Lees-Miller, JD., Hammersley, J., & Davenport, N. (2009). *Ride sharing in Personal Rapid Transit capacity planning*.
<http://hdl.handle.net/1983/1414>

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/>

Ride Sharing in Personal Rapid Transit Capacity Planning

John Lees-Miller¹, John Hammersley² and Nick Davenport²

¹Department of Engineering Mathematics, University of Bristol, Queen's Building, University Walk, Clifton, Bristol BS8 1TR. Email: enjdlm@bristol.ac.uk

²Advanced Transport Systems Ltd, Unit B3, Ashville Park, Short Way, Thornbury, Bristol BS35 3UU. Email: johnhammersley@atsltd.co.uk; ndavenport@atsltd.co.uk

ABSTRACT

Personal Rapid Transit (PRT) systems are designed so that passengers usually travel together only by choice, but strangers may choose to share a vehicle at peak times, when the system is near capacity. By predicting whether and to what extent this *ride sharing* will occur, PRT planners can better estimate the impact on system capacity and passenger experience. This paper develops a model for ride sharing based on queueing theory and applies it to explain the relationships between vehicle occupancy, passenger queue length and passenger waiting time. The effects of multiple destinations, passengers who are unwilling to share and passengers arriving in preformed parties are considered. A case study is provided to show how the model can be applied to a simple point-to-point system; in this case study it appears possible to reduce the size of the vehicle fleet by at least 30%, while still maintaining a high level of service for passengers during peak times.

1. INTRODUCTION

A Personal Rapid Transit (PRT) system provides on-demand, non-stop transportation using compact, computer-guided vehicles running on a dedicated network of guideways. In normal operation, each vehicle carries an individual passenger or a small party traveling together by choice; each party (an individual is a party of one) travels directly from their origin to their destination, without sharing with other parties, stopping or changing vehicles. However, during peak times, the number of vehicles required to provide one vehicle per party may be prohibitively large. In this paper, we consider *ride sharing*, in which several parties may choose to share a vehicle.

Previous work indicates that ride sharing can greatly reduce the number of vehicles needed to provide an acceptable level of service during peak times. Johnson (2005) reports that peak capacity for a given fleet size is roughly doubled, using a model with a single origin station and several equally likely destinations. He also discusses the passenger experience in the origin station and describes a station management strategy that facilitates ride sharing. Johnson's ride sharing model does not explicitly represent the passenger arrival process; instead, immediately after a passenger is served, a new passenger arrives to replace him, thus maintaining a queue of constant length. While this is analytically convenient, it is difficult to justify, and it limits the utility of the model for PRT planning, because the passenger arrival process is a

crucial input in the planning process. Andréasson (2005) gives a good overview of the operational issues created by ride sharing, including the implications for passenger safety and security. He reports a similar increase in capacity in a full system simulation for a single case study. However, the paper does not explore how these results may be generalized to other systems. Also, both authors assume that all passengers who have the opportunity to share will choose to do so, which is a potentially misleading assumption.

In this paper, we develop improved models for ride sharing in simple networks, discuss aspects of PRT system design and operation in the context of these models, and show how these models can be used in capacity planning. Section 2 explains a ride sharing model based on queueing theory, and section 3 shows how to use this model to explain the effect of ride sharing on system capacity; it also compares our results to those in the literature. A discussion of the effects of passengers who are unwilling to share follows in section 4, and section 5 explores the effects of larger non-separable parties on system capacity. Section 6 is a case study that shows how the models in this paper can be applied to a simple point-to-point system; it also deals briefly with the questions of how to operate stations to facilitate ride sharing and how to account for demand that changes with time.

2. A QUEUEING THEORY MODEL FOR RIDE SHARING

Consider a system with one origin and N destinations, where all passengers are traveling from the origin to one of the N destinations. When $N = 1$, this models a system of two stations or regions with dominant tidal demand from one to the other, like the point-to-point system studied in section 6. When $N > 1$, the model might describe traffic from a transit hub to several buildings, for example.

