Quantifying the Benefits of Ride Sharing

Working Paper

Author(s):
Ruch, Claudio; Lu, ChengQi; Sieber, Lukas; Frazzoli, Emilio

Publication date:
2019

Permanent link:
https://doi.org/10.3929/ethz-b-000367142

Quantifying the Benefits of Ride Sharing

Claudio Ruch, ChengQi Lu, Lukas Sieber, and Emilio Frazzoli, *Senior Member, IEEE*

***Abstract*—In unit-capacity mobility-on-demand systems the vehicles transport only one travel party at a time whereas in ride-sharing mobility-on-demand systems, a vehicle may transport different travel parties at the same time, e.g, if paths are partially overlapping. We evaluate both the efficiency-related benefits and drawbacks of ride sharing for mobility on demand in detail with a transportation simulator under the influence of different fleet operational policies and transportation scenarios. For this purpose, an open-source simulation environment is introduced which is capable of evaluating a large class of operational policies for ride-sharing mobility-on-demand systems. Then, the impact of ride sharing on efficiency and service level is assessed for several benchmark operational policies from the literature. First in a dense urban scenario, then on a line-shaped, rural one. In the urban scenario, ride sharing is able to reduce the vehicle miles traveled by 11% and the fleet size by 29% at the cost of 15% increased total travel time. 76 % of trips are with no more than 2 travel parties on board. In the rural case, a 28% fleet size reduction and a 12% reduction in vehicle miles traveled were achieved at the cost of 3% increased total travel time. 98 % of trips are with no more than 2 travel parties on board. We conclude that efficiency gains due to ride sharing in mobility on demand are present but limited and may not be sufficient to compensate for the drawbacks of reduced convenience, loss of privacy, and higher total travel and drive times. Furthermore, the study clearly demonstrates that systems should be designed with small 4-6 seat vehicles capable of handling occasional ride sharing rather than, e.g., larger and more costly minibuses.**

***Index Terms*—Ride Sharing | Mobility on Demand | Operational Policies**

Mobility on Demand (MoD) transportation systems promise to unite the convenience of motorized individual transport with the environmental friendliness and price of conventional public transit [1]. They are expected to increase their mode share significantly with the forthcoming introduction of the fully autonomous vehicle which may resolve some of the major shortcomings of today's MoD systems, namely inbalance [2], lack of system-wide coordination [3] and driver availability. In the most commonly described type of mobility on demand, a large fleet of vehicles, each with enough capacity for **one** travel party is used, e.g., 4 seats. Each vehicle serves one party at a time by transporting them from their origin directly to their travel destination. After dropping off the travel party, the vehicle may be dispatched to the next customer or relocated. Such systems are called *unit-capacity mobility-on-demand systems* (1MoD). In contrast to this, the vehicles used in *ride-sharing mobility-on-demand systems* (RMoD) may have larger capacity and might transport **several** travel parties at the same time. As an example, when two travel parties have a close-by origin and are traveling in a similar direction, it is efficient to transport both simultaneously as illustrated in Figure 1.

The resulting efficiency gains were predicted to improve the efficiency of car and taxi transportation greatly [4] by reducing the fleet size, the vehicle miles traveled and associated emissions [5]. A transportation scenario must exhibit certain characteristics in order for ride sharing to have positive effects on system efficiency, namely that a large enough fraction of requests must have at least partially aligned paths and arrive to the system approximately at the same time. There are scenarios which do not have these characteristics, e.g., an early-morning situation with small arrival rates. However, the current literature suggests that the benefits of ride sharing can be harvested in most cases. An example are the theoretically derived estimations provided in [4]. Based on macroscopic parameters including the area of the city, the average traffic speed, the total trip arrival rate and the maximum delay tolerated due to ride sharing, the results suggest that only relatively small arrival rates may be necessary to share almost 100 % of trips. For San Francisco, it is estimated that approximately 97% of trips can be shared with a delay of no more than 5 [min] at an arrival rate of 1,200 trips per hour. Generalizing the results to other cities, a potential to share almost 100% of trips in Amsterdam, Newcastle, Paris, Prague, Rome and Santiago with an arrival rate of 8.5 trips per hour and square kilometer is recognized.

The authors are with the Institute of Dynamic Systems and Control (IDSC), ETH Zurich.

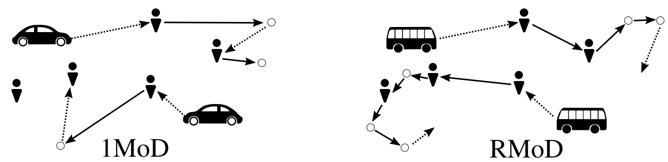

Fig. 1: In a 1MoD system, at every time only one travel party is transported. For RMoD systems, several parties can be in a vehicle at the same time.

