



Wilson, RE., & Lees-Miller, JD. (Accepted/In press). *Sampling for Personal Rapid Transit Empty Vehicle Redistribution*.  
<http://hdl.handle.net/1983/1760>

Early version, also known as pre-print

[Link to publication record in Explore Bristol Research](#)  
PDF-document

## University of Bristol - Explore Bristol Research

### General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:  
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

1 **Sampling for Personal Rapid Transit Empty Vehicle Redistribution**

2

3 John D. Lees-Miller

4 Department of Engineering Mathematics,

5 University of Bristol, Queen's Building, University Walk,

6 Bristol, BS8 1TR, United Kingdom

7 enjdlm@bristol.ac.uk

8 (corresponding author)

9

10 R. Eddie Wilson

11 Department of Engineering Mathematics,

12 University of Bristol, Queen's Building, University Walk,

13 Bristol, BS8 1TR, United Kingdom

14 RE.Wilson@bristol.ac.uk

15 Tel.: +44(0) 117 331 5627, Fax.: +44(0) 117 331 5606

16

17 6,217 words + 4 figures

**1 ABSTRACT**

2 A Personal Rapid Transit (PRT) system uses compact, computer-guided vehicles running on  
3 dedicated guideways to carry individuals or small groups directly between pairs of stations. PRT  
4 vehicles operate on demand, when a passenger requests service at his/her origin station. Because  
5 the number of trips requested from a station need not equal the number of trips ending there,  
6 some vehicles must move empty to balance the flows. The empty vehicle redistribution (EVR)  
7 problem is to decide which empty vehicles to move, either reactively, in response to known  
8 requests, or proactively, in anticipation of future requests. This paper develops a new algorithm  
9 for the EVR problem called Sampling and Voting (SV). SV chooses reactive movements using a  
10 simple nearest-neighbor rule, and it chooses proactive movements by generating an ensemble of  
11 possible sequences of future passenger requests, solving a deterministic optimization problem for  
12 each sequence individually, and then finding the empty vehicle movements that are common  
13 among the sequences. Moving vehicles proactively is essential for providing low passenger  
14 waiting times. The new SV algorithm is tested in simulation with several case study systems, and  
15 it produces significantly lower passenger waiting times than existing EVR algorithms. Variants  
16 of the SV method developed here for PRT are also applicable to conventional taxi systems and  
17 emergency response systems.  
18

## 1. INTRODUCTION

Personal Rapid Transit (PRT) is an emerging urban transport mode. The world's first PRT system is currently undergoing final testing at Heathrow Airport in London, England (1). The basic idea of a PRT system is that it provides on-demand, non-stop travel with compact, computer-guided vehicles running on a dedicated network of unidirectional guideways. Each PRT vehicle carries either an individual or a small party traveling together by choice. The vehicle begins its trip *on demand*, when a passenger arrives at a PRT station and makes a *request* for service. Once the passenger is ready to depart, his vehicle takes the quickest path to the chosen destination station: it does not stop at intermediate stations to let other passengers on or off. Hence, a PRT system is similar to a taxi system, except that PRT vehicles are constrained to start and end their trips at stations.

In general, the number of trips requested from a particular station need not equal (on average or instantaneously) the number of trips ending at that station, so that some PRT vehicles must move *empty*, that is, without passengers. The question of which vehicles to move, and where to move them, is known as the empty vehicle redistribution (EVR) problem.

In (2), we used a fluid limit to analyze the EVR problem in terms of the overall capacity of the PRT system – that is, the maximum demand for travel that can be met without diverging queues of waiting passengers. This fluid limit approach is useful for benchmarking proposed EVR algorithms in terms of capacity, but it has two important limitations:

1. It yields only average flows of empty vehicles and does not suggest an operational algorithm for the movement of vehicles at an individual level.
2. In practice, PRT systems that are presently planned will run far short of the capacity limit, and the main objective will be to minimize passenger waiting time, either at the mean or at some percentile; the fluid-limit cannot be used for this analysis.

The focus of this paper is on operational EVR algorithms, which must move individual empty vehicles in response to passenger requests or in anticipation of future requests. It is assumed that the average request rate between each pair of stations is known from historical data in the form of an origin-destination *demand matrix*, but that the individual requests are revealed only while the system is operating. Historically, EVR algorithms have been based on decision rules (3, 4) to be executed when a new passenger arrives or a vehicle becomes empty. For example, if a station is short of empty vehicles, a vehicle is chosen according to the rules and sent there; this may simply be the nearest empty vehicle at a station that is not itself short of vehicles (5), or it may be determined by solving a transportation problem (6).

This paper describes a new EVR algorithm, here called the Sampling and Voting (SV) algorithm. SV moves empty vehicles in anticipation of future passenger requests by analyzing artificial ensembles that represent possible sequences of future passenger requests over a given finite horizon. Each sequence defines a deterministic optimization problem whose (approximate) solution suggests a plan of empty vehicle movements. Features of these plans which are common across the ensemble are then extracted to determine which empty vehicles should actually be moved. Passenger waiting times with the proposed SV algorithm are significantly lower than those for existing methods, in simulation tests.

The terminology and approach used here are based on the literature for vehicle routing problems (VRPs). A classical example (7) of a VRP is the distribution of goods from a warehouse to a set of stores; each day, stores request deliveries for the next day, routes are planned over night, and vehicles are dispatched on these routes in the morning. The planning

1 objective is typically to minimize the cost of delivering the goods. This problem is *static*,  
2 because all customer requests are known when the routes are being planned. A *dynamic* VRP (8,  
3 9) arises if 'same-day' requests are to be served, because routes must be modified while the  
4 system is operating. When probabilistic information about same-day requests is known, the  
5 problem is said to be *stochastic*.

6 The EVR problem is both dynamic (because requests arrive as the system is operating)  
7 and stochastic (because the mean rates are known from the demand matrix). The corresponding  
8 static problem, in which all passenger requests are known in advance, arises in the SV algorithm,  
9 because each sequence in the ensemble generated by SV defines an instance of the static EVR  
10 problem. The static problem also provides a means of benchmarking EVR algorithms in terms of  
11 passenger waiting time, because performance in the *perfect information* situation always matches  
12 or exceeds that where the information is uncertain.

13 The sampling approach that we adopt has been successfully applied to several dynamic  
14 VRP variants related to the EVR problem (10–12). The use of the static problem to benchmark  
15 the performance of algorithms on the dynamic problem is known as *competitive analysis*, and  
16 this has also been done with dynamic VRPs (13, 14). The most relevant VRP variants are the  
17 *multivehicle truckload pickup and delivery* problems, where 'truckload' means that each vehicle  
18 can serve at most one request at a time (15), and 'pickup and delivery' means that each request  
19 has a particular origin and destination (14, 16). The objective in most VRPs is to minimize travel  
20 cost, but waiting times may also be considered (10); the objective of minimizing average  
21 passenger waiting time is known as a *minimum latency* objective (17, 18). Some dynamic  
22 problems with minimum latency objectives have also been studied (19, 20). However, none of  
23 the VRP variants described above is a perfect match for the EVR problem.

