

A Control-Oriented Model for Mobility on Demand Systems

Original

A Control-Oriented Model for Mobility on Demand Systems / Calafiore, Giuseppe; Bongiorno, Christian; Rizzo, Alessandro. - ELETTRONICO. - (2018). (Intervento presentato al convegno IEEE Conference on Decision and Control tenutosi a Miami, FL, USA nel Dec. 17-19, 2018).

Availability:

This version is available at: 11583/2715554 since: 2018-10-23T10:56:49Z

Publisher:

IEEE

Published

DOI:

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2018 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

A Control-Oriented Model for Mobility on Demand Systems

Giuseppe Carlo Calafiore, Christian Bongiorno, Alessandro Rizzo

Abstract—In this paper, we propose a control-oriented model for mobility-on-demand systems (MOD). The system is first described through dynamical stochastic state-space equations in discrete time, and then suitably simplified in order to obtain a control-oriented model, on which a control strategy based on Model Predictive Control (MPC) is devised. The control strategy aims at maintaining the average number of vehicles at stations within prescribed bounds. Relevant features of the proposed model are: *i)* the possibility of considering stochasticity and heterogeneity in the system parameters; *ii)* a state space structure, which makes the model suitable for implementation of effective parameter identification and control strategies; and *iii)* the possibility of weighting the control effort, leading to control solutions that may trade off efficiency and cost. Simulation results on a synthetic network corroborate the validity of our approach under several operational conditions.

I. INTRODUCTION

Mobility on Demand (MOD) systems are becoming pervasive in cities of any size. As of December 2016, bike- and car-sharing programs had been adopted by more than 1,000 cities worldwide [1]. The concept behind a MOD system is straightforward: a user requires a vehicle, picks it up from a designated location, executes the trip, and finally drops off the vehicle at her/his destination. MOD systems can be station-based, with vehicles parked at fixed locations (stations), or free floating, with vehicles parked with no constraints, at the user's wish.

As required to every service provider, a MOD system should be designed to meet the customer demand. In fact, such a demand is extremely heterogeneous, due to several factors, such as the time of the day, the season, commuting patterns, up-hill or down-hill stations (for bikes), and so on [2], [3]. The impossibility of meeting customer demand is usually caused by a lack of vehicles at some locations and in a corresponding surplus of vehicles somewhere else. This issue can be mitigated through the implementation of repositioning policies, also called *rebalancing*.

Rebalancing strategies are typically obtained as the solution of optimization problems and are typically executed during time periods where traffic is low, especially at night. This activity, called static rebalancing, assumes that vehicles are not used by customers during repositioning operations. Rebalancing in bike-sharing systems is usually executed by trucks able to displace high volumes of bikes, even within relatively long distances [4], [5], [6], [7], [8], [9],

[10]. Dynamic repositioning, on the other hand, assumes that customers are traveling while rebalancing operations occur, and the effects of such travels are not negligible. This repositioning is usually performed with smaller vehicles and/or over shorter distances [11], [12]. Users may also be involved in system rebalancing through incentives [13], [14].

Optimization algorithms generating repositioning strategies are based on suitable models of the MOD system. Several modeling techniques have been proposed in the literature, mostly based on statistical and data-driven approaches, to account for the stochasticity of the system under exam [15], [16], [17], [18], [19].

In this paper, we propose a novel control-oriented model for a station-based MOD system. In its general formulation, it is a dynamical model with stochastic state variables in discrete time. The model accounts for heterogeneity and stochasticity in the MOD system. It quantifies the vehicle flows from stations, and accounts for stochasticity in customer demand and traveling times. A linear and steady-state version of the model is then derived, and an MPC-like technique [20] is applied to control the expected values of the state variables, representing the expected quantity of vehicles at each station. The control problem yields a constrained optimization problem, with the objective of maintaining the number of vehicles within prescribed bounds at each station. Our approach is assessed through simulations on a synthetic MOD system.

The paper is structured as follows. In Section II, we present our model and the related assumptions, which lead to a simplified control-oriented model. In Section III, an MPC-based control algorithm to obtain rebalancing vehicle flows is presented. In Section IV, simulation results are illustrated and commented. Finally, Section V summarizes our conclusions and future work.

II. A CONTROL MODEL OF THE MOD SYSTEM

We model a MOD system as a network composed by nodes, representing the vehicle stations, and links between nodes, representing the vehicles' routes. The set \mathcal{S} of station nodes is composed of N vehicle stations, and the set \mathcal{L} is composed of N^2 links between any two station nodes in \mathcal{S} . A link (i, j) between departure station i and destination station j does not necessarily represent a specific physical route, rather the ensemble of all routes that are typically travelled by customers moving from i to j .