Several attempts have been made to quantify the benefits of ride sharing in simulation and field studies. A survey of approximately 30,000 carpool commuters in the 1970s estimated a 23% reduction of vehicle miles traveled due to ride sharing [6]. In a study on the impacts of taxi sharing in New York City, it was concluded that ride sharing would reduce cumulative trip length by 40% or more [7]. A simulation study for Prague [8] concluded that ride sharing can decrease the total miles traveled by a mobility-on-demand fleet by 65% without prolonging the travel times by more than 10 minutes. In [9] a slightly more conservative estimate of 13% savings in vehicle miles traveled (VMT) at 25% more trips served is made. The positive impact of ride sharing on mobility on demand is also subject of other studies, e.g., [10], [11], [12], [13]. Indeed, instances of real-world implementations of ride-sharing schemes exist: jitneys or dollar-vans are small buses that (often informally) operate in larger cities and transport passengers on temporarily fixed routes which are adjusted as needed. As an example, the informal but successful jitney system in New York City serves about 120,000 riders daily [14]. Another example presented in [15] are shared-ride taxis in Little Rock, Arkansas that were serving more than 1.7 million people per year in 1977. Ride sharing schemes were also introduced to the market by large ride-hailing companies, e.g., UberPOOL.

Yet, it is unclear whether the benefits are sufficient to outweigh the drawbacks both from a user and from an operator perspective in large-scale and general settings. Ride sharing increases drive times for passengers. Even at reduced total travel times, this might be perceived as a nuisance. The loss of privacy inherent to ride sharing may also be a concern, e.g., from a safety viewpoint [16]. Taking this into account, it is an open question whether price reductions are necessary to compensate for these drawbacks and if the savings in operational cost are sufficiently large to offset them. For the operator, the larger vehicles necessary to provide ride sharing increase operational costs, which can only be compensated with high enough fraction of shared trips. At least in some cases, it is questionable if such a sharing rate can be achieved. As an example, a simulation study for Austin, Texas revealed that at a mode share of 11% in a small central and dense area of the city, only 4.83% of vehicle miles traveled were ride-shared. The study also reported a reduction of vehicle miles travelled by only 4.5 to 8.7 % [17]. Another aspect that makes the discussion more complex is the possible competition of ride sharing and public transit, e.g., bus lines. In some cases, ride-sharing services were judged as unfair competition to public transportation by local authorities. According to this viewpoint, the (private) schemes only serve the profitable fraction of trips at the expense of conventional public transportation. Some studies support this view, e.g., the work on San Francisco carpoolers presented in [18] estimated that 75% of new carpoolers in the San Francisco Bay area have used public transit previously and only 10% were previously using a privately owned vehicle.The viewpoint that the applicability of ride sharing is limited is supported by a history of discontinued implementation attempts. An example is microtransit, a formal version of the jitney services. Low utilization and sharing rates have been reported for several of these services, see, e.g., [19]. The Group Ride Vehicle project in New York was intended to replace closed public transit connections with ride-sharing taxis but the service was discontinued only months after it was started [14]. Some ride-sharing services were discontinued shortly after introduction, partly due to low match rates [15]. Car pooling of private vehicles has remained on low levels of about 9% ever since it peaked in the 1970s at 20% [15].

The current situation of ambiguity makes it necessary to conduct more quantitative studies in order to compare the trade-off of costs and benefits of ride sharing for realistic transportation scenarios. The state of knowledge is characterized by two shortcomings. First, the operational policies used to control fleets of mobility-on-demand services were shown to be an important factor that cannot be neglected [20]. Many existing studies do not take this into account as they are based on simplified assumptions or heuristic fleet control policies. Second, ride sharing must be analyzed and compared coherently across different transportation scenarios including detailed road network representations and using high-resolution demand profiles. Our work aims at filling these gaps and at providing a tool to the research community to quantitatively assess ride-sharing mobility-on-demand systems.

Contributions: Our contributions are threefold. First, we provide a general interface for ride-sharing operational policies in AMoDeus [21]. The open-source interface is described in a reduced but complete manner that allows implementation and comparison of operational policies in advanced mesoscopic transportation simulation. Second, based on this interface, we have implemented four state-of-the-art operational policies from the literature. We apply these operational policies to an urban and a rural mobility-on-demand scenario and quantitatively assess the impact of ride-sharing on efficiency and service level under static demand. *Organization:* We explain the simulation environment, ride-sharing interface and the implemented operational policies in Section I. Then, we present the simulation setup in Section II. The urban and rural case studies are presented in Sections III and IV, respectively. Finally, a conclusion is drawn in Section V.

I. Ride-Sharing Operational Policies

The key objective of our research is the quantitative analysis of (ride-sharing) mobility-on-demand systems. The principal classes of tools suitable for this task are (i) custom-made simulators such as the one used in [22], (ii) comparable real system instances, e.g., taxi schemes and (iii) general transportation simulators, e.g., [23], [24] and [25]. Custom-made simulators generally lack the ability to be extended and they are of limited use when results should be compared within the research community. Real systems naturally come with high cost and limited range of possible experiments. For this reason, our research is based on the class of general, open-source transportation simulators which are used for multiple purposes by a wide research community. They provide high accuracy results and are validated by a user community of transportation scientists. Within the class, there are several subcategories. Microscopic transportation simulators such as [25] model traffic dynamics and the dynamics of cars themselves. This level of detail is unnecessary for our scope. Macroscopic transportation simulators such as [26] on the other hand will exhibit very fast simulation times but come with more limitations. For this reason, our work is based on the category of mesoscopic simulators that allow just enough granularity to capture the processes that determine the service level and efficiency of mobility-on-demand systems. Building on existing, optimized traffic simulators allows a high network resolution with respect to the number of trips: in this work, we have a ratio of trips to network links of ≈ 0.1 while for various existing approaches this value is substantially higher, e.g., ≈ 42 in [7]. In order to provide a general and user-friendly interface between transportation research and algorithmic research, AMoDeus [21] was created as an add-on to the mesoscopic transportation simulator MATSim [23] to emulate mobility-on-demand systems with advanced fleet operational policies and algorithms. For 1MoD systems, the

