Flexibility and coordination in on-demand transportation: from ride-sharing to micromobility
Jacquillat 11/18

Jason Luo

Part I. Literature (for further reading about ride sharing)

Cordeau, JF., Laporte, G. The dial-a-ride problem: models and algorithms. Ann Oper Res 153, 29–46 (2007). https://doi.org/10.1007/s10479-007-0170-8

Dimitris Bertsimas, Patrick Jaillet, Sébastien Martin (2019) Online Vehicle Routing: The Edge of Optimization in Large-Scale Applications. Operations Research 67(1):143-162. https://doi.org/10.1287/opre.2018.1763

Jean-François Cordeau, (2006) A Branch-and-Cut Algorithm for the Dial-a-Ride Problem. Operations Research 54(3):573-586. https://doi.org/10.1287/opre.1060.0283

Pallottino, S., Scutellà, M.G. (1998). Shortest Path Algorithms In Transportation Models: Classical and Innovative Aspects. In: Marcotte, P., Nguyen, S. (eds) Equilibrium and Advanced Transportation Modelling. Centre for Research on Transportation. Springer, Boston, MA. https://doi.org/10.1007/978-1-4615-5757-9_11

S. Ma, Y. Zheng and O. Wolfson, "Real-Time City-Scale Taxi Ridesharing," in IEEE Transactions on Knowledge and Data Engineering, vol. 27, no. 7, pp. 1782-1795, 1 July 2015, doi: 10.1109/TKDE.2014.2334313.

Part II. Recent News

Admin. “Washington Selects Operator for Rideshare Driver Resource Center.” Insurance Journal, 20 Oct. 2022, https://www.insurancejournal.com/news/west/2022/10/24/691216.htm.

Holt, Kris. “Lyft Brings Shared Rides Back to More Cities, Including San Francisco.” Engadget, 6 May 2022, https://www.engadget.com/lyft-shared-rides-san-francisco-san-jose-denver-las-vegas-atlanta-covid-19-183352620.html.

Holt, Kris. “NYC Will Enforce Uber and Lyft Driver Pay Increases by the End of the Year.” Engadget, 15 Nov. 2022, https://www.engadget.com/nyc-uber-and-lyft-driver-pay-increase-205833433.html.

Part III. Questions and Comments

Robert Freund: I have to say as an optimizer, Of course I couldn't help but be fascinated by the stuff related to implications and what kind of algorithms were used…what about penalties for cancellations? On the issue of flexible pick ups, I was in Manhattan just a couple of weeks ago, and I would imagine flexible pickups are much more valuable there because of the one way street configurations. I'd love to see the carbon implications at the end.

Prof. Alex Jacquillat: Thank you so much. Rob. Your point on Manhattan is very important, and actually there's a lot of data engineering that went into the last case study that I presented, because here, we are not evolving into the Manhattan distance(despite the name), we're not evolving into the Euclidean space and continuous space, but we have a very discrete network, and actually we used OpenStreetMap data in order to calibrate travel. You know all the things one way streets, travel times, and all of these things to actually capture this thing. So there is a lot of things that are going under the hood to make this as realistic as possible, and that's also one of the reasons why a little bit of flexibility can provide a lot of benefits. As I said at the beginning, on average, the walking distance is one hundred meters. So think about it as like two blocks. Not a lot. But that goes a long way in avoiding these detours because of one way streets and these types of things, so that's a great comment.

Jinhua Zhao: There is asymmetry between pick up and a drop off right? So you mentioned that they are asymmetrical in the benefits point of view, but also in the coordination cost point of view. We also noticed that pick up requires much higher cost or effort of coordination than drop off, drop off of this kind of trivial right, but pick up involves navigation, how to find each other, find a reasonable safe place to park, etc. Also from the consumer point of view waiting at home versus waiting at some street corner may be very different things. The sense of safety, the sense that I can do something useful. This also relates to the fact that for many Americans, if you are a public transit user, walking a little bit and then waiting is the norm. But many Americans will never have had the notion of public transit as they have been driving through to go everywhere. That's kind of a major culture shift on this right? So what are your remarks all about this: the asymmetry and also the behavior.

Prof. Alex Jacquillat: Yeah, that's a really good line of questions. So talking a little bit about the asymmetry. I think that you are entirely right. That's a good read of the results at the end. That's also a good read of the practical implementation challenges. The motivation for this work, and a lot of my research in general is that digitization is enabling more coordination. Therefore we have the ability to use all of the real time sensing data in order to actually power that coordination. And then the question that I'm asking is how to do that. And it's actually not that trivial to really understand where to coordinate and how to coordinate and we are surfing a wave of digitization and developing the right optimization method for this, and showing when it's most beneficial. I think that your point on behaviors is extremely important, and there are a couple of points that I would like to make here. Transit users or not, it comes at an inconvenience to walk, especially when we start looking at November weather and it being cold outside. So I think that there are two ways in which these results can be used. Number one is that right now the platforms have a very good understanding of who is more price sensitive versus who is more service quality sensitive in their user base. Actually, the users can select whether to use the service or not, so you can actually just use customer buy in as a proxy for who's the population and their interests, and then you can use the algorithm to actually power that. Number two, I also think that the model that we have can be used to answer other questions such as when is it really valuable? Maybe you have a couple of pick ups for which, working one hundred meters can actually mean so much in terms of system efficiency. We are sure that my model does recommend a little bit of walking. But honestly, there is an almost optimal solution where the customer doesn't walk, and so maybe another use of the model would be to say, “What if you actually reduce the amount of walking by a factor of two…because it's an inconvenience, and you want to minimize that inconvenience?” So this is really where not only the development of the model, but the use of the model can have an impact.

Michael Leong: What if we had a transportation system where rides are only requested between predetermined stops at every one hundred meters and the customer could choose between these discrete points instead of in a continuous space? Would that make it easier to optimize and also make it easier to operate and regulate?