24 The paper is laid out as follows. In section 2, we introduce the mathematical  
25 representation that we use as a simplified model of PRT operations, and we introduce two basic  
26 EVR algorithms (one for the dynamic case, one for the static case) that will be incorporated into  
27 the SV algorithm, which is described in detail in section 3. Then section 4 compares SV to other  
28 EVR algorithms in simulations with two test networks. Finally, section 5 presents conclusions  
29 and sets the scene for future work.

## 30 **2. PRT SYSTEM MODEL AND BASIC EVR ALGORITHMS**

31 We now develop a mathematical representation of PRT operations and use it to introduce two  
32 basic EVR algorithms: the Bell and Wong (21) Nearest-Neighbors algorithm (BWNN), which  
33 operates in a dynamic context, and the Static Nearest-Neighbors algorithm (SNN), which  
34 operates in a static context. There are two objectives. Firstly, the BWNN and SNN algorithms  
35 are components of the more complicated Sampling and Voting (SV) algorithm that we develop  
36 in section 3. Secondly, they are in themselves useful EVR algorithms which can be used to  
37 benchmark and evaluate other EVR algorithms.

38 Since we are mainly concerned with modeling passenger waiting time and how it may be  
39 minimized with the use of EVR algorithms, we develop a mathematical model which simplifies  
40 and neglects some of the details of real-world PRT system operations. These simplifications have  
41 two flavors:

- 42 1. Congestion on the guideway is ignored, so that travel times between stations are  
43 constant – thus vehicles may always depart upon the quickest path without any  
44 delay due to slot-booking processes (22, pp. 92–94) etc.  
45

- 1           2.       PRT system processes at stations are simplified. For example, we neglect  
 2                   spontaneous ride sharing (23), and we neglect passenger loading and unloading  
 3                   times. Finally, we assume that the maneuvering order of individual vehicles  
 4                   within stations is unconstrained. See (2) for a more detailed discussion of these  
 5                   points.

6           These assumptions lead to a PRT system model that is essentially the urban taxi model of  
 7 (21). Let  $S$  denote the set of stations and  $K$  denote the set of vehicles, indexed  $1, 2, \dots, n_K$ . Let  $D_{ij}$   
 8 denote the demand in requests per unit time between stations  $i \in S$  and  $j \in S$ , with  $D_{ij} = 0$  if  $i =$   
 9  $j$ , and let  $T(i, j)$  be the travel time from station  $i$  to station  $j$ , with  $T(i, j) = 0$  if  $i = j$ .

10          Each vehicle in the system maintains a list of stations which it must visit in order. Each  
 11 *trip* thus specified by consecutive elements of the list is either *occupied* (carrying passengers,  
 12 thus satisfying a request) or *empty* (without passengers, thus specified by the EVR algorithm).  
 13 Each vehicle's list evolves in time: as each trip is completed, the head of the list is deleted, and  
 14 more stations are added to the tail of the list as new occupied or empty movements are assigned.  
 15 If a vehicle completes all of the trips in its list, it becomes *idle* and remains at the last station of  
 16 its final movement waiting for new trips.

17          The key discrete events are passenger *requests*. These occur when a passenger (or a  
 18 preformed group of passengers who intend to travel together in a single vehicle) arrive at station  
 19  $i$  and request a vehicle to take them to station  $j$ . If an empty vehicle is idle at station  $i$ , then  
 20 clearly it should be assigned to the new request. Otherwise, the allocation of a vehicle is more  
 21 complicated: it may so happen that there are already vehicles en route to station  $i$ , but if not, a  
 22 new empty vehicle trip is required. However, the ideal situation is to move idle vehicles  
 23 *proactively*, in anticipation of future requests, so that future passengers do not have to wait.

24          In this model, the chief principle is that trips are completed in order, and vehicles are  
 25 never rerouted from either empty or occupied movements that have previously been assigned to  
 26 them. Particularly when the system is busy, rerouting might be beneficial, but it introduces  
 27 further complications, such as the "indefinite deferment" of requests at outlying stations (9); as  
 28 such, models that allow rerouting are the subject of ongoing research. Because the model used  
 29 here does not allow rerouting, only the tail of each vehicle's list is important: specifically, we  
 30 need to consider only the destination station  $d[k] \in S$  of the last trip assigned to each vehicle  
 31  $k \in K$  and the time  $a[k]$  at which this trip will be completed. When a vehicle becomes idle, we  
 32 simply freeze these values so that  $a[k]$  denotes the (past) time at which the vehicle became idle  
 33 and  $d[k]$  denotes the station at which it is idle. When a new trip is appended to a vehicle's list,  
 34  $d[k]$  is updated, and  $a[k]$  is recomputed according to when the new trip will be completed.  
 35

### 36 **Bell and Wong Nearest Neighbors (BWNN) Algorithm**


37 This algorithm is the simplest of several algorithms presented by (21). It does not make  
 38 proactive moves, but instead only moves empty vehicles in response to requests. Specifically,  
 39 when a request is received at time  $t$  with origin  $i$  and destination  $j$ , the vehicle

$$k^* = \underset{k}{\operatorname{argmin}} [\max\{0, a[k] - t\} + T(d[k], i)] \quad (1)$$

40 is assigned. In intuitive terms, this choice is locally optimal, because it minimizes the time that  
 41 the new request waits before a vehicle is made available to serve it. Here the terms on the right-  
 42 hand side incorporate the time until vehicle  $k$  becomes idle and the empty vehicle trip time to  
 43 reach the new request's origin station,  $i$ . This formulation includes the cases where there are  
 44 vehicles that are either already idle or will become idle at station  $i$ . To complete the specification,

1 a tie-breaking rule is required for when the argmin in (1) is non-unique: for simplicity we choose  
 2 the minimum such vehicle index  $k$ .