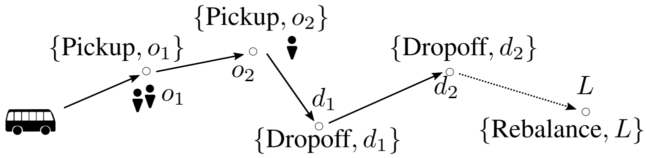

Fig. 2: A vehicle is commanded to execute a ride-sharing schedule.

fleet operational policy was shown to be a decisive factor determining the system service level and efficiency [20]. In this work, we extend the scope of our research to ride-sharing operational policies for which we developed a generic interface for evaluation in realistic transportation scenarios.

Operational Policy Simulation Interface: The underlying mathematical problem that needs to be solved when designing ride-sharing mobility-on-demand operational policies is NP-complete [9]. Proposed solution methods from the literature are based on heuristics [17], efficient setup and maintenance of data structures [9], on mixed-integer linear programming [11], [27], or on combinations of these approaches. As the decision space is considerably larger in the RMoD than in the 1MoD case, the problem is even more complex to solve. In the 1MoD case, a vehicle has only two possibilities for actions: it can either pick up a travel party and transport it to the destination or it can reposition to another location in the operational domain, a process referred to as rebalancing, or repositioning. In the ride-sharing mobility-on-demand case, the order of pickups, dropoffs and eventual rebalancing actions can be arranged in many different ways. Specifically, a schedule of future pickups and dropoffs as well as rebalance type trips has to be maintained for every vehicle at all times as illustrated in Figure 2. These schedules must satisfy constraints, namely that every pickup has to be followed by exactly one dropoff and that the number of travel parties in the vehicle must not exceed its capacity at any point in time. Our open-source interface allows to implement a large class of ride-sharing policies and evaluate them in advanced transportation scenarios, thereby leveraging the underlying simulation engine MATSim [23] and executes all commands as specified. The implementation allows schedules to be changed as long as picked up travel parties are also dropped off and it also allows to modify the routing logic of vehicles if desired by the user. For this work, we have used fastest-route routing. Based on the interface described in the previous paragraphs, we have implemented several state of the art operational policies which are described next.

High Capacity Shared Autonomous Mobility-on-Demand Algorithm: (HCRS) [11] The principle of this well formalized operational policy for the ride-sharing mobility-on-demand problem is the relatively efficient exploration of all possible options with integer linear programming. HCRS composes so-called trips out of two elements: (1) the vehicle, including the set of passengers on board and (2) the set of passengers the vehicle is scheduled to pick up. During the exploration step, trips are constrained to a maximum wait time and a maximum increase of the total travel time with respect to the shortest path. These constraints also apply to all passengers who have already boarded cars. After exploration, the cost of each possible trip is computed as the sum of the total travel delay of all the passengers involved in it. A passenger may also be deliberately ignored by the system at a high cost. The solution that minimizes the sum of costs of chosen trips is then computed with an integer linear program. After each assignment, if there are both ignored requests and idling vehicles at the same time, the algorithm will send the idling vehicles driving towards these requests minimizing the sum of Euclidean distances of these assignments.

T-Share Policy: The T-Share policy presented in [9] was one of the first contributions exploring the efficiency gains that result from sharing single-use taxis by multiple travel parties. Specifically, the algorithm assumes that a unit-capacity mobility-demand system is in place. Then, it acts on the collection of vehicles currently transporting a single travel party and evaluates if additional travel parties can be added to the vehicle. This process is done in two steps. First, whenever a customer transportation request arrives, an algorithm called "dual side taxi searching" is used to find a set of occupied vehicles which could potentially host the transportation request by adapting the current schedule. This set is of smallest possible cardinality, i.e., the iterative algorithm terminates as soon as it finds a nonempty set. Then, for this set of potential taxis the algorithm "insertion feasibility check" determines the vehicle and schedule insertion permutation that can handle the additional request with minimal additional mileage while respecting constraints on latest pickup and latest arrival of the passenger. In this work, the T-share operational policy was extended in order to allow adding a travel party to a vehicle also in the case when more than one travel party is already in the vehicle.

Dynamic Ride Sharing Strategy: (DRSS) The operational policy presented in [17] builds on previous work for the 1MoD case [28]. In each time step, for all unassigned travel parties a suitable vehicle is identified and if possible bindingly assigned to a request. Ride-sharing assignments are searched with highest priority and subject to five conditions: the total travel time of the scheduled parties may increase by at most 20%. The remaining travel time of the scheduled parties must not increase by more than 40%. The total travel time of the unassigned party may increase by no more than the greater value of 3 minutes and 20% compared to the expected travel time of a direct drive. The unassigned party must be picked up within the next five minutes. The total planned time to serve all travel parties must be shorter than the time to serve the scheduled parties plus the time to serve the unassigned party individually. If no valid ride sharing possibility can be found, the *second priority* is to assign the closest free vehicle to a travel party. Idle vehicles are rebalanced using a so called block-balance method, in which the local imbalance of free vehicles and demand between adjacent cells is used to compute rebalancing commands.