Prof. Alex Jacquillat: It's three levels of answers. The first answer is yes, great idea. Let's do it. The second level is that I have some results that I didn't show today. Let me first tell you how we implemented the Manhattan problem. The Manhattan problem again doesn't involve continuous space. We have a few candidate stopping locations. What do we consider at this point? Are they simply intersections and hotels? So whenever you have a hotel you can have a curve, and you can stop here, and otherwise we're just saying you stop by the intersections, and you are assigned to an intersection, so we could replace this network by what you're suggesting. And actually what we find is that the density of the candidate stops does not have a huge impact on the efficiency of operations, meaning that if the city doesn't want to have that at every intersection but maybe every other intersection, I don't think that you would leave a lot of money on the table, so that in that that's a positive result, right? The positive result is that you could design a few well chosen bus stops that would be just nicer to wait at and that could actually make the system more appropriate to this type of service. However, the other answer is that the one thing where I think you would be leaving money on the table is if you let the customers choose where they want to be picked up at. Imagine that. So just to take a simple example in Manhattan, you are on Fifth Avenue, on 45th Street. You would decide whether you're going to 43 or 47. But in my system the system chooses that for you. It chooses which one you would go to. And this is actually important, because once the system has that flexibility, it can optimize your ride not only for you, but actually for the other users as well, and create more efficient routes for the drivers. So I have a hybrid answer, saying that I think that city design can play a huge role toward that transition. But there needs to be some sort of centralized control by the platform in order to really maximize the efficiency.

John Moavenzadeh: What programming language or simulation tool was used to produce the results?

Prof. Alex Jacquillat: So a lot of pre-processing using Python from the OpenStreetMap data. And then for the optimization, we use a combination of Julia and C#, because we are split between MIT and Hong Kong Polytechnic University. We use different languages, but we could communicate with these programming languages.

Part IV. Summary of Memos.

Themes from Other Memos

1. Potentially using these ideas to complement the rest of the transportation industry instead of as a competitor. For example, we can use and fund such ridesharing services to help alleviate the burden on public transport on areas where having multiple buses run at high frequency would be infeasible, or to serve as the last mile of existing public transport systems
2. Would the benefit gained by walking be significantly reduced in areas with lower density and less restrictions such as one-way streets?
3. Could this potentially cause harm to existing transportation infrastructures by serving as too attractive of an alternative to public transport?

My Reflection

This week we had the pleasure of listening to Alex Jacquillat present his work on ride sharing. His presentation focused on how to optimize ride sharing by having customers walk a short distance. Professor Jacquillat presented many interesting techniques for solving this version of the problem. First of all, he presented his findings on how he was able to reduce the single-stop optimization problem to a problem in only one dimension. In addition, the optimal solution to this problem would not have the vehicle wait for the passenger for any amount of time because the vehicle moves faster. Secondly, he talked about how he was able to solve the multi-stop problem by solving a sequence of the single-stop optimization problems as a dynamic program if the norm used was differentiable or the L1 norm. He then presented his findings on how these techniques would cause a large increase in profits for ridesharing companies, who could then pass some of these profits along to the consumer. After listening to his talk, some points were left unclear to me. The topic of the distance that riders are willing to walk was addressed at the beginning of the talk, but it is unclear to me whether a large portion of riders would be willing to share a ride in the first place. Having too small a pool of riders could clearly cause some issues, and this is one possible effect of having too many riders be unwilling to share a ride. Lastly, when Professor Jacquillat presented his slide on the comparison of solutions between the standard CPLEX solution and his new algorithm, I noticed that many of the objective values were strictly better with his new algorithm than with the CPLEX solution. It makes sense as to why CPLEX is slower than his new algorithm given that CPLEX is a general solver that doesn’t utilize the problem structure, but why was it unable to solve any of the problems to anywhere close to optimality?

This seminar was the most technical of all the seminars given thus far in the semester, and it is unfortunate that we did not have time to cover the usage of the Benders cut and column generation or additional details about the applications to microtransit. This was my favorite seminar so far, and I would love to have Professor Jacquillat back next semester to give further insight on his work.

Part V. Other Information

MANAGEMENT
SLOAN SCHOOL

ROUTING OPTIMIZATION WITH VEHICLE-CUSTOMER COORDINATION

Alexandre Jacquillat

Kai Wang

Shuaian Wang

Wei Zhang

Overall research agenda

Air traffic management

On-demand mobility

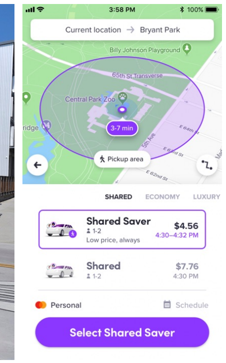

Logistics decarbonization

Societal applications

Vehicle-customer coordination (VCC)

- “Walking products”: flexibility in pickup/ dropoff locations
- How to coordinate vehicles and customers, with which impact?

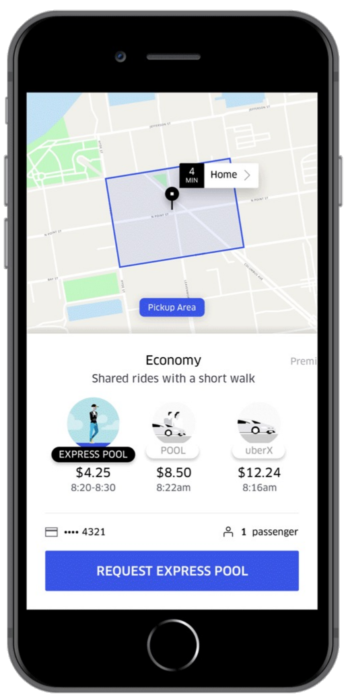

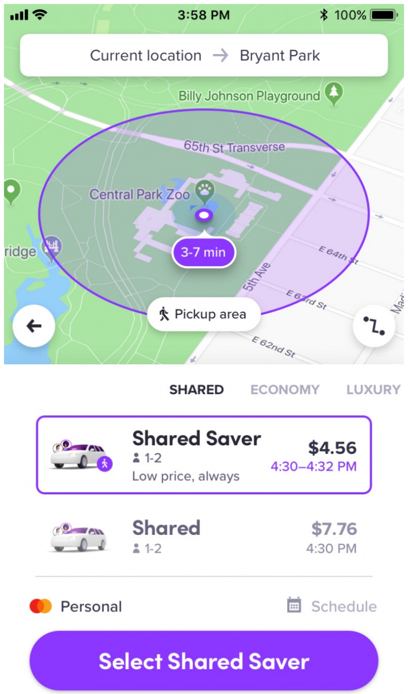

Coordinated operations

- Downstream: smoother vehicle operations
- Upstream: how to re-optimize overall routing operations?

→ Which customers to visit, in which sequence, with which vehicles, in which locations and at what times?