Parties arrive at the origin station bound for destination i according to a Poisson process with rate λ_i , in parties / hour. Assuming these N arrival processes are independent, the aggregate arrival process is also a Poisson process, with rate $\lambda = \lambda_1 + \dots + \lambda_N$. Upon arriving at the origin station, passengers queue in first-in-first-out order, each waiting for a vehicle to serve them. There are s vehicles in the fleet, each of which can carry up to C parties with the same destination. Any vehicle can serve any destination, but it serves only one destination on a given trip; when a vehicle becomes available, the first party in the queue determines its destination, and up to $C - 1$ other waiting parties with the same destination can board. The vehicle then takes d hours to serve the group and return to the origin station; these service times could vary between destinations, but for simplicity we fix them all at d . Note that a vehicle cannot leave the origin when empty; it must wait for at least one party to board. The following approximations are implicit in this model; we revisit some of them later on in the sections indicated.

- A1. The service time d is approximately deterministic because it is dominated by the vehicle round trip transit time, from the origin to the destination and back to the origin; the true transit time also includes stochastic terms for passenger loading and unloading, and for delays due to network congestion, but these are less important when the origin and destination are reasonably far apart.

- A2. While the capacity of a vehicle is a constant number of passengers, the number of *parties* it can carry depends on the number of passengers per party, which is stochastic. For simplicity, we scale the mean passenger demand and the vehicle capacity into parties; if each vehicle seats 4 passengers, and we expect less than 1.33 passengers per party, we set $C = 3$ parties and scale the demand λ appropriately. This is only approximately correct (see section 5).
- A3. The system capacity is limited by the number of vehicles available. Another limiting factor that we do not consider is the station throughput at the origin; this is mainly a function of the number of berths in the station, so we effectively assume that the origin station is large.
- A4. The total party arrival rate λ is constant over time (but see section 6).
- A5. Once the first party in the queue has determined a vehicle's destination, parties with the same destination can share the vehicle, regardless of their position in the queue (but see section 6).
- A6. All parties who can share will choose to do so (but see sections 4 and 6).

More formally, our model is known as an $M^N/D^C/s$ queueing system, in the notation of Cromie and Chaudhry (1976) and Huang (2001), which is based on the standard Kendall notation. The M^N refers to the Markovian (Poisson) arrival process with N destinations. The D^C refers to the deterministic service times and *bulk service* rule, where each vehicle has capacity C . The s denotes the number of servers; that is, we treat each vehicle as a server.

To our knowledge, there are no useful analytical results for the performance measures of the $M^N/D^C/s$ queueing system, in the literature. Cromie and Chaudhry (1976) give useful analytical results for many performance measures of the $M^1/M^C/s$ queueing system, in which service times are Markovian, rather than deterministic. While there is some variation in the service times, which we have neglected, using a Markovian service model introduces far more variation than is desirable; this is why we have not chosen an $M^N/M^C/s$ queueing system as the basis for our analysis. Tijms (2006) gives useful approximations for the $M^1/D^1/s$ system, but ride sharing is not allowed when $C = 1$. Even these analytical results are only suitable for computer calculation; we use them to validate the statistical properties of our simulations when $N = 1$ and $C = 1$. Huang et al. (2001) derive analytical results for an $M^N/M^C/s$ queueing system, in the context of semiconductor manufacturing, but again they assume Markovian services, and they use a 'largest batch first' service discipline that is not appropriate for our application. The value of the $M^N/D^C/s$ model is as a theoretically sound starting point for further extensions. We rely on Monte Carlo simulation to obtain quantitative data on our models, but we note that these models are well-suited to computer implementation, so this is not an onerous limitation. In all of our figures, each point is the mean of ten runs of one million seconds each, unless otherwise noted.