Extended Demand Supply: (Ext-DS) This policy is an extension of the 1MoD policy presented in [29] with a ride-sharing heuristic. For the unit-capacity assignment, an over-supply and an under-supply case are distinguished with more available vehicles than open requests or vice versa. Then the set with higher cardinality is iterated in random order for each vehicle (or request) and the closest request (or vehicle) is assigned. The ride-sharing extension additionally searches for every vehicle if there are open requests within a radius of no more than 0.62 miles with a deviation of their destination of not more than 5 degree from the on-board passengers. These requests are also picked up by the travelling vehicle.

II. Simulation Setup

In this study, we focus on the comparison of costs and benefits of ride sharing with respect to efficiency, specifically on the reduction of fleet size and on vehicle miles traveled. We track the occupancy of each vehicle to assess what vehicle sizes are actually required to enable ride sharing, the maximum capacity is not limited. For this analysis, we have chosen a fixed demand profile, which is served by a mobility-on-demand system of varying fleet size operating on a detailed road network guided by different operational policies. Next to the conventional metrics of total journey time, sharing rate and vehicle miles traveled, we also record the occupancy of each vehicle. For every scenario, we compare the four ride-sharing operational policies and additionally a basic 1MoD policy. In order to have a globally applicable reference not specific to any particular scenario, we have chosen the Global Bipartite Matching (GBM) policy presented in [21] which minimizes the Euclidean distance between locations of open requests and available vehicles by solving a bipartite matching problem with the Hungarian method [30]. The solution is recomputed at every time step to take into account possible improvements to the assignment.

III. The Case of Urban Mobility on Demand

An often cited use case for ride-sharing mobility on demand are dense urban areas. As a representation of such a case, we have analyzed taxi traces recorded in the city of San Francisco presented in [31] and created a reproducible mobility-on-demand scenario with the same requests [21]. The scenario contains a total of 16,439 requests which are served by the mobility-on-demand system on a road network with a total of 153,327 roads. A series of simulations with different fleet sizes and operational policies yields the resulting mean total travel times shown in Figure 3, the sharing rates in Figure 4 and the vehicle miles traveled in Figure 5.

For this case, we identify 350 vehicles as a sensible operation point with ≈ 4:46 [min] average wait time as an increased fleet size will only result in minor service level improvements. The resulting efficiency gains for this case are summarized in Table I. An RMoD operation with 250 vehicles would result in an acceptable increase in total travel time of 15% and yield 29% in fleet size reduction and 11% in reduction of vehicle miles traveled. The occupancy of vehicles during the day is shown in Figure 6. Interestingly, 40 % of requests are served with a single travel party and only 24 % of requests have more than one other travel party on board during a trip section.

An important consideration, especially for the urban case, is the density of requests. How close together in space and time must travel requests be in order for ride sharing to pay off? To ensure that the number of requests in this urban scenario is high enough we have taken the constant ratio of total request to vehicles of $\frac{16439}{350} \approx 47$ and evaluated the sharing rate for different numbers of requests. The resulting graph shown in Figure 7 reveals that the critical density to harvest the benefits of ride sharing was reached and that further increase of request density would only lightly increase system efficiency.

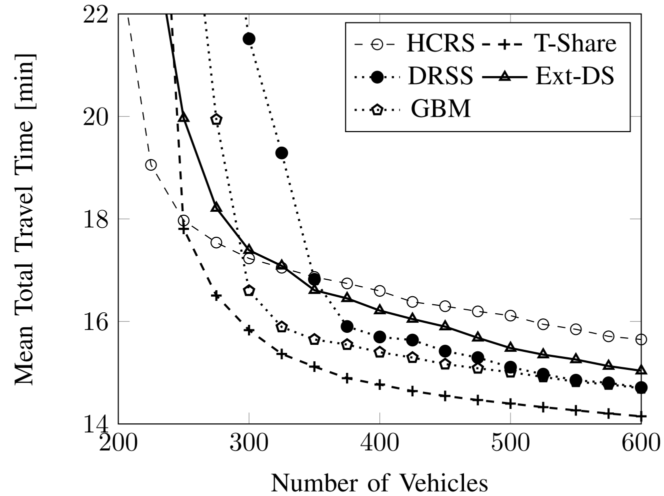

Fig. 3: Mean total travel time for the urban mobility-on-demand case.

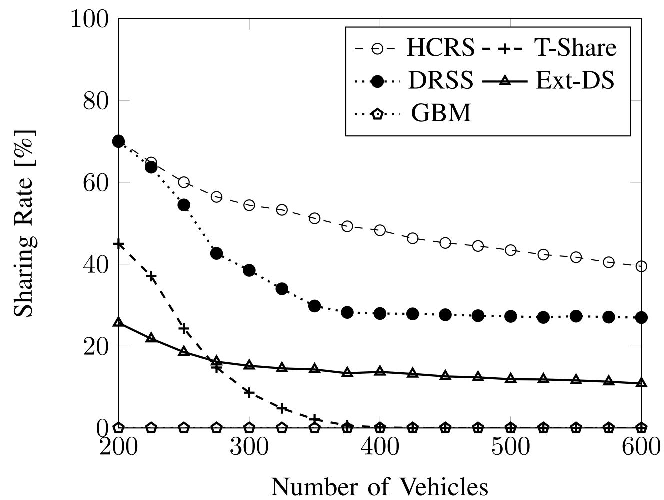

Fig. 4: Sharing rate for the urban mobility-on-demand case.