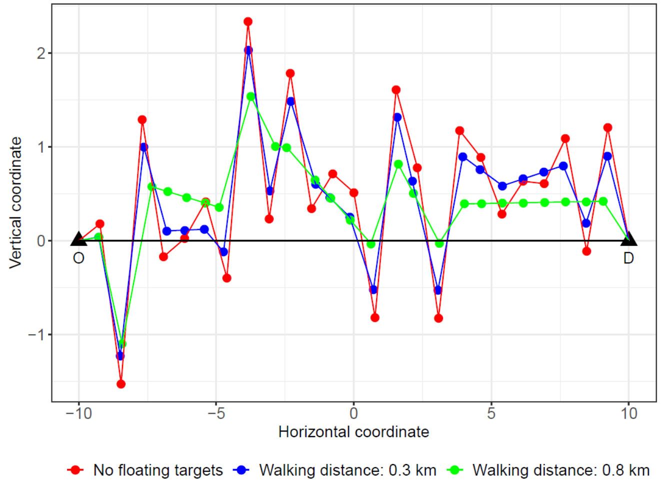

(a) Euclidean distance

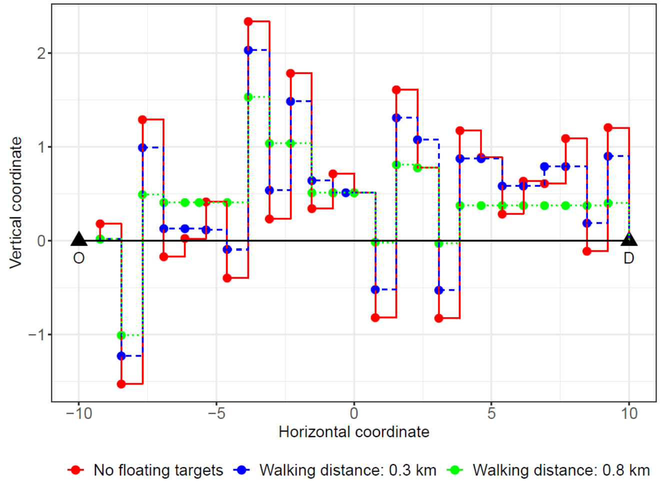

(b) Manhattan distance

Ultimate objective: DAR–VCC

- Large-scale optimization formulation
  - MISOCO: ℓ_2-norm
  - MILO: ℓ_1-norm, discrete network
- Weak polyhedral structure due to routing and coordination complexities

$$
\begin{aligned}
&Q^P_{ri} = Q^P_{r,i-1} + \sum_{j \in \mathcal{P} \cup \mathcal{D}} q_j x_{rij} \quad \forall i \in \mathcal{N},\ \forall r \in \mathcal{V} \\
&\overline{T}_j \geq v_{ri} - N^V_j (1 - x_{r,i,j+n}) \quad \forall i \in \mathcal{N}, \forall j \in \mathcal{P},\ \forall r \in \mathcal{V}
\end{aligned}
$$

$$
\begin{aligned}
&\|M^V_{ri} - M_j\| \leq N^W (1 - x_{rij}) \quad \forall i \in \mathcal{N},\ \forall j \in \mathcal{P},\ \forall r \in \mathcal{V} \\
&\|M^V_{ri} - G_j\| \leq N^W (1 - x_{rij}) \quad \forall i \in \mathcal{N},\ \forall j \in \mathcal{D},\ \forall r \in \mathcal{V} \\
&v_{ri} \geq \frac{\|M_j - H_j\|}{S_j} - N^T_j (1 - x_{rij}) \quad \forall i \in \mathcal{N},\ \forall j \in \mathcal{P},\ \forall r \in \mathcal{V}
\end{aligned}
$$

$$
\begin{aligned}
\max \quad & \sum_{r \in \mathcal{V}} \sum_{j \in \mathcal{P}} \sum_{i \in \mathcal{N}} g_j x_{rij} - \sum_{r \in \mathcal{V}} c \cdot v_{r,2n+1}, \\
\text{s.t.} \quad & M^V_{r0} = O \quad \forall r \in \mathcal{V}, \\
& M^V_{r,2n+1} = \sum_{j \in \mathcal{P}} x_{r,1,j} D + \left(1 - \sum_{j \in \mathcal{P}} x_{r,1,j}\right) O \quad \forall r \in \mathcal{V}, \\
& \|M^V_{ri} - M_j\| \leq N^W (1 - x_{rij}) \quad \forall i \in \mathcal{N},\ \forall j \in \mathcal{P},\ \forall r \in \mathcal{V}, \\
& \|M^V_{ri} - G_j\| \leq N^W (1 - x_{rij}) \quad \forall i \in \mathcal{N},\ \forall j \in \mathcal{D},\ \forall r \in \mathcal{V}, \\
& v_{ri} \geq v_{r,i-1} + \frac{\|M^V_{ri} - M^V_{r,i-1}\|}{\overline{S}} \quad \forall i \in \mathcal{N} \cup \{2n+1\},\ \forall r \in \mathcal{V}, \\
& v_{ri} \geq \frac{\|M_j - H_j\|}{S_j} - N^T_j (1 - x_{rij}) \quad \forall i \in \mathcal{N},\ \forall j \in \mathcal{P},\ \forall r \in \mathcal{V}, \\
& \|M_j - H_j\| \leq W_j \quad \forall j \in \mathcal{P}, \\
& \sum_{j \in \mathcal{P} \cup \mathcal{D}} x_{rij} \geq \sum_{j \in \mathcal{P} \cup \mathcal{D}} x_{r,i+1,j} \quad \forall i \in \mathcal{N} \setminus \{2n\},\ \forall r \in \mathcal{V}, \\
& \sum_{j \in \mathcal{P} \cup \mathcal{D}} x_{rij} \leq 1 \quad \forall i \in \mathcal{N},\ \forall r \in \mathcal{V}, \\
& \sum_{r \in \mathcal{V}} \sum_{i \in \mathcal{N}} x_{rij} \leq 1 \quad \forall j \in \mathcal{P} \cup \mathcal{D}, \\
& \sum_{i \in \mathcal{N}} x_{rij} = \sum_{i \in \mathcal{N}} x_{r,i,j+n} \quad \forall j \in \mathcal{P},\ \forall r \in \mathcal{V}, \\
& \sum_{i \in \mathcal{N}} i x_{rij} \leq \sum_{i \in \mathcal{N}} i x_{r,i,j+n} \quad \forall j \in \mathcal{P},\ \forall r \in \mathcal{V}, \\
& Q^P_{ri} = Q^P_{r,i-1} + \sum_{j \in \mathcal{P} \cup \mathcal{D}} q_j x_{rij} \quad \forall i \in \mathcal{N},\ \forall r \in \mathcal{V}, \\
& Q^P_{ri} \leq Q \quad \forall i \in \mathcal{N},\ \forall r \in \mathcal{V}, \\
& \overline{T}_j \geq v_{ri} - N^V_j (1 - x_{r,i,j+n}) \quad \forall i \in \mathcal{N}, \forall j \in \mathcal{P},\ \forall r \in \mathcal{V}, \\
& \overline{T} \geq v_{r,2n+1},\ \forall r \in \mathcal{V}, \\
& x_{rij} \in \{0,1\} \quad \forall i \in \mathcal{N},\ \forall j \in \mathcal{P} \cup \mathcal{D},\ \forall r \in \mathcal{V}, \\
& Q^P_{ri} \geq 0 \quad \forall i \in \mathcal{N},\ \forall r \in \mathcal{V}, \\
& v_{ri} \geq 0 \quad \forall i \in \mathcal{N} \cup \{2n+1\},\ \forall r \in \mathcal{V}, \\
& M^V_{ri} \in \mathbb{R}^2 \quad \forall i \in \mathcal{N} \cup \{2n+1\},\ \forall r \in \mathcal{V}, \\
& M_j \in \mathbb{R}^2 \quad \forall j \in \mathcal{P}.
\end{aligned}
$$