We characterize link (i, j) by its state $v_{ij}(t)$, representing the cardinality of the set $\mathcal{V}_{ij}(t)$ of all the vehicles *en route* from i to j at time t , and by the (random) fraction $\tilde{q}_{ij}(t, \delta) \in$

The authors are with the Department of Electronics and Telecommunications, Politecnico di Torino, Torino, Italy. G. Calafiore is also affiliated with IEIT-CNR, Italy. A. Rizzo is also affiliated with the Office of Innovation, New York University Tandon School of Engineering, Brooklyn NY, USA. [christian.bongiorno, giuseppe.calafiore, alessandro.rizzo]@polito.it.

$[0, 1]$ of the $v_{ij}(t)$ vehicles that reach their destination j within the time interval $(t, t + \delta]$.

We characterize node i by its state $z_i(t)$, denoting the number of vehicles parked (and hence available for pick up) at station i at time t , and by the instantaneous mean rate $\mu_i(t) \in \mathbb{R}^+$ of random service requests that arrive at station i at time t . Analogously, $1/\mu_i(t)$ describes the mean of the random inter-arrival time of service requests at station i . The station's throughput $\lambda_i(t)$ is instead the mean rate at which vehicles depart from station i . In reality, it always holds that $\lambda_i(t) \leq \mu_i(t)$, since not all the service requests may be fulfilled, due to the fact that there may exist station-empty periods in which no vehicles are available at the station, and hence no departure is possible from the station, even if demand from customers exists. This issue is further discussed in Section II-A.4. Node state $z_i(t)$ is bounded as $0 \leq z_i(t) \leq c_i$, where c_i is the maximum number of vehicles that can be simultaneously parked at the station, i.e., the station capacity. Both extremes of the bounding interval are not desirable for efficient operation of the station, since vehicle cannot be picked up ($z_i(t) = 0$) or dropped off ($z_i(t) = c_i$). The goal of rebalancing strategies is indeed to keep the number of vehicles at the station within the prescribed limits.

If a customer request is generated at station i at time t , and if station i has a vehicle available, then the customer takes that vehicle and starts a trip towards a destination station j . We model the selection of the destination via a set of (possibly time-varying) *routing probabilities*, i.e., we assume that, at time t and for each station i , there exist probabilities $p_{ij}(t) \in \mathbb{R}^+$, with $\sum_j p_{ij}(t) = 1$, such that a generic customer departing from i chooses destination j with probability $p_{ij}(t)$.




Fig. 1. Example of network modeling a MOD system with three stations S_1, S_2, S_3 , and corresponding itinerary links I_{ij} .

A. Simplifying assumptions for a “control” model

We next develop a *control-oriented model* for the MOD system. Contrary to a *simulation model*, whose aim is to be an extremely accurate proxy of reality, our model is devised with the objective of allowing efficient design of suitable control strategies (e.g., rebalancing policies) for the system. A good control model should be simple enough to allow for effective synthesis of the control law, although this may come at the expense of some approximation. In the end,

however, the performance of the control law should be tested and evaluated on the real system, or on a proxy of it (i.e., on a simulation model). With this in mind, we next present the main simplifying assumptions used to construct our control-oriented model of the MOD system.

1) *Piece-wise constant parameters*: The model discussed so far has naturally time-varying parameters. Previous analyses performed on logged service data suggest that system parameters can be conveniently approximated as piece-wise constant functions [15], [16]. Motivated by this observation, here we focus on the case of constant parameters, without losing generality. Stochasticity, on the other hand, is kept on the fraction of vehicles traveling from one station to the other, as detailed afterwards.

2) *Exponential inter-departure times*: Departures from each station i follow a counting process with instantaneous rate $\lambda_i(t)$. For simplicity, we shall assume specifically that they form a Poisson process of rate $\lambda_i(t)$, although this latter assumption is not critical for our developments. Since we assume that each vehicle departing from i at t chooses destination j with probability $p_{ij}(t)$, we have that the vehicles departing from i with destination j at time t also follow a Poisson process with rate $p_{ij}(t)\lambda_i(t)$.

