

Rebalancing the Rebalancers: Optimally Routing Vehicles and Drivers in Mobility-on-Demand Systems

Stephen L. Smith

Marco Pavone

Mac Schwager

Emilio Frazzoli

Daniela Rus

Abstract—In this paper we study rebalancing strategies for a mobility-on-demand urban transportation system blending customer-driven vehicles with a taxi service. In our system, a customer arrives at one of many designated stations and is transported to any other designated station, either by driving themselves, or by being driven by an employed driver. When some origins and destinations are more popular than others, vehicles will become unbalanced, accumulating at some stations and becoming depleted at others. This problem is addressed by employing rebalancing drivers to drive vehicles from the popular destinations to the unpopular destinations. However, with this approach the rebalancing drivers themselves become unbalanced, and we need to “rebalance the rebalancers” by letting them travel back to the popular destinations with a customer. In this paper we study how to optimally route the rebalancing vehicles and drivers so that the number of waiting customers remains bounded while minimizing the number of rebalancing vehicles traveling in the network and the number of rebalancing drivers needed; surprisingly, these two objectives are aligned, and one can find the optimal rebalancing strategy by solving two decoupled linear programs. We determine the minimum number of drivers and minimum number of vehicles needed to ensure stability in the system. Our simulations suggest that, in Euclidean network topologies, one would need between 1/3 and 1/4 as many drivers as vehicles.

I. INTRODUCTION

In this paper we study vehicle routing algorithms for a novel model of urban transportation system, which involves blending customer-driven vehicles with a taxi service. Our proposed car-share system is an example of a Mobility-on-Demand (MOD) system, and aims at providing urban dwellers with the tailored service of a private automobile, while utilizing limited urban land more efficiently (e.g., by minimizing the automobiles that sit unused) [1]. In our system, a customer arrives at one of many designated stations and is transported to any other designated station, either by driving themselves, or by being driven by an

This research was supported by the Future Urban Mobility project of the SMART Center, with funding from Singapore’s National Research Foundation; by ONR grant N000140911051; and by NSF grant EFR10735953.

S. L. Smith is with the Department of Electrical and Computer Engineering, University of Waterloo, Waterloo ON, N2L 3G1 Canada (stephen.smith@uwaterloo.ca).

M. Pavone is with the Department of Aeronautics and Astronautics, Stanford University, Stanford, CA 94305, USA (pavone@stanford.edu).

M. Schwager is with the Department of Mechanical Engineering and the Division of Systems Engineering, Boston University, Boston, MA 02215, USA, (schwager@bu.edu).

E. Frazzoli is with the Laboratory for Information and Decision Systems, Aeronautics and Astronautics Department, Massachusetts Institute of Technology, Cambridge, MA 02139, USA (frazzoli@mit.edu).

D. Rus is with the Computer Science and Artificial Intelligence Laboratory, Electrical Engineering and Computer Science Department, Massachusetts Institute of Technology, Cambridge, MA 02139, USA (rus@csail.mit.edu).

employed driver. In a typical one way car-share system (e.g. Car2Go) it has been observed empirically [2], and shown analytically [3], that vehicles become unbalanced, accumulating at popular destinations and becoming depleted at less popular ones. Our proposed system addresses this problem by employing rebalancing drivers to drive vehicles from the popular destinations to the unpopular destinations. However, with this approach the rebalancing drivers themselves become unbalanced, and hence we need to “rebalance the rebalancers” by letting them travel back to the popular destinations with a customer. In such a trip, the rebalancing driver operates the vehicle as a taxi, driving the customer to their desired destination. The system is illustrated in Fig. 1. The main difficulty in such a system, and the focus of this paper, is how to determine the rebalancing trips and the taxi trips in order to minimize wasted trips, while providing the best possible customer experience.

Specifically, the contribution of this paper is twofold: we study routing algorithms for the MOD system illustrated in Fig. 1 that (1) minimize the number of rebalancing vehicles traveling in the network, (2) minimize the number of drivers needed, and (3) ensure that the number of waiting customers remains bounded. Second, leveraging our analysis, we determine the relation between the minimum number of drivers needed and the minimum number of vehicles needed to ensure stability in the system; these relations would provide a system designer with essential structural insights to develop business models. Interestingly, our simulations suggest that, in Euclidean network topologies, one would need between 1/3 and 1/4 as many drivers as vehicles, and that this fraction decreases to about 1/5 if one allows up to 3-4 drivers to take a trip with a customer.

This paper builds upon the previous work of the authors in designing optimal rebalancing policies for MOD systems leveraging *autonomous operation* of the vehicles [4], [3], i.e., without the need of human drivers. On the contrary, the system proposed in this paper would use technology that is available today (i.e., by employing human drivers instead of autonomous cars), and our findings are readily applicable to *existing* one-way car-share systems, which already employ drivers to rebalance cars using heuristic methods [2]. Furthermore, by comparing the results in this paper with those in [4], one can quantitatively assess the relative benefits of “hi-tech” autonomous MOD systems versus “low-tech” driver-based MOD systems. The problem addressed in this paper has also many characteristics in common with the well-known Dynamic Traffic Assignment (DTA) problem [5], [6], [7], [8]. The key difference between rebalancing in MOD systems

and the DTA problem is that in the former the optimization is over the empty vehicle trips (i.e., the rebalancing trips) rather than the passenger carrying trips.

