Using a multi-agent simulation tool to estimate the car-pooling potential

Author(s):
Dubernet, Thibaut Jean Pierre; Rieser-Schüssler, Nadine; Axhausen, Kay W.

Publication date:
2013

Permanent link:
https://doi.org/10.3929/ethz-b-000052402

Using a multi-agent simulation tool to estimate the car-pooling potential

Date of submission: 2012-07-12

Thibaut Dubernet (corresponding author)
Institute for Transport Planning and Systems (IVT), ETH Zurich, CH-8093 Zurich
phone: +41-44-633 68 65
fax: +41-44-633 10 57
thibaut.dubernet@ivt.baug.ethz.ch

Nadine Rieser-Schüssler
Institute for Transport Planning and Systems (IVT), ETH Zurich, CH-8093 Zurich
phone: +41-44-633 30 85
fax: +41-44-633 10 57
nadine.rieser@ivt.baug.ethz.ch

Kay W. Axhausen
Institute for Transport Planning and Systems (IVT), ETH Zurich, CH-8093 Zurich
phone: +41-44-633 39 43
fax: +41-44-633 10 57
axhausen@ivt.baug.ethz.ch

Words: 4996 words + 6 figures + 4 tables = 7496 word equivalents

ABSTRACT

It is a general trend in transportation planning to try to minimize the negative externalities of the transport system as a whole, such as noise or pollutant emissions. One of the ways to achieve this is to reduce the number of cars on the roads, for instance by increasing car occupancy. This paper focuses on evaluating the potential of this possibility. The factors influencing this potential are manifold: behavioral, structural (number of potential matches), organisational (quality of available services to meet co-travelers)...

In previous studies, mainly the behavioral and organisational factors were analyzed. This paper focuses on the structural factor. To do so, the highly detailed daily plans generated by the multi-agent microsimulation software MATSim are searched for potential matches. Information about the potential matches is used to assess the feasibility of carpooling. In particular, it is shown that when considering only structural factors, it is possible to group most of the car trips into two-person car-pools.

The results of the analysis lead to the conclusion that there is no structural obstacle to carpooling development, and thus that the causes of the low share of this mode is to search in both the behavioral and organisational factors.

INTRODUCTION

It is a general trend in transportation planning to try to minimize the negative externalities of the transport system as a whole, such as noise or pollutant emissions. One of the ways to achieve this is to reduce the number of cars on the roads. This can be achieved by different means: efficient public transport, usage of individual non-motorised modes, such as bike or walk, or sharing of a vehicle for all or a fraction of a trip.

Assessing the potential of this latter possibility, known as carpooling or ride sharing, is challenging. By carpooling, one usually understands a formal service allowing passengers and drivers to meet, contrary to the more general case where individuals may share their vehicle with relatives and acquaintances. Several factors can modify greatly the potential for such a mode. The most obvious is probably the behavioral factor: carpooling can only become mainstream if individuals are willing to share their car with other, possibly unknown individuals. Second would come the organisational aspect: provided that individuals are willing to share their rides, how to put them efficiently in contact? Finally comes a structural aspect: the potential of such services is highly dependent on the number of possible matches available to an individual for a given trip.

This structural factor is perhaps the most rigid of the three: while services can be improved and the attitude toward carpooling can change to some extent, it is less likely that travel patterns change enough to lead to significant changes in the number of carpooling possibilities in a short period of time. Thus, the study of this factor can be used to determine an upper bound for the carpooling share. However, to the knowledge of the authors, this factor never was studied in isolation, and thus its relative importance to the more flexible factors is unclear. Due to the fact that this structural factor is highly dependent on travel patterns, the usage of an activity-based travel simulation tool, which produces such patterns for a population of agents, is a natural choice.

In this paper, the highly detailed daily plans produced by the MATSim microsimulation software are used to determine an upper bound for the carpooling share in the Zurich area in Switzerland. We start by reviewing different approaches used in past studies to assess the potential of car pooling. We then briefly present the MATSim software, as well as the definition of an acceptable match employed. We finish with the results from the analysis and an outlook.