3) *Densities for link arrival proportions*: As previously discussed, we model the transit of vehicles through the (i, j) link by assuming that, at each given t and given δ , only a (random) fraction $\tilde{q}_{ij}(t, \delta) \in [0, 1]$ of the $v_{ij}(t)$ vehicles reach their destination j within the time interval $(t, t + \delta]$. We let $q_{ij}(t, \delta) \doteq \mathbb{E}\{\tilde{q}_{ij}(t, \delta)\} \in [0, 1]$, and we assume that $\tilde{q}_{ij}(t, \delta)$ is statistically independent from $v_{ij}(t)$. We have verified that this assumption is plausible as long as the traffic rates are high enough, which is the situation in which rebalancing policies are mostly required. The characterization of $\tilde{q}_{ij}(t, \delta)$ should be done via a specific statistical analysis of log data. Here, we simplify the model assuming that $\tilde{q}_{ij}(t, \delta) = q_{ij}(\delta)$ is not stochastic and is independent from time.

4) *No blocking*: We assume, for the sole purpose of the control-oriented model, that stations have unlimited capacity, i.e., $c_i = +\infty, \forall i$. Also, we assume that demand is always satisfied. In reality, a station's state $z_i(t)$ remains bounded in $[0, c_i]$ at all times. In our control-oriented model we allow the state variable $z_i(t)$ to go beyond the boundaries, but we penalize out-of-boundary behavior in the control design phase.

B. The MOD control model

We let $\delta = \Delta/n_s$, where n_s is a positive integer, and Δ is the time period defined in Section II-A.1, during which the system parameters are assumed to be constant. Under the assumptions of Section II-A, we define the following quantities:

- $d_{ij}(t + \delta)$ is the number of vehicles driven by users that depart from i with destination j in the time interval $(t, t + \delta]$. According to the assumption in Section II-

A.2, $d_{ij}(t + \delta)$ has a Poisson probability mass

$$\text{Prob}\{d_{ij}(t + \delta) = k\} = \frac{1}{k!} (p_{ij}(t) \lambda_i(t) \delta)^k e^{-p_{ij}(t) \lambda_i(t) \delta},$$

for $k = 0, 1, \dots$ (1)

- $r_{ij}(t)$ is the deterministic number of “control” vehicles (i.e., vehicles used for rebalancing purposes) that are moved from i to j as dictated by the rebalancing control strategy. We assume that δ is sufficiently small with respect to the average link transit times τ_{ij} , so that there is a (practically) zero probability that any of the $d_{ij}(t + \delta)$ or of the $r_{ij}(t)$ vehicles reaches its destination by time $t + \delta$.
- $a_{ij}(t + \delta)$ is the number of vehicles, among the ones in $\mathcal{V}_{ij}(t)$, that reach the j -th station by time $t + \delta$. According to the assumption in Section II-A.3, the count $a_{ij}(t + \delta)$ can be written equivalently in terms of the *random proportion* $\tilde{q}_{ij}(t, \delta)$ as

$$a_{ij}(t + \delta) = \tilde{q}_{ij}(t, \delta) v_{ij}(t). \quad (2)$$

We observe that we are implicitly allowing $a_{ij}(t + \delta)$ to be real valued.

For $i, j = 1, \dots, N$, straightforward conservation arguments and equation (2) yield the discrete time model that describe the system behavior

$$v_{ij}(t + \delta) = (1 - \tilde{q}_{ij}(t, \delta)) v_{ij}(t) + d_{ij}(t + \delta) + r_{ij}(t) \quad (3)$$

$$z_j(t + \delta) = z_j(t) + \sum_i \tilde{q}_{ij}(t, \delta) v_{ij}(t) - \sum_h (d_{jh}(t + \delta) + r_{jh}(t)). \quad (4)$$

The system above is a linear discrete time stochastic one in the z_j and v_{ij} state variables, with stochastic inputs given by the d_{ij} departures, and control inputs given by the rebalancing departures r_{ij} . Given initial conditions, Eqs. (3)-(4) can be used to predict forward behavior of the system's state. In the following, we will derive the dynamics of the expected value of the model state variables, which will be used to design the MPC-based controller.