The paper is structured as follows. In Section II we present a fluid model for our system, and we formally state the rebalancing problem. In Section III we (i) study the well-posedness and equilibria of the model; (ii) determine the minimum number of vehicles and drivers needed to meet the customer demand; and (iii) show that the system is indeed locally stable. In Section IV we show how to optimally route the rebalancing vehicles and drivers to keep the number of waiting customers is bounded while minimizing the number of rebalancing vehicles traveling in the network and the number of rebalancing drivers needed. In Section V we study the relation between the minimum number of drivers needed and the minimum number of vehicles needed. Due to space constraints, we refer to our technical note for proofs [].

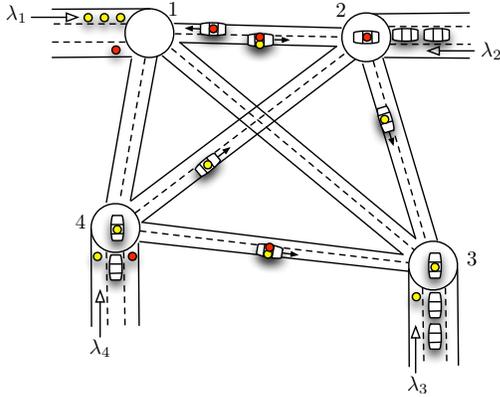


Fig. 1. At each station there are three queues: customers (yellow dots), drivers (red dots), and vehicles (small car icons). There are three types of car trips between stations: A customer can drive a car; a customer can be driven by a driver; or, a driver can drive a car to rebalance.

TABLE I
DESCRIPTION OF NOTATION FOR STATION i

	Definition
c_i	number of customers at station i
v_i	number of vehicles at station i
r_i	number of drivers at station i
λ_i	rate of arrival of customers at station i
μ_i	departure rate from station i
T_{ij}	travel time from station i to station j
p_{ij}	fraction of customers at station i destined for station j
α_{ij}	rate of rebalancing vehicles from station i to station j
γ_i	$\sum_j \alpha_{ij}$
β_{ij}	rate of rebalancing drivers from station i to station j
f_{ij}	fraction of customers traveling from i to j willing to use taxis
$H(\cdot)$	Heaviside function

II. MODELING THE MOBILITY-ON-DEMAND SYSTEM

In our prior work [3] we proposed a fluid model for mobility-on-demand systems and formulated a policy to optimally rebalance vehicles assuming that they could operate

autonomously. In this paper we consider rebalancing the vehicles through the use of dedicated personnel that are employed to drive the vehicles. In this section we extend the fluid model in [3] to capture the later scenario.

Basic model: The model in [3] can be formalized as follows. Consider a set of n stations, $\mathcal{N} = \{1, \dots, n\}$, defined over an extended geographical area (see Figure 1). Since the model is a fluid approximation, the number of customers, vehicles, and drivers are represented by real numbers. Customers arrive at station i at a constant rate $\lambda_i \in \mathbb{R}_{>0}$. The number of customers at station i at time t is $c_i(t) \in \mathbb{R}_{\geq 0}$, and the number of vehicles waiting idle at station i at time t is $v_i(t) \in \mathbb{R}_{\geq 0}$. The total number of vehicles in the system is $V \in \mathbb{R}_{>0}$. The fraction of customers at station i whose destination is station j is p_{ij} (where $p_{ij} \in \mathbb{R}_{\geq 0}$, $p_{ii} = 0$, and $\sum_j p_{ij} = 1$). The travel time from station i to station j is $T_{ij} \in \mathbb{R}_{\geq 0}$. When there are both customers and vehicles at station i (i.e., $c_i(t) > 0$ and $v_i(t) > 0$), then the rate at which customers (and hence vehicles) leave station i is μ_i ; when, instead, $c_i(t) = 0$ but $v_i(t) > 0$ the departure rate is λ_i . A necessary condition for the total number of customers at station i to remain bounded is that $\mu_i \geq \lambda_i$; we will assume $\mu_i > \lambda_i$ throughout the paper (the case $\mu_i = \lambda_i$ can be addressed with techniques similar to the ones introduced in this paper and is omitted).

From [3], we showed that a station is in need of rebalancing if $-\lambda_i + \sum_{j \neq i} \lambda_j p_{ji} \neq 0$. This can be easily understood by noting that λ_i is the rate at which vehicles leave station i , while $\sum_{j \neq i} \lambda_j p_{ji}$ is the rate at which vehicles arrive at station i . In what follows we assume that

$$-\lambda_i + \sum_{j \neq i} \lambda_j p_{ji} \neq 0 \quad \text{for all } i \in \mathcal{N},$$

and thus each station is in need of rebalancing. We comment further on this assumption in Remark III.3.