3 Figure 1 illustrates the handling of vehicle request lists and the operation of the BWNN  
 4 algorithm. It also shows that passenger waiting times could be reduced if future requests were  
 5 known, as is the case with the SNN algorithm described below.  
 6



7  
 8  
 9  
 10  
 11  
 12  
 13  
 14  
 15  
 16  
 17  
 18  
 19  
 20  
 21  
 22

**FIGURE 1** Illustration of the BWNN algorithm. There are four off-line stations (labeled A – D) in a ring and two vehicles (labeled 1 and 2). Traffic flow is counter-clockwise. (a) Vehicle 1 is initially moving to station B, and vehicle 2 is idle at station A. (b) When a request for travel from C to D is received, vehicle 1 is assigned, because it gives a smaller waiting time than vehicle 2. The new request is appended to vehicle 1's request list; this requires an empty vehicle trip (dashed line) from B to C and an occupied trip (solid line) from C to D. Note that, while vehicle 1 stops at station B and station C (filled circles), it does not become *idle* at either station, because it has not finished with its request list. (c) However, vehicle 1 does become idle at D, in this case, because no further requests are assigned to it. (d) When another request is received from C to A, vehicle 2 is assigned, and it begins an empty trip to C. Vehicle 2 was idle at A, so it could have moved to C proactively, if the request from C had been anticipated; this would have reduced the passenger's waiting time.

### 1 **Static Nearest Neighbors (SNN) Algorithm**

2 We now modify BWNN to create an algorithm for the *static* EVR problem, in which all requests  
 3 are known in advance. Let  $R$  be a set of requests, indexed  $1, 2, \dots, n_R$ , and for each request  $r \in R$   
 4 let  $i_r$  and  $j_r$  be the origin and destination respectively, and let  $t_r$  be the time at which the request is  
 5 received, with  $t_1 < t_2 < \dots < t_{n_R}$ . Like the BWNN algorithm, we step through the requests in  
 6 ascending order in  $r \in R$ , but in contrast to BWNN there is no notion of future or past because  
 7 all of the requests are known simultaneously. For request  $r$  we assign the vehicle

$$k^* = \underset{k}{\operatorname{argmin}} \max\{0, a[k] + T(d[k], i_r) - t_r\} \quad (2)$$

8 and update  $a[k^*]$  and  $d[k^*]$  accordingly. Like the BWNN rule (1), the SNN rule (2) chooses a  
 9 vehicle  $k$  to minimize the request's waiting time, but with the difference that the empty trip for  
 10 vehicle  $k$  can begin as early as  $a[k]$ , even if the request is received much later. In other words, the  
 11 SNN algorithm can move vehicles in advance of requests so that passengers are not kept waiting.

12 Ties in (2) are common, because there are often several vehicles that could reach  $i_r$  before  
 13 the request at  $t_r$ , and all such vehicles would give zero waiting time. To break ties, we first try to  
 14 select the vehicle with the minimum empty travel time  $T(d[k], i_r)$ . However, if (for example)  
 15 there is more than one vehicle inbound to station  $i_r$ , they will all have zero empty time and so a  
 16 further tie-breaking rule is required. To this end, we select the vehicle with latest arrival time  
 17  $a[k] + T(d[k], i_r)$  at  $i_r$ . The idea is that this choice allows better moves as one progresses through  
 18 the list of requests, because the vehicles that arrive earlier at  $i_r$  have more flexibility to serve  
 19 requests further down the list. Finally, if these measures are equal, we break ties by choosing the  
 20 minimum such vehicle index  $k$ .

21 With perfect information about future arrivals, SNN tends to yield average waiting times  
 22 less than those for the dynamic case; quantitative comparison follows in section 4. However,  
 23 because the routes produced by the SNN algorithm are not provably optimal, other algorithms  
 24 may outperform SNN for generic patterns of passenger requests. In separate unpublished work,  
 25 we have found that standard methods for proving the optimality of solutions to the static EVR  
 26 problem have not been effective for usefully large instances. It has, however, been found (17, 18,  
 27 20) that nearest neighbor heuristics often produce good solutions for other minimum latency  
 28 routing problems.

29

### 30 **3. NEW EVR ALGORITHM: SAMPLING AND VOTING (SV)**

31 The main challenge addressed by the proposed SV algorithm is to move idle vehicles proactively  
 32 in a dynamic context. The BWNN algorithm moves vehicles reactively in a dynamic context, and  
 33 the SNN algorithm is effective at proactive movements in a static context, in which all future  
 34 requests are known. The approach taken here aims to combine the best features of these two  
 35 algorithms.

36 When a new request is received at time  $t$ , SV assigns a vehicle using the BWNN  
 37 algorithm. Immediately after a vehicle has been assigned to serve the request, SV may then move  
 38 idle vehicles proactively. To decide which idle vehicles to move, an *ensemble* of  $n_E$  possible  
 39 sequences of  $n_R$  future requests each is generated from the demand matrix. Each sequence in the  
 40 ensemble, together with the current state of the system, defines an instance of the static EVR  
 41 problem. Each of these instances is solved approximately using the SNN algorithm, and each  
 42 resulting solution prescribes a sequence of empty vehicle trips, which constitutes 'advice' on  
 43 which idle vehicles the system should actually move. However, because each solution is  
 44 computed for a (probably) different sequence of requests, they may offer conflicting advice.



1 To determine which action should actually be taken, a *voting* system is used. The system  
 2 adopted here is that at most one idle vehicle at each station may be moved. So, each solution  
 3 casts one vote on the best destination (as defined below) for an idle vehicle at each station  $i$  with  
 4 idle vehicles; note that it may vote for  $i$ , which means that it votes *not* to move any idle vehicles  
 5 from  $i$  at this decision point. If the destination with the most votes is not  $i$ , an idle vehicle at  $i$  is  
 6 selected (breaking ties on minimum vehicle index) and moved.

7 The destinations to vote for are determined as follows. The solution for each static EVR  
 8 instance maps each input request to an empty vehicle trip, so each solution yields an ordered  
 9 sequence of exactly  $n_R$  empty trips. Empty trip  $p \in \{1, \dots, n_R\}$  is described by the tuple  
 10  $(i_p, j_p, k_p)$ , where  $k_p \in K$  is the index of the vehicle used for trip  $p$ , and  $i_p \in S$  and  $j_p \in S$  are  
 11 the stations at which the empty trip begins and ends, respectively. Trips with  $i_p = j_p$  are trivial,  
 12 in that no actual empty vehicle movement is required. For each station  $i$  with idle vehicles, let  
 13  $K_i = \{k \in K : d[k] = i \text{ and } a[k] \leq t\}$  be the set of vehicles that are currently idle at  $i$ , and hence  
 14 eligible to be moved at the current decision point. The following rules then determine the  
 15 destination to vote for.

- 16 i. vote for  $i$  if all vehicles in  $K_i$  were used for trips  $p$  with  $j_p = i$ , or
- 17 ii. vote for  $j_p$  for the first trip  $p$  with  $k_p \in K_i$  and  $j_p \neq i$ , if one exists, or
- 18 iii. vote for  $j_p$  for the first trip  $p$  with  $i_p = i$  and  $j_p \neq i$ , if one exists, or
- 19 iv. vote for  $i$ .

