# From Car to Fleet: The Compositionality of Optimal Control

Antonio Terpin[1†], Nicolas Lanzetti[1†], and Florian Dörfler[1]

***Abstract*— We study the optimal control of fleets of identical particles (e.g., robots, autonomous cars, etc.), which we capture *macroscopically* as probability measures. First, we study *fleet-to-particle* optimal control via dynamic programming in probability spaces. Second, we investigate its compositionality properties: Perhaps surprisingly, we show that in many cases of practical interest we can find the optimal solution in a *particle-to-fleet* fashion, combining two ingredients: (i) the solution of dynamic programming of each particle and (ii) the solution of an optimal transport problem. Intuitively, this means that the "low-level control of the particles" (how to reach the destination?) and "fleet-level control" (who goes where?) are decoupled. Beside its practical relevance, this work opens the field for a rigorous investigation on the compositionality of optimal control problems via category theory.**

## I. Introduction

This work deals with the following fundamental question: How to optimally control a (possibly very large) fleet of identical particles? This problem has application across various domains, including robotic coordination [10], [13], [7], [21], mobility systems [25], [24], [18], social networks [2], [15], and biological models [19], [22].

In general, one can adopt two different approaches: a *particle-to-fleet* approach and *fleet-to-particle* approach. In fleet-to-particle approaches, one *jointly* designs the control strategies of all particles. This approach is of course very principled, but suffers from the size of the fleet: For very large scale systems, the analysis and computation of (optimal) control strategies becomes intractable. In particle-to-fleet approaches, instead, one first solves the optimal control (OC) problem for every particle as if it was the only one, and then *composes* these solutions to obtain a control strategy for the fleet. This approach is particularly attractive for various reasons. First, one can effectively deploy the rich theory for the control of single agents. Second, the approach mildly suffers from the size of the fleet: One only needs to combine the solutions of the individual particles, but does not need to jointly design their control strategies. Unfortunately, these benefits are at the price of generally sub-optimal (from the fleet perspective) control strategies.

In this work, we study the interplay between these two approaches via OC and optimal transport (OT). Notably, in many settings of practical interest they coincide: We can compose the optimal control strategies of the individual particles to obtain an optimal control strategy for the fleet.

†: Equal contribution. This research was supported by Swiss National Science Foundation under the NCCR Automation (grant 51NF40_180545).

[1]Antonio Terpin, Nicolas Lanzetti, and Florian Dörfler are with the Automatic Control Laboratory, ETH Zürich, Zürich, Switzerland, {aterpin,lnicolas,dorfler}@ethz.ch

While the steering of probability measures has received much attention in continuous time [12], [9], [11], there is a lack of results in the discrete time settings, where the analysis is limited to integrator and linear dynamics [4], [5], [6]. In these works the approach is fleet-to-particle. Conversely, in [16], [17] the approach is particle-to-fleet: First, an OC problem for each particle yields a cost for every initial and terminal state; second, an OT problem gives the control strategy for the fleet. Yet, the optimality at a fleet level is not assessed: Even if the control laws for the individual particles are optimal, the macroscopic behaviour might be sub-optimal.

Our main contributions are (i) the rigorous study of discrete-time OC in probability spaces via dynamic programming (DP) (i.e., fleet-to-particle approach), and (ii) the characterization of conditions under which the particle-to-fleet and fleet-to-particle approaches coincide.

## II. Background in optimal transport

Let $\mathcal{X}$ be a Polish space (i.e., separable completely metrizable topological space, such as $\mathbb{R}^n$ or any finite space). The space of all Borel probability measures over $\mathcal{X}$ is denoted by $\mathcal{P}(\mathcal{X})$. The *pushforward* of $\mu$ via $f : \mathcal{X} \to \mathcal{X}$ is defined by $(f_{\#}\mu)(A) = \mu(f^{-1}(A))$ for all Borel sets $A \subset \mathcal{X}$. Given a non-negative lower semi-continuous function $c : \mathcal{X} \times \mathcal{X} \to \mathbb{R}_{\geq 0}$ (called transport cost) and two probability measures $\mu, \nu \in \mathcal{P}(\mathcal{X})$, the *OT cost* between $\mu$ and $\nu$ is defined by

$$C(\mu, \nu) \coloneqq \inf_{\gamma \in \Gamma(\mu,\nu)} \int_{\mathcal{X}\times\mathcal{X}} c(x_1, x_2) \mathrm{d}\gamma(x_1, x_2), \qquad (1)$$