Rebalancing vehicles: In order to rebalance the number of vehicles $v_i(t)$ at each station, vehicles without customers will be driven between stations using hired human drivers. The number of drivers waiting at station i is $r_i(t) \in \mathbb{R}_{\geq 0}$ and the total number of drivers in the system is $R \in \mathbb{R}_{>0}$. In order to send a vehicle without a customer on a rebalancing trip from station i to station j , there must be a driver present at station i . We let $\alpha_{ij} \in \mathbb{R}_{\geq 0}$ denote the rate at which we send vehicles from station i to station j when vehicles and drivers are available at station i . The total rate at which station i sends vehicles without customers is $\gamma_i := \sum_j \alpha_{ij}$, where $\alpha_{ii} = 0$. We let α denote the matrix with entries given by α_{ij} . These trips are shown in Figure 1 as vehicles with red dots in them.

Rebalancing drivers: Finally, we must rebalance the drivers in the network, as they will tend to accumulate at some stations and become depleted at others. This is done as follows. If a driver would like to make a trip from station i to station j , it can drive a car for a customer on a trip from i to j , thereby acting as a taxi driver for that trip. This allows the driver to make the journey from station i to station j by “hitching a ride” on a passenger-carrying trip, but without

negatively affecting the customer experience. We quantify this using two sets of variables. The variables $\beta_{ij} \in \mathbb{R}_{\geq 0}$ give the rate at which drivers are sent from station i to station j when there are idle drivers available at station i . We let β denote the matrix with entries given by β_{ij} and assume $\beta_{ii} = 0$.

The quantities $f_{ij} \in (0, 1]$ give the fraction of customers making the trip from station i to j that would be willing to use the taxi mode of service on their trip. The remaining fraction of customers $1 - f_{ij}$ would prefer to drive themselves on their trip. Thus, f_{ij} imposes a constraint on the largest value of β_{ij} . In what follows we assume that the f_{ij} are such that there are enough customer trips available to rebalance the drivers. In Proposition III.4 we give a necessary and sufficient condition on the f_{ij} such that this is true. These trips are shown in Figure 1 as vehicles containing red and yellow dots. The notation is summarized in Table I.

We are now ready to write the differential equations governing the evolution of the number of vehicles, customers, and drivers at each station. In order to write the expressions more compactly, we introduce the following notation:

$$v_i := v_i(t), \quad c_i := c_i(t), \quad r_i := r_i(t), \\ v_j^i := v_j(t - T_{ji}), \quad c_j^i := c_j(t - T_{ji}), \quad r_j^i := r_j^i(t - T_{ji}).$$

(In other words, v_j^i denotes the number of vehicles that *were* present at station j , specifically T_{ji} time units *prior* to the current time.) Then, we can write the customer dynamics at station i as

$$\dot{c}_i = \begin{cases} \lambda_i, & \text{if } v_i = 0, \\ 0, & \text{if } v_i > 0 \text{ and } c_i = 0, \\ \lambda_i - \mu_i, & \text{if } v_i > 0 \text{ and } c_i > 0. \end{cases}$$

Defining the Heaviside function as

$$H(x) := \begin{cases} 1, & \text{if } x > 0, \\ 0, & \text{otherwise,} \end{cases}$$

the customer dynamics can be written as

$$\dot{c}_i = \lambda_i(1 - H(v_i)) + (\lambda_i - \mu_i)H(c_i)H(v_i).$$

The rate of change of vehicles at station i can be written as the sum of four components:

- 1) the rate at which customer-carrying vehicles depart station i :

$$\begin{cases} 0, & \text{if } v_i = 0 \\ -\lambda_i, & \text{if } v_i > 0 \text{ and } c_i = 0, \\ -\mu_i, & \text{if } v_i > 0 \text{ and } c_i > 0, \end{cases}$$

which can be written more compactly as $-\lambda_i H(v_i) + (\lambda_i - \mu_i)H(c_i)H(v_i)$;

- 2) the rate at which customer-carrying vehicles arrive at station i :

$$\sum_{j \neq i} p_{ji} \left(\lambda_j H(v_j^i) - (\lambda_j - \mu_j)H(c_j^i)H(v_j^i) \right);$$

- 3) the rate at which vehicles without a customer (rebalancing vehicles) depart station i , given by $-\gamma_i H(v_i)H(r_i)$;

- 4) the rate at which vehicles without a customer (rebalancing vehicles) arrive at station i , given by $\sum_{j \neq i} \alpha_{ji} H(v_j^i)H(r_j^i)$.