C. The expected state dynamics

Observing that $\sum_h d_{jh}(t + \delta)$ is Poisson with parameter $\sum_h p_{jh} \lambda_j \delta = \lambda_j \delta$, denoting with an over bar expected quantities (i.e., $\bar{v}_{ij}(t) \doteq \mathbb{E}\{v_{ij}(t)\}$, $\bar{z}_j(t) \doteq \mathbb{E}\{z_j(t)\}$, etc.), and recalling that $\tilde{q}_{ij}(t, \delta)$ and $v_{ij}(t)$ are assumed to be independent, we can write the evolution of the expected value of the state equations in (3)-(4) as

$$\bar{v}_{ij}(t + \delta) = (1 - q_{ij}(t, \delta)) \bar{v}_{ij}(t) + p_{ij}(t) \lambda_i(t) \delta + r_{ij}(t) \quad (5)$$

$$\bar{z}_j(t + \delta) = \bar{z}_j(t) + \sum_i q_{ij}(t, \delta) \bar{v}_{ij}(t) - \lambda_j(t) \delta - \sum_h r_{jh}(t). \quad (6)$$

Equations (5)-(6) constitute a linear discrete-time deterministic dynamical system in the expected state variables $\bar{z}_j(t)$

and $\bar{v}_{ij}(t)$, with inputs given by the mean departure rates $\lambda_i(t)$, and control inputs given by the rebalancing departures $r_{ij}(t)$. With both constant parameters and rebalancing inputs, i.e., $r_{ij}(t) = r_{ij}$, $\forall t$, since $1 - q_{ij}(\delta) < 1$, it can be proved that system (5)-(6) admits a steady-state behavior, denoted with $\bar{v}_{ij}^{(ss)}$ and $\bar{z}_j^{(ss)}$.

III. CONTROL OF THE EXPECTED DYNAMICS

We next discuss a Model Predictive Control (MPC) approach for the control of the expected system's dynamics described by Eqs. (5)-(6). We fix a time horizon $T = n_h \delta$, where n_h is a positive integer. We denote with $R(t + k\delta)$ the $N \times N$ matrix whose (i, j) -th element is $r_{ij}(t + k\delta)$, and let

$$\mathcal{R} \doteq \{R \in \mathbb{R}^{N,N} : R \geq 0, \text{ and } R_{ii} = 0, i = 1, \dots, N\}.$$

The cost function we consider is the total rebalancing effort over the considered time horizon $J_T = \sum_{k=0}^{n_h-1} \|R(t + k\delta)\|_1$, where we define $\|R\|_1 \doteq \sum_j \sum_i |R_{ij}|$. The control goal is to maintain the expected states $\bar{z}_j(t + k\delta)$ within given limits $[1, c_j]$, at all times, while minimizing the rebalancing effort. Formally, we solve

$$\begin{aligned} \min_{R(t), \dots, R(t + (n_h-1)\delta) \in \mathcal{R}} \quad & \sum_{k=0}^{n_h-1} \|R(t + k\delta)\|_1 \\ \text{s.t.:} \quad & \bar{z}_j(t + k\delta) \in [1, c_j], \quad \text{for} \\ & j = 1, \dots, N, \text{ and } k = 1, \dots, n_h, \end{aligned}$$

where $\bar{z}_j(t + k\delta)$ is given by the recursion in (5)-(6), initialized with given initial conditions $\bar{z}_j(t)$, $\bar{v}_{ij}(t)$, $i, j = 1, \dots, N$.

Imposing strict feasibility for the state limits $\bar{z}_j(t + k\delta) \in [1, c_j]$ may result in infeasibility, or in high rebalancing effort. A more flexible approach is therefore to consider a tradeoff between rebalancing effort and constraint satisfaction, by introducing slack variables $s_j(t + k\delta)$. We first rewrite the state constraint as $\left| \bar{z}_j(t + k\delta) - \frac{c_j + 1}{2} \right| \leq \frac{c_j - 1}{2}$, and then relax the problem to

$$\begin{aligned} \min \quad & \sum_{k=0}^{n_h-1} \sum_j s_j(t + k\delta) + \gamma \|R(t + k\delta)\|_1 \\ \text{s.t.:} \quad & \left| \bar{z}_j(t + k\delta) - \frac{c_j + 1}{2} \right| \leq \frac{c_j - 1}{2} + s_j(t + k\delta), \quad (7) \\ & \text{for } j = 1, \dots, N, \text{ and } k = 1, \dots, n_h, \\ & R(t), \dots, R(t + (n_h - 1)\delta) \in \mathcal{R}, \\ & s_j(t + \delta) \geq 0, \dots, s_j(t + n_h \delta) \geq 0, \end{aligned}$$

where $\gamma \geq 0$ is a tunable tradeoff parameter.