RELATED STUDIES

In the past, various kind of studies were conducted to assess the potential of carpooling, or measure its market penetration. Those studies can be categorized as revealed preferences studies, stated preferences surveys, optimization algorithms, simulation studies, theoretical models, and platform design.

An obvious way of assessing the potential of carpooling is to examine the behavior of current users of this mode, by studying the success (or failure) of past or current implementations, or through revealed preferences surveys. Abrahamse and Keall analyzed the results of an initiative to increase the share of carpooling for commuter trips in New-Zealand *(1)*. Ferguson shows how increasing vehicle availability and decreasing fuel costs led to a decline of carpooling share for commuting in the United States in the eighties *(2)*. Burris and Winn conducted a survey to study the characteristics of the passengers of a special variant of carpooling, usually referred to as “slugging” or “casual carpooling” . This variant consists of spontaneous carpools, formed between strangers to meet the requirements of High Occupancy Vehicle lanes *(3)*. Morency used OD surveys in the Greater Montreal Area to study carpooling. An important result was that approximately 70% of the shared rides were performed between household members *(4)*. In those studies, the structural factors are "hidden" from the analyst, as it is difficult to determine to which extent individuals not choosing to car-pool do so by choice, rather than because of the unavailability of carpooling alternative.

Another kind of study is stated preferences surveys, where not-yet-users are included. Correia and Viegas conducted such a survey for the Lisbon Metropolitan Region, Portugal, to test the attractiveness of a "club" concept for a carpooling service *(5)*. Ciari and Axhausen conducted such a survey for Switzerland *(6)*. They used a Stated Choice survey to estimate a mode choice model, where "carpooling as a driver" and "carpooling as a passenger" were alternatives. However, Stated Preferences surveys for carpooling are limited, in the sense that they suppose that participants in a carpooling service will be able to find a match.

To consider this fact, using optimisation algorithms to actually compute matches is more pertinent. They can both be used to build co-travelers search services and to assess the maximum number of matches that obey some criterion in a sample population. A problem of this approach is that the complexity of the problem forbids to process large datasets. For example, the so-called "carpooling problem" has been formulated to translate the problem of identifying optimal matches into the terms of combinatorial optimisation, based on the more general "vehicle routing problem", where vehicles must be routed to serve customers while minimizing costs. In this problem, passengers and drivers are given as input, and the objective is to maximize the number of picked-up passengers while minimizing the overall cost, subject to constraints such as time windows, maximum driver detour or vehicle capacity. Even when restricting the problem to agents with the same destination (*e.g.* co-workers), the problem has non-polynomial complexity *(7)*. Fairly sophisticated heuristics were developed to handle large instances of such routing problems *(8)*, where "large instance" refers to hundreds of vehicles and thousands of customers. However, even those "large instance" solvers are unlikely to handle properly the very large instances corresponding to the population of a urban area, and thus cannot be used to study the structural factor. In the field of transportation planning, de Palma used optimization algorithms to generate carpooling matches for commuting trips on artificial instances *(9)*.

Some studies tried to combine the survey and optimisation approaches in a "simulation" setting, running a matching algorithm on a subset of a population derived using a behavioral model. For instance, Mühlethaler used the model estimated in *(6)* to identify carpooling drivers and passengers, which were then used as an input for various matching algorithms *(10)*. Such an approach is tractable only because of the low number of participants generated by the behavioral model.

Moreover, theoretical work has been done, to study the factors which could influence the usage of car pooling. For example, Huang *et al.* develop a simple model of participation in car-pools on a single road, and study the influence of toll policies *(11)*. Here, the structural factor is left unstudied, due to the unique OD and the non-consideration of departure times: it is assumed that any individual searching a match will find it.

Finally, applied research has been undertaken to design platforms and services allowing interested individuals to find passengers or drivers, without which formal car-pooling cannot be implemented. Different studies focus on different challenges. As early as in the mid seventies, work was done on designing computer software to match passengers and drivers *(12)*. Recent work focuses on the design of efficient web services *(13)*, or on the ability to provide real-time information *(14)*.

As can be seen from this analysis, the question of the importance of the structural factor is still open.

APPROACH

In order to assess the importance of the structural factor for car-pooling, we use the output of a MATSim run, in which we search for acceptable matches, given a criterion detailed below.