Methodological contributions

Models and algorithms to optimize routing operations with vehicle-customer coordination in large-scale networks

Single stop optimization (SSO-VCC)	Core subproblem: Optimization of location and time of a single stop, using geometric insights
Multi-stop optimization (MSO-VCC)	Optimization of locations and times of a sequence of stops, using an exact coordinate descent algorithm
Vehicle routing (VRP-VCC) Dial-a-ride (DAR-VCC)	Integration of MPO-VCC into routing optimization: subpath-based variables & time-space networks
Large-scale online routing optimization	Rolling algorithm for real-world ride-sharing: scalability and win-win-win outcomes

Single-stop optimization

Single stop optimization: statement

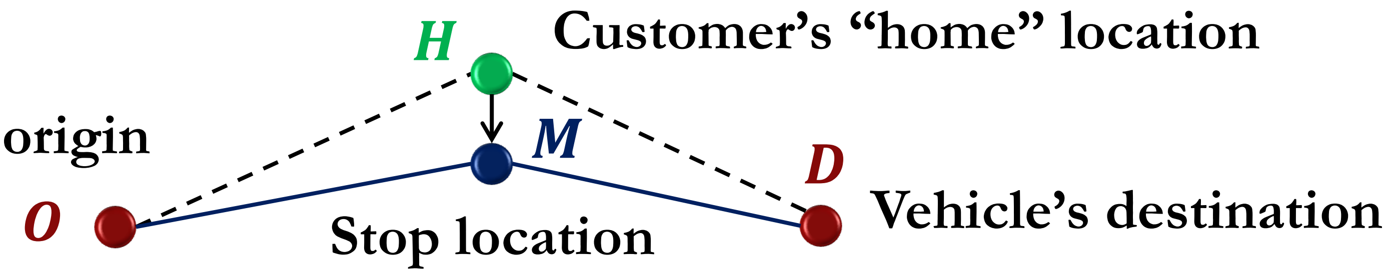

- Operations in the continuous space
- Optimizing stopping time and location
- SOCO with ℓ_2-norm; LO with ℓ_1-norm

$$
\begin{aligned}
\min\ & \left(v_M + \frac{\|D-M\|}{\overline{S}} \right), && \textbf{arrival time} \\
\text{s.t.}\ & v_M \geq \overline{v} + \frac{\|M-O\|}{\overline{S}}, && \textbf{vehicle travel time} \\
& v_M \geq \frac{\|M-H\|}{S}, && \textbf{customer travel time} \\
& \|M-H\| \leq W, && \textbf{walking distance} \\
& v_M \geq 0,\ M \in \mathbb{R}^2. && \textbf{definition domain}
\end{aligned}
$$

Single stop optimization: insights

Lemma: the vehicle does not wait at the stopping location

Lemma: (i) straight-path travel; (ii) synchronized arrivals of vehicle and customer; or (iii) binding maximum walking distance

Straight path

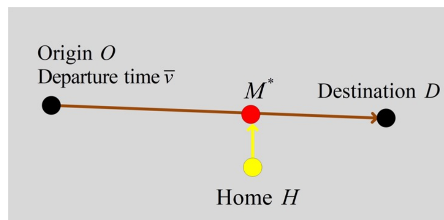

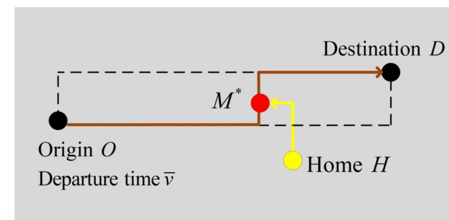

Detour

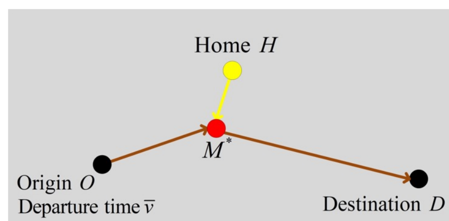

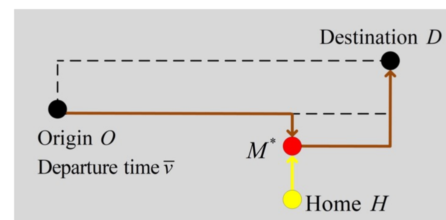

$$v^* = \min_{M \in \mathbb{R}^2} \left\{ \overline{v} + \frac{\|M-O\| + \|D-M\|}{S} \; : \; \text{s.t. } \frac{\|M-H\|}{S} \leq \overline{v} + \frac{\|M-O\|}{S}, \; \|M-H\| \leq W \right\}$$

Single stop optimization: solution

Euclidean space: SSO-VCC can be reduced from a 3D problem (time and space) to a 1D problem (walking angle), solved in $\mathcal{O}(1/\varepsilon)$