Thus, the vehicle dynamics can be written as

$$\dot{v}_i = -\lambda_i H(v_i) + (\lambda_i - \mu_i)H(c_i)H(v_i) \\ + \sum_{j \neq i} p_{ji} \left(\lambda_j H(v_j^i) - (\lambda_j - \mu_j)H(c_j^i)H(v_j^i) \right) \\ - \gamma_i H(v_i)H(r_i) + \sum_{j \neq i} \alpha_{ji} H(v_j^i)H(r_j^i),$$

Finally, the dynamics for the drivers contains four components. The first two components are identical to those of the rebalancing vehicles, given by 3) and 4) above. (This is due to the fact that each rebalancing vehicle contains a driver). The third component is the rate at which rebalancing drivers depart station i (by driving customer carrying vehicles): $-\sum_{j \neq i} \beta_{ij} H(v_i)H(r_i)$. The fourth term is the rate at which rebalancing drivers arrive at station i with a customer: $\sum_{j \neq i} \beta_{ji} H(v_j^i)H(r_j^i)$. Since drivers rebalance by driving vehicles on customer trips, we have from the customer dynamics \dot{c}_i that

$$\beta_{ij} \leq \begin{cases} f_{ij} \lambda_i p_{ij} & \text{if } c_i = 0 \\ f_{ij} \mu_i p_{ij} & \text{if } c_i > 0 \end{cases}$$

However, we will consider fixed values of β_{ij} , and since $\mu_i > \lambda_i$, we simply need to enforce the more stringent constraint $\beta_{ij} \leq f_{ij} \lambda_i p_{ij}$.

Therefore, the \dot{r}_i dynamics can be written as

$$\dot{r}_i = -\gamma_i H(v_i)H(r_i) + \sum_{j \neq i} \alpha_{ji} H(v_j^i)H(r_j^i) \\ - \sum_{j \neq i} \beta_{ij} H(v_i)H(r_i) + \sum_{j \neq i} \beta_{ji} H(v_j^i)H(r_j^i).$$

Putting everything together, we can write a set of nonlinear, time-delay differential equations describing the evolution of customers and vehicles in the system as

$$\dot{c}_i = \lambda_i(1 - H(v_i)) + (\lambda_i - \mu_i)H(c_i)H(v_i), \\ \dot{v}_i = -\lambda_i H(v_i) + (\lambda_i - \mu_i)H(c_i)H(v_i) + \\ \sum_{j \neq i} p_{ji} \left(\lambda_j H(v_j^i) - (\lambda_j - \mu_j)H(c_j^i)H(v_j^i) \right) \\ - \gamma_i H(v_i)H(r_i) + \sum_{j \neq i} \alpha_{ji} H(v_j^i)H(r_j^i), \quad (1) \\ \dot{r}_i = -\gamma_i H(v_i)H(r_i) + \sum_{j \neq i} \alpha_{ji} H(v_j^i)H(r_j^i) \\ - \sum_{j \neq i} \beta_{ij} H(v_i)H(r_i) + \sum_{j \neq i} \beta_{ji} H(v_j^i)H(r_j^i).$$

where $t \geq 0$; the initial conditions satisfy $c_i(\tau) = 0$, $v_i(\tau) = 0$, $r_i(\tau) = 0$ for $\tau \in [-\max_{i,j} T_{ij}, 0)$, $c_i(0) \in \mathbb{R}_{\geq 0}$, $v_i(0) \in \mathbb{R}_{\geq 0}$ with $v_i(0) > 0$ for at least one $i \in \mathcal{N}$, $r_i(0) \in \mathbb{R}_{\geq 0}$ with $r_i(0) > 0$ for at least one $i \in \mathcal{N}$,

and $\sum_i v_i(0) = V$ and $\sum_i r_i(0) = R$. The optimization variables α and β are constrained as follows:

$$\begin{aligned} 0 &\leq \beta_{ij} \leq f_{ij} \lambda_i p_{ij} \\ 0 &\leq \alpha_{ij}. \end{aligned}$$

The problem we wish to solve is as follows: find an *optimal* vehicle rebalancing assignment α and driver rebalancing assignment β that simultaneously

- 1) minimizes the number of rebalancing vehicles traveling in the network,
- 2) minimizes the number of drivers needed, and
- 3) ensures that the number of waiting customers remains bounded.

Note that this is a multi-objective optimization, and thus it is not clear that one can both minimize the number of rebalancing vehicles in the network and the number of drivers needed. However, it will turn out that these two objectives are aligned, and one can find an assignment (α, β) that minimizes both objectives.

III. WELL-POSEDNESS, EQUILIBRIA, AND STABILITY OF FLUID MODEL

In this section we first discuss the well-posedness of model (1) by showing two important properties, namely existence of solutions and invariance of the number of vehicles and rebalancing drivers along system trajectories. Then, we characterize the equilibria, we determine the minimum number of vehicles and drivers to ensure their existence, and we give a necessary and sufficient condition on the “user’s preference” f_{ij} such that there are enough customer trips available to rebalance the drivers. Finally, we show that rebalancing vehicles and drivers give rise to equilibria that are locally (i.e., within a neighborhood of the nominal conditions) stable.