MATSim is an open-source software which uses the activity-based approach to search for a user equilibrium on activity-travel patterns. In this approach, agents, representing individuals, are assigned daily plans, consisting of activities located in time and space, with trips between those activities. Agents get a utility from the execution of their plan, which increases with activity execution time and decreases with travel time. Agents influence the score of each other via congestion. The MATSim process searches for an equilibrium using a co-evolutionary algorithm, where each agent performs an evolutionary algorithm to improve its daily plan. The search process is as follows: trips are executed in a traffic model to obtain travel time estimates, which are used to assign a score to daily plans. Then, a fraction of the agents modify their daily plans, and the execution and scoring stages are performed again. By iteratively creating new plans from old ones or selecting a past plan based on its score, the process converges towards a steady state, used for analysis.

For this analysis, we use the output of a run for a 10% sample of the population of the Zurich metropolitan area, Switzerland. In a 10% sample, each agent actually simulated is considered representing 10 individuals in the simulated world. The synthetic population was generated using the data of the Swiss Census for the socio-demographics. Activity-trip chains from the Swiss national travel diary survey were then drawn randomly for each agent, restricting the possible chains for an agent to records of persons with similar socio-demographics. Cross-border traffic is generated from border survey data. More details about the scenario generation can be found in *(15)*. The restriction from all of Switzerland to the Zurich urban area is done by retaining only agents that pass at least once during the day through a circle of 30km radius, centered on a central Place in Zurich ("Bellevue Platz"). Taking a 10% sample of those filtered agents results in a scenario with 196,947 agents.

As stated in the introduction, our approach for assessing the carpooling potential is to identify trips which are "close enough" to be joined in a carpooling trip. In our definition, a carpooling trip consists of a driver picking up one passenger at his origin, and dropping him off at his destination, during one trip. To identify which trips can be joined into a carpooling trip, the following criteria are used:

1. A *time window width* δ: picking up and dropping off the passenger should not result in starting the trip before $t_{d,i} - \delta$ or ending it after $t_{a,i} + \delta$, where $t_{d,i}$ (*resp.* $t_{a,i}$) is the departure (*resp.* arrival) of agent i (driver of passenger). The departure time is taken from the results of the simulation. The arrival time is computed using the estimated travel time (see below), rather than the simulated arrival time, to enforce travel time consistency.
2. A *maximum detour fraction* Δ_d: the travel time of the driver trip $d_{d,\text{joint}}$, including picking up the passenger at his origin and dropping him off at his destination, should not be greater than $(1 + \Delta_d)\, d_{d,\text{init}}$, where $d_{d,\text{init}}$ is the direct travel time for the driver. Note that this constraint has an effect only when $\Delta_d\, d_{d,\text{init}} < 2\,\delta$. Thus, it mainly serves to avoid increasing the travel time for shorter trips too much, whereas for long trips, only the time windows matter.

As one of the aims of carpooling is to reduce the number of cars on the roads, only car trips are considered as possible passenger trips. It should also be emphasized that the approach only identifies potential 2-persons car-pools: identifying potential 3 or 4 persons matches would imply computing detours for all combinations of drivers and passengers, which is intractable for real world scenarios.

Using simulation allows to consider the effect of congestion in travel time estimation, using link travel times observed in the simulation. However, a problem using such time-dependent travel time estimates is that this ideally requires computing the least cost path for every travel time estimation, as the optimal path varies with departure time. The cost of such a procedure is computationally too high. To nevertheless take congestion into account, the travel times are estimated using time-dependent link travel times along the free flow shortest path. Using the free flow shortest path, rather than time-dependant shortest paths, allows to make extensive use of caching, and decreases the number of shortest path computations needed substantially. The estimates remain reasonably accurate, as shown in Figure 1, which shows the ratio of the travel time observed in the simulation with the travel times estimated with the method detailed above. As agents in the simulation may use sub-optimal routes, the estimates obtained with this method can be lower than the simulated travel times, or higher when congestion makes travel time higher on the free flow shortest path than on the route selected by the agent.

