

Robotic Load Balancing for Mobility-on-Demand Systems*

Marco Pavone[†] Stephen L. Smith[‡] Emilio Frazzoli[§] Daniela Rus[¶]

September 2, 2011

Abstract

In this paper we develop methods for maximizing the throughput of a mobility-on-demand urban transportation system. We consider a finite group of shared vehicles, located at a set of stations. Users arrive at the stations, pick-up vehicles, and drive (or are driven) to their destination station where they drop-off the vehicle. When some origins and destinations are more popular than others, the system will inevitably become out of balance: Vehicles will build up at some stations, and become depleted at others. We propose a robotic solution to this rebalancing problem that involves empty robotic vehicles autonomously driving between stations. Specifically, we develop a rebalancing policy that lets every station reach an equilibrium in which there are excess vehicles and no waiting customers and that minimizes the number of robotic vehicles performing rebalancing trips. To do this, we utilize a fluid model for the customers and vehicles in the system. We then show that the optimal rebalancing policy can be found as the solution to a linear program. We use this solution to develop a real-time rebalancing policy which can operate in highly variable environments. We verify policy performance in a simulated mobility-on-demand environment and in hardware experiments.

1 Introduction

In the past century, private automobiles have dramatically changed the concept of *personal urban mobility* by enabling fast and anytime point-to-point travel within large cities. In 2001, personal urban mobility in the US resulted in more than 3.5 trillion urban miles traveled by private cars, representing 75% of total car travel in the US (BTS, 2001). This figure, coupled with the fact that by 2030 the total population living in urban areas will jump from the current 40% to more than 60% (UN, 2007), implies that the demand for personal urban mobility will increase to formidable levels. The demand for roads and parking spaces will dramatically increase, while the availability urban land will continue to decrease. The result is that private automobiles are an unsustainable solution for the future of personal mobility in dense urban environments. To cope with this problem, a paradigm shift is emerging to replace the outdated policy of infrastructure augmentation with

*A preliminary version of this paper has appeared in the Proceedings of the 2011 Robotics: Science and Systems Conference, Pavone et al. (2011).

[†]M. Pavone is with the Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109, USA (marco.pavone@jpl.nasa.gov).

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

[§]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).

personal urban mobility leveraging robotics and automation. The challenge is to ensure the same benefits of privately-owned cars without requiring additional roads and parking spaces.

One of the leading emerging paradigms for future urban mobility systems is one-way vehicle sharing, which effectively merges private and public mobility, and directly targets the problems of parking spaces and current low vehicle utilization rates. Arguably, the most promising approach within this paradigm is represented by *Mobility-on-Demand* (MOD) systems (Mitchell et al., 2010), which provide stacks and racks of light electric vehicles at closely spaced intervals throughout a city: when a person wants to go somewhere, he/she simply walks to the nearest rack, swipes a card to pick up a vehicle, drives it to the rack nearest to his/her destination, and drops it off (see Figure 1 for an illustration of mobility-on-demand). Large-scale systems employing traditional, non-electric bicycles have already demonstrated the feasibility of mobility-on-demand in several cities throughout Europe, e.g., Paris, Lyon, Milano, Trento, Zurich (Midgley, 2009). Furthermore, experimental, limited-size, car-based MOD systems have been deployed in the past few years (Massot et al., 1999; Barth and Todd, 2001; CAR2GO, 2011) and a number of car manufacturers are currently developing two-seat electric vehicles specifically designed for MOD systems and even capable of autonomous operation (see, e.g., the General Motors EN-V prototype (GM, 2011)).

However, sharing has its drawbacks. When some origins and some destinations are more popular than others, the system will inevitably become out of balance: Vehicles will build up at some stations, and become depleted at others (this is one of the reasons why car-sharing operators such as Zipcar and Hertz allow hourly rental but *only* on a round-trip basis). In this paper we propose a *robotic solution for vehicle rebalancing* in MOD systems, whereby the shared vehicles (e.g., the General Motor’s EN-V prototype) autonomously drive from a delivery location to the next pick-up location. Rebalancing through autonomously driving vehicles has the clear potential of eliminating imbalances within the transportation network, and effectively adds another dimension to MOD systems by introducing autonomy in the design space. In the recent past, considerable efforts have been devoted to the problem of autonomous driving, and substantial progress has been made (see, for example, (Buehler et al., 2007, 2009)). However, there are virtually no tools to address the system-level problems arising at the interface between robotics and transportation science: How should one pre-position vehicles in order to anticipate future demand? Is it possible to characterize optimal, real-time rebalancing policies? How many vehicles are needed to achieve a certain quality of service (e.g., a desired average waiting time for the customers)? The purpose of this paper is to develop an approach that provides rigorous answers to such questions.

Even though rebalancing in MOD systems is an entirely new problem within the realm of transportation networks, it has many characteristics in common with the well-known Dynamic Traffic Assignment (DTA) problem (Merchant and Nemhauser, 1978; Friesz et al., 1989; Ziliaskopoulos, 2000; Peeta and Ziliaskopoulos, 2001). In this problem, one seeks to “optimize” the time varying flows on each arc of a transportation network, taking into account congestion effects along arcs and at nodes (Friesz et al., 1989). DTA models mainly differ in the methods used to capture the time-varying nature of supply and demand processes, and can be broadly divided into four categories (Peeta and Ziliaskopoulos, 2001): (i) mathematical programming, e.g., Merchant and Nemhauser (1978); Ziliaskopoulos (2000), (ii) optimal control, e.g., Friesz et al. (1989), (iii) variational inequality, e.g., Friesz et al. (1993), and (iv) simulation-based, e.g., Balakrishna et al. (2007). 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. Rebalancing in MOD systems is also related to Dynamic one-to-one Pick-up and Delivery problems (DPDPs), where dynamically-generated passengers must be transported from a pick-up site to a delivery site by a fleet of vehicles. DPDPs can be divided into three main categories (Berbeglia et al., 2010): (i) Dynamic Stacker Crane Problem, where the vehicles have unit

capacity, (ii) Dynamic Vehicle Routing Problem with Pickups and Deliveries, where the vehicles can transport more than one request, and (iii) Dynamic Dial-a-Ride Problem, where additional constraints such as time windows are considered. Excellent surveys on heuristics, metaheuristics and online algorithms for DPDPs can be found in Berbeglia et al. (2010) and Parragh et al. (2008), while analysis specifically tailored to the structural properties of transportation-on-demand systems can be found in Pavone et al. (2010). The key difference from DPDP problems is that there are a finite number of pick-up and delivery sites, the vehicles are not aware of the destination of newly arrived customers, and the optimization is over the empty vehicle trips. Finally, our problem is also related to dynamic load balancing in distributed computing systems (Cybenko, 1989; Cardellini et al., 1999). However, these problems are less constrained since the demands (i.e., jobs) do not need to wait for a “vehicle” to move across multiple processors.




Figure 1: At each station there is a queue of customers (yellow dots) and a queue of vehicles (represented by small car icons). The customer at the head of queue enters the vehicle at the head of the queue. This is shown in the circles at stations 3 and 4. Notice that at station 1 there are no vehicles, and at station 2 there are no customers. In rebalancing, we send empty vehicles from station 2 to station 1.

This paper is structured as follows. In Section 2 we present a fluid model for MOD systems, and we formally state the rebalancing problem. In Section 3 we (i) study well-posedness and equilibria of the fluid model, and we show that without rebalancing vehicles the system is unstable (i.e., at some stations the number of waiting customers will grow without bound); (ii) determine the minimum number of vehicles needed to meet the customer demand; and, (iii) show that with rebalancing vehicles the system is indeed locally stable (i.e., stable within a neighborhood of the nominal conditions). Then, in Section 4, we show how to optimally route the rebalancing vehicles across the transportation network so that stability is ensured while minimizing the number of empty vehicles traveling. The results in Sections 3 and 4 lead to a robust, real-time policy for vehicle rebalancing, which is presented in Section 5 and which is thoroughly evaluated through simulation experiments in Section 6. In Section 7 we present hardware experiments with the robotic testbed shown in




Figure 2: The MOD testbed used to implement rebalancing policies. The white rectangles represent stations (four in total), while the dashed lines represent streets. There are eight robotic vehicles providing service to “virtual” customers (simulated in a ground station laptop). Details about the hardware experiments can be found in Section 7.

Figure 2, and in Section 8 we draw our conclusions, and we present directions for future research.

A preliminary version of this paper appeared as Pavone et al. (2011). Compared to the conference version, this version presents detailed proofs of all the statements, additional remarks and examples, an expanded simulation section, and hardware experiments.

