Online Cost-Sharing Mechanism Design for Demand-Responsive Transport Systems*

Masabumi Furuhata
USC
Los Angeles, USA
k439bk439b@gmail.com

Kenny Daniel
Algorithmia Inc.
San Francisco, USA
kenny@algorithmia.io

Sven Koenig
USC
Los Angeles, USA
skoenig@usc.edu

Fernando Ordonez
Universidad de Chile
Santiago, Chile
fordon@dii.uchile.cl

Maged Dessouky
USC
Los Angeles, USA
maged@usc.edu

Marc-Etienne Brunet
McGill University
Montreal, Canada
marcetiennebrunet@gmail.com

Liron Cohen
USC
Los Angeles, USA
lironcoh@usc.edu

Xiaoqing Wang
USC
Los Angeles, USA
xiaoqinw@usc.edu

ABSTRACT

Demand-responsive transport (DRT) systems provide flexible transport services for passengers that request door-to-door rides in shared-ride mode without fixed routes and schedules. One has to design cost-sharing mechanisms for offering fare quotes to potential passengers so that all passengers are treated fairly. The main issue is how the operating costs of the DRT system should be shared among the passengers (given that different passengers cause different amounts of inconvenience to the other passengers), taking into account that DRT systems should provide fare quotes instantaneously without knowing future ride requests. We propose a novel cost-sharing mechanism, called Proportional Online Cost Sharing (POCS), that provides passengers with upper bounds on their fares immediately after their arrivals, allowing them to accept their fare quotes or drop out. We then demonstrate that POCS has attractive properties for both shuttle providers and passengers.

1. INTRODUCTION

Demand-responsive transport (DRT) systems provide flexible transport services where individual passengers request door-to-door rides by specifying their desired start and end locations. Multiple shuttles service these requests in shared-ride mode without fixed routes and schedules. DRT services are more flexible and convenient for passengers than buses since they do not operate on fixed routes and schedules, yet are cheaper than taxis due to the higher utilization of transport capacity. In the United States, DRT services are commonly used to service the transport needs of disabled and elderly citizens and have experienced rapid growth, for example, in the form of dial-a-ride paratransit services mandated under the Americans with Disabilities Act, while the National Transit Summaries and Trends report that typical DRT systems are highly subsidized.

In this paper, we propose a novel cost-sharing mechanism, called Proportional Online Cost Sharing (POCS), that provides passengers with upper bounds on their fares immediately after their arrivals, allowing them to accept their fare quotes or drop out. We then demonstrate that POCS has attractive properties for both shuttle providers and passengers. How passengers should share the operating cost in an online setting, where knowledge of future ride requests is missing, is a non-trivial problem for the following reasons: First, passengers do not submit their ride requests at the same time but should be given incentives to submit their ride requests as early as possible to allow the DRT systems more time to find routing solutions that can offer subsequent passengers lower fares due to synergies with the early ride requests, which might allow them to service more passengers. Second, passengers have different start and end locations and thus cause different amounts of inconvenience to the other passengers, which should be reflected in the fares. Finally, passengers should be quoted fares immediately after submitting their ride requests. This gives passengers certainty about the cost of service and allows the DRT system to plan routes better knowing which passengers have committed to participate. This requires DRT systems to make instantaneous and irreversible decisions despite having no knowledge of future ride requests [2].

2. ONLINE COST SHARING

In this section, we define the online cost-sharing problem for demand responsive transport (DRT) systems, provide an example, discuss existing cost-sharing mechanisms and some of their shortcomings, and finally derive a list of desirable properties for online cost-sharing mechanisms for DRT systems.

2.1 Problem Definition

DRT systems provide flexible transport services where individual passengers request door-to-door rides. Multiple shuttles service these requests without fixed routes and schedules. Passengers share shuttles. For example, after a passenger has been picked up and before it is dropped off,

*A longer version of this paper is under submission to the IEEE Transactions on Intelligent Transportation Systems. This paper was therefore submitted to this non-archival workshop purely for the exchange of ideas. Our research was supported by METRANS grant 09-19 via the California Department of Transportation and follow-up contract DTFH61-10-C-00030 from the Federal Highway Administration under the Broad Agency Announcement of Exploratory Advanced Research (EAR).

other passengers can be picked up and dropped off, resulting in a longer ride for the passenger. Passengers need to pay a share of the operating cost. Passengers arrive (that is, submit their ride requests) one after the other by specifying their desired start and end locations. The arrival time of a passenger is the time when it submits its ride request. In case the passenger decides to delay its arrival, we distinguish its truthful arrival time, which is its earliest possible arrival time, from its actual, perhaps delayed, arrival time. We assume, for simplicity, that all passengers arrive before the shuttles start to service the passengers. We also assume, without loss of generality, that exactly one passenger arrives at each time $k = 1, \ldots, t$, namely that passenger $\pi(k)$ arrives at time k under arrival order π, where an arrival order is a function that maps arrival times to passengers.

Definition 1. *For all times k and all arrival orders π with $1 \leq k$, the alpha value $\alpha_{\pi(k)}$ of passenger $\pi(k)$ quantifies the demand of its request, that is, how much of the transport resources it requests. We assume that it is positive and independent of the arrival time of the passenger.*

These assumptions are, for example, satisfied for the shortest point-to-point travel distance from the start location to the end location of a passenger, which is the quantity that we use in this paper as its alpha value.