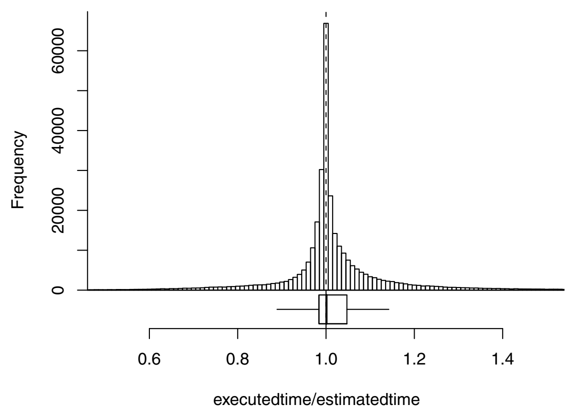

FIGURE 1 Estimated travel time vs. executed travel time

It should be emphasized that as such, the analysis does not generate an allocation of passengers to drivers: a driver can be identified as a driver for thousands of agents, and a passenger can be identified as a passenger for thousands of drivers. Moreover, each car trip is analysed as both a potential driver and a potential passenger.

The non-consideration of the combinatorial problem of finding *optimal* compatible matches allows the approach to be tractable for real-size scenarios.

However, the mere consideration of the number of acceptable matches does not allow to make conclusions about potential shares for carpooling. In particular, it does not take into account the fact that once two trips are matched, they cannot be part of another match. To take this fact into account, a simple match-generation approach is used: matches identified as acceptable are randomly selected, and all matches in which the selected driver or passenger participates are removed from the set of possibly matchable trips. The process is repeated until no further matching is possible. Repeating this random process a large number of times with different random seeds allows to get robust estimates of the feasible shares for carpooling.

The relevance of a match for the objective of reducing car traffic increases with the length of the passenger trip. However, it is also more difficult to find a potential driver for a longer passenger trip. Due to its random nature, the simple matching approach described above is likely to mainly group long driver trips with short passenger trips. To assess the importance of this factor, the above procedure can be used, but removing short trips from the set of possible passengers.

Finally, considering the fact that it is simpler to implement small scale carpooling services (for example at the firm level), the analysis was also conducted by restricting it to matches with the same origin or destination link.

RESULTS

This section presents the results from a run with $\delta = 30$min. and $\Delta_d = 0.5$. Table 1 shows some statistics about the run. It can be seen that thanks to the caching approach used for travel time estimation, computing the shortest path was actually necessary for only 0.47% of the travel time estimations. As shortest path computation is the most important part of the computation time, one can assert that without this approach, the computation time would have been around 200 times as big, making it roughly equal to 10 months.

TABLE 1 Summary statistics

Statistic	Value
# agents	196,947
# examined agents (w/ car trips)	110,468
# car trips	318,855
# car trips (> 10km.)	101,690
# matches found	510,230,774
# travel time estimates	7,382,434,602
# shortest path computations	35,013,474
runing time	37h. 16min.

Figure 2 presents the distribution of the minimal time window needed to perform each identified joint trips, that is, the time window from which a match is considered “acceptable”, given the maximum driver detour. The high number of identified trips with 0 duration time windows corresponds to trips for which the estimated time to pick up the passenger is shorter than the estimated direct time. Interestingly, the number of identified joint trips increases almost linearly with the minimum time window, before becoming relatively stable. That is, increasing a small time window leads to a greater relative increase of the number of possible joint trips than increasing an already large time window. This is probably due to the fact that a small time window decreases the possible detour: increasing a small time window does not only allow to match more departure/arrival times, but also allows the driver to search for a passenger further.

Figure 3 presents the distribution of driver detours. A quite important fraction of the identified trips correspond to “detours” of less than 1 (*i.e.* shorter travel times), due to the nature of the travel time estimation, as stated above. Half the identified joint trips correspond to increases in the driver's travel time which are lower than 20%.