IV. RESULTS

We now validate the proposed approach with numerical results obtained on three synthetic transportation networks comprising $N = 5$ stations. Due to the synthetic nature of the case study, we consider an adimensional time span $[t_s, t_f] = [0, 1]$. An adimensional time-step is set to $\delta = 0.02$, such that the discrete time index k spans the time interval as $t = k\delta$, with $k = 0, 1, \dots, n_h$. Consequently, in our simulations it results $n_h = 50$. We perform the MPC on the expected dynamics by solving problem (7) along a time-horizon equal

to the whole simulation time, i.e., considering $t \in [0, 1]$. At each time-step t , problem (7) is solved and the control commands $\{R(t), R(t + \delta), \dots, R(t_f - \delta)\}$ are derived for every time step within the considered time-horizon. The control commands are functions of the expected values of the stochastic parameters $\{d_{ij}(t), d_{ij}(t + \delta), \dots, d_{ij}(t_f)\}$. Then, only the control command corresponding to the current time-step $R(t)$ is applied to the dynamics (Eqs. (3)-(4)), after a rounding of the obtained values. We observe that the time-horizon considered to solve problem (7) shrinks at each time step, since we do not predict the system dynamics after $t = t_f$. The number of vehicles arriving from i to j are selected via a stochastic rounding [21] of $q_{ij}(\delta) v_{ij}(t - \delta)$. All the presented simulations are realized using a Monte Carlo method, averaging the obtained results over 1,000 independent trials.

In the first set of simulations we consider the system in a steady-state condition, with time-invariant parameters. Specifically, we set the routing matrix to

$$p(t) = \begin{bmatrix} 0.0 & 0.25 & 0.25 & 0.25 & 0.25 \\ 0.7 & 0.00 & 0.10 & 0.10 & 0.10 \\ 0.7 & 0.10 & 0.00 & 0.10 & 0.10 \\ 0.7 & 0.10 & 0.10 & 0.00 & 0.10 \\ 0.7 & 0.10 & 0.10 & 0.10 & 0.00 \end{bmatrix}.$$

This choice mimics a system comprising a station that is mostly “attractive” to the others, while traffic from the other stations is uniformly distributed. Then, we set $\lambda(t) = \{102.9, 36.8, 36.8, 36.8, 36.8\}$, modeling the highest rate of customer requests at the most attractive station and a uniform rate of requests at the others, and such that $\sum_j \lambda_j(t) p_{ij}(t) = \lambda_i(t)$ holds. For simplicity, we set $q_{ij}(\delta) = 0.75$ for every $i \neq j$ and $q_{ii}(\delta) = 1$, independently from time. The initial value of $v_{ij}(t_s)$ is set to the corresponding steady-state value $\bar{v}_{ij}^{(ss)}$, with $r_{ij} = 0$. For all the sets of simulations, we initially assign to each station $z_j(t_s) = 7, \forall j$, vehicles, and we fix the lower and upper capacity limits to 1 and $c_j = 13, \forall j$, respectively.

To assess the performance of the control procedure, we consider the number of vehicles that exceed the capacity limits per time-step over the whole station set, and indicate it with f_E . We perform a stress test by adding an additional expected rate of departures to station 5, indicated with λ_5^+ , which takes the system out of the steady-state condition. In Fig. 2(a) and (b) we plot f_E as a function of λ_5^+ obtained through a series of Monte Carlo simulations without and with control, respectively. We observe that, even when the system is at steady-state ($\lambda_5^+ = 0$), the control procedure provides a significant improvement of the system performance. Due to the stochastic components, in fact, the uncontrolled system, nominally at steady-state, exhibits an average of $f_E = 9.3$ excess vehicles. This value reduces to $f_E = 0.6$ for the controlled system. When λ_5^+ increases, the advantage of applying the control procedure becomes apparent. Figs. 2(c) and (d) illustrate the trend over time of the number of vehicles in each station $\lambda_5^+ = 100$, for the uncontrolled and the controlled system, respectively. While

in the uncontrolled system station 5 tends to become empty, whereas the remaining station tends to fill up, the application of the control action keeps the number of vehicles within the prescribed limits. Although the control is designed on the expected value of the state variables, it consistently helps in mitigating fluctuations about such expected values. For the value $\lambda_5^+ = 100$ used to realize Figs. 2(c) and (d), in fact, we observe from Figs. 2(a) and (b) that the control action substantially decreases the performance parameter f_E from 77.6 to 1.5.