Definition 2. *For all times t and all arrival orders π with $1 \leq t$, the total cost $totalcost^t_\pi$ at time t under arrival order π is the operating cost required to service passengers $\pi(1), \ldots, \pi(t)$. We define $totalcost^0_\pi := 0$ and assume that 1) the total cost is non-decreasing over time, that is, for all times t and t' and all arrival orders π with $t \leq t'$, $totalcost^t_\pi \leq totalcost^{t'}_\pi$; and 2) the total cost at time t is independent of the arrival order of passengers $\pi(1), \ldots, \pi(t)$, that is, for all times t and all arrival orders π and π' with $1 \leq t$ and $\{\pi(1), \ldots, \pi(t)\} = \{\pi'(1), \ldots, \pi'(t)\}$, $totalcost^t_\pi = totalcost^t_{\pi'}$.*

These assumptions are, for example, satisfied for the minimal operating cost, which is the quantity that we use in this paper for the total cost. The DRT system can accommodate advanced features, such as operating times and capacities of shuttles and time constraints of passengers, as long as it can determine total costs that satisfy the assumptions. The assumptions are typically not satisfied if passengers can arrive after the shuttles have started to service passengers since the shuttle locations influence the total cost. We initially assume for simplicity in the theoretical part of this paper that the DRT system can easily calculate the total cost at any given time.

Definition 3. *For all times k and all arrival orders π with $1 \leq k$, the marginal cost $mc_{\pi(k)}$ of passenger $\pi(k)$ under arrival order π is the increase in total cost due to its arrival, that is, $mc_{\pi(k)} := totalcost^k_\pi - totalcost^{k-1}_\pi$.*

Definition 4. *For all times k and t and all arrival orders π with $1 \leq k \leq t$, the shared cost $cost^t_{\pi(k)}$ of passenger $\pi(k)$ at time t under arrival order π is its share of the total cost at time t.*

The DRT system provides a (myopic) fare quote to a passenger immediately after its arrival. The fare quoted to passenger $\pi(k)$ immediately after its arrival at time k is $cost^k_{\pi(k)}$.

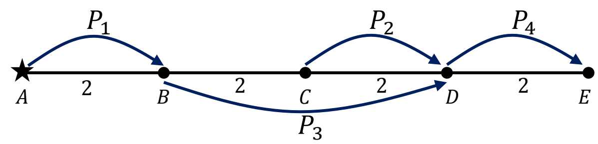

Figure 1: DRT Example 1

Table 1: DRT Values

		$k=1$ $\pi(k)=P_1$	$k=2$ $\pi(k)=P_2$	$k=3$ $\pi(k)=P_3$	$k=4$ $\pi(k)=P_4$
Alpha Value:	$\alpha_{\pi(k)}$	2	2	4	2
Total Cost:	$totalcost^k_\pi$	40	120	120	160
Marginal Cost:	$mc_{\pi(k)}$	40	80	0	40

(A fare quote of infinity means that the passenger cannot be serviced.)

Definition 5. *For all times k and all arrival orders π with $1 \leq k$, the fare limit $w_{\pi(k)}$ of passenger $\pi(k)$ is the maximum amount that it is willing to pay for its requested ride.*

Passenger $\pi(k)$ drops out and is not serviced if its fare limit $w_{\pi(k)}$ is lower than its fare quote, that is, $w_{\pi(k)} < cost^k_{\pi(k)}$. In this case, the DRT system simply pretends that the passenger never arrived, which explains why we assume, without loss of generality, that all passengers accept their fare quotes. When the passenger accepts its fare quote and is serviced, its fare is $cost^t_{\pi(k)}$ (which is not guaranteed to equal its fare quote).

2.2 Demand-Responsive Transport Example

We use the DRT example in Figure 1 to illustrate typical cost-sharing mechanisms. There is one shuttle that can transport up to four passengers and starts at the star. The shuttle incurs an operating cost of 10 for each unit of distance traveled and needs to return to its initial location. There are four passengers with arrival order $\pi(1) = P_1$, $\pi(2) = P_2$, $\pi(3) = P_3$ and $\pi(4) = P_4$. For example, Passenger P_3 requests a ride from location B to location D, as shown in Figure 1. All passengers accept all fare quotes. Table 1 shows the alpha value of each passenger, the total cost after the arrival of each passenger and the marginal cost of each passenger. For example, the alpha value of Passenger P_3 is the shortest point-to-point travel distance from its start location B to its end location D. Thus, $\alpha_{\pi(3)} = 4$. The total cost at time 3, after the arrival of Passenger P_3, is 10 times the minimal travel distance of the shuttle required to service Passengers P_1, P_2 and P_3 and return to its initial locations. Thus, $totalcost^3_\pi = 120$ since the shuttle has to drive from location A (to pick up Passenger P_1) via location B (to drop off Passenger P_1 and pick up Passenger P_3) and location C (to pick up Passenger P_2) to location D (to drop off Passengers P_2 and P_3) and to return to its initial location A. The marginal cost of Passenger P_3 is the increase in total cost due to its arrival. Thus, $mc_{\pi(3)} = totalcost^3_\pi - totalcost^2_\pi = 120 - 120 = 0$ since the total cost remains 120.

2.3 Typical Cost-Sharing Mechanisms

Online cost-sharing mechanisms determine the shared costs in an online setting, where knowledge of future arrivals

Table 2: Proportional Cost Sharing: $cost^t_{\pi(k)}$

	$k=1$ $\pi(k)=P_1$	$k=2$ $\pi(k)=P_2$	$k=3$ $\pi(k)=P_3$	$k=4$ $\pi(k)=P_4$
$t=1$	40			
$t=2$	60	60		
$t=3$	30	30	60	
$t=4$	32	32	64	32

of passengers is missing. We present typical cost-sharing mechanisms and some of their shortcomings in an online setting using the DRT example in Section 2.2.

2.3.1 Proportional Cost Sharing

One commonly used cost-sharing mechanism is *proportional cost sharing* [16, 14], where the total cost is distributed among all passengers proportionally to their alpha values, which reflects that passengers with higher demands should contribute more toward the total cost. Consequently, for all times k and t and all arrival orders π with $1 \leq k \leq t$, the shared cost of passenger $\pi(k)$ at time t under arrival order π is

$$cost^t_{\pi(k)} := totalcost^t_\pi \frac{\alpha_{\pi(k)}}{\sum_{j=1}^{t} \alpha_{\pi(j)}}.$$

Instead of distributing the total (operating) cost among all passengers, one could also distribute the operating cost of each shuttle among all passengers serviced by that shuttle, which results in identical properties for the DRT example in Section 2.2 since there is only one shuttle in the DRT example.