3. SYSTEM CAPACITY WITH RIDE SHARING

We now apply our model to explore the effects of ride sharing on *system capacity*, which is the largest number of parties that the system can serve per hour. When the

system is saturated, vehicles become available for service at rate $\mu = s / d$ vehicles per hour, and all vehicles operate at their full capacity, C , so the system capacity is μC parties per hour. That is, if the party arrival rate λ remains constant (assumption A4) at or above μC , the number of waiting parties grows without bound. So, for fixed fleet size s and service time d , increasing the vehicle capacity C results in a proportional increase in system capacity. Figure 1 shows this effect; when $C = 1$, no ride sharing is allowed, and the queue grows without bound as the arrival rate λ exceeds 110 parties/h. For $C = 2$, divergence is delayed until λ exceeds 220 parties/h. This increase in capacity is explained by an increase in mean vehicle occupancy, which approaches the vehicle capacity ($C = 2$), as λ exceeds 220 parties/h.

Figure 1: Mean queue length and mean occupancy for fixed service time and fleet size, with increasing party arrival rate ($d = 660s$, $s = 20$ vehicles).

This increase in mean vehicle occupancy requires an increase in mean queue length. When a vehicle becomes available, only those parties currently waiting in the queue can share with one another. If there are fewer than C parties (with the same destination) in the queue, then the vehicle makes that trip at less than full occupancy. The queue length fluctuates because of randomness in the arrival process, but high mean occupancy requires, on average, a standing queue. Moreover, as a consequence of Little's Law (Little 1961) for queueing systems, the mean party waiting time is directly proportional to the mean queue length; so, using larger vehicles increases system capacity at the cost of increased passenger waiting time. The degree to which ride sharing can increase capacity in practice thus depends on how much additional waiting time the passengers will accept; we return to this subject in section 6.

Next, we consider systems with more than one destination and compare these results with the existing results in the literature (Johnson 2005). For simplicity, we assume that the demand is split evenly among the N destinations. Then, in a queue of a given number of parties, the number of parties that are bound for any particular destination is inversely proportional to N . Only parties with the same destination can share a

vehicle, so for larger N , a longer queue is needed to achieve a given increase in the mean vehicle occupancy, and hence the system capacity. This suggests that ride sharing is most effective when the number of destinations is small.

To quantify this, and for comparison with Johnson's results, we refer to Figure 2, which shows a linear relationship between mean passenger waiting time and the number of destinations. Johnson also finds a linear relationship between mean waiting time and the number of destinations, but for a different ride sharing model. In Johnson's model, the arrival process is chosen so that the queue length is held constant at $N + 1$ parties, in order to make the model more tractable. The mean waiting time is then $(N + 1)/(2\mu)$, in our notation; that is, the constant of proportionality is fixed at $1/(2\mu)$. Figure 2 shows that the constant of proportionality varies with the total arrival rate. In this sense, Johnson's results also hold in our model, for a limited number of arrival rates. It is also worth remarking that passenger waiting time increases considerably as the number of destinations grows; when $\lambda = 180$ parties/h and $N = 1$, passengers wait 0.4 minutes on average, but when N increases to 24, as in Johnson's paper, this increases to 4.8 minutes. This indicates that ride sharing is less helpful for such a large number of destinations.

Figure 2: Mean waiting time for fixed service rate and several total arrival rates, with increasing number of destinations ($C = 3$, $d = 660s$, $s = 20$).

Figure 3: Savings in fleet size are sensitive to the percentage of parties that are willing to share ($C = 3$, $d = 660s$, $N = 1$).

4. PASSENGER WILLINGNESS TO RIDE SHARE

All ride sharing models that we are aware of (Johnson 2005; Andréasson 2005) allow parties to choose whether to ride share with other parties. These models also assume that all parties are willing to share (assumption A6), which is potentially misleading. There are many factors that can influence whether a party is willing to share; here, we restrict our analysis to waiting time, monetary incentives, and peer pressure.