A. Well-posedness

The fluid model (1) is nonlinear, time-delayed, and the right-hand side is discontinuous. Due to the discontinuity, we need to analyze the model within the framework of Filippov solutions (see, e.g., [9]). The following proposition verifies that the fluid model is well-posed.

Proposition III.1 (Well-posedness of fluid model). *For the fluid model (1), the following hold:*

- 1) *For every initial condition, there exist continuous functions $c_i(t) : \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$, $v_i(t) : \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$, and $r_i(t) : \mathbb{R} \rightarrow \mathbb{R}_{\geq 0}$ $i \in \mathcal{N}$, satisfying the differential equations (1) in the Filippov sense.*
- 2) *The total number of vehicles and rebalancing drivers is invariant for $t \geq 0$ and is equal, respectively, to $V = \sum_i v_i(0)$ and $R = \sum_i r_i(0)$.*

B. Equilibria

The following result characterizes the equilibria of model (1). The proof is contained in [] Recall that no station is exactly balanced, and thus $-\lambda_i + \sum_{j \neq i} \lambda_j p_{ji} \neq 0$, for all $i \in \mathcal{N}$.

Theorem III.2 (Existence of equilibria). *Let $\mathcal{A} \times \mathcal{B}$ be the set of assignments (α, β) that verify the equations*

$$\sum_{j \neq i} (\alpha_{ij} - \alpha_{ji}) = D_i, \quad (2)$$

$$\sum_{j \neq i} (\beta_{ij} - \beta_{ji}) = -D_i, \quad (3)$$

for each $i \in \mathcal{N}$, where $D_i := -\lambda_i + \sum_{j \neq i} \lambda_j p_{ji}$. Moreover, let

$$V_\alpha := \sum_{i,j} T_{ij} (p_{ij} \lambda_i + \alpha_{ij}), \quad \text{and}$$

$$R_{\alpha,\beta} := \sum_{i,j} T_{ij} (\alpha_{ij} + \beta_{ij}).$$

If $(\alpha, \beta) \notin \mathcal{A} \times \mathcal{B}$, then no equilibrium exists. If $(\alpha, \beta) \in \mathcal{A} \times \mathcal{B}$, there are two cases:

- 1) If $V > V_\alpha$ and $R > R_{\alpha,\beta}$, then the set of equilibria is

$$c_i = 0, \quad v_i > 0, \quad r_i > 0 \quad \forall i \in \mathcal{N},$$

where $\sum_i v_i = V - V_\alpha$ and $\sum_i r_i = R - R_{\alpha,\beta}$.

- 2) If $V \leq V_\alpha$ or $R \leq R_{\alpha,\beta}$, then no equilibrium exists.

Remark III.3 (Balanced stations case). *We have assumed that $D_i = -\lambda_i + \sum_{j \neq i} \lambda_j p_{ji} \neq 0$ for each station i . This assumption removes the pathological case that a station is perfectly balanced and does not need any rebalancing effort. In the case that $D_i = 0$ for a station, then $r_i = 0$ becomes a valid equilibrium. Due to space constraints we have omitted a full treatment of the $D_i = 0$ case in this presentation. •*

One question remains; does there always exist an assignment $(\alpha, \beta) \in \mathcal{A} \times \mathcal{B}$ that satisfies the constraints $\alpha_{ij} \geq 0$, and $0 \leq \beta_{ij} \leq f_{ij} \lambda_i p_{ij}$ for each $i, j \in \mathcal{N}$? We call such an assignment *feasible*. It is straightforward to verify that a feasible assignment for α always exists, since the variables are constrained only to be non-negative [3]. The β variables, however, are bounded from above (that is, they have finite capacities), and thus it is not clear whether there exists a feasible β assignment. The following result gives a standard condition for the existence of a feasible assignment (see, for example [11, p. 220] and a consequence of this condition. The proof is given in []).

Proposition III.4 (Existence of a feasible assignment). *A feasible assignment (α, β) exists if and only if,*

$$-\sum_{i \in S} D_i \leq \sum_{i \in S, j \notin S} f_{ij} \lambda_i p_{ij} \quad \text{for every set } S \subseteq \mathcal{N}, \quad (4)$$

where $D_i = -\lambda_i + \sum_{j \neq i} \lambda_j p_{ji}$. As a consequence, if $f_{ij} = 1$ for all $i, j \in \mathcal{N}$, then a feasible assignment always exists.

