Routing Optimization with Vehicle-Customer Coordination

Alexandre Jacquillat[a,*], Kai Wang[b], Shuaian Wang[c], Wei Zhang[c]

[a] Sloan School of Management, MIT, Cambridge, MA, USA
[b] Heinz College, Carnegie Mellon University, Pittsburgh, PA, USA
[c] Faculty of Business, The Hong Kong Polytechnic University, Hong Kong
[*] Corresponding author

Extended abstract submitted for presentation at the 11th Triennial Symposium on Transportation Analysis conference (TRISTAN XI)
June 19-25, 2022, Mauritius Island

January 12, 2022

Keywords: vehicle-customer coordination, vehicle routing, ride-sharing, time-space network

1 INTRODUCTION

Several transportation systems leverage connected technologies and digitization to coordinate vehicles' and customers' operations. In the most prominent example, new ride-sharing services enable riders can walk to meet drivers in mutually convenient locations in exchange of a discount (e.g., Uber Express Pool, Lyft Shared Saver). Similarly, company and school buses pick up riders in a few central locations prior to traveling to a common destination. Another example is aerial refueling, in which a tanker aircraft coordinates its operations with other aircraft.

In all these examples, vehicle-customer coordination provides an extra degree of freedom to enhance the efficiency of first- and last-mile transportation. This flexibility, however, comes with challenges. At the downstream level, the immediate question is how to optimize the timing and location of each stop, based on vehicle operations and customer locations. At the upstream level, service providers need to comprehensively re-optimize routing operations to take full advantage of vehicle-customer coordination—which customers to serve, with which vehicles, and in which sequence. In turn, vehicle-customer coordination requires dedicated algorithms. The only attempt to date to solve a vehicle routing problem with vehicle-customer coordination comes from Gambella *et al.* (2018), who formulated the problem in the Euclidean space as a mixed-integer second-order cone program and developed a branch-and price algorithm. Yet, their algorithm falls short of the large-scale instances arising in many practical applications.

2 OPTIMIZATION METHODOLOGY

In response, the first goal of this paper is to develop scalable algorithms to support routing optimization with vehicle-customer coordination. We tackle a *Dial-A-Ride problem with Vehicle-Customer Coordination (DAR–VCC)*, in which customers request transportation from an origin to a destination by a deadline, and an operator optimizes vehicle-customer assignments, the sequence of pickups and dropoffs, as well as the location and time of each stop. We formulate it as a mixed-integer second-order cone program (MISOCP) in the Euclidean ℓ_2 space and as a mixed-integer linear program (MILP) in the Manhattan ℓ_1 space. Either way, off-the-shelf implementation can only solve small-scale instances. Instead, we develop a decomposition algorithm that breaks down the problem into three steps:

1. Single Stop Optimization with Vehicle-Customer Coordination (SSO– VCC). We derive geometric insights on the optimal stopping location (between two fixed locations) as a function of the vehicle's speed, the customer's speed and the customer's maximum walking distance. This reduces SSO–VCC from a three-dimensional problem (one temporal dimension and two spatial dimensions) to one-dimensional problems in the ℓ_2 space (the direction of walking, as illustrated in Figure 1a), or a closed-form system of linear equations in the ℓ_1 space (Figure 1b). Either way, SSO–VCC can be solved very efficiently.