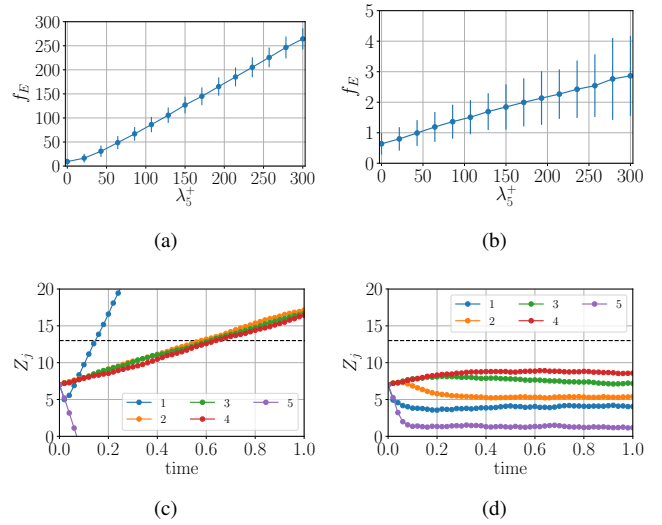


Fig. 2. Average number of excess bikes over the whole station set f_E as a function of λ_5^+ , in the uncontrolled (a) and controlled (b) system for $\gamma = 0.01$. Whiskers correspond to the standard deviation. Time evolution of $z_j(t)$ for $\lambda_5^+ = 100$ in the uncontrolled (c) and controlled (d) system. Dashed black lines indicate capacity limits, colored dotted lines refer to different stations.

The second set of simulations reflects the case of an asymmetry in the journey times. This can happen, for example, when there are stations located on a hill and, therefore, the journey time from these stations and those located downhill is asymmetric. To emulate the situation where station 1 is located uphill and the others downhill, we fix matrix q as

$$q(\delta) = \begin{bmatrix} 1.00 & 0.75 & 0.75 & 0.75 & 0.75 \\ q^* & 1.00 & 0.75 & 0.75 & 0.75 \\ q^* & 0.75 & 1.00 & 0.75 & 0.75 \\ q^* & 0.75 & 0.75 & 1.00 & 0.75 \\ q^* & 0.75 & 0.75 & 0.75 & 1.00 \end{bmatrix},$$

where q^* is a constant parameter that is tuned to emulate the effect under observation. Notably, strong asymmetries in the journey time are modeled using small values of q^* .

We also fix the initial $v_{ij}(0)$ as in the first simulation set, so that when $q^* \neq 0.75$, the system is no more in equilibrium. Figure 3(a) illustrates the time evolution of variables z_i for $q^* = 0.1$. As expected, station 1 suffers from a lack of vehicles, due to the fact that the journey time toward it is increased. The application of the control strategy mitigates this phenomenon, as illustrated in Fig. 3(b), showing that the control action drives the system toward an equilibrium condition, where the number of vehicles is within the prescribed




Fig. 3. Time evolution of $z_j(t)$ for $q^* = 0.1$, for the uncontrolled (a) and the controlled system (b). The control tradeoff parameter is set to $\gamma = 0.01$. Number of rebalancing vehicles in the controlled system, in equilibrium and with symmetric journey times, that is, $q^* = 0.75$ (c), and with asymmetric journey times, that is, $q^* = 0.1$. Performance parameter f_E (e) and rebalancing effort per time-step J_T/n_h (f) as a function of q^* in the controlled system.

limits. Figures 3(c) and (d) highlight the outcome of the control activity at the link level, that is, showing the number of rebalancing vehicles required to travel per time-step from a station to another. In particular, Fig. 3(c) illustrates the rebalancing activity computed for $q^* = 0.75$, that is, in the equilibrium condition and with symmetrical journey times. In this case, we observe a modest rebalancing activity, which is only generated to compensate the stochasticity around the system equilibrium. On the other hand, Fig. 3(d) illustrates the rebalancing activity when $q^* = 0.1$ is selected. As expected, the rebalancing algorithm dictates that a larger number of bikes should be displaced from all the stations toward station 1, while the rebalancing activity among all the remaining stations remains moderate, although greater than in the equilibrium condition. Figures 3(e) and (f) illustrate the average system performance f_E and the average rebalancing effort over time J_T/n_h as a function of the asymmetry parameter q^* , in the controlled system. The numerical results confirm the intuition that a stronger asymmetry yields a worse system performance and a higher rebalancing effort.