2 Modeling the Mobility-on-Demand System

In this paper, we use a *fluidic approach* to model mobility-on-demand systems. Our fluid model is intended to serve as an approximation of a corresponding stochastic queueing model, where customers enter the system according to a Poisson process and where travel times between stations are nondeterministic.¹ Consider n stations defined over an extended geographical area (see Figure 1; see also Figure 2 for a corresponding hardware testbed). We denote the set of stations with \mathcal{N} . In this model, the number of customers and vehicles 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 relation between V and the v_i 's will be made explicit at the end of this section). 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

¹This can be formalized by showing that the fluid model arises as the limit of a sequence of appropriately scaled queueing models. Such analysis will be presented in a forthcoming paper. In this paper, the relation between the fluidic approximation and the queueing model will be discussed through simulation in Section 6.

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 0. 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).

In order to rebalance the number of vehicles at each station, (robotic) empty vehicles will be sent between stations. The rate at which station i sends empty (i.e., rebalancing) vehicles to station j is denoted by $\alpha_{ij} \in \mathbb{R}_{\geq 0}$ and the total rate at which station i sends empty vehicles is $\gamma_i := \sum_j \alpha_{ij}$, where $\alpha_{ii} = 0$. We let α denote the matrix with entries given by α_{ij} . The notation is summarized in Figure 1.

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

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

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 empty (rebalancing) vehicles depart station i , given by $\gamma_i H(v_i)$;

4. the rate at which empty (rebalancing) vehicles arrive at station i , given by $\sum_{j \neq i} \alpha_{ji} H(v_j^i)$.

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

$$\begin{aligned} \dot{c}_i &= \lambda_i(1 - H(v_i)) + (\lambda_i - \mu_i)H(c_i)H(v_i), \\ \dot{v}_i &= -\lambda_i H(v_i) + (\lambda_i - \mu_i)H(c_i)H(v_i) + \sum_{j \neq i} p_{ji} \left(\lambda_j H(v_j^i) - (\lambda_j - \mu_j)H(c_j^i)H(v_j^i) \right) \\ &\quad - \gamma_i H(v_i) + \sum_{j \neq i} \alpha_{ji} H(v_j^i), \end{aligned} \quad (1)$$

where $t \geq 0$; the initial conditions satisfy $c_i(\tau) = 0$, $v_i(\tau) = 0$ 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 $i \in \mathcal{N}$, and $\sum_i v_i(0) = V$.

The problem we wish to solve is as follows: find an *optimal* rebalancing assignment α that minimizes the number of rebalancing vehicles traveling in the network *and* ensures that the number of waiting customers remains bounded. It will turn out (see Remark 4.1) that this objective is equivalent to minimizing the vehicle utilization rate while meeting the customer demand.

3 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 along system trajectories. Then, we characterize the equilibria, we show that without rebalancing vehicles the system, in general, does not admit any equilibrium (and indeed waiting customers will grow without bounds), and we determine the minimum number of vehicles required for stabilizability (i.e., to ensure existence of equilibria). Finally, we show that rebalancing vehicles give rise to equilibria that are locally (i.e., within a neighborhood of the nominal conditions) stable.

3.1 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., Filippov (1988)). The following proposition verifies that the fluid model is well-posed.

Proposition 3.1 (Well-posedness of fluid model). *For the fluid model (1), the following hold:*

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

Proof. To prove the first claim, it can be checked that all assumptions of Theorem II-1 in Haddad (1981) for the existence of Filippov solutions to time-delayed differential equations with discontinuous right-hand side are satisfied, and the claim follows.

To prove the second claim, let $v_{ij}(t)$, where $t \geq 0$, be the number of vehicles in-transit from station i to station j (i.e., the vehicles for which the last station visited is i and the next station they will visit is j). Clearly, $v_{ii}(t) = 0$. Now, the total number $V(t)$ of vehicles in the system at time $t \geq 0$ is given, by definition, by $V(t) = \sum_{i=1}^n v_i(t) + \sum_{i,j} v_{ij}(t)$. One can express $v_{ij}(t)$ as

$$v_{ij}(t) = \int_{t-T_{ij}}^t p_{ij} \left(\lambda_i H(v_i(\tau)) - (\lambda_i - \mu_i)H(c_i(\tau))H(v_i(\tau)) \right) + \alpha_{ij} H(v_i(\tau)) d\tau. \quad (2)$$

By applying the Leibniz integral rule, one can write

$$\begin{aligned}\dot{v}_{ij}(t) &= p_{ij} \left(\lambda_i H(v_i) - (\lambda_i - \mu_i) H(c_i) H(v_i) \right) + \alpha_{ij} H(v_i) \\ &\quad - p_{ij} \left(\lambda_i H(v_i^j) - (\lambda_i - \mu_i) H(c_i^j) H(v_i^j) \right) - \alpha_{ij} H(v_i^j).\end{aligned}$$

Therefore, one immediately obtains, for $t \geq 0$,

$$\begin{aligned}\dot{V}(t) &= \sum_{i=1}^n \dot{v}_i(t) + \sum_{i=1}^n \sum_{j=1}^n \dot{v}_{ij}(t) \\ &= \sum_{i=1}^n \sum_{j=1}^n p_{ij} \left(-\lambda_i H(v_i(t)) + (\lambda_i - \mu_i) H(c_i(t)) H(v_i(t)) \right) + \sum_{i=1}^n \sum_{j=1}^n p_{ji} \left(\lambda_j H(v_j(t - T_{ji})) \right. \\ &\quad \left. - (\lambda_j - \mu_j) H(c_j(t - T_{ji})) H(v_j(t - T_{ji})) \right) - \sum_{i=1}^n \sum_{j=1}^n \alpha_{ij} H(v_i(t)) \\ &\quad + \sum_{i=1}^n \sum_{j=1}^n \alpha_{ji} H(v_j(t - T_{ji})) + \sum_{i=1}^n \sum_{j=1}^n \dot{v}_{ij}(t) \\ &= 0.\end{aligned}$$

This proves the claim. □

3.2 Equilibria

The equilibria of model (1) are characterized by the following theorem.

Theorem 3.2 (Existence of equilibria). *Let \mathcal{A} be the set of assignments α that verify the equation*

$$\sum_{j \neq i} \alpha_{ij} - \sum_{j \neq i} \alpha_{ji} = -\lambda_i + \sum_{j \neq i} \lambda_j p_{ji}, \quad (3)$$

for each $i \in \mathcal{N}$, and let

$$V_\alpha := \sum_{ij} T_{ij} (p_{ij} \lambda_i + \alpha_{ij}).$$

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

1. If $V > V_\alpha$, then the set of equilibria is

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

where $\sum_i v_i = V - \sum_{ij} T_{ij} (p_{ij} \lambda_i + \alpha_{ij})$.

2. If $V \leq V_\alpha$, then no equilibrium exists.

Proof. To prove the theorem, we set $\dot{c}_i = 0$ and $\dot{v}_i = 0$ for all $i \in \mathcal{N}$. From the $\dot{c}_i = 0$ equations we obtain

$$\lambda_i = \lambda_i H(v_i) - (\lambda_i - \mu_i) H(v_i) H(c_i). \quad (4)$$

Since $\lambda_i < \mu_i$, the above equations have a solution only if

$$c_i = 0 \quad \text{and} \quad v_i > 0 \quad \forall i \in \mathcal{N}.$$

Then, setting $\dot{v}_i = 0$, combined with (4), we obtain

$$0 = -\lambda_i + \sum_{j \neq i} \lambda_j p_{ji} - \gamma_i H(v_i) + \sum_{j \neq i} \alpha_{ji} H(v_j),$$

where we have used the fact that in a stationary equilibrium $v_i(t)$ and $c_i(t)$ are constants.

For every equilibrium we require $v_i > 0$, and thus $H(v_i) = 1$. Therefore, a necessary condition for the existence of equilibria is that the rebalancing assignment α can be chosen such that

$$\sum_{j \neq i} \alpha_{ij} - \sum_{j \neq i} \alpha_{ji} = -\lambda_i + \sum_{j \neq i} \lambda_j p_{ji},$$

for each $i \in \mathcal{N}$. Hence, if $\alpha \notin \mathcal{A}$, no equilibrium exists and the first claim is proven.

Assume now that $\alpha \in \mathcal{A}$ and assume that $V > V_\alpha$; we now want to show that the candidate equilibria $c_i = 0$ and $v_i > 0$ for all $i \in \mathcal{N}$ are indeed valid equilibria (note that these are the only possible equilibria). Since $\alpha \in \mathcal{A}$, then the necessary condition for the existence of equilibria in equation (3) is satisfied; the only condition that yet needs to be verified is that the overall number of vehicles is sufficient to sustain the equilibrium system traffic, i.e., the equilibrium flow of transit vehicles. Indeed, when $c_i = 0$ and $v_i > 0$ for all $i \in \mathcal{N}$, the equilibrium number of transit vehicles is given by V_α (recall equation (2)). Hence, in order to satisfy the condition $v_i > 0$ for all $i \in \mathcal{N}$, one needs a number of vehicles larger than V_α , which is verified by assumption. This, together with the invariance result in Theorem 3.1, shows the second claim.