Another major factor is the station design; while this is largely below the resolution of our model, we return to it briefly in section 6.

If passengers are rational, and they act to minimize their remaining waiting time, no sharing will occur. This is because the first party in the queue must consent to sharing their vehicle, something which gives them no waiting time benefit; once they have selected their destination, they can either choose to share, in which case they incur a small extra wait due to other passengers loading, or not to share, in which case they leave as soon as the vehicle arrives. Thus, although parties further back in the queue can usually reduce their remaining waiting time by sharing, the passenger at the front of the queue has no incentive (in terms of waiting time) to allow others to share his vehicle.

There are, however, two mitigating factors. Firstly, we speculate that there is considerable peer pressure to allow sharing when in a crowded station; taking a private vehicle might be frowned upon by those left waiting in the queue. This effect can only be quantified by experiment. Secondly, the operator can adjust the fare policy to offer a monetary incentive for sharing. Suggestions include charging by vehicle rather than by person (Andréasson 2005), or giving a discount to those who are willing to share (Andréasson 2005; Johnson 2005). A more thorough analysis of such fare collection policies is required, but it is beyond the scope of this paper. We also note that some systems (in airports, for example) are operated without fares; in these systems, peer pressure is the only incentive for sharing.

While further experiments and analysis are needed to properly answer these social engineering questions, our model can be modified to provide some sensitivity analysis. We consider the effect of varying a fixed probability w that a party is willing to share; so far, w has been 100%. This fixed probability is a fairly crude approximation, because it assumes that a party's decision on whether or not to share is entirely intrinsic; in reality, it may depend in a complex way on the actions of other parties around them. For example, parties may see that a vehicle is filling up and become less willing to share, further preventing high occupancies. However, this assumption provides a reasonable starting point.

Figure 3 shows the effect of w on the number of vehicles needed to ensure that 90% of parties wait less than 60s (see also section 6). For example, when $\lambda = 180$ parties/hr and $w = 100\%$, the number of vehicles required is reduced by 46% (from 39 to 21, in the particular system under study). When w drops to 80%, the required fleet size is reduced by only 30% (to 27). We note that a small change in w when w is near 100% can significantly affect the required fleet size; that is, system capacity is quite sensitive to w . The main reason is that the probability of n parties sharing is w^n , so achieving high vehicle occupancy ($1 < n < C$) requires a disproportionately longer queue as w decreases.

5. THE EFFECT OF PARTY SIZES ON RIDE SHARING

We have so far assumed that a vehicle can always carry up to three parties (assumption A2). In reality, party sizes will vary stochastically, allowing a possible conflict between the number of passengers arriving in a new party and the number of

remaining empty seats in a vehicle. In this case, the arriving party will have to decide on whether to split up or stay together. The distribution of party sizes differs considerably between applications. For example, many of the parties in a theme park will be families, and each family would require their own vehicle; ride sharing would be less effective in this case. In most applications that the authors have considered, however, the vast majority of parties will be individuals or pairs. We now explore several relaxations of assumption A2 to assess its validity.

We consider a model in which parties arrive according to a Poisson process, but, each time a party arrives, X passengers with the same destination join the queue; here, X is a random variable taking positive integer values. This is known as a compound Poisson process (Woodward 1994). The vehicle capacity in this model is defined to be S passengers, rather than C parties. The distribution of X would be based on the actual group size data for the application under study, but here we use a parameterized distribution. For simplicity, we still assume that party sizes cannot exceed vehicle capacity (no party has more than S passengers). We also note that in many applications, parties arrive by automobile, and so the party size is limited by the capacity of a typical automobile. These considerations lead us to define X by a binomial distribution with

$$\Pr(X = x) = \binom{S-1}{x-1} p^{x-1} (1-p)^{S-x}, \quad x = 1, 2, \dots, S$$

where $p = (G - 1)/(S - 1)$ and G is the mean party size. This means that a party consists of at least one passenger, accompanied by up to $S - 1$ additional passengers; each additional passenger occurs with probability p . In general, the group size distribution could vary between destinations, but we ignore this for the sake of simplicity. The distribution of X when $G = 1.33$ is computed in Table 1.