where $\Gamma(\mu, \nu)$ is the set of all probability measures over $\mathcal{X}\times\mathcal{X}$ with marginals $\mu$ and $\nu$, often called *transport plans* [23], [20], [3]. The semantics are as follows: We seek the minimum cost to transport the probability distribution $\mu$ onto the probability distribution $\nu$, when transporting a unit of mass from $x_1$ to $x_2$ costs $c(x_1, x_2)$. Accordingly, a transport plan $\gamma \in \Gamma(\mu, \nu)$ encodes the allocation of probability mass: if $(x_1, x_2) \in \mathrm{supp}(\gamma)$, then some of the probability mass at $x_1$ is displaced to $x_2$. When the OT plan is of the form $\gamma = (\mathrm{Id} \times T)_{\#}\mu$ for some measurable $T$, then $T$ is the OT map from $\mu$ and $\nu$.

## III. Problem setting

### A. Fleets as probability measures

In agreement with mean-field approches [1], we identify a fleet of identical particles whose state takes value in $\mathcal{X}$ as a probability measure $\mu$ over $\mathcal{X}$. Accordingly, $\mu(A)$ denotes the *share* of particles whose state belongs to the set $A \subset \mathcal{X}$. This setting encompasses many cases of practical interest:

*Example 1 (Robots in a grid):* Consider $n$ robots in a grid of $W \times H$ cells, where the $k^{th}$ agent is located in the cell

$(i_k, j_k)$. Then, $\mu = \frac{1}{n}\sum_{k=1}^{n} \delta_{(i_k,j_k)}$, where $\delta_{(i,j)}$ denotes a delta measure at $(i,j)$.

Via a simple augmentation of the state space, we can also consider fleets of heterogeneous particles:

*Example 2 (Intermodality):* Consider $n$ vehicles in a transportation network, represented by $N$ routes. Denote by $i_k$ the route in which the $k^{\text{th}}$ vehicle is located, and by $j_k$ its class: bicycle, car, train, etc. Then, the fleet is described macroscopically as $\mu = \frac{1}{n}\sum_{k=1}^{n} \delta_{(i_k,j_k)}$.

The discrete-time dynamics of the fleet of identical particles usually results from a pushforward operation via a map $f : \mathcal{X} \to \mathcal{X}$; i.e., $\mu^+ = f_\# \mu$. Intuitively, this means that all particles of $\mu$ located at $x \in \mathcal{X}$ are displaced to $f(x) \in \mathcal{X}$.

*Example 3:* Consider a fleet of robots, where every robot at $x \in \mathcal{X}$ receives as input $u(x) \in U$ and evolves with the dynamics $f(x, u(x))$: The fleet evolves as $\mu^+ = f(\cdot, u(\cdot))_\# \mu$.

### B. The optimal control of fleets

In this setting, the OC problem reads

$$J(\mu_0) = \inf_{u_k \in \mathcal{U}_k} \sum_{k=0}^{N-1} G_k(\mu_k, u_k) + V_N(\mu_N), \tag{2}$$

subject to dynamics and input constraints. In contrast with traditional OC, the optimization variables do not take values in $U$ (e.g., in $\mathbb{R}^p$). Rather, they are continuous maps from the particle space to the input space; i.e., $u_k \in \mathcal{U}_k \subseteq C^0(\mathcal{X}_k, U_k)$. Consequently, (2) is an infinite-dimensional optimization problem. For the cost terms there are many options:

*Example 4:* Given a quantity $h : \mathcal{X} \times U \to \mathbb{R}_{\geq 0}$, we can consider its average value over the fleet $\mathbb{E}^\mu[h(x, u(x))]$, and its variance $\mathrm{Var}^\mu[h(x, u(x))]$. More general terms can also be considered; e.g., the OT cost $C(\mu, \nu)$ from a reference measure $\nu$, or the Kullback-Leibler divergence.

Perhaps surprisingly, we can show that the conditions for well-posedness of the DP recursion are aligned with traditional DP; see [8, §4.2] and [14].

### C. When things are compositional

Unfortunately, DP in probability spaces is often impractical: There are limited instances with analytical solution, and a computational approach is not practical due to the infinite dimensionality of (2). Nonetheless, we prove that if (i) the stage costs are of the type $G_k(\mu, u) = \mathbb{E}^\mu[g_k(x, u(x))]$ and (ii) the terminal cost is an OT problem with transport cost $v_N$ between the fleet state $\mu_N$ at the end of the horizon and a reference probability measure $\nu$, then we can solve (2) in a particle-to-fleet fashion.