Finally, by using similar arguments, one can show that when $\alpha \in \mathcal{A}$ but $V \leq V_\alpha$ no equilibrium exists. \square

Equation (3) implies that without rebalancing vehicles (i.e., when each α_{ij} is equal to zero), the system, in general, does not have equilibria. Also, it can be shown that in absence of equilibria the number of waiting customers will grow without bounds (the proof of this statement is rather straightforward—it can be obtained with a two station example—and is omitted in the interest of brevity). Hence, in general, *rebalancing vehicles are necessary to ensure equilibria and stability*.

Theorem 3.2 shows that if the set of assignments \mathcal{A} is empty, then no equilibrium can exist. The next lemma shows that, fortunately, there always exists a *rebalancing assignment* that satisfies equation (3), i.e., the set \mathcal{A} is always non-empty.

Lemma 3.3 (Existence of assignments satisfying equation (3)). *There always exists an assignment α such that equation (3) is satisfied, i.e., set \mathcal{A} is always non-empty.*

Proof. Consider the assignment in which $\alpha_{1k} := \lambda_k - \sum_{j \neq k} \lambda_j p_{jk}$, for $k \in \{2, \dots, n\}$, and all other entries in α are zero. For each $i \in \{2, \dots, n\}$, we can substitute the assignment into the LHS of (3), and verify that the constraint is satisfied. Thus, we just need to verify the assignment satisfies the constraint in (3) obtained when we set $i = 1$. Setting $i = 1$ in (3), and bringing all terms to the LHS, the constraint becomes $\sum_{k \neq 1} \alpha_{1k} + (\lambda_1 - \sum_{j \neq 1} \lambda_j p_{j1}) = 0$. Substituting the assignment into the above equation, the LHS becomes $\sum_k \left(\lambda_k - \sum_j \lambda_j p_{jk} \right) = \sum_k \lambda_k - \sum_k \sum_j \lambda_j p_{jk} = \sum_k \lambda_k - \sum_j \lambda_j \sum_k p_{jk} = 0$, since $\sum_k p_{jk} = 1$. Thus, the proposed assignment is feasible and the linear system of equations given by (3) is consistent with at least one solution. \square

We conclude this section by presenting a simple result about the minimum number of vehicles required to ensure the existence of equilibria, which is a direct consequence of Theorem 3.2 and Lemma 3.3.

Corollary 3.4 (Minimum number of vehicles for the existence of equilibria). *Model (1) admits equilibria if and only if*

$$V > \underline{V} := \min_{\alpha \in \mathcal{A}} V_\alpha.$$

3.3 Stability of Equilibria

In this section we consider the following questions: assume that the system is in equilibrium, what happens if there is a “burst” of customers arriving at the stations? Or, what happens if there is a sudden change in the number of available vehicles (e.g., because of a disruption)? In other words, we investigate the (local) *stability* of the equilibria of our model.

3.3.1 Stability of Equilibrium Sets

We consider the following notion of local stability. Let $\alpha \in \mathcal{A}$ and assume $V > V_\alpha$ (this is a necessary and sufficient condition to have equilibria, see Theorem 3.2). The (non-empty) set of equilibria

$$\mathcal{E}_\alpha := \{(\mathbf{c}, \mathbf{v}) \in \mathbb{R}^{2n} \mid c_i = 0, v_i > 0, \text{ for all } i \in \mathcal{N}, \text{ and } \sum_i v_i = V - V_\alpha\} \quad (5)$$

is locally asymptotically stable if for any equilibrium $(\underline{\mathbf{c}}, \underline{\mathbf{v}}) \in \mathcal{E}_\alpha$ there exists a neighborhood $\mathcal{B}_\alpha^\delta(\underline{\mathbf{c}}, \underline{\mathbf{v}}) := \{(\mathbf{c}, \mathbf{v}) \in \mathbb{R}^{2n} \mid c_i \geq 0, v_i \geq 0 \text{ for all } i \in \mathcal{N}, \|(\mathbf{c} - \underline{\mathbf{c}}, \mathbf{v} - \underline{\mathbf{v}})\| < \delta, \text{ and } \sum v_i = V - V_\alpha\}$ such that every evolution of model (1) starting at

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

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

Theorem 3.5 (Stability of equilibria). *Let $\alpha \in \mathcal{A}$ be a feasible assignment, and assume $V > V_\alpha$; then, the set of equilibria \mathcal{E}_α is locally asymptotically stable.*

Proof. The proof is provided in the Appendix. □

3.3.2 Stability of Single Equilibrium Points

In the previous section, we have discussed the stability of the set of equilibria \mathcal{E}_α . Here, we discuss stability of single equilibrium points. For simplicity, in this section we restrict perturbations to the number of customers only (the extension to the general case is possible but more involved).

Specifically, consider an assignment $\alpha \in \mathcal{A}$, and assume, as usual, $V > V_\alpha$. We say that an equilibrium point $(\underline{\mathbf{c}}, \underline{\mathbf{v}}) \in \mathcal{E}_\alpha$ is locally asymptotically stable if there exists a neighborhood

$$\mathcal{B}_\alpha^\delta(\underline{\mathbf{c}}, \underline{\mathbf{v}}) := \{(\mathbf{c}, \mathbf{v}) \in \mathbb{R}^{2n} \mid c_i \geq 0, v_i = \underline{v}_i \text{ for all } i \in \mathcal{N}, \text{ and } \|(\mathbf{c} - \underline{\mathbf{c}}, \mathbf{0})\| < \delta\}$$

such that every evolution of model (1) starting (note that we have redefined $\mathcal{B}_\alpha^\delta(\underline{\mathbf{c}}, \underline{\mathbf{v}})$) at the initial conditions in (6) satisfies $\lim_{t \rightarrow +\infty} \mathbf{c}(t) = \underline{\mathbf{c}}$ and $\lim_{t \rightarrow +\infty} \mathbf{v}(t) = \underline{\mathbf{v}}$.

In general, a single equilibrium point is *not* locally asymptotically stable, as the following example shows.

Example 3.6. Consider the simplest possible case where there are only two stations. Assume that $T_{ij} = T_{ji} := T$, and that $\lambda_1 < \lambda_2$. In this case, an assignment belonging to set \mathcal{A} is $\gamma_1 = \alpha_{12} = \lambda_2 - \lambda_1$ and $\gamma_2 = \alpha_{21} = 0$. Consider an arbitrary equilibrium $(\underline{\mathbf{c}}, \underline{\mathbf{v}}) \in \mathcal{E}_\alpha$. Assume that δ is small enough, so that $T_1 := c_1/(\mu_1 - \lambda_1) < T$, $T_2 := c_2/(\mu_2 - \lambda_2) < T$, $c_1 < \underline{v}_1$, $c_2 < \underline{v}_2$. Assume also that $c_1 \neq c_2$. Consider the case $T_1 < T_2$ (the other case is analogous). Then, model (1) with initial conditions as in (6) has the following evolution. In time interval $[0, T_1]$ one has

$$\begin{aligned} \dot{c}_1(t) &= \lambda_1 - \mu_1, & \dot{c}_2(t) &= \lambda_2 - \mu_2, \\ \dot{v}_1(t) &= \lambda_1 - \mu_1, & \dot{v}_2(t) &= \lambda_2 - \mu_2. \end{aligned}$$

Then, in time interval $[T_1, T_2]$ one has

$$\begin{aligned} \dot{c}_1(t) &= 0, & \dot{c}_2(t) &= \lambda_2 - \mu_2, \\ \dot{v}_1(t) &= 0, & \dot{v}_2(t) &= \lambda_2 - \mu_2. \end{aligned}$$

In time interval $[T_2, T]$ all derivatives are equal to zero. Then, in time interval $[T, T + T_1]$ one has

$$\begin{aligned} \dot{c}_1(t) &= 0, & \dot{c}_2(t) &= 0, \\ \dot{v}_1(t) &= \mu_2 - \lambda_2, & \dot{v}_2(t) &= \mu_1 - \lambda_1, \end{aligned}$$

and, finally, in time interval $[T + T_1, T + T_2]$

$$\begin{aligned} \dot{c}_1(t) &= 0, & \dot{c}_2(t) &= 0, \\ \dot{v}_1(t) &= \mu_2 - \lambda_2, & \dot{v}_2(t) &= 0. \end{aligned}$$