Operational Policy	Fleet Size	Vehicle Miles Traveled	Mean Total Travel Time
1MoD (GBM)	350	137,022 miles	15:39 min
RMoD (HCRS)	350	120,587 miles	16:52 min
RMoD (HCRS)	300	121,301 miles	17:13 min
RMoD (HCRS)	250	123,686 miles	17:58 min
RMoD (HCRS)	200	120,956 miles	23:09 min

TABLE I: Efficiency gains of RMoD compared to 1MoD under a basic global bipartite matching policy in the urban scenario.

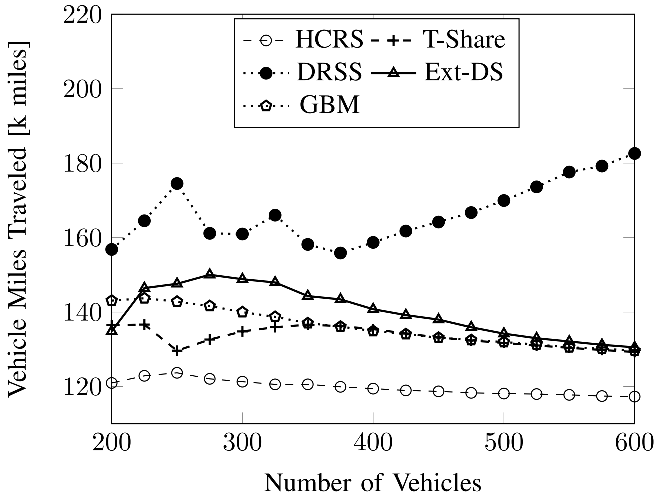

Fig. 5: Vehicle miles traveled for the urban mobility-on-demand case.

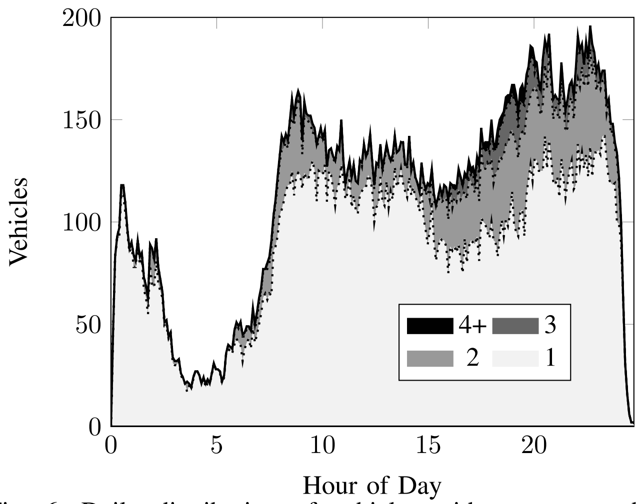

Fig. 6: Daily distribution of vehicles with respect to the number of travel parties on board. 6576 trips were unshared, 5983 shared trip segments with at most 1 other, 2685 with at most 2 other and 1194 with 3 to 6 other travel parties.

IV. Substitution of Public Transportation Lines

The potential efficiency gains of ride sharing certainly depend on the spatio-temporal distribution of the travel demand. While the study in the urban setting utilizes high request density, this case attemps to exploit the spatial distribution of requests. To explore this possibility, we have analyzed a scenario in which a train line in a rural area is substituted with a hypothetical RMoD system. The travel demand of 1,000 daily requests that is currently served by train is now served door to door by the taxi fleet. The road network has a total of 9,049 roads, for details see [32]. As in the urban case, the scenario was evaluated in simulation using [21] under all the operational policies and for different fleet sizes. The resulting mean total travel times are shown in Figure 8, the sharing rates in Figure 9 and the daily fleet distance in

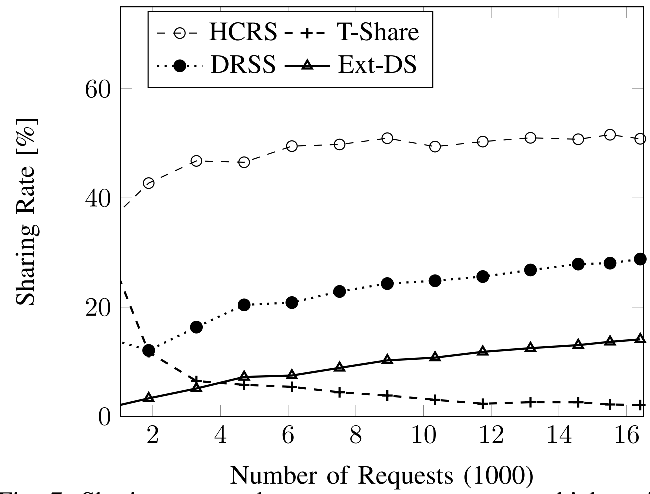

Fig. 7: Sharing rate under a constant request to vehicle ratio of 47.