Manhattan space: SSO-VCC solved via a system of linear equations

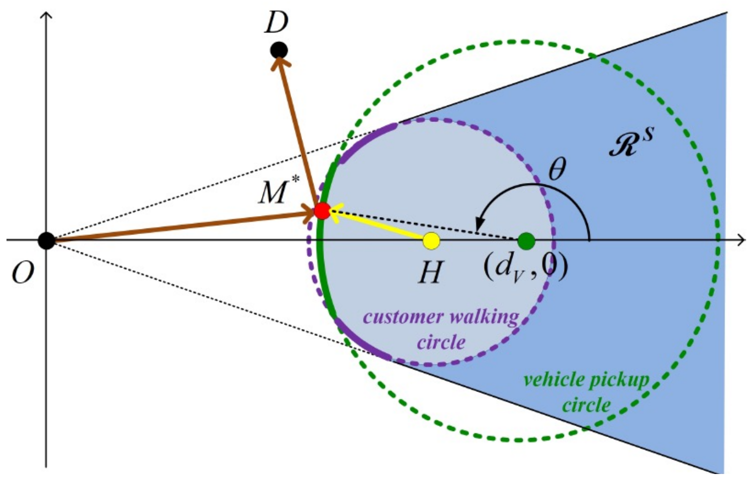

Euclidean space

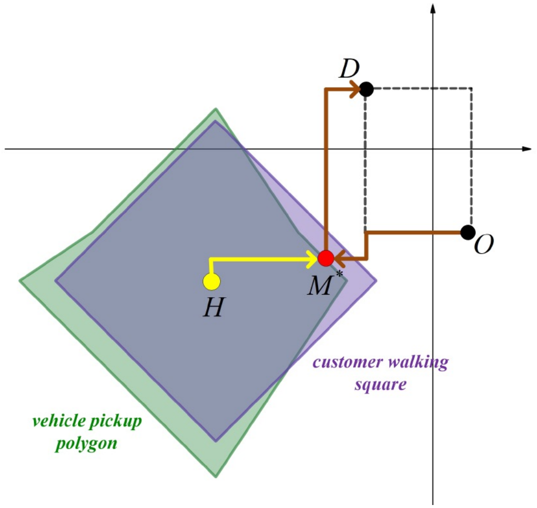

Manhattan space

Multi-stop optimization

Multi-stop optimization: statement

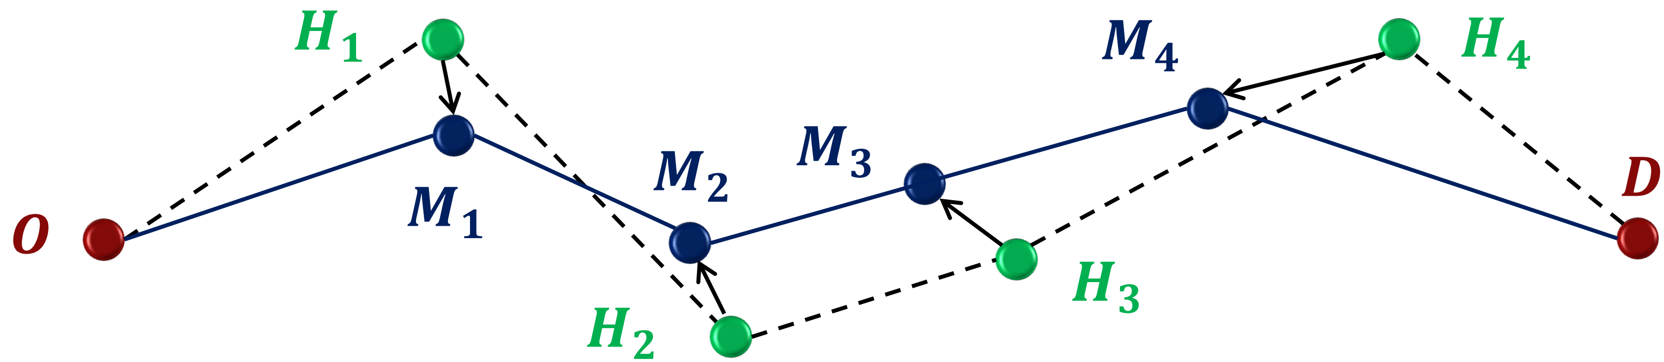

$$
\begin{aligned}
\text{[MSO–VCC]} \quad & \min\ v_{K+1}, && \textbf{arrival time}\\
& \text{s.t. } v_0 = \overline{v},\ M_0 = O,\ M_{K+1} = D, && \textbf{departure time}\\
& v_k \geq v_{k-1} + \frac{\|M_k - M_{k-1}\|}{\overline{S}} \quad \forall k = 1, \cdots, K+1, && \textbf{vehicle travel time}\\
& v_k \geq \frac{\|M_k - H_k\|}{S_k} \quad \forall k = 1, \cdots, K, && \textbf{customer travel time}\\
& \|M_k - H_k\| \leq W_k \quad \forall k = 1, \cdots, K, && \textbf{walking distance}\\
& v_k \geq 0,\ M_k \in \mathbb{R}^2 \quad \forall k = 0, \cdots, K+1. && \textbf{definition domain}
\end{aligned}
$$

Multi-stop optimization: results

- Constrained and non-separable convex optimization problem
  → Coordinate descent algorithm with feasibility and optimality guarantees
- Reformulation as a continuous-space dynamic program
  → A tailored backward induction algorithm (outer loop) that ensures the optimality of all subsequent decisions (inner loop)

$$\Pi_{k-1}(M_{k-1}, v_{k-1}) = \min_{M_k \in \mathbb{F}_k(M_{k-1}, v_{k-1})} \left\{ \frac{\|M_k - M_{k-1}\|}{\overline{S}} + \Pi_k \left(M_k, v_{k-1} + \frac{\|M_k - M_{k-1}\|}{\overline{S}} \right) \right\}$$

Theorem: If $||.||$ is differentiable or the ℓ_1-norm, MSO–VCC can be solved to optimality by solving SSO–VCC problems iteratively

Multi-stop optimization: algorithm

- Iteratively re-optimize stop k, while re-optimizing "tail" $k+1, \ldots$

Full

$$\min_{M_1 \in \mathbb{F}_1(O,\overline{v})} \left\{ \overline{v} + \frac{\|M_1 - O\|}{\overline{S}} + \Phi\left(M_1, \Xi^*(M_1, v_1)\right) \right\}, \quad \text{where } v_1 = \overline{v} + \frac{\|M_1 - O\|}{\overline{S}}$$