Table 2 shows the shared costs for the DRT example. For example, the total cost at time 3 is 120. It is distributed among all passengers that have arrived by time 3, namely Passengers P_1, P_2 and P_3, proportionally to their alpha values, namely 2, 2 and 4, respectively. Consequently, the shared cost of Passenger P_3 at time 3 and thus the fare quoted to Passenger P_3 after its arrival is $cost^3_{\pi(3)} = 60$. Similarly, the total cost at time 4 is 160. It is distributed among all passengers that have arrived at time 4, namely Passengers P_1, P_2, P_3 and P_4, proportionally to their alpha values, namely 2, 2, 4 and 2, respectively. Consequently, the shared cost of Passenger P_3 at time 4 and thus its fare is $cost^4_{\pi(3)} = 64$, implying that its fare is higher than its fare quote at time 3. This is undesirable because Passenger P_3 might accept the fare quote but not the higher fare, meaning that it will have to drop out shortly before receiving its ride and then needs to search for a last-minute alternative to using the DRT system, which might be pricy and is not guaranteed to exist. Thus, we suggest that a fare quote should be an upper bound on the fare (*immediate-response property*). We also suggest that the upper bound should be reasonably low since passengers might otherwise look for alternatives to using the DRT system, commit to one and then drop out unnecessarily. Obtaining reasonably low upper bounds can be difficult since the DRT system has no knowledge of future arrivals of passengers.

2.3.2 Incremental Cost Sharing

Another commonly used cost-sharing mechanism is *incremental cost sharing* [9], where the shared cost of each passenger is its marginal cost, which is the increase in total cost due to its arrival. Consequently, for all times k and t and all arrival orders π with $1 \leq k \leq t$, the shared cost of passenger $\pi(k)$ at time t under arrival order π is

$$cost^t_{\pi(k)} := mc_{\pi(k)}.$$

Table 3 (left) shows the shared costs for the DRT example in Section 2.2. For example, the marginal cost of Passenger P_3 is 0. Consequently, the shared cost of Passenger P_3 from its arrival at time 3 on is 0, and thus both its fare quote and fare are 0 as well. In general, incremental cost sharing satisfies the immediate-response property since the marginal costs are independent of time. The fares of Passengers P_1, P_2, P_3 and P_4 are 40, 80, 0 and 40, respectively. Thus, Passenger P_3 is a free rider, which is undesirable in general and especially in the context of the DRT example since Passenger P_3 has the highest demand, which should be reflected in the fares. Proportional cost sharing does not suffer from this problem. For the discussion below, notice that the fare per alpha value of Passenger P_1 is 20 and the one of Passenger P_3 is 0 even though Passenger P_1 arrives before Passenger P_3.

Table 3 (right) shows the shared costs for the DRT example in Section 2.2 if Passenger P_1 delays its arrival and the passengers arrive in order P_2, P_1, P_3 and P_4. Now, the shared cost of Passenger P_1 from its arrival at time 2 on is 0, and thus both its fare quote and fare are 0 as well. Thus, Passenger P_1 can reduce its fare from 40 to 0 by strategically delaying its arrival. This delay is undesirable because synergies with the early ride requests allow the DRT system to offer low fare quotes to new passengers. We therefore suggest to ensure that the best strategy of every passenger is to arrive truthfully (that is, as early as possible) because it cannot decrease its fare by delaying its arrival (*incentive-compatibility property*). Incremental cost sharing does not satisfy this property as shown above. Similarly, under incremental cost sharing, Passenger P_1 and P_2 prefers to pay the fare of Passenger P_3 rather than their own fare because Passenger P_3 enjoys a free ride due to payments of these two passengers. We therefore suggest that the fares per alpha value of passengers are never higher than those of passengers that arrive after them (*online-fairness property*).

2.4 Desirable Properties

None of the cost-sharing mechanisms discussed so far are well-suited for the DRT problem. Based on their shortcomings, we derive a list of desirable properties for online cost-sharing mechanism. Our primary objective is to design an online cost-sharing mechanism that provides incentives for passengers to arrive truthfully while satisfying basic properties of cost-sharing mechanism in general, such as fairness and budget balance.

Online Fairness: The shared costs per alpha value of passengers are never higher than those of passengers who arrive after them, that is, for all times k_1, k_2 and t and all arrival orders π with $1 \leq k_1 \leq k_2 \leq t$,

$$\frac{cost^t_{\pi(k_1)}}{\alpha_{\pi(k_1)}} \leq \frac{cost^t_{\pi(k_2)}}{\alpha_{\pi(k_2)}}.$$

Immediate Response: Passengers are provided immediately after their arrivals with (ideally low) upper bounds on their shared costs at any future time, that is, for all times k, t_1 and t_2 and all arrival orders π with $1 \leq k \leq t_1 \leq t_2$,

$$cost^{t_1}_{\pi(k)} \geq cost^{t_2}_{\pi(k)}.$$

Individual Rationality: The shared costs of passengers who accepted their fare quotes never exceed their fare limits