C. Stability of Equilibria

In this section we investigate the (local) *stability* of the equilibria of our model. We consider the following notion of local stability. Let $(\alpha, \beta) \in \mathcal{A} \times \mathcal{B}$ and assume $V > V_\alpha$ and $R > R_{\alpha,\beta}$ (this is a necessary and sufficient condition

to have equilibria, see Theorem III.2). We say that the (non-empty) set of equilibria

$$\mathcal{E}_{\alpha,\beta} := \{(\mathbf{c}, \mathbf{v}, \mathbf{r}) \in \mathbb{R}^{3n} \mid c_i = 0, v_i > 0, r_i > 0 \text{ for all } i \in \mathcal{N}, \text{ and } \sum_i v_i = V - V_\alpha \text{ and } \sum_i r_i = R - R_{\alpha,\beta}\} \quad (5)$$

is locally asymptotically stable if for any equilibrium $(\underline{\mathbf{c}}, \underline{\mathbf{v}}, \underline{\mathbf{r}}) \in \mathcal{E}_{\alpha,\beta}$ there exists a neighborhood $\mathcal{B}_{\alpha,\beta}^\delta(\underline{\mathbf{c}}, \underline{\mathbf{v}}, \underline{\mathbf{r}}) := \{(\mathbf{c}, \mathbf{v}, \mathbf{r}) \in \mathbb{R}^{3n} \mid c_i \geq 0, v_i \geq 0, r_i \geq 0 \text{ for all } i \in \mathcal{N}, \|(\mathbf{c} - \underline{\mathbf{c}}, \mathbf{v} - \underline{\mathbf{v}}, \mathbf{r} - \underline{\mathbf{r}})\| < \delta, \text{ and } \sum v_i = V - V_\alpha \text{ and } \sum r_i = R - R_{\alpha,\beta}\}$ such that every evolution of model (1) starting at

$$\begin{aligned} c_i(\tau) &= \underline{c}_i \text{ for } \tau \in [-\max_{i,j} T_{ij}, 0) \\ v_i(\tau) &= \underline{v}_i \text{ for } \tau \in [-\max_{i,j} T_{ij}, 0) \\ r_i(\tau) &= \underline{r}_i \text{ for } \tau \in [-\max_{i,j} T_{ij}, 0) \\ (\mathbf{c}(0), \mathbf{v}(0), \mathbf{r}(0)) &\in \mathcal{B}_{\alpha,\beta}^\delta(\underline{\mathbf{c}}, \underline{\mathbf{v}}, \underline{\mathbf{r}}) \end{aligned} \quad (6)$$

has a limit which belongs to the equilibrium set. In other words, $(\lim_{t \rightarrow +\infty} \mathbf{c}(t), \lim_{t \rightarrow +\infty} \mathbf{v}(t), \lim_{t \rightarrow +\infty} \mathbf{r}(t)) \in \mathcal{E}_{\alpha,\beta}$. The next theorem characterizes stability. The proof is contained in []

Theorem III.5 (Stability of equilibria). *Let $(\alpha, \beta) \in \mathcal{A} \times \mathcal{B}$ be a feasible assignment, and assume $V > V_\alpha$ and $R > R_{\alpha,\beta}$; then, the set of equilibria $\mathcal{E}_{\alpha,\beta}$ is locally asymptotically stable.*

IV. OPTIMAL REBALANCING

Our objective is to find a rebalancing assignment (α, β) that simultaneously minimizes the number of rebalancing vehicles traveling in the network and the number of rebalancing drivers needed, while ensuring the existence of (locally) stable equilibria for model (1). From the previous section, we already know that the set of assignments ensuring the existence of stable equilibria is $\mathcal{A} \times \mathcal{B}$ (provided that the total number of vehicles V and drivers R is large enough).

The time-average number of rebalancing vehicles traveling in the network is simply given by $\sum_{i,j} T_{ij} \alpha_{ij}$. In minimizing this quantity we are also minimizing the lower bound on the necessary number of vehicles V_α . The time-average number of drivers in the network is given by $\sum_{i,j} T_{ij} (\alpha_{ij} + \beta_{ij})$. In minimizing this quantity we are minimizing the lower bound on the necessary number of drivers $R_{\alpha,\beta}$.

Combining the two objectives with the existence of stable equilibria constraints in (2) and (3), we obtain the following:

$$\begin{aligned} &\text{minimize } \sum_{i,j} T_{ij} \alpha_{ij} \text{ and } \sum_{i,j} T_{ij} (\alpha_{ij} + \beta_{ij}) \\ &\text{subject to } \sum_{j \neq i} (\alpha_{ij} - \alpha_{ji}) = D_i \quad \forall i \in \mathcal{N} \\ &\quad \sum_{j \neq i} (\beta_{ij} - \beta_{ji}) = -D_i \quad \forall i \in \mathcal{N} \\ &\quad 0 \leq \alpha_{ij} \quad \forall i, j \in \mathcal{N}, \\ &\quad 0 \leq \beta_{ij} \leq f_{ij} \lambda_i p_{ij} \quad \forall i, j \in \mathcal{N}, \end{aligned}$$

where $D_i = \lambda_i + \sum_{j \neq i} \lambda_j p_{ji}$, and the optimization variables are α_{ij} and β_{ij} , where $i, j \in \mathcal{N}$. The constraints ensure that the optimization is over the set $\mathcal{A} \times \mathcal{B}$. Note, however, that this optimization can be decoupled into an optimization over α and an optimization over β . Both optimizations are minimum cost flow problems [11]. The α optimization is identical to that presented in [3]:

$$\begin{aligned} &\text{minimize } \sum_{i,j} T_{ij} \alpha_{ij} \\ &\text{subject to } \sum_{j \neq i} (\alpha_{ij} - \alpha_{ji}) = D_i \quad \forall i \in \mathcal{N} \\ &\quad \alpha_{ij} \geq 0 \quad \forall i, j \in \mathcal{N}. \end{aligned}$$