Algorithm

$$M_1^{(l)} \in \underset{M_1 \in \mathbb{F}_1(O,\overline{v})}{\arg\min} \left\{ \overline{v} + \frac{\|M_1 - O\|}{\overline{S}} + \Phi\left(M_1, \Xi^*\left(M_1^{(l-1)}, v_1^{(l-1)}\right)\right) \right\}, \text{ and } v_1^{(l)} = \overline{v} + \frac{\left\|M_1^{(l)} - O\right\|}{\overline{S}}$$

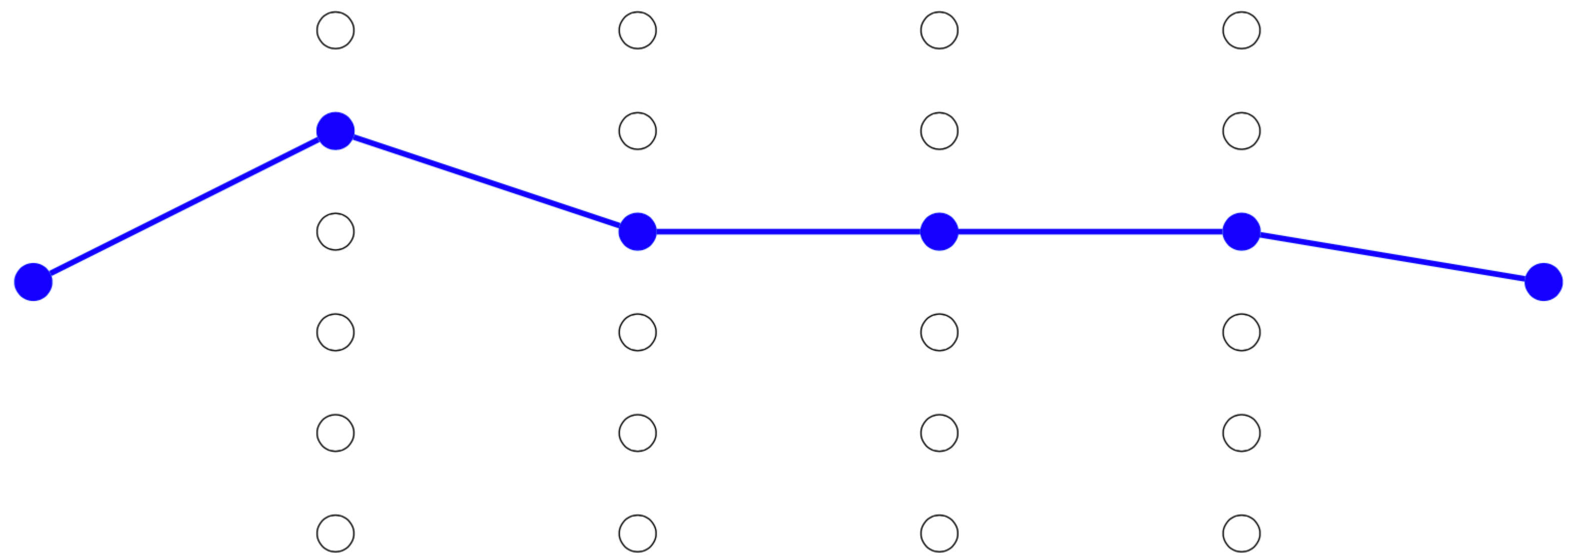

Routing optimization

Routing optimization: statements

- Vehicle routing problem (VRP–VCC)
  - Visiting customers with multiple vehicles to minimize travel time
- Offline dial a ride (DAR–VCC)
  - Transport customers from origin to destination to maximize profits
- Online dial-a-ride (O–DAR–VCC)
  - Dynamic arrivals, re-optimization every 10 s
  - NYT data: 20,000 requests per hour