Figure 10. We identify 35 vehicles (≈ 2:32 [min] average wait

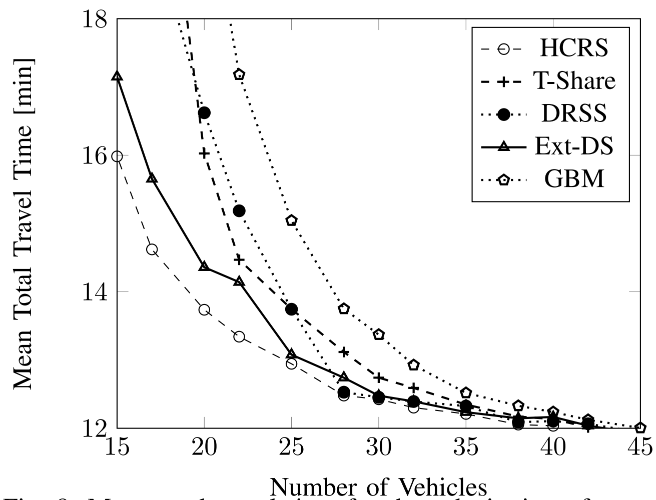

Fig. 8: Mean total travel time for the substitution of a rural public transportation line.

time) as a sensible operation point after which an additional increase in fleet size will not significantly improve service level. The resulting efficiency gains for this operation point due to ride sharing are summarized in Table II. Assuming that

Operational Policy	Fleet Size	Vehicle Miles Traveled	Mean Total Travel Time
1MoD (GBM)	35	6,447 miles	12:31 min
RMoD (HCRS)	35	5,637 miles	12:12 min
RMoD (HCRS)	25	5,649 miles	12:56 min
RMoD (HCRS)	15	5,140 miles	15:58 min
RMoD (HCRS)	10	4,365 miles	23:01 min

TABLE II: Efficiency gains of RMoD compared to 1MoD under a basic global bipartite matching policy in the rural scenario.

the stakeholders would not like to increase the total travel time of the system when introducing RMoD, they could chose an

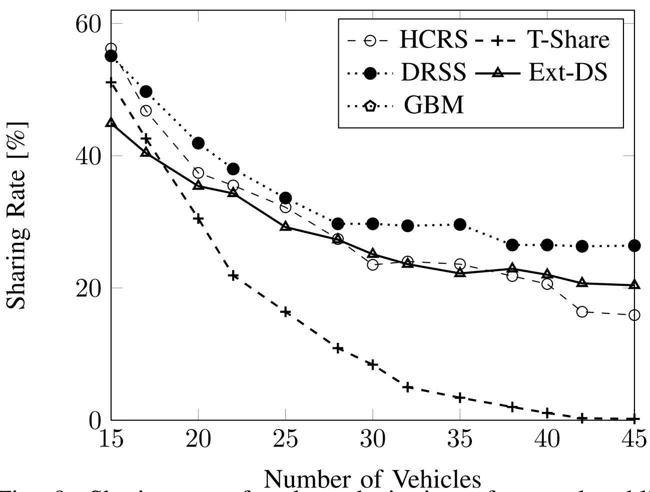

Fig. 9: Sharing rate for the substitution of a rural public transportation line.

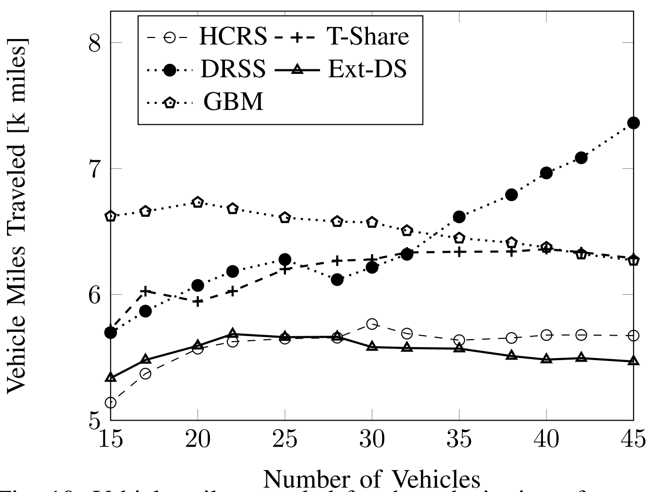

Fig. 10: Vehicle miles traveled for the substitution of a rural public transportation line.

operating point with 25 vehicles which represents a reduction of the fleet size by 28% and a reduction of the vehicle miles traveled by 12% at the cost of 3% increased mean total travel time. Figure 11 reveals that also in this case, a large percentage of 68% of trips travels with a single travel party, only 8.6% share their trip with 2 to maximum 5 other travel parties.

V. Conclusion

We have implemented a generic interface to test operational policies for ride-sharing mobility-on-demand systems in a high-fidelity transportation simulator. We have implemented four operational policies and have applied them to an urban and a rural scenario. In both cases, we have measured the efficiency gains in terms of fleet size reduction and reduction of vehicle miles traveled due to ride sharing. These efficiency gains compared to the unit-capacity mobility-on-demand case are present but of modest size. We conclude that it is debatable if these efficiency gains and associated fare reductions are sufficiently high to compensate for the potential loss of privacy and for the higher drive times. Furthermore, we identified that