In particular, the cost-to-go at stage $k$ (i.e., $J_k$) is the OT cost (1) with transport cost $j_k$ between the probability measures $\mu_k$ (state of the fleet at stage $k$) and $\nu$ (same reference probability measure as in the terminal cost), where for all $x \in \mathcal{X}_k, y \in \mathcal{X}_N$ $j_k(x, y)$ is the cost to steer a particle at stage $k$ and state $x_k$ to $y$, with stage costs $g_k$ and terminal cost $v_N$. Moreover, each $j_k$ encodes the optimal strategy for the individual particles (i.e., the solution to the DP recursion in the ground space). Therefore, we effectively have a computational compositional approach to optimally steer the fleet:

1) Perform DP for every pair of initial and terminal states $x \in \mathrm{supp}(\mu_k)$, $y \in \mathrm{supp}(\nu)$ to find the optimal input $u_{k,xy}$ and the cost-to-go $j_k(x, y)$; and
2) Find the OT map between $\mu_k$ and $\nu$ with transportation cost $j_k(x, y)$ and apply $u_k(x) = u_{k,xT_k(x)}$.

### D. Discussion

If we consider a single particle, the macroscopic description reduces to a delta ($\mu = \delta_x$), and we recover the traditional case: Let $\nu = \delta_{x^*}$, then $T_k(x) = x^*$ and the fleet solution collapses to the single particle one. This suggests that OT induces a *lifting operator* from $\mathcal{X}$ to $\mathcal{P}(\mathcal{X})$. Similarly, one can prove that whenever the terms $g_k$ and $v_N$ satisfy the required assumptions for well-posedness of DP, so do their "lifted" counterparts $G_k, V_N$ [23], [3]. Moreover, a probability space over a Polish space $\mathcal{X}$ (endowed with the narrow topology) is itself a Polish space. Thus, our results can be rephrased in the language of category theory (e.g., via functors).

## IV. An Example

We now instantiate our results in Example 1. We endow the grid world with the commutative group structure $(i, j) + (h, k) = (i + h \bmod W, j + k \bmod H)$. Each robot has five actions available ($\texttt{UP} \equiv (0, 1)$, $\texttt{RIGHT} \equiv (1, 0)$, $\texttt{DOWN} \equiv (0, H-1)$, $\texttt{LEFT} \equiv (W-1, 0)$, and $\texttt{HOVER} \equiv (0, 0)$) and moves with the integrator dynamics $f(x, u(x)) = x + u(x)$. Additionally, some cells are blocked by bushes. The goal is to drive the swarm towards a fixed final configuration $\nu$, minimizing the travelled distance and avoiding the obstacles. The OC problem can be formulated for each particle via the cost terms: (i) $g(x, u(x)) = +1$ if $f(x, u(x)) \equiv \texttt{FREE}$, and $+\infty$ if there is an obstacle or the agents goes left when in the first column (to avoid trivial solution); and (ii) $v_N(x, y) = 0$ if $x = y$, else $+\infty$. As we seek to minimize the total cost of the fleet (i.e., expected value), we can exploit the compositionality of the OC problem: We find all shortest paths, and solve an OT problem (here, a linear program); see Fig. 1.

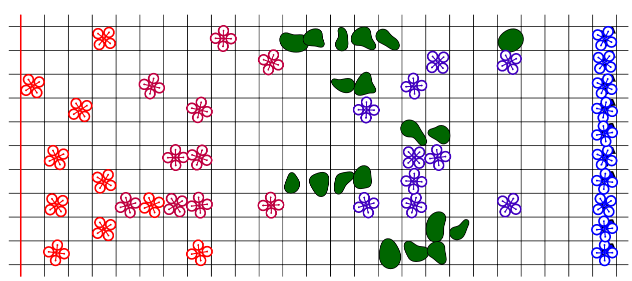

Fig. 1. Snapshots of a simulation for an instance of the Forest ride problem. The time evolution is color-coded from red (start) to blue (end).

## V. Conclusions and Future Work

We studied OC in probability spaces. Notably, in many cases of practical interest, optimal solutions are *compositional*, and result from combining optimal solutions for each particle. In future work, we would like to study the compositionality properties of OC through the formalism of category theory.

## References


[1] Giacomo Albi, Young-Pil Choi, Massimo Fornasier, and Dante Kalise. Mean field control hierarchy. *Applied Mathematics & Optimization*, 76(1):93–135, 2017.

[2] Giacomo Albi, Lorenzo Pareschi, and Mattia Zanella. On the optimal control of opinion dynamics on evolving networks. In *IFIP Conference on System Modeling and Optimization*, pages 58–67. Springer, 2015.

[3] L. Ambrosio, N. Gigli, and G. Savaré. *Gradient flows: in metric spaces and in the space of probability measures*. Birkhäuser Basel, 1 edition, 2008.