Table 3: Incremental Cost Sharing: $cost^t_{\pi(k)}$

	Truthful Arrival				Delayed Arrival			
	$k=1$ $\pi(k)=P_1$	$k=2$ $\pi(k)=P_2$	$k=3$ $\pi(k)=P_3$	$k=4$ $\pi(k)=P_4$	$k=1$ $\pi(k)=P_2$	$k=2$ $\pi(k)=P_1$	$k=3$ $\pi(k)=P_3$	$k=4$ $\pi(k)=P_4$
$t=1$	40				120			
$t=2$	40	80			120	0		
$t=3$	40	80	0		120	0	0	
$t=4$	40	80	0	40	120	0	0	40

at any future time, that is, for all times k and t and all arrival orders π with $1 \leq k \leq t$,

$$cost^t_{\pi(k)} \leq w_{\pi(k)}.$$

Budget Balance: The total cost equals the sum of the shared costs of all passengers, that is, for all times t and all arrival orders π with $1 \leq t$,

$$\sum_{j=1}^{t} cost^t_{\pi(j)} = totalcost^t_{\pi}.$$

Ex-Post Incentive Compatibility:[1] The best strategy of every passenger is to arrive truthfully, provided that all other passengers arrive truthfully as well and do not change whether they accept their fare quotes or drop out, because it then cannot decrease its shared cost by delaying its arrival, that is, for all times k_1, k_2 and t and all arrival orders π and π' with $1 \leq k_1 < k_2 \leq t$ and

$$\begin{aligned} \pi'(k) &= \begin{cases} \pi(k+1) & \text{if } k_1 \leq k < k_2 \\ \pi(k_1) & \text{if } k = k_2 \\ \pi(k) & \text{otherwise,} \end{cases} \\ cost^t_{\pi(k_1)} &\leq cost^t_{\pi'(k_2)}. \end{aligned}$$

The online fairness and ex-post incentive-compatibility properties are similar but one does not imply the other. Basically, they provide incentives for passengers to arrive truthfully. Thus, the DRT systems have more time to prepare and might also be able to offer subsequent passengers lower fares due to synergies with the early ride requests, which might allow them to service more passengers. The online-fairness property is also meant to ensure that passengers consider the fares to be fair. The immediate-response property enables DRT systems to provide fare quotes, in form of upper bounds on the fares, to passengers immediately after their arrivals despite missing knowledge of future arrivals of passengers. Thus, passengers have no uncertainty about whether they can be serviced or how high their fares will be, while the DRT systems reduce their uncertainty about passengers dropping out and can thus prepare better. Yet, the DRT system still retains some flexibility to optimize the routes and schedules after future arrivals of passengers. The budget-balance property guarantees that the sum of the fares of all passengers always equals the total cost. Thus, no profit is made and no subsidies are required.

We stated sufficient rather than necessary conditions for the properties. For example, the budget-balance property could be weakened to state that the total cost equals the sum of the shared costs of all passengers after the arrival of the last passenger. Requiring the properties to be satisfied at any time rather than only after the arrival of the last passenger simplifies the development of the online cost-sharing mechanism since they do not know in advance when the last passenger arrives.

[1]We would like the ex-post incentive-compatibility property ideally to state that the best strategy of every passenger is to arrive truthfully because it cannot decrease its shared cost by delaying its arrival. However, we impose two conditions in this paper that we hope to be able to relax in the future. The first condition is that all other passengers arrive truthfully, which, for example, rules out collusion of several passengers. In general, the literature on online-mechanism design [12] distinguishes two types of incentive compatibility, namely *dominant-strategy incentive compatibility* and *ex-post incentive compatibility*. Dominant-strategy incentive compatibility does not require the first condition, while ex-post incentive compatibility does. Dominant-strategy incentive compatibility is difficult to achieve in an online setting [12], which is why we impose the first condition in this paper. The second condition is that the other passengers do not change whether they accept their fare quotes or drop out, even though, for example, the delayed arrival of a passenger could cause the fare quotes of subsequent passengers to increase, which might make them drop out. The arrival orders with and without the delayed arrival of the passenger are then difficult to relate, which is why we impose the second condition in this paper.

3. POCS

In this section, we describe a novel online cost-sharing mechanism, called Proportional Online Cost Sharing (POCS), which satisfies the properties listed in Section 2.4, as proved in the technical report [4]. The idea behind POCS is the following: POCS partitions passengers into coalitions, where coalitions contain all passengers that arrive within given time intervals (rather than, for example, all passengers served by the same shuttle). Initially, each newly arriving passenger forms its own coalition. However, passengers can choose to form coalitions with passengers that arrive directly after them to decrease their shared costs per alpha value, which implies the online fairness, immediate response, and ex-post incentive-compatibility properties. For example, the immediate-response property is satisfied because passengers add other passengers to their coalitions only when this decreases their shared costs per alpha value and thus also their shared costs (since the alpha values are positive).

3.1 Calculation of Shared Costs

We now describe how POCS calculates the shared costs.

Definition 6. *For all times k_1, k_2 and t and all arrival orders π with $k_1 \leq k_2 \leq t$, the coalition cost per alpha value of passengers $\pi(k_1), \ldots, \pi(k_2)$ at time t under arrival order π is*

$$ccpa_{\pi(k_1,k_2)} := \frac{\sum_{j=k_1}^{k_2} mc_{\pi(j)}}{\sum_{j=k_1}^{k_2} \alpha_{\pi(j)}}.$$

Definition 7. *For all times k and t and all arrival orders π with $k \leq t$, the shared cost of passenger $\pi(k)$ at time t*

under arrival order π *is*

$$cost^t_{\pi(k)} := \alpha_{\pi(k)} \min_{k \le j \le t} \max_{1 \le i \le j} ccpa_{\pi(i,j)}.$$

3.2 Other Cost-Sharing Mechanisms

The following definition and lemma, whose proof is provided in the technical report [4], helps to understand the similarities between POCS and other cost-sharing mechanisms. It states that the shared costs per alpha value of all passengers in any coalition are always identical and equal to the coalition cost per alpha value of the coalition.