X	1	2	3	4
$\Pr(X = x)$	0.705	0.261	0.032	0.001
$\Pr(X \leq x)$	0.705	0.966	0.999	1.000

Table 1: Distribution of group size X when mean group size $G = 1.33$.

The cumulative density indicates a 97% chance of party size one or two, and our simulations indicate that this does not have a significant impact on the number of vehicles required to provide satisfactory service; assumption A2 is a reasonable approximation when $G \leq 1.33$. This is the case for most applications. For mean party sizes up to $G = 2$, results are mainly the same, but when $G = 3$ there is a 70% chance of a party with size 2 or 3, and assumption A2 significantly overestimates the potential for ride sharing.

6. POINT-TO-POINT PRT SYSTEM CASE STUDY

We now apply our model to a simple but useful PRT system that connects two locations, where we assume that there is one station in each location and that the network layout and the peak demands are given. Our objective is to determine how many vehicles are needed in order to provide an acceptable level of service. The level of service is defined in terms of the 90th percentile of the party waiting time distribution; for example, service might be acceptable when 90% of parties arriving

in the peak period wait less than 1 minute before boarding a vehicle. For each combination of peak demand and fleet size, the peak period is simulated 1000 times to build an accurate estimate of the waiting time distribution; we then choose the smallest fleet size that provides an acceptable level of service for all of the expected peak demands.

We have so far assumed that the passenger demand λ is constant (A4), but this is not usually true in peak periods. The party arrival rate will usually rise to a maximum and then fall off. It is straightforward to extend our model to capture this. We use the representative demand profiles shown in Figure 4, which were generated from Gaussian curves with “standard deviations” of 15 minutes for the AM peak and 30 minutes for the PM peak. The AM peak is 2 hours long and is sharper and higher than the PM peak, which is 3 hours long. The simulator records waiting times for all passengers arriving in peak hours, and it terminates upon serving the last passenger that arrived during the peak. Waiting times for passengers who arrive after the peak are discarded; waiting times from the first two hours are also discarded, to reduce the importance of the simulator’s initial conditions (all vehicles begin at the origin, ready to serve passengers).

Figure 4: Demand used for AM peak and PM peak simulations.

We have also assumed that parties anywhere in the queue can share with one another (assumption A5). Whether this can be achieved in practice depends on how the stations operate; our assumptions about this are as follows. Each station contains a fixed number of *berths*, at which parties can load into or unload from vehicles. Each berth has a destination selection panel, with which a party tells the system where they are traveling to. This layout is typical of stations in the ULTra PRT system, developed by Advanced Transport Systems Ltd.; it differs from the station layout in (Johnson 2005), which separates destination selection from berths, but the following discussion suggests that our layout can also facilitate ride sharing.

At low intensity, there will usually be some empty vehicles parked in the berths, waiting for passengers to arrive. Ride sharing is unlikely at low intensity, because parties will arrive, choose a berth, select their destination and then depart immediately on a waiting empty vehicle. However, at high demand there will usually be a queue of parties waiting for vehicles (section 2). We assume that the party at the head of the queue will go to a free berth, select their destination and wait there, while the other parties wait in first-in-first-out order. When a party selects their destination, they are asked whether they want to share (section 4); if they choose to share, their destination is displayed on a screen above their berth. Other parties with the same destination can then “jump the queue” to share a vehicle with that party.

It is unlikely that the station process outlined above will be perfectly efficient (assumption A5). The apparent complexity of the human factors involved suggests that more work, including experimental work, is required in this area. For now, we examine what happens when parties can only communicate with their immediate neighbors in the queue; this assumption is intended to provide a lower bound on the likely level of interaction between parties in a station. When the “neighbors only?” column in Table 2 is “Y,” the parties can only share with their neighbors; otherwise, assumption A5 is in effect.

