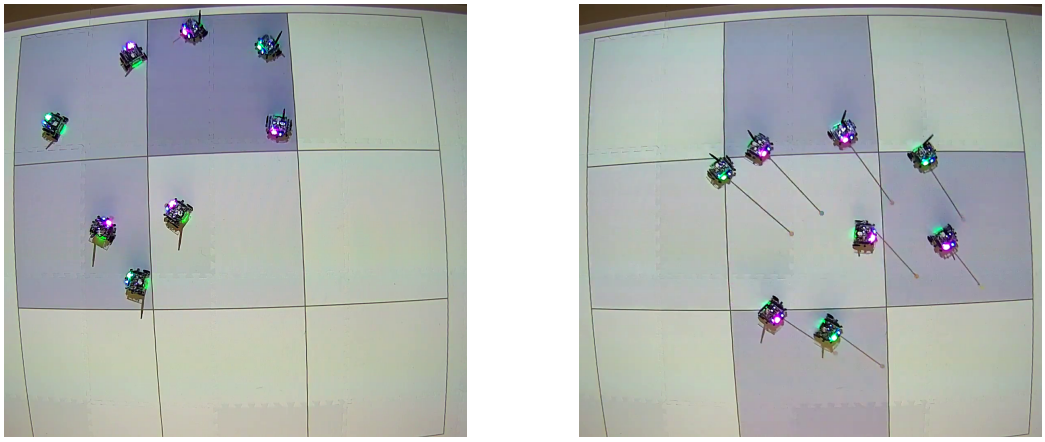

Fig. 16 Initial random distribution at $t = 0$ [s] and and at $t = 5$ [s]

VI. Conclusion

In this paper, a novel density control strategy for shaping a very-large cloud of particles into a reflective dish, using a small number of actuators is presented. Such a strategy has been successfully tested both using numerical simulations and experimentally on the Robotarium testbed. The main innovative feature is to cast the control problem in the space of probability distributions and derive an iterative single step optimal control algorithm for the swarm's density distribution. The Wasserstein distance has been introduced as a more effective cost functional for the optimal control problem. This metric, taken from optimal transport theory, can solve the optimization problem even when the actual and target distributions are disjointed, and it is also a more meaningful measure of distance between distributions than the Euclidean norms. Furthermore, a more accurate density evolution has been introduced that will eventually lead to more

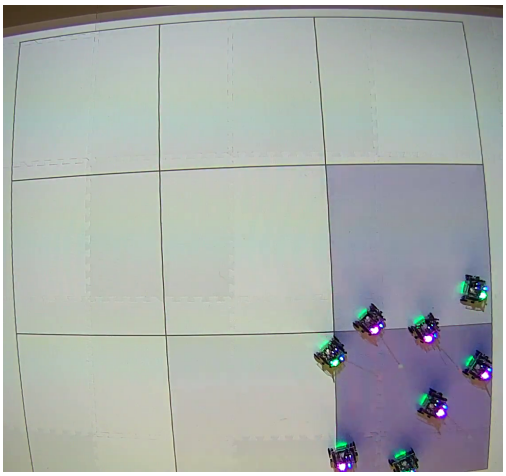

Fig. 17 Target density reached at $t = 10$ [s]

efficient and physically meaningful macroscopic models for particle swarms. It is envisioned that the algorithms and results introduced in this paper will enable a new generation of large-aperture space telescopes for the space exploration.

Funding Sources

The U.S. Government sponsorship is acknowledged. This research was carried out at the Jet Propulsion Laboratory (JPL), California Institute of Technology, under a contract with the National Aeronautics and Space Administration, during the internship of the first author sponsored by the JPL Visiting Student Research Program. The reference project for the activity was the Orbiting Rainbows concept, developed by Marco B. Quadrelli under a contract with the NASA Innovative Advanced Concepts Program NNH14ZOA001N-14NIAC_A2.

References

[1] Corbacho, V. V., Kuiper, H., and Gill, E., "Review on thermal and mechanical challenges in the development of deployable space optics," *Journal of Astronomical Telescopes, Instruments, and Systems*, Vol. 6, No. 1, 2020. https://doi.org/10.1117/1.JATIS.6.1.010902.