The third set of numerical results models time-varying flows of users that commute between the periphery and the center of a hypothetical city during daytime, with a heterogeneous use of stations. To mimic this scenario, we fix at first a set of initial heterogeneous departure rates

in stations as $\lambda(t_s) = \{100, 56.2, 31.2, 17.8, 10\}$, and the corresponding final expected departure rates as $\lambda_i(t_f) = \lambda_{N-i}(t_s) \forall i$. Then, we generate the expected departure rate profiles using $\lambda_i(t) = t \lambda_i(t_f) + (t_f - t) \lambda_i(t_s)$. Similarly, we fix a heterogeneous set of initial routing probabilities as

$$p(t_s) = \begin{bmatrix} 0.46 & 0.26 & 0.15 & 0.08 & 0.05 \\ 0.46 & 0.26 & 0.15 & 0.08 & 0.05 \\ 0.46 & 0.26 & 0.15 & 0.08 & 0.05 \\ 0.46 & 0.26 & 0.15 & 0.08 & 0.05 \\ 0.46 & 0.26 & 0.15 & 0.08 & 0.05 \end{bmatrix}$$

and the corresponding final values as $p_{ij}(t_f) = p_{i,N-j}(t_s)$. The routing probability profiles in time are obtained as $p_{ij}(t) = t p_{ij}(t_f) + (t_f - t) p_{ij}(t_s)$. We fix the initial velocity $v_{ij}(0) = 0 \forall i, j$. Finally, we choose $q_{ij}(\delta) = 0.75$ for every $i \neq j$ and $q_{ii}(\delta) = 1$, independently from time.

Figure 4(a) illustrates the evolution over time of the state variables $z_j(t)$ of the uncontrolled system. The system dynamics is characterized by two pairs of stations with an opposite behavior. Specifically, the violet and the red stations exceed the upper capacity limit during the central part of the simulation interval, whereas the orange and the blue ones exceed the lower capacity limit in the same time-window. In particular, the violet and the blue stations have the largest deviations from the initial condition. The green station works around an equilibrium condition, since it is characterized by a similar rate of departures and arrivals per time-step.




Fig. 4. Time evolution of variables $z_j(t)$ for (a) the uncontrolled system; (b) the controlled system with $\gamma = 1$; and (c) the controlled system with $\gamma = 10^{-6}$. In (d), the total number of rebalancing vehicles in the controlled system, for $\gamma = 1$ and $\gamma = 10^{-6}$.

Figure 4(b) illustrates the evolution of $z_j(t)$ over time, with a tradeoff parameter $\gamma = 1$. Such value implies a high penalization of the rebalancing action, yielding a weak control effort. This is confirmed in Fig. 4(d), which illustrates the control effort, that is, the number of balancing vehicles over time. Vehicles from the violet and red stations are directed towards the blue and orange ones by the control algorithm, whereas the green station is almost untouched.

However, the control effort is not sufficient to satisfy the lower capacity limit for the blue and the violet stations. Decreasing the tradeoff parameter to $\gamma = 10^{-6}$, we observe from Fig. 4(c) that all the capacity limits are satisfied in every station, at the expenses of a higher control effort (see Fig. 4(d)). We observe that, in this case, the control action is enforced by moving vehicles from the green station, which operates around its equilibrium, toward the blue and orange stations. We finally observe from Fig. 4(d) that the control action has a similar trend, but different absolute values, in accordance with the selection of the tradeoff parameters. In particular, our results confirm the intuition that the control action is stronger when the variability in the demand is higher.

V. CONCLUSIONS

In this work, we have defined a novel control-oriented model for MOD systems, which accounts for the inherent heterogeneity and stochasticity in the system parameters. The model is a dynamical stochastic one, and evolves in discrete time. A control strategy to achieve system rebalancing has been devised using model predictive control on the expected values of the state variables. We have validated our approach over three sets of numerical simulations on a synthetic system comprising five stations. The proposed numerical experiments aim at assessing the system performance in some fundamental aspects of a realistic urban mobility system. Notably, the heterogeneous use of the stations and the intra-day time-varying flow of users, who commute from the periphery to the city center throughout the day, have been modeled.

Our analysis shows that a control effort concentrated at the beginning and at the end of the day can be sufficient to significantly mitigate the unbalance of the stations. Furthermore, an important feature of our model is that the control effort can be tuned by a tradeoff parameter between performance and cost.

Future work will include the application of our method to a real system, using real logged service data from large bike- or car- sharing networks, and to assess the results by using different objective functions, which would possibly have different repercussions in the impact of variables in the overall performance. We will also account for heterogeneity in the travel time, and, most importantly, we will devise a control model that explicitly takes into account stochastic fluctuations of the system variables, rather than limiting the focus to their expected values.

ACKNOWLEDGEMENTS

This research is supported by Compagnia di San Paolo, Italy.