To fix the remaining parameters, we set the vehicle capacity at $C = 3$ parties (see assumption A2 and section 5) and the vehicle round trip time at $d = 660s$ (ten minutes travel plus one minute for passenger loading and unloading; see assumption A1). Table 2 shows the predicted fleet sizes for several ride sharing scenarios. The “% willing to share” column corresponds to the probability of a party sharing, as defined in section 4. We consider two possible definitions of acceptable service, one where 90% of parties wait less than 1 minute, and another where 90% of parties wait less than 3 minutes.

peak profile	% willing to share	neighbors only?	vehicles needed for "90% wait < ..."	
			60s	180s
AM	0		53	47
AM	60	Y	44	36
AM	60	N	40	32
AM	80	Y	37	28
AM	80	N	34	25
AM	100		26	18
PM	0		45	41
PM	60	Y	37	31
PM	60	N	34	27
PM	80	Y	31	24
PM	80	N	29	21
PM	100		22	15

Table 2: Fleet sizes for case study system under varying ride sharing assumptions.

First, the AM peak consistently requires more vehicles than the PM peak, so the AM peak determines the fleet size. At the “90% wait < 60s” service level, the system requires 53 vehicles if no ride sharing is allowed, but only 26 vehicles under the most optimistic ride sharing assumptions; this is a 51% reduction, which is in line with other results in the literature (section 1). If a lower service level is acceptable, the savings can be greater; at the “90% wait < 180s” service level, and under the most optimistic ride sharing assumptions, the fleet size is reduced by 62%, from 47 vehicles to 18 vehicles. This is because longer waiting times imply longer queue lengths, which in turn allow increased vehicle occupancy, as discussed in section 3.

When not all parties are willing to share, or the communication between parties in the station is more limited, the savings due to ride sharing are reduced, but still significant. Assuming that 80% of parties are willing to share, and that parties are limited to sharing with their neighbors, the fleet size required to provide the higher service level is reduced by 30%, from 53 vehicles to 37 vehicles. The fleet size required to meet the lower service level is reduced by 40%, from 47 vehicles to 28 vehicles. This is still a substantial reduction, but, as noted in section 4, these results are quite sensitive to the fraction of parties that are willing to ride share; when this drops to 60%, the corresponding reductions in fleet size are 17% and 23%.

The number of extra vehicles required because of the “neighbors only” restriction is fairly small (on the order of 10%) in the system under study, because there is only one destination and most parties are willing to share. Its effect is larger when there are more destinations; if there are two equally likely destinations, a party with a given number of neighbors is only half as likely to find a suitable party to share with. Our model indicates that for a similar system with two destinations and the AM peak demand split evenly between them, 45 vehicles are required to provide “90% wait < 60s” when 80% of parties are willing to share; this is a 15% reduction from the number required when there is no ride sharing at all. When there are multiple destinations, the station signage and layout become much more important.

7. CONCLUSIONS

The aim of this paper was to establish a suitable model to analyze the effects ride sharing has on PRT system performance, and examine how station design and passenger behavior factors should be taken into account. To this end, we developed a model for ride sharing based on queueing theory, and although the model requires a number of assumptions (see section 2), we believe it is a sound basis for analysis, and it provides an alternative to anything found already in the literature. This model was then used to explain the relationship between occupancy and queue length in the presence of ride sharing, and to demonstrate the effect increasing the number of destinations has on these relationships, comparing our results with those in the current literature.

A crucial issue seemingly ignored in previous studies is the *willingness* of passengers to rideshare; in both Johnson (2005) and Andréasson (2005) it is assumed that all parties are perfectly willing to share. As discussed in section 4, if all parties behave rationally and seek to minimize their waiting time, no ridesharing will occur as it is

the decision of the party at the head of the queue whether to share or not, and they get no benefit from doing so. Whilst incentives such as peer pressure and monetary savings may increase the likelihood of ride sharing occurring, as the effect of unwillingness to share on the beneficial effects of ridesharing is quite pronounced (see figure 3), one must take this issue into account in any analysis.