DEFINITION 8. *For all times* k_1, k_2 *and* t *and all arrival orders* π *with* $k_1 \le k_2 \le t$, *a coalition* (k_1, k_2) *at time* t *is a group of passengers* $\pi(k_1), \dots, \pi(k_2)$ *with*

$$\frac{cost^t_{\pi(k)}}{\alpha_{\pi(k)}} = \frac{cost^t_{\pi(k_1)}}{\alpha_{\pi(k_1)}}$$

for all times k *with* $k_1 \le k \le k_2$ *and the preceeding equality not holding for all times* k *with* $(k = k_1 - 1$ *or* $k = k_2 + 1)$ *and* $1 \le k \le t$.

LEMMA 1. *The shared cost per alpha value of any passenger in any coalition at any time equals the coalition cost per alpha value of the coalition, that is, for all times* k_1, k, k_2 *and* t *and all arrival orders* π *with* $1 \le k_1 \le k \le k_2 \le t$ *such that* (k_1, k_2) *is a coalition at time* t,

$$\frac{cost^t_{\pi(k)}}{\alpha_{\pi(k)}} = ccpa_{\pi(k_1,k_2)}.$$

Lemma 1 implies that POCS is a combination of proportional and incremental cost sharing. The sum of the marginal costs of all passengers in any coalition ("the total cost of all passengers in the coalition") at time t is distributed among all passengers in the coalition proportionally to their alpha values since, for all times k_1, k, k_2 and t and all arrival orders π with $k_1 \le k \le k_2 \le t$ such that (k_1, k_2) is a coalition at time t,

$$\begin{aligned} cost^t_{\pi(k)} &\overset{Lem.1}{=} \alpha_{\pi(k)} ccpa_{\pi(k_1,k_2)} \\ &\overset{Def.6}{=} \alpha_{\pi(k)} \frac{\sum_{j=k_1}^{k_2} mc_{\pi(j)}}{\sum_{j=k_1}^{k_2} \alpha_{\pi(j)}} \\ &= \left(\sum_{j=k_1}^{k_2} mc_{\pi(j)} \right) \frac{\alpha_{\pi(k)}}{\sum_{j=k_1}^{k_2} \alpha_{\pi(j)}}, \end{aligned}$$

which is similar to proportional cost sharing where the total cost (of all passengers) is distributed among all passengers proportionally to their alpha values.

The sum of the shared costs of all passengers in any coalition ("the shared cost of the coalition") at time t equals the sum of the marginal costs of all passengers in the coalition ("the marginal cost of the coalition") at the same time since, for all times k_1, k_2 and t and all arrival orders π with $k_1 \le k_2 \le t$ such that (k_1, k_2) is a coalition at time t,

$$\begin{aligned} \sum_{j=k_1}^{k_2} cost^t_{\pi(j)} &\overset{Lem.1}{=} ccpa_{\pi(k_1,k_2)} \sum_{j=k_1}^{k_2} \alpha_{\pi(j)} \\ &\overset{Def.6}{=} \frac{\sum_{j=k_1}^{k_2} mc_{\pi(j)}}{\sum_{j=k_1}^{k_2} \alpha_{\pi(j)}} \sum_{j=k_1}^{k_2} \alpha_{\pi(j)} \\ &= \sum_{j=k_1}^{k_2} mc_{\pi(j)}, \end{aligned}$$

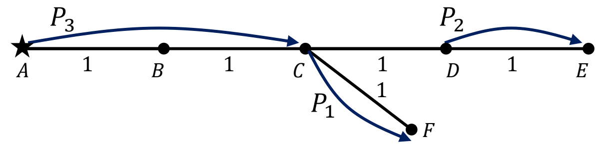

Figure 2: DRT Example 2

Table 4: POCS: $ccpa_{\pi(k_1,k_2)}$

	$k_2 = 1$ $\pi(k_2) = P_1$	$k_2 = 2$ $\pi(k_2) = P_2$	$k_2 = 3$ $\pi(k_2) = P_3$	$k_2 = 4$ $\pi(k_2) = P_4$
$k_1 = 1$ $\pi(k_1) = P_1$	20	30	15	16
$k_1 = 2$ $\pi(k_1) = P_2$		40	13 1/3	15
$k_1 = 3$ $\pi(k_1) = P_3$			0	6 2/3
$k_1 = 4$ $\pi(k_1) = P_4$				20

which is similar to incremental cost sharing. where the shared cost of a passenger is its marginal cost. It also implies the budget-balance property since summing over all passengers in all coalitions is identical to summing over all passengers and the sum of the marginal costs of all passengers equals the total cost.

3.3 Illustration

Table 4 shows the coalition costs per alpha value for the DRT example in Section 2.2. The coalition costs per alpha value are used to calculate the shared costs, shown in Table 5. The shared costs, in turn, are used to calculate the shared costs per alpha value, shown in Table 6, by dividing the shared costs by the alpha values, shown in Table 1. For example, at time 4, Passengers P_1, P_2 and P_3 form a coalition (since their shared costs per alpha value are equal), and Passenger P_4 forms a coalition by itself. The sum of the marginal costs of the three passengers in the first coalition ("the total cost of all passengers in the coalition") is 120 and is distributed among all passengers in the coalition proportionally to their alpha values, namely 2, 2 and 4, respectively. Consequently, the shared cost of Passenger P_3 at time 4 and thus its fare is $cost^4_{\pi(3)} = 60$. Table 6 shows that the shared costs per alpha value in each row are monotonically non-decreasing from left to right, corresponding to the online-fairness property. Table 5 shows that the shared costs in each column are monotonically non-increasing from top to bottom (and consequently Table 6 shows that the shared costs per alpha value have the same property), corresponding to the immediate-response property. Table 5 also shows that the sum of the shared costs in each row equals the total cost at the corresponding time, corresponding to the budget-balance property.