For $t > T + T_2$ all derivatives are identically zero. Hence, one easily obtains that $\lim_{t \rightarrow +\infty} \mathbf{c}(t) = \mathbf{0}$ and $\lim_{t \rightarrow +\infty} \mathbf{v}(t) = [v_1(0) - c_1(0) + c_2(0), v_1(0) - c_2(0) + c_1(0)] \neq [v_1(0), v_2(0)]$. Since δ is arbitrarily small, we conclude that, in general, single equilibrium points in \mathcal{E}_α are not locally asymptotically stable. •

The above example shows that, in general, single equilibrium points in \mathcal{E}_α are *not* locally asymptotically stable; this implies that, to make individual equilibrium points locally asymptotically stable, an additional feedback term is needed. Specifically, assume that $(\underline{\mathbf{0}}, \underline{\mathbf{v}}^d) \in \mathcal{E}_\alpha$ is a desired equilibrium point (e.g., $v_i^d = (V - V_\alpha)/n$ for each $i \in \mathcal{N}$). Then, one can make $(\underline{\mathbf{0}}, \underline{\mathbf{v}}^d)$ locally asymptotically stable by adding to the vehicles' dynamics the feedback term

$$-H(v_i(t) - v_i^d) + \sum_{j \neq i} \frac{1}{n-1} H(v_j(t - T_{ji}) - v_i^d). \quad (7)$$

That is, each station i sends away empty vehicles at a rate of 1 when the current number of vehicles $v_i(t)$ exceeds the desired number of vehicles v_i^d . This rate is in addition to the rebalancing vehicles, which are sent at rate $\gamma_i = \sum_j \alpha_{ij}$. The normalizing constant $1/(n-1)$ has the interpretation that the extra rebalancing vehicles are equally likely sent to the other stations. The next theorem shows that under model (1) with the additional feedback term (7) on vehicle's dynamics the equilibrium point $(\underline{\mathbf{0}}, \underline{\mathbf{v}}^d)$ is locally asymptotically stable.

Theorem 3.7 (Stability of single equilibrium points). *Let $\alpha \in \mathcal{A}$ be a feasible assignment, assume $V > V_\alpha$, and let $(\mathbf{0}, \mathbf{v}^d) \in \mathcal{E}_\alpha$ be a desired equilibrium point. Under model (1), with the additional feedback term (7) on the vehicles' dynamics, the equilibrium point $(\mathbf{0}, \mathbf{v}^d)$ is locally asymptotically stable.*

Proof. The proof is provided in the Appendix. □

4 Optimal Rebalancing

Our objective is to find a rebalancing assignment α that minimizes the number of empty vehicles traveling in the network and ensures 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} (provided that the total number of vehicles V is large enough). Hence, we are left with the problem of finding the rebalancing assignment in \mathcal{A} that minimizes the number of empty vehicles traveling in the network.

The time-average number of rebalancing vehicles is simply given by $\sum_{i,j} T_{ij} \alpha_{ij}$. Note that in minimizing this quantity, we are also minimizing the lower bound on the necessary number of vehicles \underline{V} from Corollary 3.4. Combining this objective with the existence of stable equilibria (i.e., the constraints in (3)), we see that α should be chosen as the solution to the following minimum cost flow problem (Korte and Vygen, 2007):

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

where $\mathcal{N} := \{1, \dots, n\}$. The two constraints ensure that the optimization is over the set \mathcal{A} .

From Lemma 3.3, this linear program is feasible, and thus an optimal solution α^* exists. The rebalancing policy is then given by sending empty vehicles from station i to station j at a rate of α_{ij}^* .

Remark 4.1 (Minimizing the vehicle utilization rate). *Our stated objective is to minimize the number of rebalancing vehicles traveling in the network. For a given rebalancing assignment in \mathcal{A} , the time-average number of rebalancing vehicles is given by $\sum_{i,j} T_{ij} \alpha_{ij}$. The time-average number of vehicles carrying customers is given by $\sum_{i,j} T_{ij} p_{ij} \lambda_i$ and thus the time-average number of vehicle trips (empty and customer carrying) is $V_\alpha := \sum_{i,j} T_{ij} (p_{ij} \lambda_i + \alpha_{ij})$. Therefore, in minimizing the number of rebalancing trips we also minimize the total number of vehicles on the road, or equivalently, the total vehicle utilization rate. •*

The above optimization gives optimal values for the α_{ij} 's. However, when the travel times T_{ij} satisfy the triangle inequality there is additional structure in the optimization that can be leveraged. Consider the flow of vehicles when no re-balancing is performed. We can divide the stations into two sets; the set S consisting of stations with a *surplus* of vehicles, and the set D consisting of stations with a *deficit* vehicles. The sets are given by

$$S := \left\{ i \in \mathcal{N} \mid \lambda_i < \sum_j \lambda_j p_{ji} \right\}, \tag{8}$$

$$D := \mathcal{N} \setminus S. \tag{9}$$

The following lemma characterizes the structure of the optimal solution when the travel times satisfy the triangle inequality.

Lemma 4.2 (Re-balancing structure). *Assume that the travel times satisfy the triangle inequality. In an optimal re-balancing solution, empty vehicles are sent from station i to station j (i.e., $\alpha_{ij}^* > 0$) only if $i \in S$ and $j \in D$.*

Proof. Let us begin by showing the following fact about an optimal assignment α^* . If for some $i, j \in \mathcal{N}$, we have $\alpha_{ij}^* > 0$, then $\alpha_{jk}^* = 0$ for every $k \in \mathcal{N}$. To prove this, suppose by way of contradiction, that there exist three stations a, b , and c , such that $\alpha_{ab}^*, \alpha_{bc}^* > 0$. Now, let $\epsilon := \min\{\alpha_{ab}^*, \alpha_{bc}^*\} > 0$, and let us define a new assignment $\bar{\alpha}$ such that $\bar{\alpha}_{ab} := \alpha_{ab}^* - \epsilon$, $\bar{\alpha}_{bc} := \alpha_{bc}^* - \epsilon$, $\bar{\alpha}_{ac} := \alpha_{ac}^* + \epsilon$, and $\bar{\alpha}_{ij} = \alpha_{ij}^*$ for all other entries. Note that since α^* is feasible, so is $\bar{\alpha}$. That is, it satisfies the constraints in (3). However, letting $C(\alpha) = \sum_{i,j} \alpha_{ij} T_{ij}$, we see that

$$\begin{aligned} C(\alpha^*) - C(\bar{\alpha}) &= T_{ab}(\alpha_{ab}^* - \bar{\alpha}_{ab}) + T_{bc}(\alpha_{bc}^* - \bar{\alpha}_{bc}) + T_{ac}(\alpha_{ac}^* - \bar{\alpha}_{ac}) \\ &= \epsilon(T_{ab} + T_{bc} - T_{ac}) > 0, \end{aligned}$$

where the final inequality comes from the fact that the travel times T_{ij} satisfy the triangle inequality. But, this implies that strictly fewer empty trips occur with $\bar{\alpha}$ than with α^* , a contradiction.

Now, consider the set $S := \{i \in \mathcal{N} \mid \lambda_i < \sum_j \lambda_j p_{ji}\}$, and a station $i \in S$. Since $i \in S$, an assignment α^* satisfies the constraint $\sum_{j \neq i} (\alpha_{ij}^* - \alpha_{ji}^*) = -\lambda_i + \sum_{j \neq i} \lambda_j p_{ji}$, only if $\alpha_{ij}^* > 0$ for some j . But, we have shown above that if $\alpha_{ij}^* > 0$, then $\alpha_{jk}^* = 0$ for every $k \in \mathcal{N}$. Thus, $j \notin S$, implying that $j \in D$. Hence we have that the stations are partitioned into two sets S and D , such that $\alpha_{ij}^* > 0$ only if $i \in S$ and $j \in D$. \square

The above lemma allows us to rephrase the optimization as follows:

$$\begin{aligned} &\text{minimize} && \sum_{i \in S, j \in D} T_{ij} \alpha_{ij} \\ &\text{subject to} && \sum_{j \in D} \alpha_{ij} = -\lambda_i + \sum_{j \neq i} \lambda_j p_{ji} && \forall i \in S \\ &&& \sum_{i \in S} \alpha_{ij} = \lambda_j - \sum_{i \neq j} \lambda_i p_{ij} && \forall j \in D \\ &&& \alpha_{i,j} \geq 0 && \forall i \in S, j \in D. \end{aligned} \tag{10}$$

Notice that the above optimization has fewer variables, and fewer constraints than the first optimization. We can visualize this new optimization as a continuous matching problem in a bipartite graph, with vertex sets S and D .

We conclude this section by discussing some important aspects of our problem setup.