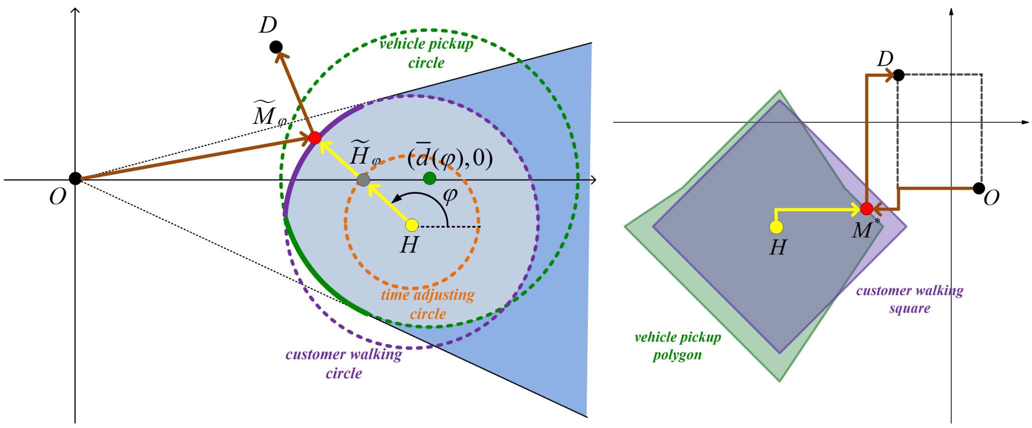

(a) *Euclidean space.* (b) *Manhattan space.*

Figure 1 – *Geometric representation of SSO–VCC (O: vehicle origin; D; vehicle destination; H: customer home; M: optimized stopping location).*

2. Multiple Stop Optimization with Vehicle-Customer Coordination (MSO– VCC). This problem involves optimizing the locations and times of multiple stops in a given sequence. We propose a tailored coordinate descent scheme that optimizes one stop at a time, thus decomposing MSO–VCC into a sequence of SSO–VCC problems. We prove that, in the ℓ_1 space and the ℓ_2 space, this algorithm terminates at a globally optimal solution of the MSO–VCC—a constrained and non-separable optimization problem (Theorem 1). This algorithm provides our core optimization engine, which solves the MSO–VCC much faster than off-the-shelf implementation.

Theorem 1 *In a metric space armed with a differentiable norm (e.g., the ℓ_2-norm) or the ℓ_1-norm, MSO–VCC can be solved to optimality by solving SSO–VCC problems iteratively.*

3. Dial-A-Ride with Vehicle-Customer Coordination (DAR– VCC). We propose a new time-space network representation for dial-a-ride problems where "empty vehicles" flow from node to node and arcs represent customer-serving trips. This structure captures vehicle capacities and time windows into the network itself. We develop an algorithm combining dynamic programming to generate candidate trips (which embeds MSO–VCC to optimize the time and location of each stop) and a time-space network optimization to select trips (see Figure 2). This algorithm provides a new dial-a-ride methodology combining set partitioning and time-space formulations.

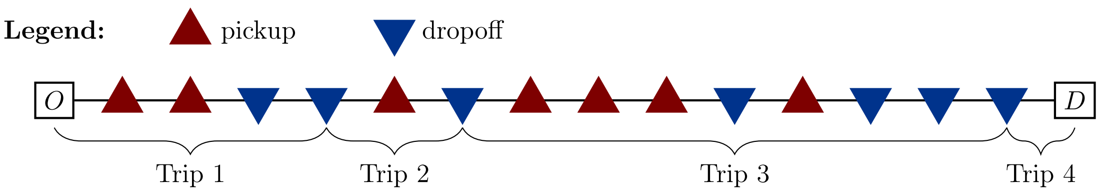

Figure 2 – *A vehicle path combines several customer-serving trips via time-space network optimization.*

Computational results show that this algorithm significantly outperform state-of-the-art MISOCP and MILP benchmarks, both in terms of solution quality and computational times (Table 1). Indeed, MISOCP or MILP implementation can only solve instances with up to 6–8 customers, returns a very loose optimality gap with 10–50 customers, and does not even find a feasible solution with more than 25–50 customers. We also propose a CPLEX-based heuristic that achieves better scalability and consistently returns feasible solutions—albeit, rarely terminating in 1 hour. In comparison, our algorithm can handle medium-scale instances with up to 50 customers in seconds and large-scale instances with up to 200 customers in minutes. To show the robustness of these results, we also solve the *Vehicle Routing Problem with Vehicle-Customer Coordination (VRP–VCC)* from Gambella *et al.* (2018) by embedding MSO–VCC into a dynamic programming algorithm. Our approach provides Pareto improvements as compared to the branch-and-price benchmark: higher-quality solutions in shorter runtimes.