20 Rule (i) votes to leave all idle vehicles at  $i$  if they were all needed there. If an idle vehicle  
 21 at  $i$  was moved to another station, rule (ii) votes to perform this trip. If none of the idle vehicles  
 22 were moved, as is common when demand is light, rule (iii) looks at all trips for a hint at where it  
 23 should send one of these idle vehicles. If there were no trips from station  $i$ , rule (iv) leaves them  
 24 where they are.

25 The concept of planning with an ensemble is very general, and many variants on the SV  
 26 algorithm are possible, such as (a) using algorithms other than BWNN and SNN to assign  
 27 vehicles to requests; (b) running the ensemble generation and voting system at different decision  
 28 points; (c) using different voting systems; (d) using different rules to choose destinations from  
 29 each sequence in the ensemble. The SV algorithm described here gave the lowest mean  
 30 passenger waiting times among several such variants that we have tested in simulation.

#### 31 32 **4. RESULTS**

33 In this section, the proposed algorithm is evaluated in simulations of two case study systems. The  
 34 required input data are a network, an origin-destination demand matrix and a fleet size. The main  
 35 outputs are passenger waiting times and empty vehicle use. The steady state distributions of  
 36 these outputs are estimated by running long simulations with the demand matrix held constant  
 37 for each run. Passenger requests from station  $i$  to station  $j$  are generated from a Poisson process  
 38 with rate  $D_{ij}$ . For convenience, passenger arrival and travel times are rounded to the nearest  
 39 integer second.

40 The observed waiting times strongly depend on how ‘heavy’ the demand is, relative to  
 41 the network size and the available vehicle fleet. This is made precise using the *intensity* measure  
 42 from (2). Given a network and a demand matrix, the method in (2) computes the minimum  
 43 number of vehicles needed to serve the demand based on the required average flows of both  
 44 occupied and empty vehicles. The intensity of the demand on the system is defined as the ratio of

1 the number of vehicles needed to the number of vehicles actually available ( $n_K$ ). When the  
 2 intensity is larger than one, passenger requests are arriving faster than they can possibly be  
 3 served, regardless of the EVR algorithm used, because there are not enough vehicles. The focus  
 4 here is therefore on demands with intensity less than one. To measure the effects of intensity, an  
 5 initial demand matrix is scaled up and down to produce a family of demand matrices with the  
 6 same spatial distribution but different intensities between zero and one.

7 The input data used here are the ‘Grid’ and ‘Corby’ networks and their corresponding  
 8 demand matrices from (2) (for the Grid network, the demand matrix with dispersion parameter  $\theta$   
 9 = 0.01 is used). The fleet size is set at  $n_K = 200$  vehicles. Intensity one corresponds to a total  
 10 demand of 1414 requests/hour for the Corby system and 2035 requests/hour for the Grid system.  
 11 Simulations have been conducted on nine systems in total, with between 15 and 60 stations,  
 12 between 50 and 600 vehicles, and total demand at intensity one between 360 and 5050  
 13 requests/hour, and the results and conclusions are consistent with those presented here.

14 For comparison, another EVR algorithm, here called the Surplus / Deficit (SD) algorithm,  
 15 is also evaluated. It is an algorithm for the dynamic EVR problem that moves vehicles  
 16 proactively. The general approach in SD is similar to several other published EVR algorithms  
 17 (4); it is most similar to that of (5). Each station  $i$  has an associated *call time*  $l_i$ , which is the  
 18 cumulative average of all previous empty vehicle trip times to that station (simulations with  
 19 exponential moving averages gave the same results). The *surplus* of vehicles at station  $i$  is the  
 20 number of inbound vehicles (with  $d[k] = i$ ) minus the expected number of requests over the call  
 21 time, namely  $l_i \sum_j D_{ij}$ . When a new request is received, a vehicle is assigned using BWNN.

22 Immediately afterward, SD may move idle vehicles proactively, as follows. For each station  $i$   
 23 with idle vehicles, in descending order by number of idle vehicles, if the surplus of vehicles at  $i$   
 24 is greater than or equal to one, an idle vehicle at  $i$  is sent to the nearest station with surplus less  
 25 than zero (if any). Additionally, when a vehicle becomes idle at station  $i$ , the above actions are  
 26 taken for station  $i$  only.