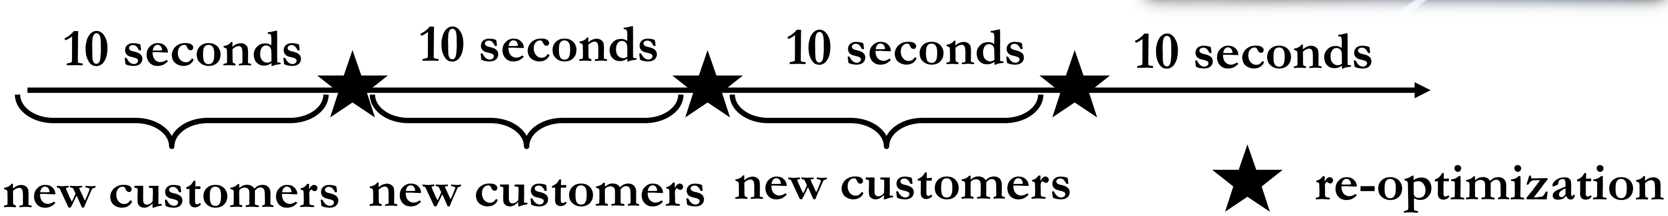

Subpath-based time-space model

Route-based formulation

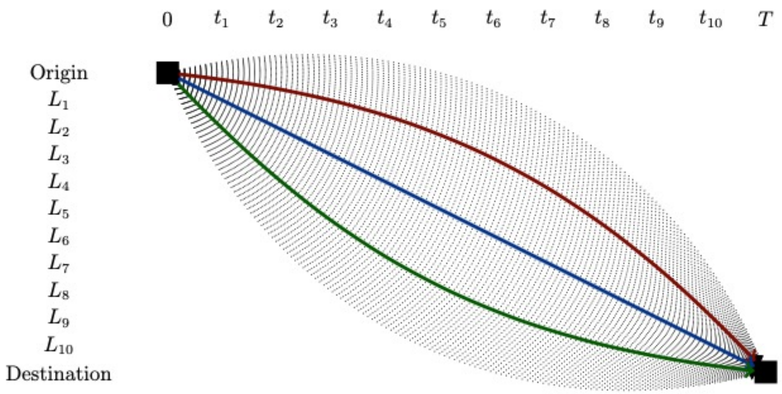

Arc-based formulation

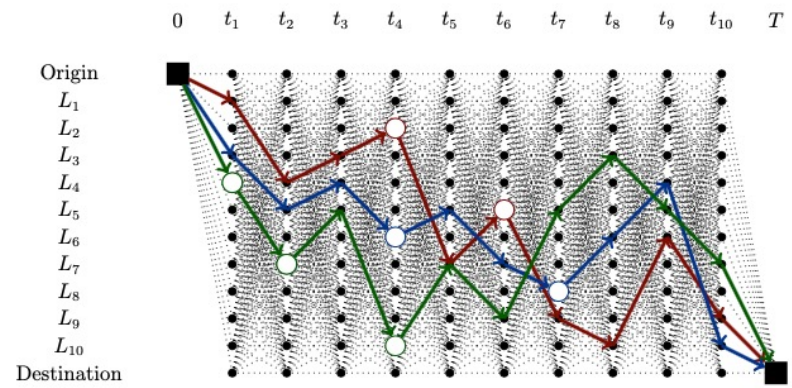

Subpath-based formulation

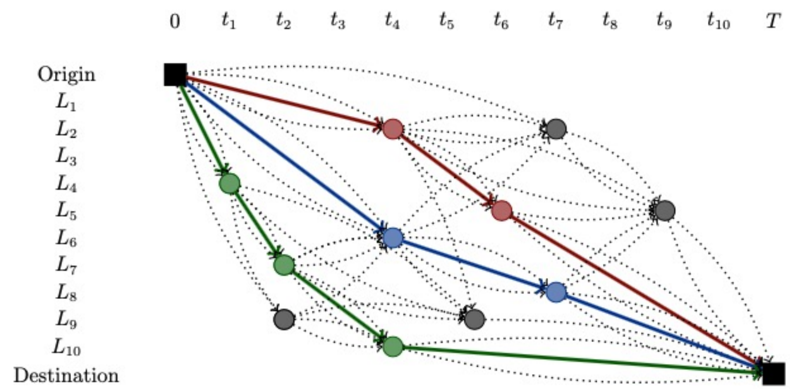

Offline dial-a-ride: subpath model

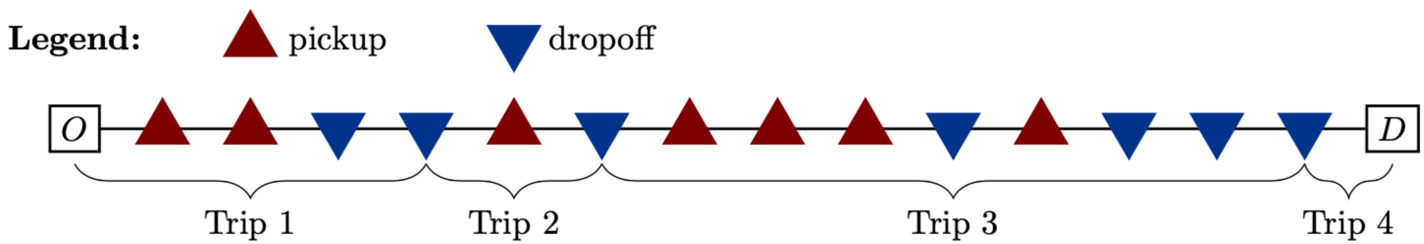

Path optimization [MSO-VCC] → **Subpath generation [dynamic program]** → **Route optimization [time-space network]**

- Node: "empty point" for vehicles
- Arc: customer-carrying trip (subpath)

→ Combination of set partitioning principles and time-space principles

$$
\begin{aligned}
\max \; & \sum_{u \in \mathcal{U}} g^u z^u, \\
\text{s.t.} \; & \sum_{u \in \mathcal{U}_s^+} z^u - \sum_{u \in \mathcal{U}_s^-} z^u = \begin{cases} 0 & \forall s \neq s^O, s^D \\ \Lambda & \text{for } s = s^O \\ -\Lambda & \text{for } s = s^D \end{cases} \\
& \sum_{u \in \mathcal{U}} a^{uj} z^u \leq 1, \; \forall j \in \mathcal{P}, \\
& 0 \leq \Lambda \leq |\mathcal{V}|, \\
& z^u \in \{0,1\}, \; \forall u \in \mathcal{U}.
\end{aligned}
$$

Offline dial-a-ride: sample results

	CPLEX		Our algorithm	
# customers	Solution	CPU (s)	Solution	CPU (s)
5	18.3	66	18.3	<1
7	10.7	>3,600	16.9	<1
8	9.6	>3,600	24.9	<1
10	2.2	>3,600	17.7	<1
15	11.3	>3,600	42.5	<1
20	7.0	>3,600	41.3	<1
25	4.6	>3,600	46.4	<1
50	-	>3,600	63.2	340
75	-	>3,600	62.6	246
100	-	>3,600	57.8	180
125	-	>3,600	49.5	126
150	-	>3,600	57.6	390
175	-	>3,600	56.9	678
200	-	>3,600	63.1	2,843

Small-scale instances:
much faster than CPLEX

Medium-scale instances:
better and faster than CPLEX

Large-scale instances:
optimal solution
no solution with CPLEX

Offline dial-a-ride: optimization edge

- Small downstream benefits from vehicle-customer coordination
- Instead, benefits stem from re-optimization of "upstream" routing operations in view of added operating flexibility

	No VCC		Fix Seq. + VCC		Fix Cust. + VCC		Our algorithm	
$\lvert\mathcal{V}\rvert$	Profit	% change	Profit	% change	Profit	% change	Profit	% change
1	17.4	(base)	17.5	(+0.4%)	17.5	(+0.6%)	23.4	(+34.4%)
3	50.9	(base)	51.1	(+0.5%)	51.1	(+0.6%)	67.5	(+32.7%)
5	82.0	(base)	82.4	(+0.5%)	82.4	(+0.5%)	105.8	(+29.0%)
10	149.3	(base)	150.0	(+0.4%)	150.1	(+0.5%)	188.0	(+25.9%)
15	201.6	(base)	201.8	(+0.1%)	204.8	(+1.6%)	260.0	(+28.9%)
20	247.6	(base)	248.0	(+0.2%)	253.3	(+2.3%)	321.3	(+29.8%)
25	284.0	(base)	284.7	(+0.3%)	292.7	(+3.1%)	369.2	(+30.0%)
30	308.8	(base)	310.0	(+0.4%)	321.0	(+4.0%)	401.7	(+30.1%)