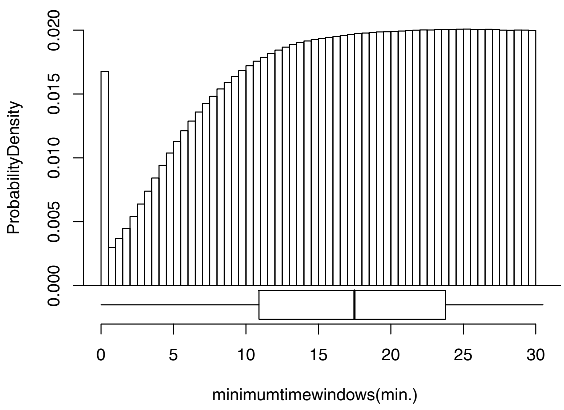

FIGURE 2 Distribution of the minimal time window, $\Delta_d = 0.5$

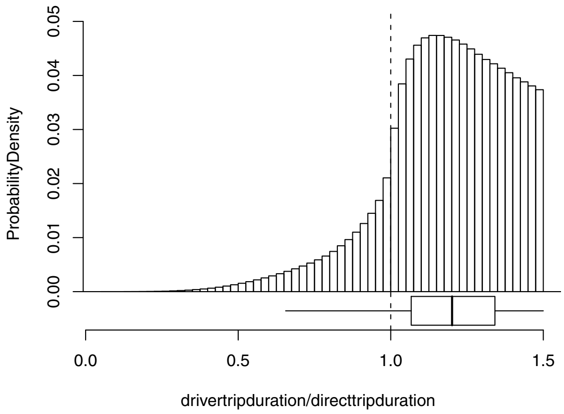

FIGURE 3 Distribution of the driver detours, $\delta = 30$min.

For the further analysis, the driver detour is limited to the more reasonable value of 0.15 (*i.e.* increases in the driver travel time are no more than 15%), and the time window to 15min..

Table 2 shows a summary of the number of potential matches found for the different criteria. The numbers of matches found in the search process are multiplied by 100: as each agent in the sample is considered as representing 10 individuals, one match in the search process represents all the possible combinations of those 10 represented passengers with those 10 represented drivers. After restricting the matches to trips with the same origin or destination link, less than 0.3% of the potential matches remain. The numbers remain high: even for the most restrictive conditions, each trip can be part of more than 13 joint trips, on average. When no restriction on origin or destination is imposed, for each trip, on average, more than 6,000 joint trips are possible.

TABLE 2 Number of matches for different conditions

	δ = 30min., Δ_d = 0.5		δ = 15min., Δ_d = 0.15	
	# matches	avg. # per trip	# matches	avg. # per trip
all matches	51,023,077,400	16002.0	19,911,749,700	6244.8
same origin	121,249,500	38.0	425,673	13.4
same destination	119,723,900	37.5	427,275	13.4

Table 3 shows a summary of the proportions of examined trips for which at least one possible driver or one possible passenger trip were identified, for the different settings. The matching process were run 300 times with different random seeds for each setting. The table displays the maximum and average fraction of car trips which have been matched as passengers.

For the less restrictive conditions, almost all car trips were identified as possible passengers or drivers (only 7 trips out of 318,855 could not be either passenger or driver). As passengers do not suffer any modification of the travel time, but only of the departure/arrival time, for all conditions more trips are identified as possible passengers than as possible drivers. Even for the most restrictive conditions, more than 60% of the trips could be identified as possible drivers or passengers, and more than 23% could actually be matched with a driver. For the condition δ = 30min. Δ_d = 0.5 , more than 45% could actually be matched with a driver. In other words, counting drivers and passengers, more than 90% of the trips could be matched.

To assess the impact of a long trip on the likelihood to find a passenger, the process was also run considering only trips of more than 10km as possible passengers. Table 4 presents the number of such trips which have been matched with a driver, both as a fraction of the number of car trips of less than 10km and of the overall number of car trips. When restricting the matching process to passenger trips of more than 10km, without OD restriction, the share of the considered potential passenger trips which are actually matched is higher than when also considering shorter passenger trips. Due to the fact that no restriction was imposed on drivers trips, dropping potential passenger trips of less than 10km does not modify the number of potential driver trips for the remaining passenger trips. However, the number of potential passenger trips per driver trip is lower. Thus, for each passenger trip, the probability that a given driver trip pertains to the choice set when the random matching is performed is higher, allowing for an increase in the likelihood of a match for those trips. Of course, this increase in the likelihood to be matched for long trips does result in a higher overall likelihood only if the likelihood of a match does not decreases to quickly with trip length. The fact that this increase in the overall likelihood is observed indicates that the potential of carpooling for long trips is real.