[4] Efstathios Bakolas. Optimal covariance control for discrete-time stochastic linear systems subject to constraints. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, pages 1153–1158, 2016.

[5] Efstathios Bakolas. Covariance control for discrete-time stochastic linear systems with incomplete state information. In *Proceedings of the American Control Conference*, 2017.

[6] Isin M Balci and Efstathios Bakolas. Covariance control of discrete-time gaussian linear systems using affine disturbance feedback control policies. *arXiv preprint arXiv:2103.14428*, 2021.

[7] Giuseppe Belgioioso, Dominic Liao-McPherson, Mathias Hudoba de Badyn, Saverio Bolognani, John Lygeros, and Florian Dörfler. Sampled-data online feedback equilibrium seeking: Stability and tracking. In *Proc. 60th Conference on Decision and Control*, 2021.

[8] Dimitri P. Bertsekas. *Abstract Dynamic Programming*. Number June. Athena Scientific, 2 edition, 2022.

[9] Benoît Bonnet. *Optimal control in Wasserstein spaces*. PhD thesis, Ecole Doctorale Mathématiques et Informatique de Marseille, 2019.

[10] Francesco Bullo, Jorge Cortés, and Sonia Martínez. *Distributed control of robotic networks: A mathematical approach to motion coordination algorithms*. Kindle Direct Publishing, 2022.

[11] Giulia Cavagnari and Antonio Marigonda. Attainability property for a probabilistic target in Wasserstein spaces. *Discrete and Continuous Dynamical Systems - Series A*, 41(2), 2020.

[12] Yongxin Chen, Tryphon T. Georgiou, and Michele Pavon. Optimal transport in systems and control. *Annual Review of Control, Robotics, and Autonomous Systems*, 4(1), 2021.

[13] Silvia Ferrari, Greg Foderaro, Pingping Zhu, and Thomas A Wettergren. Distributed optimalcontrol of multiscale dynamical systems: A tutorial. *IEEE Control Systems Magazine*, 36(2):102–116, 2016.

[14] Willem K Klein Haneveld. On the behavior of the optimal value operator of dynamic programming. *Mathematics of Operations Research*, 5(2):308–320, 1980.

[15] Elizabeth Y Huang, Dario Paccagnan, Wenjun Mei, and Francesco Bullo. Assign and appraise: Achieving optimal performance in collaborative teams. *IEEE Transactions on Automatic Control*, 2022.

[16] Mathias Hudoba de Badyn, Erik Miehling, Dylan Janak, Behçet Açıkmeşe, Mehran Mesbahi, Tamer Başar, John Lygeros, and Roy Smith. Discrete-time linear-quadratic regulation via optimal transport. *arXiv preprint arXiv:2109.02347*, 2021.

[17] Vishaal Krishnan and Sonia Martínez. Distributed online optimization for multi-agent optimal transport. *arXiv preprint arXiv:1804.01572*, 2019.

[18] Nicolas Lanzetti, Maximilian Schiffer, Michael Ostrovsky, and Marco Pavone. On the interplay between self-driving cars and public transportation. *arXiv preprint arXiv:2109.01627*, 2021.

[19] Suzanne Lenhart and John T. Workman. *Optimal Control Applied to Biological Models*. Chapman and Hall/CRC, 2007.

[20] Filippo Santambrogio. *Optimal Transport for Applied Mathematicians*. Number May. Springer, 2015.

[21] Antonio Terpin, Sylvain Fricker, Michel Perez, Mathias Hudoba de Badyn, and Florian Dörfler. Distributed feedback optimisation for robotic coordination. *arXiv preprint arXiv: 2109.14486*, 2021.

[22] Peter J Thomas, Mette Olufsen, Rodolphe Sepulchre, Pablo A Iglesias, Auke Ijspeert, and Manoj Srinivasan. Control theory in biology and medicine. *Biological Cybernetics*, 113(1):1–6, 2019.

[23] Cédric Villani. *Optimal Transport: Old and New*. Springer, Berlin, Heidelberg, 1 edition, 2007.

[24] Gioele Zardini, Nicolas Lanzetti, Andrea Censi, Emilio Frazzoli, and Marco Pavone. Co-design to enable user-friendly tools to assess the impact of future mobility solutions. *arXiv preprint arXiv:2008.08975*, 2020.

[25] Gioele Zardini, Nicolas Lanzetti, Marco Pavone, and Emilio Frazzoli. Analysis and Control of Autonomous Mobility-on-Demand Systems. *Annual Review of Control, Robotics, and Autonomous Systems*, 5(1), 2021.