The β optimization then looks as follows:

$$\begin{aligned} &\text{minimize } \sum_{i,j} T_{ij} \beta_{ij} \\ &\text{subject to } \sum_{j \neq i} (\beta_{ij} - \beta_{ji}) = -D_i \quad \forall i \in \mathcal{N} \\ &\quad 0 \leq \beta_{ij} \leq f_{ij} \lambda_i p_{ij} \quad \forall i, j \in \mathcal{N}. \end{aligned}$$

The α optimization is an uncapacitated minimum cost flow problem and thus is always feasible. In Proposition III.4 we give conditions on the f_{ij} fractions in order for the β optimization to be feasible.

The rebalancing policy is then given by solving the two minimum cost flow problems to obtain solutions α_{ij}^* and β_{ij}^* . We then send empty rebalancing vehicles (along with drivers) from station i to station j at a rate of α_{ij}^* (when vehicles and drivers are available at station i). We send drivers on customer-carrying vehicles from i to j at a rate of β_{ij}^* (when customers and vehicles are available at station i).

V. SIMULATIONS

In this section we study the relation between the minimum number of drivers needed for stability $R_{\alpha,\beta}$ and the minimum number of vehicles needed V_α from Theorem III.2. To evaluate these quantities, we need to generate sample data consisting of arrival rates λ_i at each station i , customer destination probabilities p_{ij} , travel times between stations T_{ij} , and the fraction of customers f_{ij} traveling from i to j that are willing to be driven by a driver. We generate this data as follows: We uniformly randomly place n stations in a 100×100 environment, and calculate the travel times T_{ij} as the Euclidean distance between stations. We uniformly randomly generate the arrival rates λ_i on the interval $[0, 0.05]$ arrivals per time unit. Similarly we uniformly randomly generate the destination probabilities p_{ij} such that they are nonnegative and $\sum_j p_{ij} = 1$ for each station i . Finally, we assume that $f_{ij} = 1$ for each pair of stations in order to avoid issues with feasibility.

To solve the optimizations in Section IV for the optimal assignment $(\alpha^*, \beta^*) \in \mathcal{A} \times \mathcal{B}$, we use the freely available SeDuMi (Self-Dual-Minimization) toolbox.

Figure 2 shows results for numbers of stations ranging from 10 up to 200. For each number of stations we generate 20 random problem instances of the form described above.

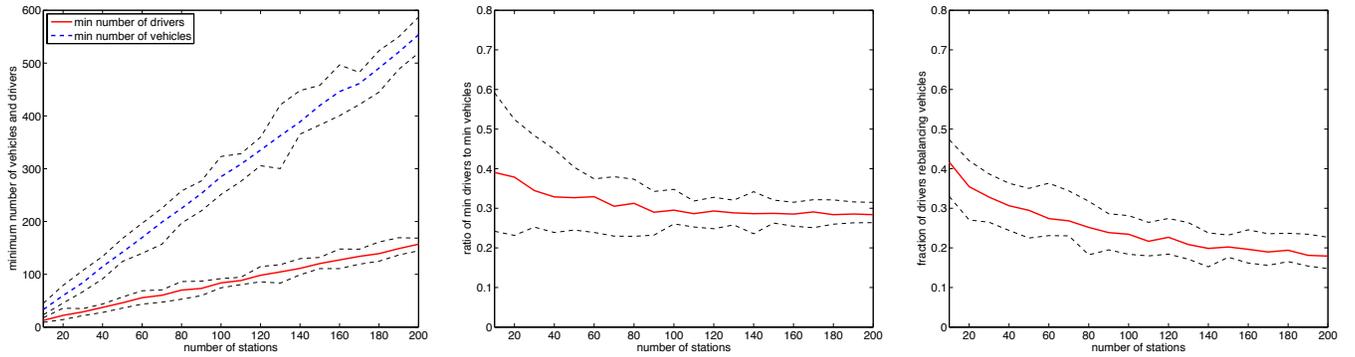


Fig. 2. Left figure: The minimum number of vehicles and drivers. Middle figure: The ratio between the minimum number of drivers and number of vehicles. Right figure: The fraction of drivers that are performing vehicle rebalancing trips. For each fixed number of stations, 20 trials were performed. Thick lines show the mean of the 20 trials while thin dashed lines show the maximum and minimum over the trials.