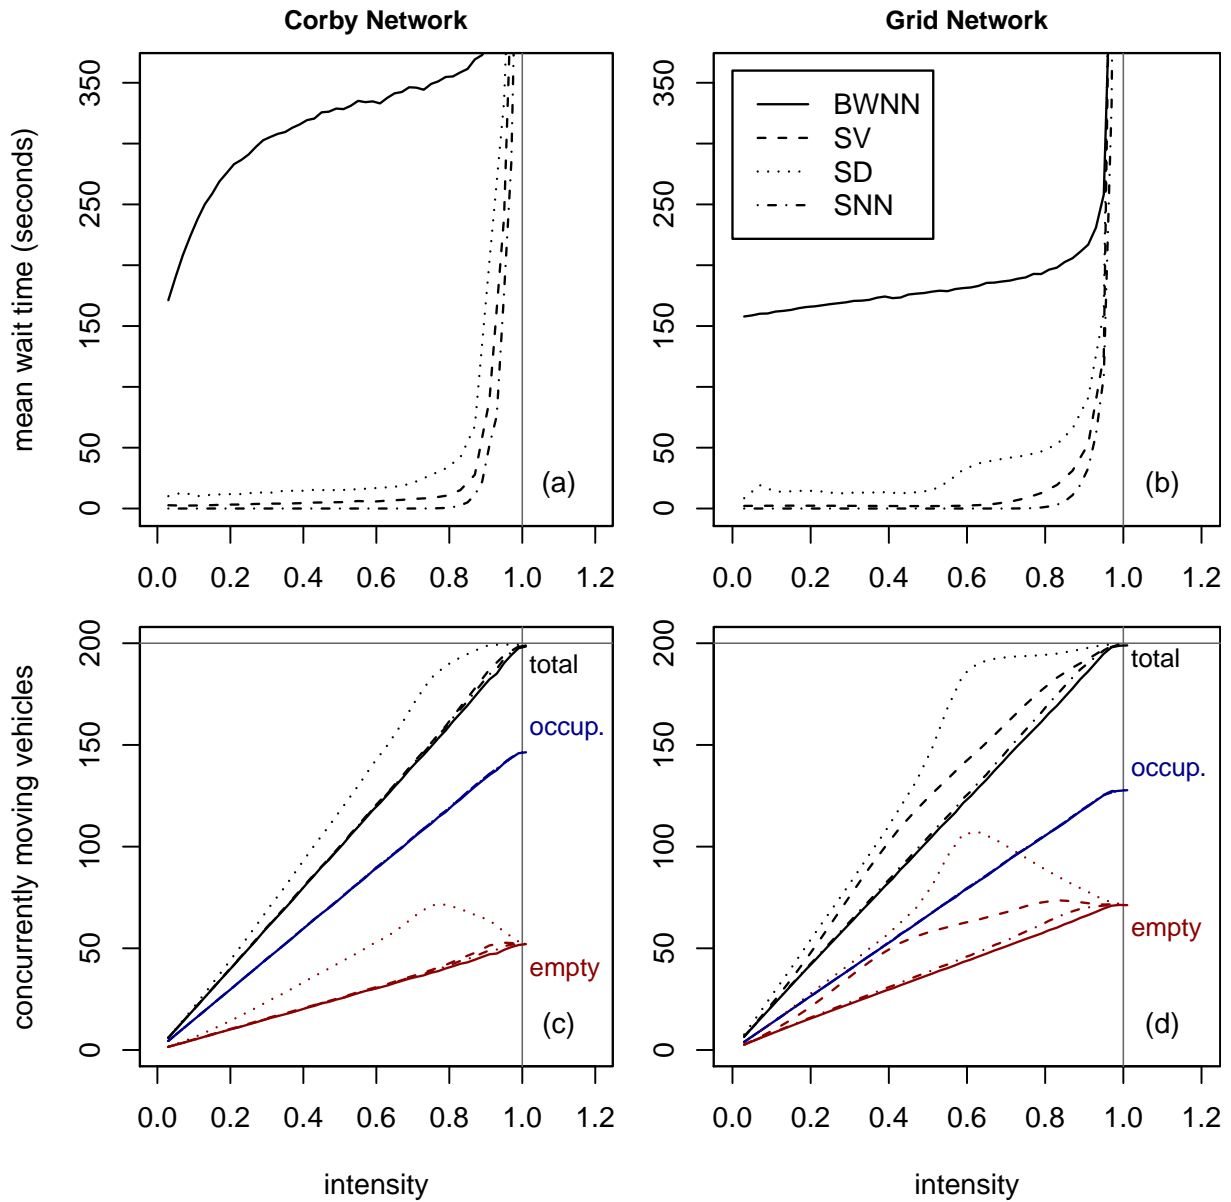
## 28 Comparison of Algorithms

29 Figure 2(a) compares the mean waiting times observed for the four algorithms on the Corby  
 30 system. An important observation is that waiting times increase rapidly as intensity approaches  
 31 one, regardless of which EVR algorithm is used, as is expected based on the definition of  
 32 intensity. However, the SV algorithm produces the lowest mean waiting times at all intensities  
 33 tested. In practice, we are most interested in the system's performance at around 70% to 90%  
 34 intensity, because in this range the system is well-utilized, but acceptably low passenger waiting  
 35 times may still be obtained. At intensity 0.8, for example, mean waiting times are 355s for  
 36 BWNN, 41s for SD and 15s for SV, so SV reduces mean waiting times by 96% from BWNN and  
 37 by 63% from SD. The relative reduction decreases as intensity increases, however, and in fact  
 38 the SD and SV algorithms become increasingly similar to the BWNN algorithm at higher  
 39 intensities, because there are fewer idle vehicles to redistribute, as shown in Figure 2(c). Figure  
 40 2(c) also shows that the reduction in passenger waiting times comes from a modest increase in  
 41 the average number of moving empty vehicles, or equivalently in empty vehicle travel time. The  
 42 largest increase occurs at intensity 0.91, and this is from 47 concurrently moving empty vehicles  
 43 with BWNN to 51 with SV (out of 200 vehicles). With perfect information about future arrivals,  
 44 SNN finds routes with average waiting times less than those for the dynamic case, as expected;  
 45 at intensity 0.8, the mean waiting time for SNN is 3s. Only mean waiting times are reported, but

1 the ranking of the four algorithms is the same at the 90th percentile of the waiting distribution; at  
2 intensity 0.8, 90% of passengers wait less than 51s with SV and less than 106s with SD.

3 Results for the Grid network are qualitatively similar, as shown in Figure 2(b, d).  
4 However, it is notable that for all three algorithms the number of concurrently moving vehicles  
5 reaches its maximum ( $200 = n_K$ ) at around intensity 0.95, which is below the theoretical  
6 maximum (intensity 1). It thus appears that nearest neighbor algorithms of the type studied here  
7 do not deliver maximum possible throughput; it is not yet known whether there is any practical  
8 algorithm that does.

9 Computation time for the SV algorithm is larger than that for the other algorithms, but  
10 SV is still fast enough for real time use with the case study networks. The mean computing time  
11 per passenger request for the SV results in Figure 2 is 0.04s (user plus system CPU time on an  
12 Intel Xeon E5405 at 2.0GHz with Red Hat g++ 4.1.2 and -O3). The largest system so far tested  
13 with SV has 60 stations and 600 vehicles. When  $n_E = 50$  sequences and  $n_R = 750$  generated  
14 requests / sequence, SV produces lower waiting times than the SD and BWNN methods on this  
15 system, and it uses 1.2s of computing time per passenger request. For this system, the demand at  
16 intensity one is 5050 requests/hour, or 0.7s per request, so the *sequential* SV algorithm used for  
17 testing is not fast enough for real time use at high intensity, in this case. However, SV is easy to  
18 parallelize, because each of the  $n_E$  sequences can be generated and processed independently; for  
19 example, real elapsed time per request could in this case be reduced to near 0.6s if two  
20 processors were used. The main limit on the scalability of the SV algorithm is thus the time to  
21 process a single sequence. The SNN minimization (2) takes  $O(n_K)$  time, so processing a single  
22 sequence takes  $O(n_K n_R)$  time.  
23