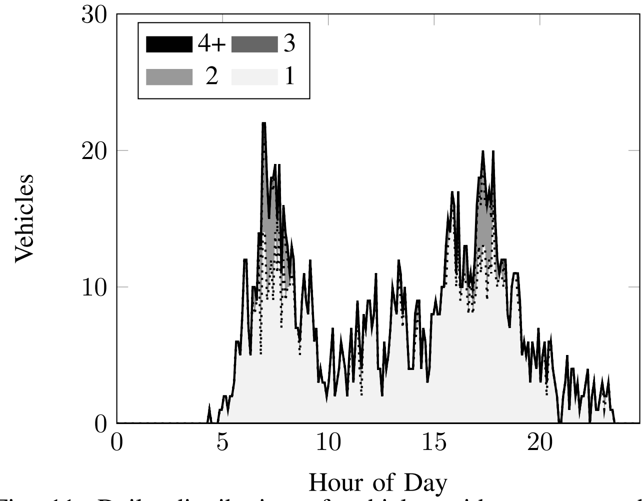

Fig. 11: Daily distribution of vehicles with respect to the number of travel parties on board. 678 trips were unshared, 236 shared trip segments with at most 1 other, 61 with at most 2 other and 25 with 3 to 5 other travel parties.

for a majority of shared trips, at most 3 travel parties share a vehicle. We see this as a clear sign that small autonomous vehicles designed for 4-6 people would be preferential to larger and more costly minibuses. This suggests that the design of mobility systems should aim for an efficient combination of unit-capacity mobility on demand with small vehicles and high-capacity transit links such as subways or trains. We aim to continue research in three directions. First, we plan to implement and compare additional operational policies. Second, we aim to theoretically evaluate the difference of high-capacity public transportation systems, e.g., subways to ride-sharing mobility on demand. Specifically, we aim to understand if there are cases in which ride-sharing mobility on demand would be advantageous both in comparison to such high-capacity public transit systems and unit-capacity mobility-on-demand systems. Finally, we plan to explore other applications of ride sharing, specifically to employ ride-sharing as a tool to relieve congestion.

References

[1] K. Spieser, K. Treleaven, R. Zhang, E. Frazzoli, D. Morton, and M. Pavone, "Toward a systematic approach to the design and evaluation of automated mobility-on-demand systems: A case study in singapore," in *Road vehicle automation*. Springer, 2014, pp. 229–245.

[2] M. Pavone, S. Smith, E. Frazzoli, and D. Rus, "Load balancing for Mobility-on-Demand systems," in *Robotics: Science and Systems VII*, 2011.

[3] C. Ruch, S. Richards, and E. Frazzoli, "The value of coordination in one-way mobility-on-demand systems," *IEEE Transactions on Network Science and Engineering*, 2019.

[4] R. Tachet, O. Sagarra, P. Santi, G. Resta, M. Szell, S. Strogatz, and C. Ratti, "Scaling law of urban ride sharing," *Scientific reports*, vol. 7, p. 42868, 2017.

[5] N. Fellows and D. Pitfield, "An economic and operational evaluation of urban car-sharing," *Transportation Research Part D: Transport and Environment*, vol. 5, no. 1, pp. 1–10, 2000.

[6] A. Cohen and S. Shaheen, *Planning for Shared Mobility*, 2018.

[7] P. Santi, G. Resta, M. Szell, S. Sobolevsky, S. H. Strogatz, and C. Ratti, "Quantifying the benefits of vehicle pooling with shareability networks," *Proceedings of the National Academy of Sciences*, vol. 111, no. 37, pp. 13 290–13 294, 2014.