REFERENCES

- [1] R. Meddin and P. DeMaio. The bike-sharing world map. <http://www.bikesharingworld.com/>, 2018.
- [2] P. Borgnat, P. Abry, P. Flandrin, C. Robardet, J.-B. Rouquier, and E. Fleury. Shared bicycles in a city: A signal processing and data analysis perspective. *Advances in Complex Systems*, 14(03):415–438, 2011.
- [3] R. Nair, E. Miller-Hooks, R. C. Hampshire, and A. Bušić. Large-scale vehicle sharing systems: analysis of Vélib’. *International Journal of Sustainable Transportation*, 7(1):85–106, 2013.
- [4] Z. Yang, J. Hu, Y. Shu, P. Cheng, J. Chen, and T. Moscibroda. Mobility modeling and prediction in bike-sharing systems. In *Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services*, pages 165–178. ACM, 2016.
- [5] D. Chemla, F. Meunier, and R. W. Calvo. Bike sharing systems: Solving the static rebalancing problem. *Discrete Optimization*, 10(2):120–146, 2013.
- [6] M. Dell’Amico, E. Hadjicostantinou, M. Iori, and S. Novellani. The bike sharing rebalancing problem: Mathematical formulations and benchmark instances. *Omega*, 45:7–19, 2014.
- [7] M. Dell’Amico, M. Iori, S. Novellani, T. Stützel, et al. A destroy and repair algorithm for the bike sharing rebalancing problem. *Computers & Operations Research*, 71:149–162, 2016.
- [8] G. Erdoğan, M. Battarra, and R. W. Calvo. An exact algorithm for the static rebalancing problem arising in bicycle sharing systems. *European Journal of Operational Research*, 245(3):667–679, 2015.
- [9] A. Faghhi-Imani, R. Hampshire, L. Marla, and N. Eluru. An empirical analysis of bike sharing usage and rebalancing: Evidence from Barcelona and Seville. *Transportation Research Part A: Policy and Practice*, 97:177–191, 2017.
- [10] T. Raviv, M. Tzur, and I. A. Forma. Static repositioning in a bike-sharing system: models and solution approaches. *EURO Journal on Transportation and Logistics*, 2(3):187–229, 2013.
- [11] J. Schuijbroek, R. C. Hampshire, and W.-J. Van Hoeve. Inventory rebalancing and vehicle routing in bike sharing systems. *European Journal of Operational Research*, 257(3):992–1004, 2017.
- [12] J. Pfrommer, J. Warrington, G. Schildbach, and M. Morari. Dynamic vehicle redistribution and online price incentives in shared mobility systems. *IEEE Transactions on Intelligent Transportation Systems*, 15(4):1567–1578, 2014.
- [13] A. Singla, M. Santoni, G. Bartók, P. Mukerji, M. Meenen, and A. Krause. Incentivizing users for balancing bike sharing systems. In *AAAI*, pages 723–729, 2015.
- [14] C. Fricker and N. Gast. Incentives and redistribution in homogeneous bike-sharing systems with stations of finite capacity. *Euro journal on transportation and logistics*, 5(3):261–291, 2016.
- [15] G. C. Calafiore, F. Portigliotti, and A. Rizzo. A network model for an urban bike sharing system. *IFAC-PapersOnLine*, 50(1):15633–15638, 2017.
- [16] G. C. Calafiore, F. Portigliotti, A. Rizzo, and C. Novara. A flow optimization approach for the rebalancing of mobility on demand systems. In *IEEE Conference on Decision and Control*, pages 1–6, 2017.
- [17] D. Zeng, J. Xu, J. Gu, L. Liu, and G. Xu. Short term traffic flow prediction based on online learning SVR. In *Power Electronics and Intelligent Transportation System, 2008. PEITS’08. Workshop on*, pages 616–620. IEEE, 2008.
- [18] L. Zhao and F.-Y. Wang. Short-term traffic flow prediction based on ratio-median lengths of intervals two-factors high-order fuzzy time series. In *Vehicular Electronics and Safety, 2007. ICVES. IEEE International Conference on*, pages 1–7. IEEE, 2007.
- [19] I. Kaysi, M. E. Ben-Akiva, and H. Koutsopoulos. *An integrated approach to vehicle routing and congestion prediction for real-time driver guidance*, volume 1408. Transportation Research Board, 1993.
- [20] E. F. Camacho and C. B. Alba. *Model predictive control*. Springer Science & Business Media, 2013.
- [21] S. Gupta, A. Agrawal, K. Gopalakrishnan, and P. Narayanan. Deep learning with limited numerical precision. In *International Conference on Machine Learning*, pages 1737–1746, 2015.