1  
2  
3  
4  
5  
6  
7  
8  
9  
10

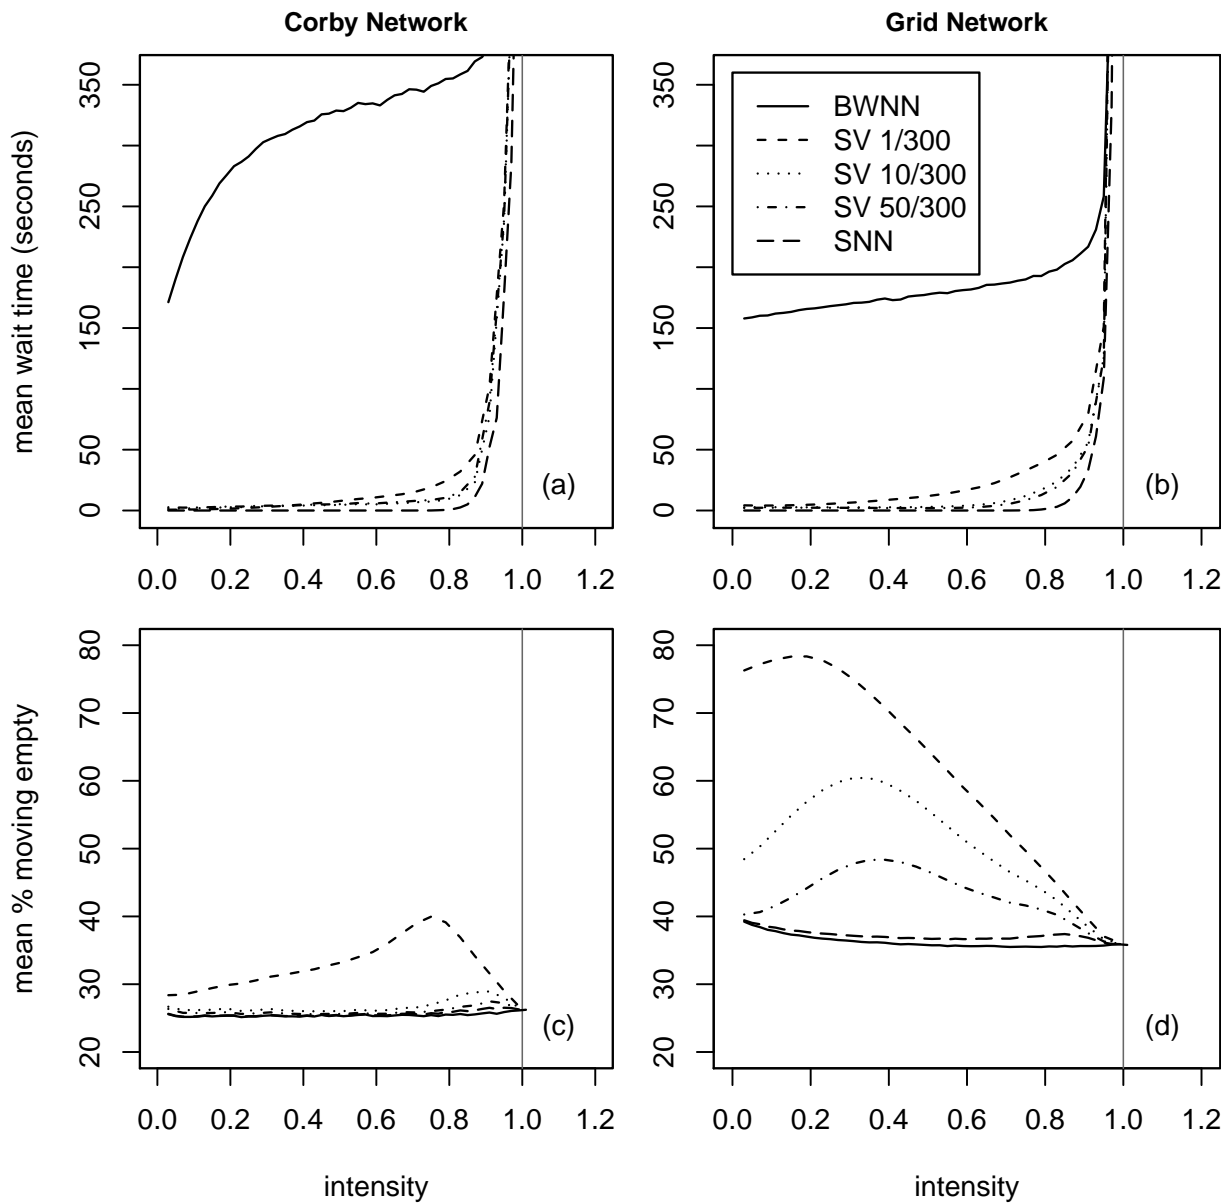
**FIGURE 2** Mean passenger waiting times (*a, b*) and vehicle fleet utilization (*c, d*) for the Grid and Corby networks for the four algorithms. The proposed SV algorithm moves idle vehicles in anticipation of future requests, which reduces waiting times significantly below the BWNN baseline and the SD algorithm. The SNN algorithm operates with perfect information about future requests in order to estimate how much further waiting times might be reduced. Here there are  $n_E = 50$  sequences, each with  $n_R = 300$  requests for SV. Each point is averaged over 10 independent runs of 50,000 simulated passengers each.

### 1 **Effect of Number and Length of Ensemble Sequences on SV**

2 The SV algorithm has two free parameters. Figure 3 shows the effects of varying the number of  
3 sequences  $n_E$ , and Figure 4 shows the effects of varying the number of requests,  $n_R$ , in each  
4 sequence.