Table 1 – *Average computational results for DAR–VCC (3 vehicles, Euclidean distance).*

		CPLEX					CPLEX-based heuristic			Our algorithm	
Norm	# cus.	Sol.	UB	Opt. gap	Gap vs. alg.	CPU (s)	Sol.	Gap vs. alg.	CPU (s)	Sol.	CPU (s)
ℓ_2	5	18.3	18.3	0%	0%	66	18.3	0%	110	18.3	<1
	10	4.5	33.5	652%	439%	>3,600	15.2	58%	>3,600	24.0	<1
	15	6.4	56.3	787%	568%	>3,600	18.6	128%	>3,600	42.4	<1
	20	7.0	72.2	937%	493%	>3,600	18.3	126%	>3,600	41.3	<1
	25	4.6	94.4	1,936%	901%	>3,600	20.2	130%	>3,600	46.4	<1
	50	-	-	-	-	-	17.9	253%	>3,600	63.2	340
	100	-	-	-	-	-	15.8	266%	>3,600	57.8	180
	200	-	-	-	-	-	15.1	317%	>3,600	63.1	2,843
ℓ_1	5	21.9	21.9	0%	0%	1	21.9	0%	3	21.9	<1
	10	29.5	40.2	36%	1%	>3,600	26.8	11%	3,100	29.8	<1
	15	38.1	65.7	73%	17%	>3,600	33.8	31%	>3,600	44.4	<1
	20	30.7	90.1	193%	48%	>3,600	30.2	50%	>3,600	45.4	<1
	25	20.6	112.3	445%	160%	>3,600	35.4	51%	>3,600	53.5	<1
	50	7.0	212.1	2,946%	919%	>3,600	29.0	145%	>3,600	71.0	96
	100	-	-	-	-	-	33.2	98%	>3,600	65.9	90
	200	-	-	-	-	-	30.7	138%	>3,600	72.9	1,939

3 ONLINE IMPLEMENTATION

Our second goal is to support real-world operations with vehicle-customer coordination, with a focus on ride-sharing services where the platform can request customers to walk to the pickup location or from the dropoff location. We define an *Online Dial-A-Ride with Vehicle-Customer Coordination (O–DAR– VCC)*. Consistent with practice and the literature, we assume that operations proceed via batching and optimization, that is, customer requests are aggregated and matched to vehicles every few seconds (Yan *et al.*, 2020, Bertsimas *et al.*, 2019, Ashlagi *et al.*, 2019). At each epoch, our algorithm optimizes service to "new" customers, service to "backlogged" customers, and vehicle repositioning. It also captures vehicles' and customers' operations in complex road networks featuring one-way streets, traffic congestion, and a handful of possible stopping locations (as opposed to continuous operations in the ℓ_1 or ℓ_2 space). Using real-world data from New York City, we show that our algorithm consistently terminates in seconds, thus enabling its real-time implementation in very large-scale networks of operations.