A factor which appears to have a much smaller effect is the arrival party size; although larger, non-separable parties reduce the mean vehicle occupancy, this reduction is only significant when the mean party size approaches three. Thus under our assumption of less than 1.33 passengers per party (assumption A2), this effect is negligible.

In the case study of section 6, our models were applied to a point-to-point system to determine the required fleet size to provide an acceptable level of service. In order to more realistically approximate peak period behavior, we dropped the assumption of a constant demand (A4) and instead used the two profiles shown in figure 4, representing AM and PM peaks. The simulation results presented in table 2 reveal that it is the sharper and higher AM peak which determines the fleet size, and under most optimistic ride sharing assumptions, we find a 51% reduction in the number of vehicles required at the “90% wait < 60s” service level, consistent with the findings in other literature (section 1).

What our results also show, however, is the reduction in savings one obtains if some passengers are unwilling to share, or the station isn’t properly designed to promote ridesharing. At the same service level, but only assuming 80% of parties are willing to share, and that parties are limited to sharing with their neighbors, the reduction in fleet size drops to 30% (from 51%), and if the willingness is further reduced to 60%, the saving on vehicles is only 17%. Generating an environment which encourages passengers to rideshare at busy times is thus very important for it to be effective in allowing for smaller fleet sizes.

Facilitating the passenger’s ability to rideshare also plays a crucial role, as the final analysis of section 6 demonstrated; for a station with two equally likely destinations, a willingness to share of 80%, and neighbors only interactions, the fleet size was only reduced by 15% (rather than 30% in the single destination case).

Thus the optimistic projections of a 50% reduction in fleet size requirements due to ride sharing need to be tempered by the observations that such a figure makes potentially unrealistic assumptions about passenger behavior and station design. In order to achieve a benefit anywhere close to this figure when there are multiple destinations, station design (signage and layout) needs to be carefully considered so as to both facilitate and provide sufficient incentives for ride sharing in PRT.

ACKNOWLEDGEMENTS

For useful discussions and feedback, the authors wish to thank Prof. Martin Lowson and Phil Smith of Advanced Transport Systems Ltd, and Dr. R. E. Wilson of Bristol University Department of Engineering Mathematics. This work was partly funded by the CityMobil Sixth Framework Programme for DG Research Thematic Priority 1.6, Sustainable Development, Global Change and Ecosystems, Integrated Project, Contract Number TIP5-CT-2006-031315.

REFERENCES

- Andréasson, I. (2005). "Ride-sharing on prt." *Automated People Movers*.
- Cromie, M. V. and Chaudhry, M. L. (1976). "Analytically explicit results for the queueing system M/Mx/C with charts and tables for certain measures of efficiency." *Operational Research Quarterly (1970-1977)*, 27(3), 733-745.
- Huang, M.-G., Chang, P.-L., and Chou, Y.-C. (2001). "Analytic approximations for multiserver batch-service workstations with multiple process recipes in semiconductor wafer fabrication." *Semiconductor Manufacturing, IEEE Transactions on*, 14(4), 395-405.
- Johnson, R. E. (2005). "Doubling personal rapid transit capacity with ridesharing." *Transit: Intermodal Transfer Facilities, Rail, Commuter Rail, Light Rail, and Major Activity Center Circulation Systems*, 1930, 107-112.
- Little, J. D. C. (1961). "A proof for the queueing formula: $L = \lambda W$." *Operations Research*, 9(3), 383-387.
- Tijms, H. (2006). "New and old results for the M/D/c queue." *AEU – International Journal of Electronics and Communications*, 60(2), 125-130.
- Woodward, M. E. (1994). "Modelling with discrete-time queues." *IEEE Computer Society Press, Loughborough*.