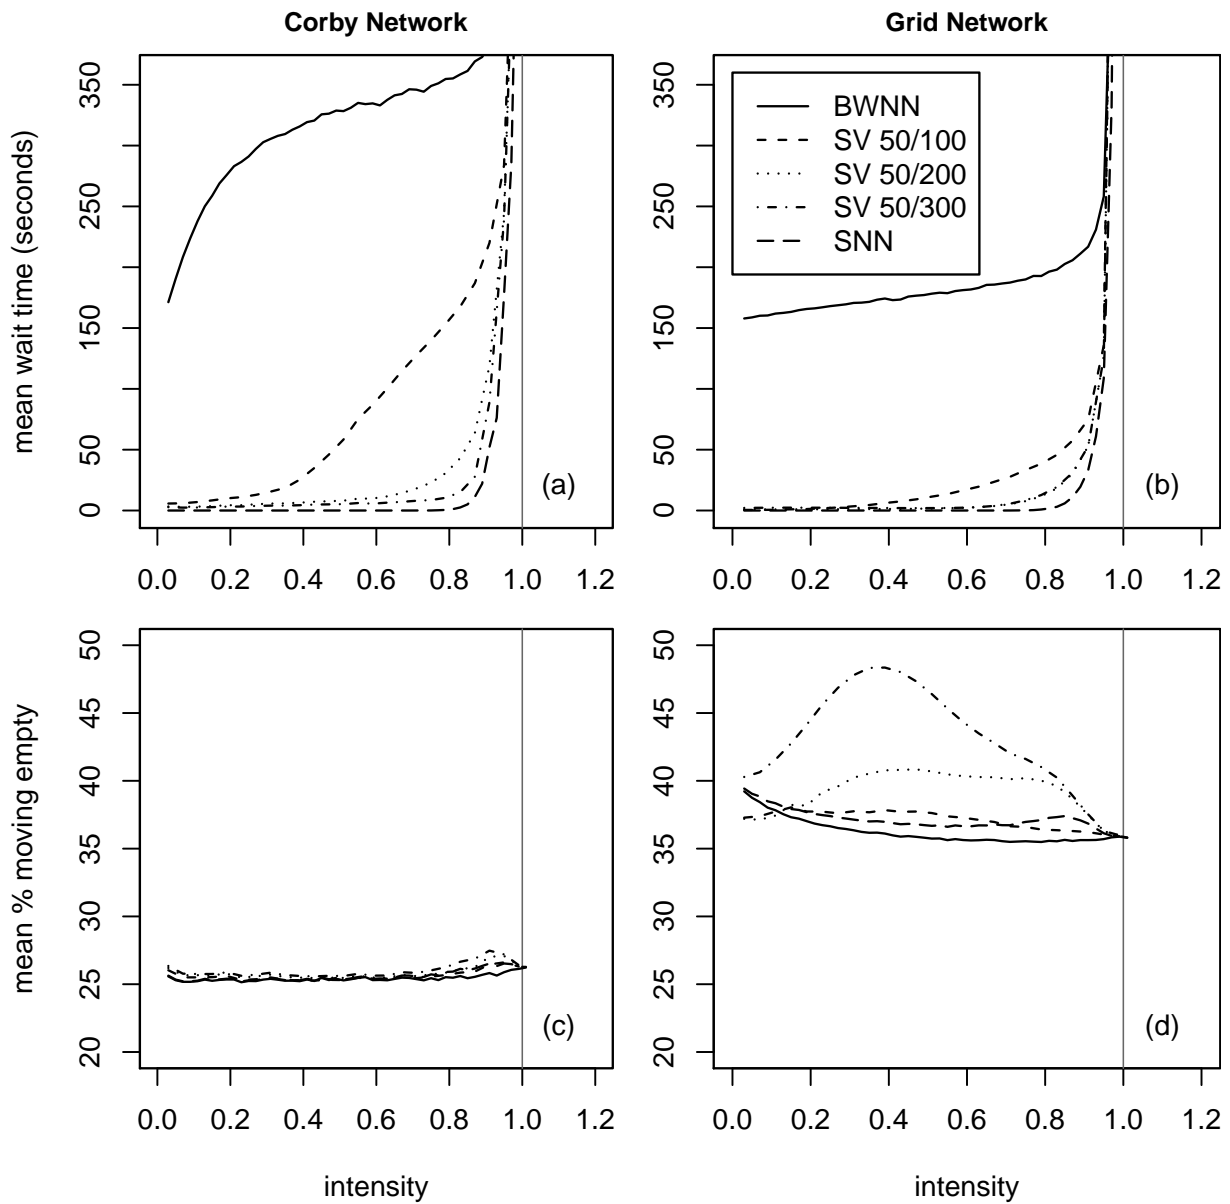
5 Figures 3(*a, b*) show that waiting times are not very sensitive to the number of sequences  
6 used. On the Corby system, a single sequence is sufficient to reduce waiting times well below the  
7 BWNN baseline. However, Figures 3(*c, d*) show that using more sequences produces some  
8 further reductions in waiting time and also significant reductions in empty vehicle travel. For  
9 example, on the Grid system at intensity 0.4 with a single sequence, 65% of moving vehicles are  
10 empty on average, but this falls to 45% when 50 sequences are used. This is because a small  
11 ensemble is less likely to be representative of the actual demand, and also because the actions  
12 recommended by a small ensemble are more likely to change at each decision point. So, while a  
13 small ensemble will sometimes make good recommendations that reduce passenger waiting  
14 times, it is also likely to make bad recommendations that result in wasted empty vehicle trips.

15 Figures 4(*a, b*) show that waiting times decrease for all intensities as  $n_R$  increases. A  
16 sequence with more requests generates more occupied and empty vehicle trips, which makes it  
17 more likely that the sequence will vote on what to do with an idle vehicle, rather than just  
18 leaving it idle. This results in more empty vehicle movement, but also lower waiting times. For  
19 example, at intensity 0.8 on the Grid network, increasing  $n_R$  from 100 to 200 increases the  
20 fraction of moving vehicles that are empty from 36% to 40%, and it decreases mean passenger  
21 waiting times from 46s to 18s.



1  
2  
3  
4  
5  
6  
7  
8

**FIGURE 3** Effect of the number of sequences ( $n_E$ ) on mean passenger waiting times (*a, b*) and the fraction of moving vehicles which are empty (*c, d*). Even a single sequence ( $n_E = 1$ ) is sufficient to significantly reduce waiting times below the BWNN baseline, but using more sequences reduces the amount of empty vehicle travel required. Here  $n_R = 300$  requests per sequence. Each point is averaged over 10 independent runs of 50,000 simulated passengers each.



1  
2  
3  
4  
5  
6  
7  
8

**FIGURE 4** Effect of sequence length ( $n_R$ ) on mean waiting times (*a*, *b*) and the fraction of moving vehicles which are empty (*c*, *d*). Longer sequences (with more requests) reduce passenger waiting times at the cost of modest increases in empty vehicle running. Here  $n_E = 50$  sequences. Each point is averaged over 10 independent runs of 50,000 simulated passengers each.

## 1 5. CONCLUSIONS

2 The Sampling and Voting (SV) algorithm developed in this paper is a new algorithm for solving  
3 the dynamic Empty Vehicle Redistribution (EVR) problem for Personal Rapid Transit systems.  
4 Variants of SV might also be applied to manage fleets of conventional taxis or emergency  
5 response vehicles. Average passenger waiting times must diverge as the demand approaches  
6 capacity, regardless of the EVR algorithm used. However, the SV algorithm allows the system to  
7 maintain a good service level through a larger part of the intensity range than other algorithms in  
8 the literature, albeit with some additional computational cost.

9 Simulation results show that moving idle vehicles proactively is key to providing low  
10 waiting times, because this is what the SV algorithm does that the baseline Bell and Wong  
11 Nearest Neighbors (BWNN) algorithm does not. This entails some risk of making wasted empty  
12 vehicle trips, but, in the case studies considered here, SV yields large relative reductions (up to  
13 96%) in waiting time with smaller relative increases in empty vehicle travel time, particularly at  
14 high intensities (around 22% at intensity 0.8 on the Grid network, for example). This increase in  
15 empty vehicle travel contributes to network congestion, which is not modeled here. In a system  
16 operating near its congested limit, other algorithms (or variants of SV that do model congestion)  
17 might perform better than SV as described here.