[8] D. Fiedler, M. Čertickỳ, J. Alonso-Mora, and M. Čáp, "The impact of ridesharing in mobility-on-demand systems: Simulation case study in prague," in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2018, pp. 1173–1178.
[9] S. Ma, Y. Zheng, and O. Wolfson, "T-share: A large-scale dynamic taxi ridesharing service," in *Data Engineering (ICDE), 2013 IEEE 29th International Conference on*. IEEE, 2013, pp. 410–421.
[10] S. Ma, Y. Zheng, O. Wolfson *et al.*, "Real-time city-scale taxi ridesharing." *IEEE Trans. Knowl. Data Eng.*, vol. 27, no. 7, pp. 1782–1795, 2015.
[11] J. Alonso-Mora, S. Samaranayake, A. Wallar, E. Frazzoli, and D. Rus, "On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment," *Proceedings of the National Academy of Sciences*, vol. 114, no. 3, pp. 462–467, 2017.
[12] B. Yang, S. Ren, E. F. Legara, Z. Li, E. Y. Ong, L. Lin, and C. Monterola, "Phase transition in taxi dynamics and impact of ridesharing," *arXiv preprint arXiv:1801.00462*, 2018.
[13] D. O. Santos and E. C. Xavier, "Taxi and ride sharing: A dynamic dial-a-ride problem with money as an incentive," *Expert Systems with Applications*, vol. 42, no. 19, pp. 6728–6737, 2015.
[14] D. A. King and E. Goldwyn, "Why do regulated jitney services often fail? evidence from the new york city group ride vehicle project," *Transport Policy*, vol. 35, pp. 186–192, 2014.
[15] S. Shaheen, *Shared Mobility: The Potential of Ridehailing and Pooling*. Washington, DC: Island Press/Center for Resource Economics, 2018, pp. 55–76.
[16] A. Sanguinetti, K. Kurani, and B. Ferguson, "Is it ok to get in a car with a stranger? risks and benefits of ride-pooling in shared automated vehicles," 2019.
[17] D. J. Fagnant and K. M. Kockelman, "Dynamic ride-sharing and fleet sizing for a system of shared autonomous vehicles in austin, texas," *Transportation*, vol. 45, no. 1, pp. 143–158, 2018.
[18] S. A. Shaheen, N. D. Chan, and T. Gaynor, "Casual carpooling in the sf bay area: Understanding user characteristics, behaviors, and motivations," *Transport Policy*, vol. 51, pp. 165–173, 2016.
[19] StreetsblogNYC, "Ford's chariot vans are mostly empty," https://nyc.streetsblog.org/2018/08/02/fords-chariot-vans-are-mostly-empty/, 2018.
[20] S. Hörl, C. Ruch, F. Becker, E. Frazzoli, and K. W. Axhausen, "Fleet control algorithms for automated mobility: A simulation assessment for zurich," *Transportation Research. Part C, Emerging Technologies*, 2018.
[21] C. Ruch, S. Hörl, and E. Frazzoli, "Amodeus, a simulation-based testbed for autonomous mobility-on-demand systems," in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2018, pp. 3639–3644.
[22] K. Treleaven, M. Pavone, and E. Frazzoli, "Asymptotically optimal algorithms for One-to-One pickup and delivery problems with applications to transportation systems," *IEEE Trans. Automat. Contr.*, vol. 58, no. 9, pp. 2261–2276, 2013.
[23] A. Horni, K. Nagel, and K. W. Axhausen, *The multi-agent transport simulation MATSim*. Ubiquity Press London, 2016.
[24] M. Adnan et al, "Simmobility: A multi-scale integrated agent-based simulation platform," in *95th Annual Meeting of the Transportation Research Board Forthcoming in Transportation Research Record*, 2016.
[25] D. Krajzewicz, G. Hertkorn, C. Rössel, and P. Wagner, "Sumo (simulation of urban mobility)-an open-source traffic simulation," in *Proceedings of the 4th middle East Symposium on Simulation and Modelling (MESM20002)*, 2002, pp. 183–187.
[26] H. J. Payne, "Freflo: A macroscopic simulation model of freeway traffic," *Transportation Research Record*, no. 722, 1979.
[27] J. Alonso-Mora, A. Wallar, and D. Rus, "Predictive routing for autonomous mobility-on-demand systems with ride-sharing," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 3583–3590.
[28] D. J. Fagnant, K. M. Kockelman, and P. Bansal, "Operations of shared autonomous vehicle fleet for austin, texas, market," *Transportation Research Record: Journal of the Transportation Research Board*, no. 2536, pp. 98–106, 2015.
[29] J. Bischoff and M. Maciejewski, "Simulation of city-wide replacement of private cars with autonomous taxis in berlin," *Procedia computer science*, vol. 83, pp. 237–244, 2016.
[30] H. W. Kuhn, "The hungarian method for the assignment problem," *Naval research logistics quarterly*, vol. 2, no. 1-2, pp. 83–97, 1955.
[31] M. Piorkowski, N. Sarafijanovic-Djukic, and M. Grossglauser, "A parsimonious model of mobile partitioned networks with clustering," in *2009 First International Communication Systems and Networks and Workshops*. IEEE, 2009, pp. 1–10.
[32] C. Ruch, L. Sieber, S. Hörl, K. W. Axhausen, and E. Frazzoli, "Autonomous mobility-on-demand providing superior public transportation in rural areas," *Under Review*, 2019.

Claudio Ruch received his MSc degree from ETH Zürich in the interdisciplinary "Robotics, Systems and Control" program in 2014. At the time he worked on rebalancing algorithms for bicycle sharing schemes. After working internationally in industrial projects he is now a graduate student at the Institute for Dynamic Systems and Control. His research, at the intersection of robotics and transportation, is focused on optimizing the fleet operation of mobility-on-demand systems.

ChengQi Lu war born in Beijing, China and has received his MSc degree from ETH Zürich in 2019. His principal interest of research are congestion-aware operating strategies for mobility-on-demand systems. Using such strategies, he aims to reduce the congestion of mobility in large cities. His newest results indicate that with good fleet control, coordinated mobility-on-demand systems can substantially reduce traffic congestion in large cities. In his free time, ChengQi Lu enjoys hiking in the Swiss Alps.

Lukas Sieber received his MSc degree from ETH Zürich in the interdisciplinary "Robotics, Systems and Control" program in 2019. Since then, has has worked on several transprotation related research projects with the Swiss federal railways SBB. Together with other researchers, he was able to show in his work that mobility-on-demand systems might be a very attractive option for rural areas that can increase the service level and at the same time substantially reduce costs compared to existing public transit connections. Before focusing on mobility related topics, Lukas has worked as a data analyst and in the field of power grid automation.

Emilio Frazzoli is a Professor of Dynamic Systems and Control at ETH Zurich, as well as co-founder and CTO of nuTonomy Inc.. He received a Laurea degree in Aerospace Engineering from the University of Rome, "Sapienza", Italy, in 1994, and a Ph. D. degree from the Department of Aeronautics and Astronautics of the Massachusetts Institute of Technology, in 2001. Before joining ETH in 2016, he held faculty positions at the University of Illinois, Urbana-Champaign, the University of California, Los Angeles, and at the Massachusetts Institute of Technology. He was the recipient of a NSF CAREER award in 2002, of the IEEE George S. Axelby award in 2015, and of the IEEE Kiyo Tomiyasu award in 2017. His current research interests focus primarily on autonomous vehicles, mobile robotics, and transportation systems.