[2] Quadrelli, M. B., Basinger, S., Arumugam, D., and Swartzlander, G., "NIAC Phase II Orbiting Rainbows: Future Space Imaging with Granular Systems," NASA Innovative Advanced Concepts (NIAC) NNH14ZOA001N, February 2017.

[3] Quadrelli, M. B., Basinger, S., and Sidick, E., "Unconventional imaging with contained granular media," *Unconventional and Indirect Imaging, Image Reconstruction, and Wavefront Sensing 2017*, Vol. 10410, International Society for Optics and Photonics, 2017, p. 104100W. https://doi.org/10.1117/12.2272617.

[4] Quadrelli, M. B., Basinger, S., Swartzlander, G., and Arumugam, D., "Dynamics and control of granular imaging systems," *AIAA SPACE 2015 Conference and Exposition*, 2015, p. 4484. https://doi.org/10.2514/6.2015-4484.

[5] Quadrelli, M. B., and Ius, P., "Modeling and Simulation of Trapping Mechanisms of Granular Media In Space," *AIAA/AAS Astrodynamics Specialist Conference*, 2016, p. 5572. https://doi.org/10.2514/6.2016-5572.

[6] Basinger, S. A., Palacios, D., Quadrelli, M. B., and Swartzlander Jr, G. A., "Optics of a granular imaging system (ie "orbiting rainbows")," *UV/Optical/IR Space Telescopes and Instruments: Innovative Technologies and Concepts VII*, Vol. 9602, International Society for Optics and Photonics, 2015, p. 96020E. https://doi.org/10.1117/12.2185699.

[7] Krishnan, V., and Martínez, S., "Distributed optimal transport for the deployment of swarms," *2018 IEEE Conference on Decision and Control (CDC)*, IEEE, 2018, pp. 4583–4588. https://doi.org/https://10.1109/CDC.2018.8619816.

[8] Rossi, F., Bandyopadhyay, S., Wolf, M., and Pavone, M., "Review of multi-agent algorithms for collective behavior: a structural taxonomy," *IFAC-PapersOnLine*, Vol. 51, No. 12, 2018, pp. 112–117. https://doi.org/10.1016/j.ifacol.2018.07.097.

[9] Bandyopadhyay, S., Chung, S.-J., and Hadaegh, F. Y., "A probabilistic Eulerian approach for motion planning of a large-scale swarm of robots," *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2016, pp. 3822–3829. https://doi.org/10.1109/IROS.2016.7759562.

[10] Demir, N., Eren, U., and Açıkmeşe, B., "Decentralized probabilistic density control of autonomous swarms with safety constraints," *Autonomous Robots*, Vol. 39, No. 4, 2015, pp. 537–554. https://doi.org/10.1007/s10514-015-9470-z.

[11] Açıkmeşe, B., and Bayard, D. S., "Probabilistic swarm guidance for collaborative autonomous agents," *2014 American control conference*, IEEE, 2014, pp. 477–482. https://doi.org/10.1109/ACC.2014.6859358.

[12] Açikmeşe, B., and Bayard, D. S., "A Markov chain approach to probabilistic swarm guidance," *2012 American Control Conference (ACC)*, IEEE, 2012, pp. 6300–6307. https://doi.org/10.1109/ACC.2012.6314729.

[13] Brockett, R., and Khaneja, N., "On the stochastic control of quantum ensembles," *System theory*, Springer, 2000, pp. 75–96. https://doi.org/10.1007/978-1-4615-5223-9_6.

[14] Zlotnik, A., and Li, S., "Synthesis of optimal ensemble controls for linear systems using the singular value decomposition," *2012 American Control Conference (ACC)*, IEEE, 2012, pp. 5849–5854. https://doi.org/10.1109/ACC.2012.6315297.

[15] Chen, Y., *Modeling and control of collective dynamics: From Schrödringer bridges to Optimal Mass Transport*, Ph. D. Dissertation, Department of Mechanical Engineering, University of Minnesota, 2016.