3.4 Ex-Post Incentive Compatibility

We use the DRT example in Figure 2 to illustrate that POCS does not satisfy the ex-post incentive-compatibility property if the second condition (namely that the other passengers do not change whether they accept their fare quotes or drop out) is removed. There is one shuttle that can transport up to four passengers and starts at the star. The shuttle incurs an operating cost of 10 for each unit of distance traveled and needs to return to its initial location. There are three passengers. Passengers P_1 and P_3 accept all fare quotes, while Passenger P_2 accepts all fare quotes up to 60. Assume that the passengers arrive in order P_1, P_2 and P_3. First, Passenger P_1 arrives, receives a fare quote of 60 and

Table 5: POCS: $cost^t_{\pi(k)}$

	$k=1$ $\pi(k)=P_1$	$k=2$ $\pi(k)=P_2$	$k=3$ $\pi(k)=P_3$	$k=4$ $\pi(k)=P_4$
$t=1$	40			
$t=2$	40	80		
$t=3$	30	30	60	
$t=4$	30	30	60	40

Table 6: POCS: $cost^t_{\pi(k)}/\alpha_{\pi(k)}$

	$k=1$ $\pi(k)=P_1$	$k=2$ $\pi(k)=P_2$	$k=3$ $\pi(k)=P_3$	$k=4$ $\pi(k)=P_4$
$t=1$	20			
$t=2$	20	40		
$t=3$	15	15	15	
$t=4$	15	15	15	20

accepts it. Second, Passenger P_2 arrives, receives a fare quote of 50 and accepts it. Third, Passenger P_3 arrives, receives a fare quote of 50 and accepts it. In the end, Passengers P_1, P_2 and P_3 are serviced with fares of 25, 25 and 50, respectively. Now assume that Passenger P_1 delays its arrival, and the passengers arrive in order P_2, P_3 and P_1. First, Passenger P_2 arrives, receives a fare quote of 80 and drops out since the fare quote exceeds its fare limit of 60. Second, Passenger P_3 arrives, receives a fare quote of 40 and accepts it. Third, Passenger P_1 arrives, receives a fare quote of 20 and accepts it. In the end, Passengers P_1 and P_3 are serviced with fares of 20 and 40, respectively. Thus, Passenger P_1 managed to decrease both its fare quote and fare by delaying its arrival since this caused Passenger P_2 to drop out.

4. EXPERIMENTAL ANALYSIS

We have proved that POCS satisfies five properties that make DRT systems more attractive to both shuttle providers and passengers, provided that our assumptions are satisfied. For example, Definition 2 assumes that the total cost satisfies two properties that hold for the minimal operating cost, which is therefore the quantity that we have used so far for the total cost. Calculating the minimal operating cost is typically an NP-hard problem and thus time-consuming. However, DRT systems need to calculate the minimal operating cost every time a ride request is submitted, which would prevent them from operating in real-time. We thus present an experimental study with a transport simulation where the DRT system uses a heuristic to compute a low operating cost that is not guaranteed to be minimal [10]. In this case, the assumption in Definition 2 that the total cost is independent of the arrival order of passengers (which implies that the decisions of passengers to accept their fare quotes or drop out and thus also their fare quotes themselves do not depend on the arrival order of passengers) is not satisfied. This assumption is used (only) to prove that POCS satisfies the ex-post incentive-compatibility property. We thus investigate whether the best strategy of every passenger remains to arrive truthfully, for example because the likelihood of transport capacity still being available tends to decrease over time.

4.1 Transport Simulator

Our transport simulator first generates a given number of shuttles and passengers. Each shuttle is characterized by its capacity, start location, end location, operating time window and operating cost for each unit of distance traveled. Each passenger is characterized by its truthful arrival time, start location, end location, pick-up time window, drop-off time window and fare limit. The settings of our simulator are slightly more general than what we have used in the DRT examples because operating time windows of shuttles and pick-up and drop-off time windows of passengers are taken into account. The transport simulator then simulates each passenger. Once a passenger is assigned to a shuttle, it is never re-assigned to a different shuttle, which makes it possible to calculate the marginal cost of a passenger as the lowest operating cost increase of adding the passenger to any shuttle, but is also a reason why the total cost (which equals the sum of the operating costs of all shuttles) is not guaranteed to be equal to the minimal operating cost or to be independent of the arrival order of the passengers. When a new passenger submits a ride request, the transport simulator requests from each shuttle the operating cost increase from adding the passenger to all passengers previously assigned to it, selects a shuttle with the lowest operating cost increase and then uses POCS to calculate a fare quote for the passenger under the assumption that the passenger is assigned to the selected shuttle. If the fare limit of the passenger is lower than this fare quote, then the passenger drops out, and the transport simulator does not service it. Otherwise, the passenger accepts the fare quote, and the transport simulator adds it to all passengers previously assigned to the selected shuttle and then updates the shared costs of all passengers assigned to the shuttles.

Each shuttle has to calculate its route, schedule and operating cost increase (or, equivalently, operating cost) when adding a new passenger to all passengers previously assigned to it. The shuttle maintains an itinerary for all passengers assigned to it - in the form of a sequence of locations, namely its start location, its end location and the start and end locations of all passengers assigned to it. It calculates its travel distance as the shortest travel distance needed to visit all locations in the order given in its itinerary, and it calculates its operating cost as the product of its travel distance and its operating cost for each unit of distance traveled. Determining an itinerary for the new passenger and all passengers previously assigned to it that minimizes its operating cost is time-consuming. The shuttle therefore uses a non-optimal scheduling method [17, 11], which is another reason why the total cost is not guaranteed to equal the minimal operating cost and not guaranteed to be independent of the arrival order of passengers. In the construction phase of the scheduling method, the shuttle uses a cheapest-insertion method to construct a (feasible) itinerary by inserting the start and end locations of the new passenger into the cached itinerary for the passengers previously assigned to it. In the subsequent improvement phase of the scheduling method, the shuttle uses tabu search [7, 13, 5, 6], a form of hill climbing, to improve the itinerary from the construction phase.