4.1 Optimality with respect to Generalized Feedback Control Laws

In our model setup we have considered a restricted class of rebalancing control laws (specifically, *constant* rebalancing rates). However, at least when the travel times satisfy the triangle inequality, the solution that we derived is optimal within the general class of causal feedback control laws, in the sense that no other causal control law can achieve stability with a smaller average number of rebalancing trips. Indeed, without rebalancing, vehicles will accumulate at stations belonging to set S at an average rate $\sum_j (p_{ji} \lambda_j - \lambda_i)$. Hence, to ensure stability, every causal control law needs to

“pump” into the system empty vehicles at an average rate at least as large as $\sum_{i \in S} \sum_j (p_{ji} \lambda_j - \lambda_i)$. When the T_{ij} ’s satisfy the triangle inequality, this is exactly the rate at which empty vehicles are “pumped” into the system by the optimal rebalancing control law derived in this section (see the constraints in the optimization problem 10).

4.2 Provision of Door-to-Door Service

An interesting extension of a MOD system is to allow the robotic vehicles to pick up some of the customers directly at their travel origins (for example, their houses) and deliver them at their travel destinations (for example, their offices). This aspect could be modeled in our framework by assigning to each station i all locations in the city that are closer to i than any other station j . This forms a partition of the city into n dominance regions, one for each station. Then, a customer in dominance region i would be serviced in the following manner: 1) A vehicle travels from station i to customer i ’s location; 2) the vehicle drives the customer to its destination in region j ; and 3) the vehicle travels to the corresponding station j . The travel time T_{ij} would then represent the average time to complete the three legs of this trip. (In Section 6.3 we demonstrate the effect of uncertain travel times on the performance of the MOD system).

5 Adaptive Real-time Rebalancing Policy

Until now, the policies have required knowledge of the arrival rates λ_i , and the destination distribution p_{ij} . In this section we introduce a policy that does not require any *a priori* information. The idea is to repeatedly solve the optimization introduced in Section 4, but using the current distribution of customers and vehicles. Let us define the number of vehicles *owned* by a station i to be the number of vehicles at that station, given by $v_i(t)$, plus the number of vehicles in transit to the station, given by $\sum_j v_{ji}(t)$: $v_i^{\text{own}}(t) := v_i(t) + \sum_j v_{ji}(t)$. Note that by definition, $\sum_i v_i^{\text{own}}(t) = V$ at all times $t \geq 0$.

Now, if station i has $c_i(t)$ customers and $v_i(t)$ vehicles, then $\min\{c_i(t), v_i(t)\}$ vehicles will leave the station to serve the waiting customers. The excess vehicles at station i is then $v_i^{\text{excess}}(t) := v_i^{\text{own}}(t) - c_i(t)$. These are the vehicles that station i currently has available to send to other stations in need. Thus, the total number of excess vehicles in the system is $\sum_i v_i^{\text{excess}}(t) = V - \sum_i c_i(t)$. Note that this number may be negative. We would like to split these excess vehicles among the n stations according to some desired distribution. That is, for each station i we have a desired number of vehicles $v_i^d(t)$, such that $\sum_i v_i^d(t) \leq V - \sum_i c_i(t)$ for all t . As an example, we may choose that the excess vehicles are evenly split among the stations so that

$$v_i^d(t) = \left\lfloor \frac{V - \sum_j c_j(t)}{n} \right\rfloor \quad \text{for each station } i, \quad (11)$$

where we take the floor to obtain an integer. Our goal is to have $v_i^{\text{excess}}(t) \geq v_i^d(t)$ for all t .

Let us define an optimization horizon $t_{\text{hor}} > 0$. At time instants kt_{hor} , where k is a non-negative integer, we rebalance the excess vehicles by solving the following optimization:

$$\begin{aligned} & \text{minimize} && \sum_{i,j} T_{ij} \text{num}_{ij} \\ & \text{subject to} && v_i^{\text{excess}}(t) + \sum_{j \neq i} (\text{num}_{ji} - \text{num}_{ij}) \geq v_i^d(t) \quad \forall i \in \mathcal{N}, \\ & && \text{num}_{ij} \in \mathbb{N} \quad \forall i, j \in \mathcal{N}, i \neq j. \end{aligned}$$

Note that this is an integer linear program, where num_{ij} is the number of vehicles that station i will send to station j . It can be written in the form $\{\min cx \mid Ax \geq b, x \in \mathbb{N}\}$. However, it can easily be verified that the constraint matrix A is *totally unimodular* (Korte and Vygen, 2007). In addition, the vector b contains integer entries $v_i^{\text{excess}}(t) - v_i^d(t)$. Therefore, we can relax the integer constraints to $\text{num}_{ij} \geq 0$, and solve the problem as a linear program $\{\min cx \mid Ax \geq b, x \geq 0\}$. The resulting solution will necessarily have integer values, and thus will also be the optimal solution to the integer linear program.

Therefore, we can efficiently rebalance the system every t_{hor} time units without knowledge of λ_i or p_{ij} . Each time the optimization is solved, we simply send num_{ij} rebalancing vehicles from station i to station j . In the next section we will characterize the performance of this policy in simulation. For future work, we are looking at modelling customer arrivals as a stochastic process and then analyzing the theoretical performance of this policy.

6 Simulation Results

We have developed a simulation environment in MATLAB[®] for testing rebalancing policies. An example of a 12 station environment is shown in Figure 3. In this environment customers arrive *stochastically* at each station i according to a Poisson process with parameter λ_i . Each customer’s destination is sampled from the distribution p_{ij} . Since the evolution of the system is stochastic, we perform several trials, and then compute statistics in order to characterize a policy’s performance. For each policy, we solve the necessary linear programs using the freely available SeDuMi (Self-Dual-Minimization) toolbox. Simulations were run on a laptop computer with a 2.66 GHz dual core processor and 4 GB of RAM.

This section presents simulation results for an environment with 50 randomly distributed stations. The environment is shown in Figure 4 for 65 vehicles. The travel times T_{ij} between stations are given by the Euclidean distance. However, in Section 6.3 we investigate the effects of uncertainty in travel times on the system performance. The dimensions of the environment are 20 by 20 (dimensionless) units, and each vehicle moves 0.2 units per time step. The arrival rates at each station were randomly selected, as was the destination density p_{ij} . Using the necessary condition on the number of vehicles in Theorem 3.4, we obtain that $V > 58.8$ for any stabilizing policy. A simulation study for a smaller 12 station environment was presented in the preliminary work Pavone et al. (2011).

6.1 Real-time Rebalancing Policy

Figure 5 summarizes the performance of the real-time rebalancing policy of Section 5 with $v_i^d(t)$ as shown in (11). The left figure shows the number of waiting customers, and the number of in-transit vehicles as a function of the optimization horizon t_{hor} . The total number of vehicles is $V = 120$. Each data point is the mean of 20 independent trials, where each trial starts from an initial condition of equally distributed vehicles, and no customers, and runs for 5000 time steps. Error bars show the standard deviation over the 20 trials. We see that as the optimization horizon increases, the number of in-transit vehicles decreases, but the number of waiting customers increases. Thus, the optimal choice of t_{hor} is a trade-off between the cost of performing rebalancing trips, and the wait-time of customers.

The right figure shows the stability of the real-time rebalancing policy. Each data point shows the mean of 20 trials. In each trials, we start the system with 2000 customers, and run the system for 15,000 time steps. We then look at the difference between the initial number of customers, and the time-average present over the last 1000 time steps of the trial. From the simulation, we can see




Figure 3: The simulation environment for load balancing in mobility-on-demand systems. At each station, the grey bar shows the number of vehicles, and the blue bar shows the number of waiting customers. Blue vehicles are carrying customers, while grey vehicles are performing rebalancing trips. The first frame shows the initial condition, with 24 vehicles and 3 customers. The second and third frames show the system as it evolves.




Figure 4: The 50 station simulation environment. At each station, the grey bar shows the number of vehicles, and the blue bar shows the number of waiting customers. Blue vehicles are carrying customers, while grey vehicles are performing rebalancing trips.

that approximately 65 vehicles is the threshold for stability in this example. With $V \geq 65$ we have stability. This compares quite well with the necessary condition of $V > 58.8$.

6.2 Policy Comparison

Finally, we compare the performance of three policies: 1) the real-time rebalancing policy with $t_{\text{hor}} = 100$ and $v_i^d(t)$ as shown in (11); 2) the basic policy from the fluid model, where we send a constant rate of α_{ij} vehicles between stations; and 3) the fluid model policy with feedback discussed in Section 3.3.2. In this third policy we set the desired number of vehicles at each station to be $v_i^d = \lceil V/n \rceil$. It should be noted that the choice of v_i^d is not an equilibrium point in \mathcal{E}_α since $\sum_i v_i^d > V - V_\alpha$. However, in practice we obtain better performance with this choice than we do in choosing an equilibrium point such as $v_i^d = (V - V_\alpha)/n$. The reason for is that in the fluid policy with feedback, a station i sends away additional vehicles as soon as $v_i(t) > v_i^d$. When the arrival rate of customers is stochastic, it frequently occurs that there is a temporary increase in customer arrivals, and thus a station requires additional vehicles that it just sent away. By setting v_i^d to a value larger than the equilibrium, we see a reduction in both the number of rebalancing trips and in the number of waiting customers.