[16] De Badyn, M. H., Eren, U., Açikmeşe, B., and Mesbahi, M., "Optimal mass transport and kernel density estimation for state-dependent networked dynamic systems," *2018 IEEE Conference on Decision and Control (CDC)*, IEEE, 2018, pp. 1225–1230. https://doi.org/10.1109/CDC.2018.8619808.

[17] Villani, C., *Optimal transport: old and new*, Vol. 338, Springer Science & Business Media, 2008. https://doi.org/10.1007/978-3-540-71050-9.

[18] Bandyopadhyay, S., Chung, S.-J., and Hadaegh, F. Y., "Probabilistic swarm guidance using optimal transport," *2014 IEEE Conference on Control Applications (CCA)*, IEEE, 2014, pp. 498–505. https://doi.org/"10.1109/CCA.2014.6981395".

[19] Rubner, Y., Tomasi, C., and Guibas, L. J., "The earth mover's distance as a metric for image retrieval," *International journal of computer vision*, Vol. 40, No. 2, 2000, pp. 99–121. https://doi.org/10.1023/A:1026543900054.

[20] Kolouri, S., Park, S., Thorpe, M., Slepčev, D., and Rohde, G. K., "Transport-based analysis, modeling, and learning from signal and data distributions," *arXiv preprint arXiv:1609.04767*, 2016.

[21] Yang, Y., and Engquist, B., "Analysis of optimal transport and related misfit functions in full-waveform inversion," *Geophysics*, Vol. 83, No. 1, 2018, pp. A7–A12. https://doi.org/10.1190/geo2017-0264.1.

[22] S.Bandyopadhyay, and Quadrelli, M., "Optimal Transport Based Control of Granular Imaging System in Space," *International Workshop on Satellite Constellations and Formation Flying 17-17*, 2017.

[23] Zhang, K., and Zhang, Y., "Optimal reconfiguration with collision avoidance for a granular spacecraft using laser pressure," *Acta Astronautica*, Vol. 160, 2019, pp. 163 – 174. https://doi.org/10.1016/j.actaastro.2019.04.021.

[24] Brockett, R., "Notes on the control of the Liouville equation," *Control of partial differential equations*, Springer, 2012, pp. 101–129. https://doi.org/10.1007/978-3-642-27893-8_2.

[25] Monge, G., "Mémoire sur la théorie des déblais et des remblais," *Histoire de l'Académie Royale des Sciences de Paris*, 1781.

[26] Kantorovitch, L., "On the translocation of masses," *Management science*, Vol. 5, No. 1, 1958, pp. 1–4.

[27] Ambrosio, L., "Lecture notes on optimal transport problems," *Mathematical aspects of evolving interfaces*, Springer, 2003, pp. 1–52. https://doi.org/10.1007/978-3-540-39189-0_1.

[28] Snow, M., and Van lent, J., "Monge's Optimal Transport Distance for Image Classification," *arXiv preprint arXiv:1612.00181*, 2016.

[29] Sinigaglia, C., *Granular Imager: Modeling and Control*, Master's Thesis, Department of Mechanical Engineering, Polytechnic University of Milan, 2019. URL https://dartslab.jpl.nasa.gov/References/pdf/other/2019-granularimager-Sinigaglia.pdf.

[30] Jones, E., Oliphant, T., Peterson, P., et al., "SciPy: Open source scientific tools for Python," 2001. URL http://www.scipy.org/.

[31] Yu, J., Vishwanathan, S., Günter, S., and Schraudolph, N. N., "A quasi-Newton approach to nonsmooth convex optimization problems in machine learning," *The Journal of Machine Learning Research*, Vol. 11, 2010, pp. 1145–1200. https://doi.org/https://10.1145/1390156.1390309.

[32] Pickem, D., Glotfelter, P., Wang, L., Mote, M., Ames, A., Feron, E., and Egerstedt, M., "The robotarium: A remotely accessible swarm robotics research testbed," *2017 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2017, pp. 1699–1706. https://doi.org/10.1109/ICRA.2017.7989200.

[33] Pickem, D., Lee, M., and Egerstedt, M., "The GRITSBot in its natural habitat-a multi-robot testbed," *2015 IEEE International conference on robotics and automation (ICRA)*, IEEE, 2015, pp. 4062–4067. https://doi.org/10.1109/ICRA.2015.7139767.