18 The Static Nearest Neighbors (SNN) algorithm gives an approximate solution to the static  
19 EVR problem, in which there is perfect information of all future passenger requests. We have  
20 thus used the SNN algorithm to give a rough estimate of the potential for further waiting time  
21 reductions for the dynamic EVR problem, in which only the mean request rates are known. For  
22 example, such reductions might be obtained by allowing empty vehicles to be rerouted, or by the  
23 use of empty vehicle sidings (or *car barns* (4)), which act as buffers for empty vehicles. SV (as  
24 currently defined) cannot deal with empty vehicle sidings, because they are effectively stations  
25 with no requests, which thus receive no vehicles. Finally, improvements may be possible in the  
26 details of the sampling algorithm, as outlined in section 3.

27  
28



1 **REFERENCES**

- 2 1. ULTra PRT. ULTra at London Heathrow Airport.  
3 [www.ultraprt.com/applications/existing-systems/heathrow](http://www.ultraprt.com/applications/existing-systems/heathrow). Accessed July 31, 2010.
- 4 2. Lees-Miller, J. D., J. C. Hammersley and R. E. Wilson. Theoretical Maximum  
5 Capacity as a Benchmark for Empty Vehicle Redistribution in Personal Rapid Transit. In  
6 *Transportation Research Record: Journal of the Transportation Research Board*. No. 2146,  
7 Transportation Research Board of the National Academies, Washington, D.C., 2010.
- 8 3. Andréasson, I. Vehicle Distribution in Large Personal Rapid Transit Systems. In  
9 *Transportation Research Record: Journal of the Transportation Research Board*. No. 1451,  
10 Transportation Research Board of the National Academies, Washington, D.C., 1994, pp. 95–99.
- 11 4. Anderson, J. E. Control of personal rapid transit systems. *Journal of Advanced*  
12 *Transportation*, vol. 32, no. 1, 1998.
- 13 5. Andréasson, I. Quasi-Optimum Redistribution of Empty PRT Vehicles. In *Automated*  
14 *People Movers VI: Creative Access for Major Activity Centers*, W. J. Sproule, E. S. Neumann  
15 and S. W. Lynch, eds. American Society of Civil Engineers, Reston, VA, April 1998, pp. 541–  
16 550.
- 17 6. Andréasson, I. Reallocation of empty personal rapid transit vehicles en route. In  
18 *Transportation Research Record: Journal of the Transportation Research Board*. No. 1838,  
19 Transportation Research Board of the National Academies, Washington, D.C., 2003, pp. 36–41.
- 20 7. Toth, P. and D. Vigo. *The vehicle routing problem*. SIAM monographs on discrete  
21 mathematics and applications, Society for Industrial and Applied Mathematics, 2002.
- 22 8. Gendreau, M. and J.-y. Potvin. Dynamic Vehicle Routing and Dispatching. In *Fleet*  
23 *management and logistics*, T. G. Crainic and G. Laporte, eds., Kluwer, Boston, chap. 5, 1998,  
24 pp. 115–126.
- 25 9. Psaraftis, H. N. Dynamic vehicle routing: Status and prospects. *Annals of Operations*  
26 *Research*, vol. 61, no. 1, December 1995, pp. 143–164.
- 27 10. Yang, J., P. Jaillet and H. Mahmassani. Real-Time Multivehicle Truckload Pickup  
28 and Delivery Problems. *Transportation Science*, vol. 38, no. 2, May 2004, pp. 135–148.
- 29 11. Bent, R. W. and P. Van Hentenryck. Scenario-Based Planning for Partially Dynamic  
30 Vehicle Routing with Stochastic Customers. *Operations Research*, vol. 52, no. 6, 2004, pp. 977–  
31 987.
- 32 12. Hvattum, L. M., A. Løkketangen and G. Laporte. Solving a Dynamic and Stochastic  
33 Vehicle Routing Problem with a Sample Scenario Hedging Heuristic. *Transportation Science*,  
34 vol. 40, no. 4, November 2006, pp. 421–438.
- 35 13. Mitrovic-Minic, S., R. Krishnamurti and G. Laporte. Double-horizon based heuristics  
36 for the dynamic pickup and delivery problem with time windows. *Transportation Research Part*  
37 *B: Methodological*, vol. 38, no. 8, September 2004, pp. 669–685.
- 38 14. Berbeglia, G., J.-F. Cordeau and G. Laporte. Dynamic pickup and delivery problems.  
39 *European Journal of Operational Research*, vol. 202, no. 1, April 2010, pp. 8–15.
- 40 15. Powell, W. B. A Stochastic Formulation of the Dynamic Assignment Problem, with  
41 an Application to Truckload Motor Carriers. *Transportation Science*, vol. 30, no. 3, August  
42 1996, pp. 195–219.
- 43 16. Berbeglia, G., J.-F. Cordeau, I. Gribkovskaia and G. Laporte. Static pickup and  
44 delivery problems: a classification scheme and survey. *TOP*, vol. 15, no. 1, July 2007, pp. 131.
- 45 17. Fischetti, M., G. Laporte and S. Martello. The Delivery Man Problem and Cumulative  
46 Matroids. *Operations Research*, vol. 41, no. 6, 1993, pp. 1055–1064.

- 1           18. Larsen, A., O. B. G. Madsen and M. M. Solomon. The A Priori Dynamic Traveling  
2 Salesman Problem with Time Windows. *Transportation Science*, vol. 38, no. 4, 2004, pp. 459–  
3 472.
- 4           19. Bertsimas, D. J. and D. S. Levi. A New Generation of Vehicle Routing Research:  
5 Robust Algorithms, Addressing Uncertainty. *Operations Research*, vol. 44, no. 2, 1996, pp. 286–  
6 304.
- 7           20. Swihart, M. R. and J. D. Papastavrou. A stochastic and dynamic model for the single-  
8 vehicle pick-up and delivery problem. *European Journal of Operational Research*, vol. 114, no.  
9 3, May 1999, pp. 447–464.
- 10          21. Bell, M. G. H. and K. I. Wong. A Rolling Horizon Approach to the Optimal  
11 Dispatching of Taxis. In *Transportation and traffic theory : Flow, dynamics and human*  
12 *interaction : proceedings of the 16th International Symposium on Transportation and Traffic*  
13 *Theory*, H. S. Mahmassani, ed. Elsevier, Amsterdam, July 2005, pp. 629–648.
- 14          22. Irving, J. H., H. Bernstein, C. L. Olson and J. Buyan. *Fundamentals of Personal*  
15 *Rapid Transit*. Lexington Books, D.C. Heath and Company, 1978.
- 16          23. Lees-Miller, J., J. Hammersley and N. Davenport. Ride Sharing in Personal Rapid  
17 Transit Capacity Planning. In *Automated People Movers 2009*, R. R. Griebenow, ed.  
18