From a practical standpoint, our results suggest that vehicle-customer coordination can provide significant improvements in routing operations, with an average profit increase of 3–4%. For a system of the size of Manhattan, this represents an estimated gain of $55,000 daily, or $20M annually. Moreover, the profit increase stems from two sources: (i) cost savings at the "downstream" level, by meeting customers in more convenient locations, and (ii) profit improvements at the "upstream" level, by serving more customers or higher-margin customers. In fact, the second source has a much larger contribution to the profit increase than the first one (85–95% vs. 5–15%). These results suggest that the main benefits of vehicle-customer coordination do not stem from downstream adjustments in vehicle routes; rather, most of the gains stem from comprehensively re-optimizing "upstream" operations: which customers to serve, in which sequence and with which vehicles. We also consider several discount schemes through which the operator can share the benefits of vehicle-customer coordination with customers—as is the case in practice. By sharing the benefits equally between the operator and customers, we obtain a solution such that (i) the operator's profit increases by 1.7%, (ii) 4.6% extra customers receive a service, (iii) 13% of customers are requested to walk to the pickup location and 3% from the dropoff location (96.5 meters on average) and offered a 4.6% discount, and (iv) vehicle miles traveled are reduced by 11.0%. This solution represents a win-win-win outcome: higher operating profits, better customer level of service, and smaller environmental footprint.

Table 2 – *Average performance (14 days between 12/1/2019 and 12/14/2019; 90-minute window between 5:00 PM and 6:30 PM; around 16,000 customer requests per day; 2,000 vehicles).*

	No ride-sharing		Ride-sharing			
Metric	No VCC	VCC:PU	No VCC	VCC:PU	VCC:DO	VCC:PU–DO
Profit	$133,997	$139,041	$139,780	$143,967	$140,507	$144,583
Profit increase	(base)	3.8%	(base)	3.0%	0.5%	3.4%
Downstream contribution	(base)	12.4%	(base)	15.5%	9.6%	14.8%
Upstream contribution	(base)	87.6%	(base)	84.5%	90.4%	85.2%
Revenue	$139,635	$144,302	$144,494	$148,196	$145,181	$148,767
Revenue increase	(base)	3.3%	(base)	2.6%	0.5%	3.0%
Cost	$5,639	$5,261	$4,714	$4,229	$4,674	$4,184
Cost per request	$0.48	$0.43	$0.39	$0.33	$0.38	$0.33
Cost decrease	(base)	6.7%	(base)	10.3%	0.8%	11.2%
Acceptance rate	71%	75%	75%	78%	75%	78%
Pickups away from origins	—	10%	—	13%	—	13%
Dropoffs away from destinations	—	—	—	—	3%	3%
Vehicle miles traveled (total)	46,367	43,503	38,995	35,031	38,677	34,683
Vehicle miles traveled (per request)	4.01	3.58	3.21	2.77	3.16	2.73
Max discount (uniform)	(base)	3.5%	(base)	2.8%	0.5%	3.2%
Max discount (targeted)	(base)	18.7%	(base)	9.6%	8.4%	9.2%
Max discount (prorated, per 100m)	(base)	$2.8	(base)	$1.4	$1.2	$1.2

"VCC:PU", "VCC:DO", "VCC:PU–DO": vehicle-customer coordination for pickups, dropoffs and both.
"Downstream contribution": share of the profit increase from downstream cost savings.
"Upstream contribution": share of the profit increase from upstream change in customer mix.
Max. discount: largest discounts for the operator to break even, under various schemes.

References

Ashlagi, Itai, Burq, Maximilien, Dutta, Chinmoy, Jaillet, Patrick, Saberi, Amin, & Sholley, Chris. 2019. Edge weighted online windowed matching. *Pages 729–742 of: Proceedings of the 2019 ACM Conference on Economics and Computation.*

Bertsimas, Dimitris, Jaillet, Patrick, & Martin, Sébastien. 2019. Online vehicle routing: The edge of optimization in large-scale applications. *Operations Research*, **67**(1), 143–162.

Gambella, Claudio, Naoum-Sawaya, Joe, & Ghaddar, Bissan. 2018. The Vehicle Routing Problem with Floating Targets: Formulation and Solution Approaches. *INFORMS Journal on Computing*, **30**(3), 554–569.

Yan, Chiwei, Zhu, Helin, Korolko, Nikita, & Woodard, Dawn. 2020. Dynamic pricing and matching in ride-hailing platforms. *Naval Research Logistics*, **67**(8), 705–724.