4.2 Experiment 1

In Experiment 1, we demonstrate that passengers have an incentive to arrive truthfully since their fare quotes and fares tend to increase as their arrival times increase. Thus, it is more likely that they accept their fare quotes and are serviced for low fares if they arrive as early as possible. We perform 10,000 simulations with the transport simulator in a grid city of size 11×11 (that is, with 121 locations) and

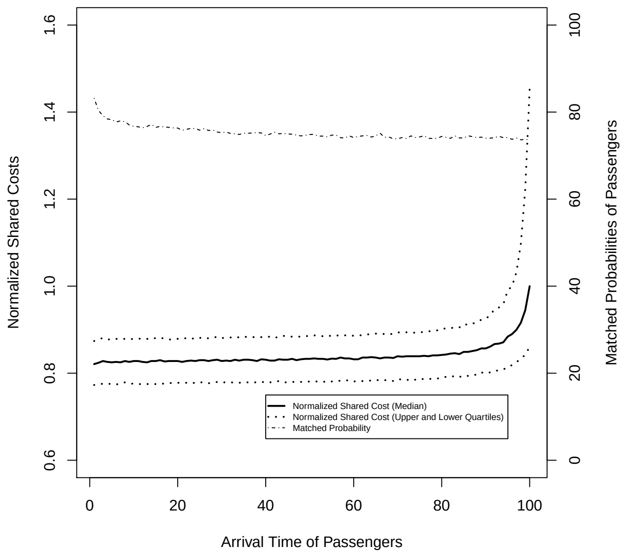

Figure 3: Results of Experiment 1

report average results. There are 25 shuttles that can each transport up to 10 passengers and operate the same hours from dawn (time 101) to dusk (time 1440). We assume that passengers submit their requests before dawn (the departure time of the shuttles) because otherwise the marginal costs depend on their arrival times. We also assume that shuttles have sufficient time to service all passengers before dusk. The shuttles start at a depot in the center of the city. Each shuttle incurs an operating cost of 1 for each unit of distance traveled and needs to return to its initial location at dusk. There are 100 passengers that all arrive truthfully one at a time (that is, their arrival times range from time 1 to time 100). The start location of 20 percent of the passengers is the depot. The start locations of the other passengers and the end locations of all passengers are randomly selected from all locations with uniform probability. The pick-up and drop-off time windows are identical for each passenger but might be different from passenger to passenger. Their lower bounds are dawn, and the differences between their upper and lower bounds are randomly selected from being 2.5 to 3.0 times higher than their alpha values (that is, the shortest point-to-point travel distance from their start location to their end location). Thus, passengers do not have tight schedules, resulting in low fare quotes. The fare limits of passengers are randomly selected from being 1.5 to 3.0 times higher than their alpha values. Thus, passengers have high fare limits. For both of these reasons, the fare quotes often do not exceed the fare limits. Many passengers therefore accept their fare quotes and are serviced.

Figure 3 shows the probability that passengers accept their fare quotes ("Matched Probabilities of Passengers") as a function of their arrival times k, that is, the percentage of simulations with $cost^{k}_{\pi(k)} \leq w_{\pi(k)}$. The probability that passengers accept their fare quotes is around 75 percent. It decreases as their arrival times increase (since their fare quotes tend to increase as their arrival times increase) but only very slowly. Figure 3 also shows the fares per alpha value of all passengers that accepted their fare quotes ("Normalized Shared Costs") as a function of their arrival times k, that is, $cost^{100}_{\pi(k)}$ averaged over all simulations with $cost^{k}_{\pi(k)} \leq w_{\pi(k)}$. The fares per alpha value of passengers increase as their arrival times increase (as suggested by the online fairness property) but only very slowly. The only exception is the sharp increase for arrival times close to 100 since passengers that arrive then can no longer share their costs with a high number of passengers that arrive after them.

4.3 Experiment 2

The definition of ex-post incentive compatibility states that the best strategy of every passenger is to arrive truthfully, provided that all other passengers arrive truthfully as well and do not change whether they accept or decline their fare quotes, two assumptions that are not guaranteed to be satisfied in practice. We have already shown in Section 3.4 that POCS does not satisfy the ex-post incentive-compatibility property if the second condition is removed. In Experiment 2, we therefore evaluate how likely it is that passengers can decrease their fares by delaying their arrivals if the second condition is removed. Experiment 2 is similar to Experiment 1, except that we distinguish four scenarios with different flexibilities of shuttles and passengers and use experimental parameters that decrease the scale of the experiment since each simulation is now more time-consuming. We perform 1,000 simulations with the transport simulator in a grid city of size 5×5 and report average results. Each simulation consists of at most 45 runs in addition to a run where Passengers $P_1 \ldots P_{10}$ arrive truthfully in order $P_1 \ldots P_{10}$, namely runs where all passengers arrive truthfully except that Passenger P_i delays its arrival and arrives only immediately after Passenger P_j for all i and j with $1 \leq i < j \leq 10$ where Passenger P_i accepts its fare quote when all passengers arrive truthfully. There are either 2 or 10 shuttles (for two scenarios) that can each transport up to 3 passengers, operate the same hours from dawn to dusk and start at a depot in the center of the city. Each shuttle incurs an operating cost of 1 for each unit of distance traveled and needs to return to its initial location at dusk. There are 10 passengers that arrive one at a time (that is, their arrival times range from time 1 to time 10) before the shuttles start to service them. The start and end locations of all passengers are randomly selected from all locations with uniform probability. The pick-up and drop-off time windows are identical for each passenger but might be different from passenger to passenger. Their lower bounds are dawn, and the differences between their upper and lower bounds are either 3.0 or 4.0 times (for two scenarios) higher than their alpha values. The fare limits of passengers are 3.0 times higher than their alpha values.