On the contrary, when restricting the search to trips with the same origin or destination, the share of potential passenger trips actually matched is lower than when also considering shorter passenger trips. Due to the already low number of possible matches, the decrease in the likelihood to find a match due to the trip length is not compensated by the higher likelihood of a driver trip to be available.

Of course, the potential may vary with time of day. Figure 4 and Figure 5 show the distribution of the number of identified possible trips as a passenger or a driver, per examined trip, as a function of the time of day. It can be seen that the number of potential matches remains

TABLE 3 Proportions of trips with matches

		δ = 30min., Δ_d = 0.5	δ = 15min., Δ_d = 0.15
all matches	driver	83.09%	78.12%
	passenger	99.85%	99.54%
	both	82.94%	77.74%
	either	100.00%	99.93%
	matched pass.: avg.	45.17%	43.68%
	matched pass.: max.	45.23%	43.89%
same origin	driver	45.45%	34.60%
	passenger	65.75%	46.52%
	both	30.07%	16.72%
	either	81.13%	64.41%
	matched pass.: avg.	30.78%	23.92%
	matched pass.: max.	30.84%	23.97%
same destination	driver	44.80%	33.92%
	passenger	65.48%	46.15%
	both	29.02%	16.22%
	either	80.95%	63.85%
	matched pass.: avg.	30.40%	23.67%
	matched pass.: max.	30.47%	23.73%

TABLE 4 Proportion of trips with passenger matches (trips longer than 10km)

		δ = 30min., Δ_d = 0.5		δ = 15min., Δ_d = 0.15	
relative to trips of:		> 10km	all lengths	> 10km	all lengths
all matches	avg.	50.38%	16.07%	49.38%	15.75%
	max.	50.49%	16.10%	49.50%	15.79%
same origin	avg.	26.07%	8.31%	16.46%	5.25%
	max.	26.17%	8.35%	16.52%	5.27%
same destination	avg.	25.66%	8.18%	16.31%	5.20%
	max.	25.73%	8.21%	16.37%	5.22%

quite stable during day time. In addition to the already-mentioned difference in the number of possible trips as a passenger or a driver, one can see that whereas the average values are similar, the variability in the possible number of trips as a passenger is much higher than for trips as a driver. The reason is probably the fact that the detour imposed to the driver limits the flexibility on departure/arrival times, as was pointed out before.

Figure 6 presents the average number of identified possible trips as a passenger or a driver, per examined trip, as a function of the time of day, when restricting the search to trips with the same origin or destination link. As expected, the number of identified possible joint trips is

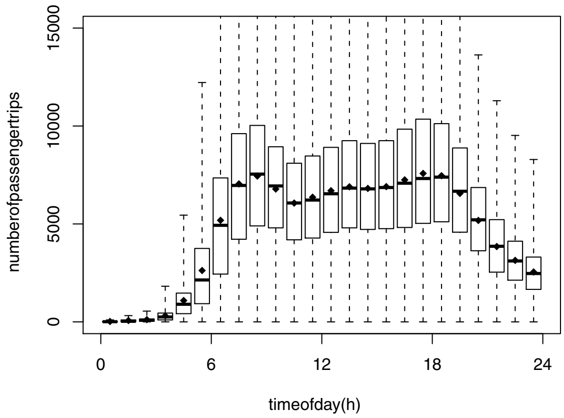

FIGURE 4 Number of potential trips as a passenger as a function of time of day

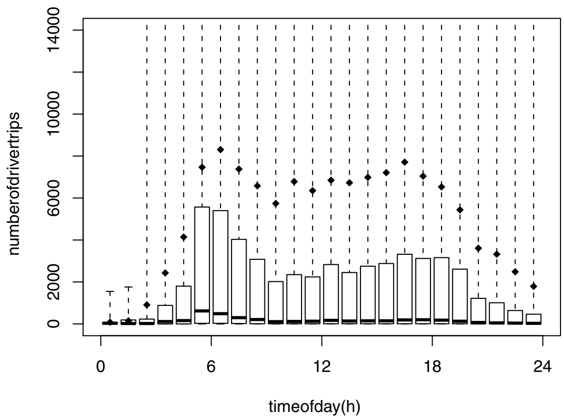

FIGURE 5 Number of potential trips as a driver as a function of time of day

higher in the morning for trips with the same destination (going to work), and in the afternoon for trips with the same origin (back from work). The average number of matches per examined trip is lower when restricting to the same origin.