For each policy, Figure 6 (left) shows the number of waiting customers as a function of the total number of vehicles V , on a log scale. Figure 6 (right) shows the number of in-transit vehicles as a function of V . Each data point is the mean of 20 independent trials, and each trial consists of 5000 time steps. In each trial we compute mean number of waiting customers, and the mean number of in-transit vehicles over the last 2000 time steps of the trial. Note that a finite number of waiting customers does not necessarily indicate stability. We see from Figure 5 that at least 65 vehicles are needed for stability.

From the left figure we see that the real-time rebalancing policy has the best performance in




Figure 5: Analysis of the real-time rebalancing policy. Left figure: The number of customers or vehicles as a function of the optimization horizon t_{hor} . Right figure: Stability as a function of the number of vehicles V . Positive numbers indicate instability, while negative numbers indicate stability.

terms of number of waiting customers. The basic fluid policy performs quite poorly. This is due to the stochastic fluctuations in the customer arrival rates, and their chosen destinations. The fluid policy with feedback performs adequately, but consistently has over 50% more waiting customers than the real-time rebalancing policy.

From the right figure we see that the fluid model policy sends the fewest rebalancing vehicles. Thus, to minimize rebalance cost without regard to customer satisfaction (i.e., wait times), the fluid model policy performs best. The real-time rebalancing policy sends fewer vehicles than the fluid policy with feedback. Thus, the real-time policy outperforms the feedback policy in terms of waiting customers, and in-transit vehicles. Finally, the number of in-transit vehicles for the fluid policy with feedback temporarily drops when the number of vehicles exceeds 100. This occurs because the desired number of vehicles at a station was chosen as $v^d = \lceil V/n \rceil$. Since there are $n = 50$ stations, when the number of vehicles is increased from 100 to 101, the value v^d changes from 2 to 3. This shift dramatically affects the number of in-transit vehicles. However, we can see that it does not have as large of an effect on the number of waiting customers. Thus, one may be able to further improve the performance of the fluid policy with feedback by setting v^d to a value larger than $\lceil V/n \rceil$. We plan to explore this further in our future work.

6.3 Uncertainty in Travel Times

Finally, in this section we demonstrate the effects of nondeterministic travel times on the performance of the real-time rebalancing policy. In the above simulations, each vehicle moves a fixed (nominal) distance of $\text{step}_{\text{nom}} = 0.2$ in each time-step. Thus, the travel time T_{ij} is given by the distance from station i to station j , divided by step_{nom} . To simulate nondeterministic travel times, we add a uniform random variable with support on $[-a, a]$, denoted $U(-a, a)$, to each step of each vehicle:

$$\text{step}_{\text{rand}} = \text{step}_{\text{nom}} + U(-a, a).$$

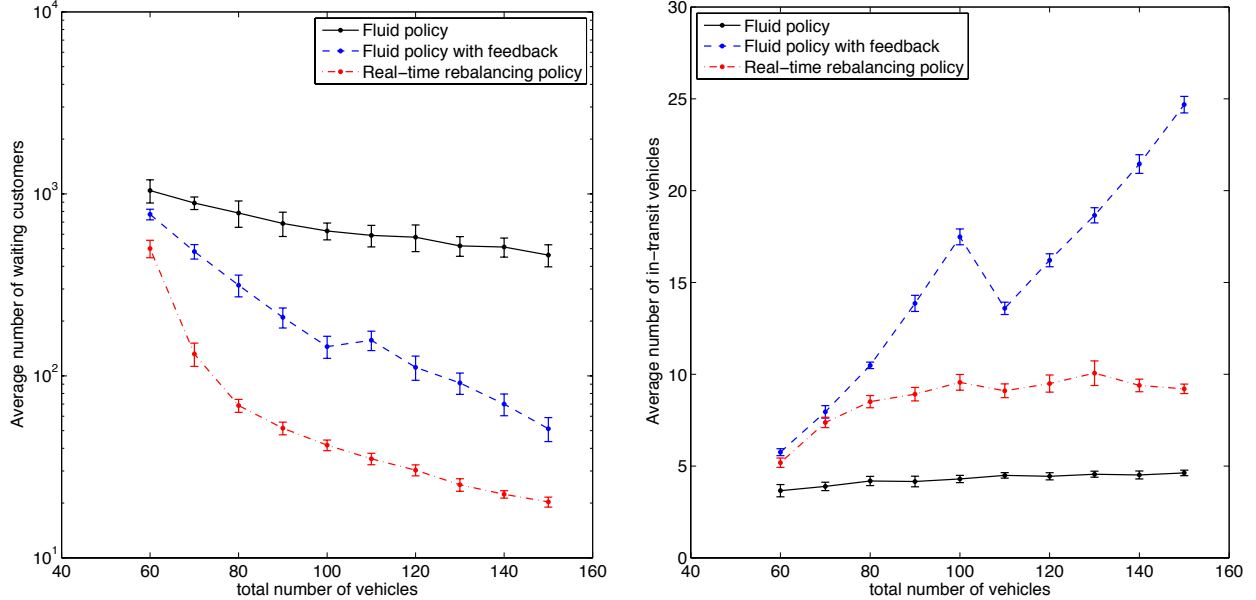


Figure 6: The performance of the three different policies. Left figure: The time-average number of customers waiting in the system. Note that the vertical axis is a log-scale. Right figure: The time-average number of vehicles performing rebalancing trips per unit of time.

The expected step size at each time-step remains equal to step_{nom} , but as a increases, the uncertainty in the travel times increases. However, the expected time to travel from station i to station j remains approximately equal to the deterministic time T_{ij} . By varying a we can see the effect of uncertain travel times on the performance of the policy. We consider large perturbations in which $a > \text{step}_{\text{nom}}$ to simulate large deviations from the expected travel time. Note that this has a non-intuitive physical interpretation that the vehicle may take a step in the opposite direction of its destination station. However, the travel time distribution remains centered around T_{ij} , and thus it is just used to increase the variance.

Figure 7 shows the average number of waiting customers and the average number of rebalancing vehicles as a function of a . The Monte Carlo simulation is performed for the same 50 station environment as in the previous sections. The number of vehicles is set to 120. Each data point is the mean of the 20 independent trials, and each trial consists of 5000 time steps. In each trial we compute the mean number of waiting customers and the mean number of in-transit vehicles over the last 2000 time steps of the trial.

From Figure 7 we can see that the effect of uncertain travel times on the average number of customers in the system is very small. We see a slight trend to higher wait-times (and large standard deviations), but the trend is not statistically significant even when the maximum perturbation a is 48 times the nominal step size. Even less change is noticed in the average number of rebalancing vehicles shown on the right side of Figure 7. Thus, as long as the mean travel times between stations are known accurately, uncertainty in the travel times does not appear to have a significant effect on the system performance. Of course, if the mean travel time is not accurately known, we would expect that the performance would degrade.




Figure 7: The performance of the real-time rebalancing policy as a function of uncertainty on travel-times. The nominal vehicle step size is 0.2, and a zero mean uniform random variable $U(-a, a)$ is added to the step size at each time step. The value of a is shown on the x -axis of the plot, ranging from 0 to 2.4. Left figure: The time-average number of customers waiting in the system. Right figure: The time-average number of vehicles performing rebalancing trips.

7 Hardware Experiments

We have also validated the real-time rebalancing policy through hardware experiments involving eight ground robots and four stations. The robots are iRobot iCreate platforms, each outfitted with a netbook computer (see Figure 8). The testing area consists of a $2\text{m} \times 2\text{m}$ square with a station at each corner and the surrounding area for parking queues (see Figure 8). Localization was performed by using a Vicon motion capture system, which allows for exact localization of the vehicles. Collision avoidance between vehicles were handled by using a first-in-first-out rule. Each ground robot communicates via WiFi with a ground station laptop, where the real-time rebalancing policy is run. Specifically, the real-time rebalancing policy has been implemented with $v_i^d(t)$ as shown in (11) and an horizon $t_{\text{hor}} = 100$ s.

We ran eight trials, each starting at the initial configuration shown in Figure 8. Each trial lasts approximately 17 minutes, and a robot can travel from one station to another in roughly two minutes. For each run, “virtual” customers arrive stochastically at each station $i \in \{1, 2, 3, 4\}$ according to a Poisson process with parameter λ_i . Both the arrival rates at each station and the destination density p_{ij} were randomly selected.