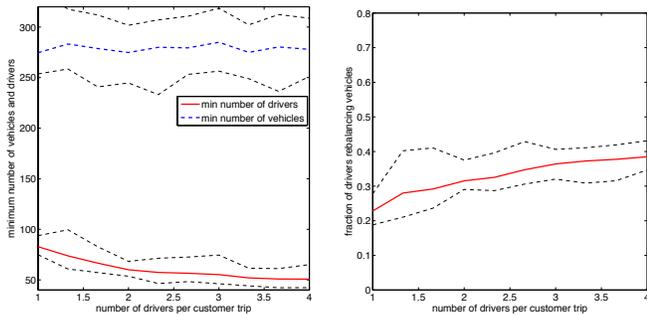


Fig. 3. Increasing the number of drivers per customer trip for 100 station problems. Left figure: The minimum number of vehicles and drivers. Right figure: The fraction of drivers that are performing vehicle rebalancing trips.

The thick line in each plot shows the mean over the 20 trials while the thin dashed lines show the maximum and minimum values. The left figure shows how V_{α^*} and R_{α^*,β^*} vary with the number of stations. The middle figure shows the ratio $R_{\alpha^*,\beta^*}/V_{\alpha^*}$ as a function of the number of stations. We can see that we need between $1/3$ and $1/4$ as many drivers as we do vehicles. The right figure shows the ratio between the minimum number of rebalancing vehicles in transit and the number of drivers. This gives a measure of the fraction of drivers that are driving rebalancing vehicles. It is interesting to note that this ratio is quite low, reaching approximately $1/5$ for 200 stations.

One way to increase the fraction of drivers performing vehicle rebalancing is to allow multiple drivers to take a trip with a customer. This allows drivers to take more efficient routes back to stations that are in need of drivers. In our model it corresponds to setting $f_{ij} > 1$. This is explored in Figure 3 where we range f_{ij} from 1 to 4 for 20 problem instances on 100 stations. We see that as f_{ij} increases from 1 to 4, the number of drivers decreases from approximately 80 to 50, and the fraction of drivers performing vehicle rebalancing increases from under $1/4$ to nearly $2/5$.

VI. CONCLUSIONS

In this paper we studied the problem of rebalancing the rebalancers in a mobility-on-demand system, which blends

customer-driven vehicles with a taxi service. For a fluid model of the system, we showed that the optimal rebalancing policy can be found as the solution of two linear programs. Also, we showed that in Euclidean network topologies one would need between $1/3$ and $1/4$ as many drivers as vehicles, and that this fraction decreases to about $1/5$ if one allows up to 3-4 drivers to take a trip with a customer. These results could have an immediate impact on existing one-way car-sharing systems such as Car2Go.

For future work we plan to analyze a stochastic queueing model and study the time-varying case whereby the system's parameters change periodically (thus modeling the day/night variations). Also, we plan to develop real-time rebalancing policies that do not require any a priori information, and to incorporate uncertainty in the travel times, time windows for the customers, and capacity constraints for the roads.

REFERENCES

- [1] W. J. Mitchell, C. E. Borroni-Bird, and L. D. Burns, *Reinventing the Automobile: Personal Urban Mobility for the 21st Century*. Cambridge, MA: The MIT Press, 2010.
- [2] CAR2GO, "CAR2GO Austin. Car Sharing 2.0: Great Idea for a Great City," Tech. Rep., 2011.
- [3] M. Pavone, S. L. Smith, E. Frazzoli, and D. Rus, "Robotic load balancing for mobility-on-demand systems," *International Journal of Robotics Research*, vol. 31, no. 7, pp. 839–854, 2012.
- [4] —, "Load balancing for mobility-on-demand systems," in *Robotics: Science and Systems*, Los Angeles, CA, June 2011.
- [5] D. K. Merchant and G. L. Nemhauser, "Optimality Conditions for a Dynamic Traffic Assignment Model," *Transportation Science*, vol. 12, no. 3, pp. 200–207, 1978.
- [6] T. L. Friesz, J. Luque, R. L. Tobin, and B. W. Wie, "Dynamic Network Traffic Assignment Considered as a Continuous Time Optimal Control Problem," *Operations Research*, vol. 37, no. 6, pp. 893–901, 1989.
- [7] A. K. Ziliaskopoulos, "A Linear Programming Model for the Single Destination System Optimum Dynamic Traffic Assignment Problem," *Transportation Science*, vol. 34, no. 1, pp. 37–49, 2000.
- [8] S. Peeta and A. Ziliaskopoulos, "Foundations of Dynamic Traffic Assignment: The Past, the Present and the Future," *Networks and Spatial Economics*, vol. 1, pp. 233–265, 2001.
- [9] A. F. Filippov, *Differential Equations with Discontinuous Righthand Sides*, ser. Mathematics and its Applications. Dordrecht, The Netherlands: Kluwer Academic Publishers, 1988, vol. 18.
- [10] G. Haddad, "Monotone trajectories of differential inclusions and functional differential inclusions with memory," *Israel Journal of Mathematics*, vol. 39, no. 1-2, pp. 83 – 100, 1981.
- [11] B. Korte and J. Vygen, *Combinatorial Optimization: Theory and Algorithms*, 4th ed., ser. Algorithmics and Combinatorics. Springer, 2007, vol. 21.