The results from this section show that what we called the “structural” factor for carpooling does not severely limit the feasibility of carpooling. Randomly matching trips, with a maximum time window of 15min. and a maximum driver detour corresponding to 15% of the driver’s initial travel time, lead to more than 87% of the car trips being part of a carpooling trip. Moreover, the results show that even when restricting the matches to trips with the same origin or the same destination, more than 47% of the car trips could still be part of a carpooling trip. This indicates a real potential for small-scale services, *e.g.* at the firm level. Such services are easier to implement, and it is known that the probability of accepting a carpooling match increases if the potential co-traveler is a colleague or acquaintance, which is likely at the firm level.

Also, the influence of the length of a trip on the likelihood to find a match was tested. Running the random matching process only for passenger trips of more than 10km did not result

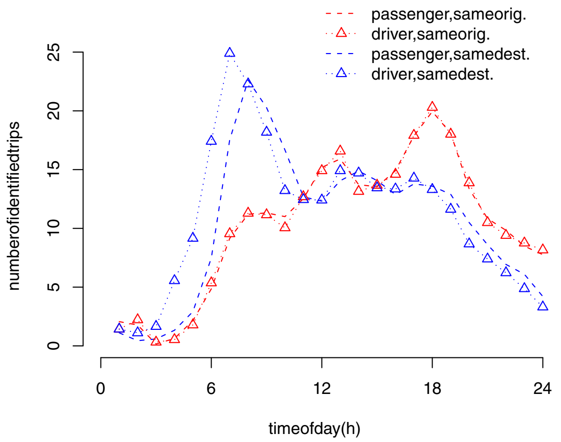

FIGURE 6 Number of potential joint trips as a function of time of day

in a lower likelihood for passengers to find a driver. Matching such trips in carpooling trips has a higher impact on the overall traveled distance than matching shorter trips, and this results indicate that this impact could actually be significant.

CONCLUSIONS

This paper focuses on studying the structural feasibility of carpooling in the area of Zurich, Switzerland. To do so, the output of a MATSim run was used as a source for detailed trip data, in which all matches fulfilling a maximum time window and maximum detour time were identified.

The results of the analysis show that high shares of trips are "close enough" to be performed jointly rather than individually: for the retained condition, more than 87% of the trips could be performed in two-persons car pools, rather than individually.

The case of trips where individuals have the same origin or destination was also investigated. It was shown that even with such restriction, more than 47% of the trips could be performed in two-persons car-pools. Such solutions are easier to implement, for example at the firm level. Such implementations also make easier to limit the number of platforms existing to find matches for a given OD, making it easier to find a match if it exists. However, with such limitations, the average number of possible drivers per passenger decreases dramatically, dropping from more than 6,000 to less than 20. This makes this kind of solutions much less robust facing a driver deciding to opt-out.

The influence of the trip length on the likelihood to find a driver was also explored. Not considering short trips in the matching process actually allowed to increase the likelihood to find a match for the remaining trips, due to the availability of drivers which would otherwise have been matched with shorter passenger trips.

An obvious limitation of the approach, even when accepting to leave behavioural or organisational factors out of the analysis, is the non-consideration of the tour constraint: the fact that a passenger can be driven from home to work does not mean that he could be driven back from work. To cope with this problem, the process could be extended to identify possibilities of return.

These results indicate that the most limiting factor in the development of carpooling is the behavioral factor, the most prominent characteristics of which probably being the willingness to be alone in one's car and the fear of a loss in flexibility.

To get a feel of the influence of this factor, the impact of the matches found in this paper on the generalized cost could be analyzed. The generalized cost could be derived from a behavioral model such as the one estimated in *(6)*. However, this model is a mode choice model, which predicts the choice of the carpooling mode, without considering the choice of the co-traveler. Its usage to choose between alternative co-travelers may produce inconsistent results, due to what is sometimes called the "red bus/blue bus paradox" *(16)*.