Figure 9 summarizes the performance of the rebalancing policy in the first run, where the total number of vehicles is *larger* than the minimum number of required vehicles (which, according to Theorem 3.4, is approximately equal to 4). Specifically, Figure 9 (left) shows the total number of waiting customers, while Figure 9 (right) shows the total number of vehicles waiting idle at the stations. One can note that the number of waiting customers stays bounded and is almost zero, which is in accordance with our theoretical results.

Figure 9 summarizes the performance of the rebalancing policy in the eighth run, where the total number of vehicles is *smaller* than the minimum number of required vehicles (which, according




Figure 8: View of hardware testbed with eight robots. Rectangles represent stations (four in total), while dashed lines represent “streets.” Each trial starts with two robots per station.

to Theorem 3.4, is approximately equal to 10). Specifically, Figure 10 (left) shows the total number of waiting customers, while Figure 10 (right) shows the total number of vehicles waiting idle at the stations. One can note that the number of waiting customers keeps growing, which is in accordance with our theoretical findings.

8 Conclusions

In this paper we studied the problem of rebalancing a mobility-on-demand system to efficiently transport customers in an urban environment. We proposed a robotic solution to rebalancing that involves empty robotic vehicles autonomously driving between stations. For a fluid model of the system, we showed that the optimal rebalancing policy can be found as the solution to a linear program. Under this policy, every station reaches an equilibrium in which there are excess vehicles and no waiting customers. We used this solution to develop a real-time rebalancing policy that can operate under stochastic customer demand. For future work we plan to analyze the stochastic queueing model and characterize the performance of the real-time rebalancing policy. We also plan to enrich our model by including uncertainty in the travel times, time windows for the customers, and capacity constraints for the roads. Finally, we are interested in using dynamic pricing to provide incentives for customers to perform rebalancing trips themselves.

Acknowledgements

The authors would like to acknowledge the contribution of Edward Mugica to the hardware experiments. This research was initiated while M. Pavone and S. L. Smith were at the Massachusetts




Figure 9: Performance of the rebalancing policy in run 1. Left figure: total number of waiting customers. Right figure: total number of vehicles waiting idle at the stations. The minimum number of vehicles required for stabilizability is $V \geq 4$. The number of waiting customers stays bounded.

Institute of Technology.

Funding

This research was supported by the Future Urban Mobility project of the Singapore-MIT Alliance for Research and Technology (SMART) Center, with funding from Singapore’s National Research Foundation; by the Office of Naval Research [grant number N000140911051]; and by the National Science Foundation [grant number EFRI0735953].

References

- R. Balakrishna, M. Ben-Akiva, and H. N. Koutsopoulos. Offline Calibration of Dynamic Traffic Assignment: Simultaneous Demand-and-Supply Estimation . *Transportation Research Record: Journal of the Transportation Research Board*, 2003(1):50–58, 2007.
- M. Barth and M. Todd. User behavior evaluation of an intelligent shared electric vehicle system. *Transportation Research Record: Journal of the Transportation Research Board*, 1760:145–152, 2001. URL <http://dx.doi.org/10.3141/1760-19>.
- G. Berbeglia, J. F. Cordeau, and G. Laporte. Dynamic pickup and delivery problems. *European Journal of Operational Research*, 202(1):8 – 15, 2010.
- BTS. Highlights of the 2001 National Household Travel Survey. Technical report, Bureau of Transportation Statistics, 2001.
- M. Buehler, K. Iagnemma, and S. Singh, editors. *The 2005 DARPA Grand Challenge: the Great Robot Race*. Number 36 in Springer Tracts in Advanced Robotics. Springer, 2007.
- M. Buehler, K. Iagnemma, and S. Singh, editors. *The DARPA Urban Challenge: Autonomous Vehicles in City Traffic*. Number 56 in Springer Tracts in Advanced Robotics. Springer, 2009.




Figure 10: Performance of the rebalancing policy in run 8. Left figure: total number of waiting customers. Right figure: total number of vehicles waiting idle at the stations. The minimum number of vehicles required for stabilizability is $V \geq 10$. The number of waiting customers grows without bound.

CAR2GO. CAR2GO Austin. Car Sharing 2.0: Great Idea for a Great City. Technical report, 2011.

V. Cardellini, M. Colajanni, and P. S. Yu. Dynamic load balancing on Web-server systems. *IEEE Internet Computing*, 3(3):28 – 39, 1999.

George Cybenko. Dynamic load balancing for distributed memory multiprocessors. *Journal of Parallel and Distributed Computing*, 7(2):279 – 301, 1989.

A. F. Filippov. *Differential Equations with Discontinuous Righthand Sides*, volume 18 of *Mathematics and its Applications*. Kluwer Academic Publishers, Dordrecht, The Netherlands, 1988.

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*, 37(6):pp. 893–901, 1989.

T. L. Friesz, D. Bernstein, T. E. Smith, R. L. Tobin, and B. W. Wie. A Variational Inequality Formulation of the Dynamic Network User Equilibrium Problem. *Operations Research*, 41(1): 179–191, 1993.

GM. EN-Vs Impress Media at Consumer Electronics Show. Technical report, General Motors Company, 2011.

G. Haddad. Monotone trajectories of differential inclusions and functional differential inclusions with memory. *Israel journal of mathematics*, 39(1-2):83 – 100, 1981.

B. Korte and J. Vygen. *Combinatorial Optimization: Theory and Algorithms*, volume 21 of *Algorithms and Combinatorics*. Springer, 4 edition, 2007.

M.-H. Massot, J.-F. Allouche, E. Bénéjam, and M. Parent. Praxitèle: Preliminary results from the Saint-Quentin station-car experiment. *Transportation Research Record: Journal of the Transportation Research Board*, 1666:125–132, 1999. URL <http://dx.doi.org/10.3141/1666-15>.

D. K. Merchant and G. L. Nemhauser. Optimality Conditions for a Dynamic Traffic Assignment Model. *Transportation Science*, 12(3):200–207, 1978.

- P. Midgley. The role of smart bike-sharing systems in urban mobility. Technical report, 2009.
- W. J. Mitchell, C. E. Borroni-Bird, and L. D. Burns. *Reinventing the Automobile: Personal Urban Mobility for the 21st Century*. The MIT Press, Cambridge, MA, 2010.
- S. N. Parragh, K. F. Doerner, and R. F. Hartl. A survey on pickup and delivery problems. *Journal fur Betriebswirtschaft*, 58(2):81 – 117, 2008.
- M. Pavone, K. Treleaven, and E. Frazzoli. Fundamental Performance Limits and Efficient Policies for Transportation-On-Demand Systems. In *IEEE Conf. Decision and Control*, pages 5622–5629, Atlanta, GA, 2010.
- M. Pavone, S. L. Smith, E. Frazzoli, and D. Rus. Load balancing for mobility-on-demand systems. In *Robotics: Science and Systems*, Los Angeles, CA, June 2011.
- S. Peeta and A. Ziliaskopoulos. Foundations of Dynamic Traffic Assignment: The Past, the Present and the Future. *Networks and Spatial Economics*, 1:233–265, 2001.
- UN. World Urbanization Prospects: The 2007 Revision Population Database. Technical report, United Nations, 2007.
- A. K. Ziliaskopoulos. A Linear Programming Model for the Single Destination System Optimum Dynamic Traffic Assignment Problem. *Transportation Science*, 34(1):37–49, 2000.

Appendix

Proof of Theorem 3.5. Consider an equilibrium $(\underline{\mathbf{c}}, \underline{\mathbf{v}}) \in \mathcal{E}_\alpha$ (note that $\underline{\mathbf{c}} = \mathbf{0}$ by Theorem 3.2). We now prove that every evolution of model (1) starting at

$$\begin{aligned}
c_i(\tau) &= 0 \text{ for } \tau \in [-\max_{i,j} T_{ij}, 0) \\
v_i(\tau) &= \underline{v}_i \text{ for } \tau \in [-\max_{i,j} T_{ij}, 0) \\
(\mathbf{c}(0), \mathbf{v}(0)) &\text{ such that } 0 \leq c_i(0) < v_i(0) \text{ for all } i \in \mathcal{N}, \text{ and } \sum v_i(0) = V - V_\alpha,
\end{aligned} \tag{12}$$

has a limit which belongs to the equilibrium set. The claim of the theorem will then be an easy consequence of this statement.