Online optimization

Online dial-a-ride: algorithm

- Batching and optimization
  - Very large scale: 10,000+ requests/hour
  - Dynamic customer arrivals

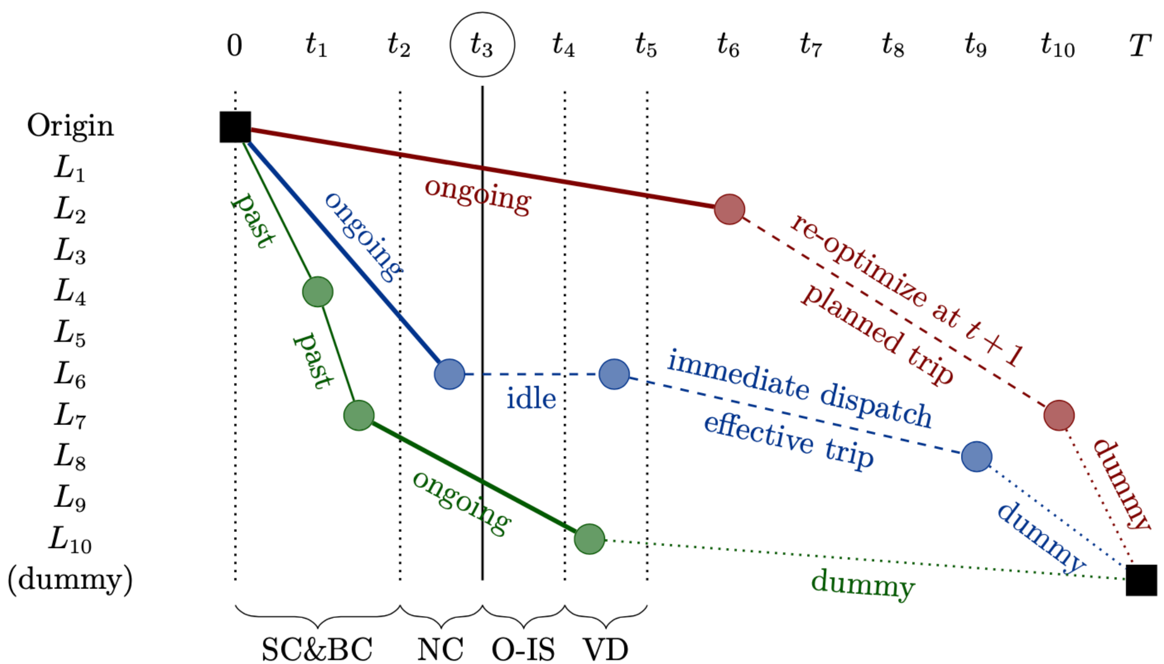

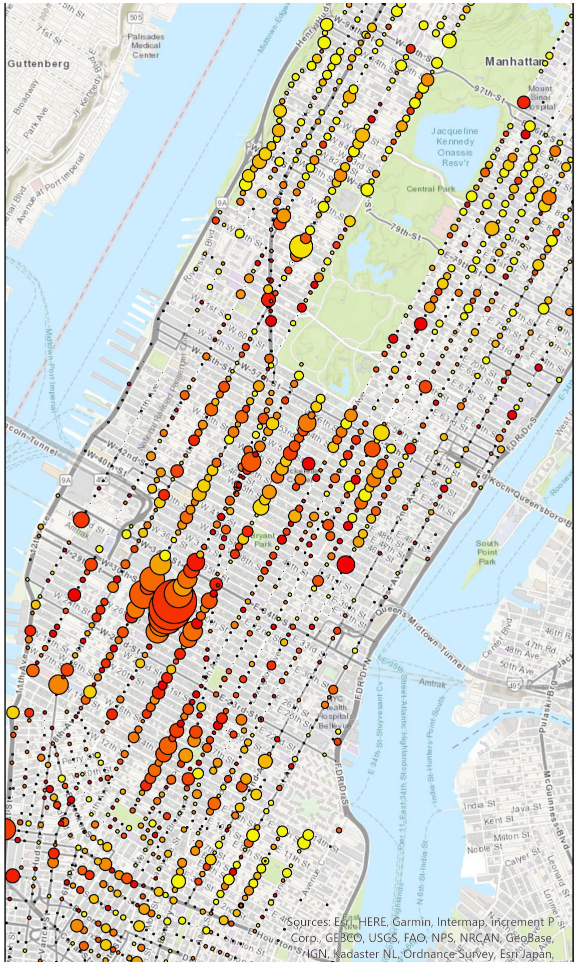

Online dial-a-ride: sample results

	No ride-pooling		Ride-pooling			
Metric	No VCC	VCC:PU	No VCC	VCC:PU	VCC:DO	VCC:PU–DO
Profit	$137,888	$144,378	$144,662	$152,697	$144,757	$153,714
Profit increase	(base)	4.8%	(base)	5.6%	0.0%	6.2%
Revenue increase	(base)	4.4%	(base)	4.9%	0.0%	5.4%
Cost decrease	(base)	3.6%	(base)	13.1%	2.0%	15.4%
Downstream contribution	(base)	10.2%	(base)	13.7%	1.2%	14.2%
Upstream contribution	(base)	89.8%	(base)	86.3%	98.8%	85.8%
VCC (pickups)	—	9%	—	13%	—	13%
VCC (dropoffs)	—	—	—	—	3%	4%
Pooled trips	—	—	18%	23%	19%	26%
Acceptance rate	75%	80%	79%	84%	79%	85%
Max discount	(base)	24.1%	(base)	18.0%	−0.2%	16.0%
VMT	56,803	54,471	46,705	40,619	45,720	39,612

Win-win-win outcomes with profit sharing: profitability (+3%), customer service (+6%), and environmental footprint (-15%)

Summary

Models and algorithms to optimize routing operations with vehicle-customer coordination in large-scale networks

Single stop optimization (SSO-VCC)	Core subproblem: Optimization of location and time of a single stop, using geometric insights
Multi-stop optimization (MSO-VCC)	Optimization of locations and times of a sequence of stops, using an exact coordinate descent algorithm
Vehicle routing (VRP-VCC) Dial-a-ride (DAR-VCC)	Integration of MPO-VCC into routing optimization: subpath-based variables & time-space networks
Large-scale online routing optimization	Rolling algorithm for real-world ride-sharing: scalability and win-win-win outcomes

Thank you