Table 7 shows, for each scenario, both the number of runs and the probabilities that passengers who delay their arrivals improve (since their fares decrease), do not change (since their fares remain unchanged) or worsen (since either their fare quotes increase sufficiently for them to drop out or - in case they do not drop out - their fares increase) their situations. Experiment 2 demonstrates that passengers have an incentive to arrive truthfully since, in all scenarios, the probability that passengers who delay their arrivals improve their situations is lower than 20 percent while the probability that they worsen their situation is higher than 50 percent. Exper-

Table 7: Results of Experiment 2

Scenario	Number of Shuttles	Time Window	Number of Runs	Situation Improves	No Change	Situation Worsens: Not Dropping Out	Situation Worsens: Dropping Out
1	2	3.0	32,808	11%	32%	24%	33%
2	2	4.0	37,259	15%	31%	39%	15%
3	10	3.0	36,955	16%	31%	51%	2%
4	10	4.0	37,990	17%	29%	51%	3%

iment 2 does not measure one advantage of passengers who delay their arrivals, namely the situation when passengers originally dropped out since their fare quotes exceeded their fare limits and by delaying their arrivals improve their fare quotes so they no longer drop out. Also, Experiment 2 assumes that passengers delay their arrivals randomly (rather than strategically) due to missing knowledge of future arrivals of passengers. The probability that the situation for passengers who delay their arrivals worsens is zero if passengers are able to delay their arrivals strategically since they can always decide to arrive truthfully instead, in which case their situations do not change. We thus expect the probability that their situations improve to increase.

5. CONCLUSIONS

In this paper, we determined properties of cost-sharing mechanisms that we believe make demand-responsive transport systems attractive to both shuttles and passengers, namely online fairness, immediate response, individual rationality, budget balance and ex-post incentive compatibility. We then proposed a novel cost-sharing mechanism, called Proportional Online Cost Sharing (POCS), that has these properties. Overall, POCS is a first step towards addressing some of the problems raised by the missing knowledge of future arrivals of passengers, which differentiates our research from previous research [3, 15, 8, 1]. However, some issues remain to be addressed by more advanced online cost-sharing mechanisms, including integrating more complex models of passengers, shuttles and transport environments. Our current simplifying assumptions include, for example, that the availability of shuttles does not change unexpectedly, that all passengers arrive before the shuttles start to service passengers, that fares depend only on the ride requests and no other considerations (for example, that DRT systems do not face competition), that all passengers evaluate their trips uniformly according to the criteria quantified by the alpha values (for example, that all passengers consider travel time to be equally important), that DRT systems provide fare quotes to passengers without predicting future arrivals of passengers (for example, that DRT systems service hard-to-accommodate passengers even though these passengers increase the shared costs of subsequent passengers and might make subsequent passengers drop out), that passengers try to decrease their fares only by delaying their arrivals (rather than, for example, by colluding with other passengers or entering fake ride requests under false names) and that passengers honor their commitments (for example, that passengers do not change ride requests, cancel them or show up late).

6. REFERENCES

[1] R. Anderson and A. Claus. Cost allocation in transportation systems. *Southern Economic Journal*, pages 793–803, 1976.

[2] J. Brenner and G. Schäfer. Online cooperative cost sharing. In T. Calamoneri and J. Diaz, editors, *Algorithms and Complexity*, volume 6078, pages 252–263. Springer, 2010.

[3] M. Frisk, M. Göthe-Lundgren, K. Jörnsten, and M. Rönnqvist. Cost allocation in collaborative forest transportation. *European Journal of Operational Research*, 205(2):448–458, 2010.

[4] M. Furuhata, L. Cohen, and S. Koenig. Using auctions to allocate transportation requests for demand responsive transit systems. Technical Report 09-19, Metrans Transportation Center, University of Southern California, Los Angeles, USA, 2013.

[5] M. Gendreau. An introduction to tabu search. In F. Glover and G. Kochenberger, editors, *Handbook of Metaheuristics*, volume 57 of *International Series in Operations Research and Management Science*, pages 37–54. Springer, 2003.

[6] M. Gendreau and J. Potvin. Tabu search. In E. Burke and G. Kendall, editors, *Search Methodologies*, pages 165–186. Springer, 2005.

[7] F. Glover, E. Taillard, and D. de Werra. A user's guide to tabu search. *Annals of Operations Research*, 41(1):1–28, 1993.

[8] J. Lemaire. An application of game theory: Cost allocation. *Astin Bulletin*, 14(1):61–81, 1984.

[9] H. Moulin. Incremental cost sharing: Characterization by coalition strategy-proofness. *Social Choice and Welfare*, 16(2):279–320, 1999.

[10] H. Moulin. Cost sharing in networks: Some open questions. *International Game Theory Review*, 15(2):1–10, 2013.

[11] W. Nanry and J. Barnes. Solving the pickup and delivery problem with time windows using reactive tabu search. *Transportation Research Part B*, 34(2):107–121, 2000.

[12] D. Parkes. Online mechanisms. In N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 16, pages 411–439. Cambridge University Press, 2007.

[13] C. Reeves. *Modern Heuristic Techniques for Combinatorial Problems*. Blackwell, 1993.

[14] Y. Sprumont. Ordinal cost sharing. *Journal of Economic Theory*, 81(1):126–162, 1998.

[15] F. Tohmé and T. Sandholm. Coalition formation processes with belief revision among bounded-rational self-interested agents. *Journal of Logic and Computation*, 9(6):793–815, 1999.

[16] Y. Wang and D. Zhu. Ordinal proportional cost sharing. *Journal of Mathematical Economics*, 37(3):215–230, 2002.

[17] Z. Xiang, C. Chu, and H. Chen. A fast heuristic for solving a large-scale static dial-a-ride problem under complex constraints. *European Journal of Operational Research*, 174(2):1117–1139, 2006.