We start by observing the following fact. Assume that $v_i(\tau) > 0$ for all $\tau \in [-\max_{i,j} T_{ij}, t]$, then at time t the differential equations read $\dot{c}_i(t) = (\lambda_i - \mu_i)H(c_i(t))$, for all $i \in \mathcal{N}$; recalling that $-\lambda_i + \sum_{j \neq i} \lambda_j p_{ji} - \gamma_i + \sum_{j \neq i} \alpha_{ji} = 0$,

$$\begin{aligned}
\dot{v}_i(t) &= -\lambda_i + (\lambda_i - \mu_i)H(c_i) + \sum_{j \neq i} p_{ji} \left(\lambda_j - (\lambda_j - \mu_j)H(c_j^i) \right) - \gamma_i + \sum_{j \neq i} \alpha_{ji} \\
&= (\lambda_i - \mu_i)H(c_i) - \sum_{j \neq i} p_{ji} (\lambda_j - \mu_j)H(c_j^i) \\
&\geq (\lambda_i - \mu_i)H(c_i), \quad \text{for all } i \in \mathcal{N}.
\end{aligned}$$

Since $v_i(\tau) > 0$ for all $\tau \in [-\max_{i,j} T_{ij}, 0]$, and since $v_i(0) > c_i(0)$ for all $i \in \mathcal{N}$, we conclude that no $v_i(t)$ can reach the value 0 *before* the corresponding number of customers $c_i(t)$ has reached the value 0. However, once $c_i(t)$ reaches the value 0 (after a time interval $c_i(0)/(\mu_i - \lambda_i)$), the time

derivative $\dot{v}_i(t)$ is larger than or equal to zero. This implies that when the initial conditions satisfy (12), then $v_i(t) > 0$ for all $t \geq 0$.

Since $v_i(t) > 0$ for all $t \geq 0$, and since this implies that $\dot{c}_i(t) = (\lambda_i - \mu_i)H(c_i(t))$ for all $i \in \mathcal{N}$ and $t \geq 0$, we conclude that all $c_i(t)$ will be equal to zero for all $t \geq T' := \max_i c_i(0)/(\mu_i - \lambda_i)$. Then, for $t \geq T' + \max_{ij} T_{ij} =: T''$ the differential equations become: $\dot{c}_i(t) = 0$, $\dot{v}_i(t) = 0$.

Collecting the results obtained so far, we have that $\lim_{t \rightarrow +\infty} c_i(t) = 0$ for all $i \in \mathcal{N}$. Moreover, since $\dot{v}_i(t) = 0$ for all $t \geq T''$, the limit $\lim_{t \rightarrow +\infty} v_i(t)$ exists. Finally, one has $v_i(t) = v_i(0) + \int_0^t \dot{v}_i(\tau) d\tau \geq v_i(0) + \int_0^t \dot{c}_i(\tau) d\tau = v_i(0) + c_i(t) - c_i(0)$. Since $v_i(0) > c_i(0)$, we conclude that $\lim_{t \rightarrow +\infty} v_i(t) > 0$. Thus any solution with initial conditions (12) has a limit which belongs to \mathcal{E}_α (the property $\lim_{t \rightarrow +\infty} \sum v_i(t) = V - V_\alpha$ is guaranteed by the invariance property in Proposition 3.1 and the assumption $\sum v_i(0) = V - V_\alpha$).




Figure 11: The relation $0 \leq c_i < v_i$, and the definition of the radius ψ_i .

Let $\psi_i := \underline{v}_i \sin \frac{\pi}{4}$ (see Figure 11), and let $\psi_{\min} := \min_i \psi_i$. Then, from the definitions of ψ_i and ψ_{\min} , it follows that if one chooses $\delta = \psi_{\min}$, then any solution of model (1) with initial conditions satisfying (6) has a limit which belongs to the equilibrium set. This concludes the proof. \square

Proof of Theorem 3.7. The proof is similar to the proof of Theorem 3.5. Assume perturbations in the number of customers \underline{c} such that $0 \leq c_i(0) < v_i^d = v_i(0)$ for all $i \in \mathcal{N}$.

We start by observing the following fact. Assume that $v_i(\tau) > 0$ for all $\tau \in [-\max_{i,j} T_{ij}, t]$, then at time t the differential equations read $\dot{c}_i(t) = (\lambda_i - \mu_i)H(c_i(t))$, for all $i \in \mathcal{N}$; recalling that $-\lambda_i + \sum_{j \neq i} \lambda_j p_{ji} - \gamma_i + \sum_{j \neq i} \alpha_{ji} = 0$, one obtains

$$\begin{aligned} \dot{v}_i &= -\lambda_i + (\lambda_i - \mu_i)H(c_i) + \sum_{j \neq i} p_{ji} \left(\lambda_j - (\lambda_j - \mu_j)H(c_j^i) \right) - \gamma_i + \sum_{j \neq i} \alpha_{ji} \\ &\quad - H(v_i - v_i^d) + \sum_{j \neq i} \frac{1}{n-1} H(v_j^i - v_i^d) \\ &= (\lambda_i - \mu_i)H(c_i) - \sum_{j \neq i} p_{ji} (\lambda_j - \mu_j)H(c_j^i) - H(v_i - v_i^d) + \sum_{j \neq i} \frac{1}{n-1} H(v_j^i - v_i^d) \\ &\geq (\lambda_i - \mu_i)H(c_i) + \begin{cases} -1 & \text{if } v_i > v_i^d, \\ 0 & \text{if } v_i \leq v_i^d, \end{cases} \end{aligned}$$

for all $i \in \mathcal{N}$. Since $v_i(\tau) > 0$ for all $\tau \in [-\max_{i,j} T_{ij}, 0]$, and since $v_i(0) = v_i^d > c_i(0)$ for all $i \in \mathcal{N}$, we conclude that no $v_i(t)$ can reach the value 0 before the corresponding number of customers $c_i(t)$ has reached the value 0. However, once $c_i(t)$ reaches the value 0 (after a time interval $c_i(0)/(\mu_i - \lambda_i)$), the time derivative $\dot{v}_i(t)$ is larger than or equal to zero whenever $v_i(t) \leq v_i^d$. This implies that under the assumptions on the initial conditions, then $v_i(t) > 0$ for all $t \geq 0$.

Since $v_i(t) > 0$ for all $t \geq 0$, and since this implies that $\dot{c}_i(t) = (\lambda_i - \mu_i)H(c_i(t))$ for all $i \in \mathcal{N}$ and $t \geq 0$, we conclude that all $c_i(t)$ will be equal to zero for all $t \geq T' := \max_i c_i(0)/(\mu_i - \lambda_i)$. Then, for $t \geq T' + \max_{ij} T_{ij} =: T''$ the differential equations become:

$$\begin{aligned}\dot{c}_i(t) &= 0, \\ \dot{v}_i(t) &= -H(v_i(t) - v_i^d) + \sum_{j \neq i} \frac{1}{n-1} H(v_j(t - T_{ji}) - v_i^d).\end{aligned}$$

Note that $\dot{v}_i(t) < 0$ if $v_i(t) > v_i^d$ and $\dot{v}_i(t) \geq 0$ if $v_i(t) \leq v_i^d$. Hence, if at some instant $t \geq T''$ the number of vehicles satisfies $v_i(t) = v_i^d$, then $v_i(\tau) = v_i^d$ for all $\tau \geq t$, in other words v_i will “slide” along the mode v_i^d for ever. Let $\mathcal{U}(t)$ be the set of indexes of stations such that $v_i(t) > v_i^d$, and $\mathcal{D}(t)$ be the set of indexes of stations such that $v_i(t) < v_i^d$.

One can show that $\mathcal{U}(t) \neq \emptyset$ if and only if $\mathcal{D}(t) \neq \emptyset$, for all $t \geq T''$. Indeed, assume by the sake of contradiction that $\mathcal{U}(t) \neq \emptyset \not\Rightarrow \mathcal{D} \neq \emptyset$. Then, since $\mathcal{D}(t) = \emptyset$, and since $c_i = 0$ and $v_i > 0$ for all $i \in \mathcal{N}$ and for all $t \geq T'$, one can write for all $t \geq T''$: $V(t) = \sum_i^n v_i(t) + \sum_{i,j} v_{ij}(t) > \sum_i v_i^d + \sum_{ij} T_{ij}(p_{ij}\lambda_i + \alpha_{ij}) = V(0)$, which is in contradiction with the fact the the total number of vehicles is invariant. The proof for the converse implication is identical.

Now, if $\mathcal{U}(t) = \emptyset$, then one immediately obtains the claim of the theorem. Assume, instead, that $\mathcal{U}(t) \neq \emptyset$; from the previous discussion, one concludes that for each $i \in \mathcal{U}(t)$, $\dot{v}_i(t) \leq -1/(n-1)$. This implies that for $t \geq T'' + (n-1) \max_{i \in \mathcal{U}(T'')} v_i(T'')$ the set $\mathcal{U}(t)$ will be empty, and thus also set $\mathcal{D}(t)$ will be empty, implying that $v_i = v_i^d$ from then on. This concludes the proof. \square