The matches generated by the approach were used to analyze the impact of the structural factor on the carpooling potential. They could also be used as an input for a simulation-based approach for the evaluation of the behavioral factor. The authors are developing an approach to simulate joint traveling behavior in MATSim *(17)*, which would allow to test the behavioral acceptance of the identified matches. Using behavioral information from previous studies, microsimulation could be used to test the importance of the organisational aspect, by simulating the process of individuals searching for a match via one or several platforms. However, such a simulation setting still doesn't exist, and would require significant improvements of the state-of-the-art.

ACKNOWLEDGEMENTS

This research was funded by Volkswagen AG as part of the project "An analysis of mode choice effects in different scenarios simulated with MATSim".

REFERENCES

1. Abrahamse, W. and M. Keall (2012) Effectiveness of a web-based intervention to encourage carpooling to work: A case study of Wellington, New Zealand, *Transport Policy*, **21**, 45–51.

2. Ferguson, E. (1997) The rise and fall of the american carpool: 1970-1990, *Transportation*, **24**, 349–376.

3. Burris, M. W. and J. R. Winn (2006) Slugging in houston-casual carpool passenger characteristics, *Journal of Public Transportation*, **9** (5) 23–40.

4. Morency, C. (2007) The ambivalence of ridesharing, *Transportation*, **34**, 239–253.

5. Correia, G. and J. M. Viegas (2011) Carpooling and carpool clubs: Clarifying concepts and assessing value enhancement possibilities through a stated preference web survey in Lisbon, Portugal, *Transportation Research Part A: Policy and Practice*, **45**, 81–90.

6. Ciari, F. and K. W. Axhausen (2011) Choosing carpooling or car sharing as a mode: Swiss stated choice experiments, *Working Paper*, **694**, IVT, ETH Zurich, Zurich.

7. Baldacci, R., V. Maniezzo and A. Mingozzi (2004) An exact method for the car pooling problem based on lagrangean column generation, *Operations Research*, **52** (3) 422–439.

8. Kytöjoki, J., T. Nuortio and M. Gendreau (2007) An efficient variable neighborhood search heuristic for very large scale vehicle routing problems, *Computers and Operations Research*, **34**, 2743–2757.

9. de Palma, A. (2012) The economics of the family in transportation and urban economics, presentation, 12th Swiss Transport Research Conference (STRC), Ascona.
10. Mühlethaler, F. (2012) Potential of car-pooling in Switzerland, paper presented at the *12th Swiss Transport Research Conference*, Ascona, May 2012.
11. Huang, H.-J., H. Yang and M. G. H. Bell (2000) The models and economics of carpools, *Annals of Regional Science*, **34** (1) 55–68.
12. Hansen, R., S. Kahne and R. Houska (1975) A car pooling system using heuristic costs, *Transportation Research*, **9** (2–3) 159–165.
13. Calvo, R. W., F. d. Luigi, P. Haastrup and V. Maniezzo (2004) A distributed geographic information system for the daily car pooling problem, *Computers and Operations Research*, **31** (13) 2263–2278.
14. Sghaier, M., H. Zgaya, S. Hammadi and C. Tahon (2010) A distributed dijkstra's algorithm for the implementation of a real time carpooling service with an optimized aspect on siblings, paper presented at the *Intelligent Transportation Systems Conference (ITSC)*, Madeira, September 2010.
15. Balmer, M., K. Meister, M. Rieser, K. Nagel and K. W. Axhausen (2008) Agent-based simulation of travel demand: Structure and computational performance of MATSim-T, paper presented at the *Innovations in Travel Modeling (ITM'08)*, Portland, June 2008.
16. Ben-Akiva, M. E. and S. R. Lerman (1985) *Discrete Choice Analysis: Theory and Application to Travel Demand*, MIT Press, Cambridge.
17. Dubernet, T. and K. W. Axhausen (2012) Including joint trips in a multi-agent transport simulation, paper presented at the *13th International Conference on Travel Behaviour Research (IATBR)*, Toronto, July 2012.