The Impact of Autonomous Trucking

A Case-Study of Ryder's Dedicated Transportation Network

Ryder System, Inc. and the Socially Aware Mobility Lab (Georgia Tech)

1 Introduction

Autonomous trucks are expected to fundamentally transform the freight transportation industry, and the enabling technology is progressing rapidly. Morgan Stanley estimates the potential savings from automation at $168 billion annually for the US alone.[1] Additionally, autonomous transportation may improve on-road safety, and reduce emissions and traffic congestion.[2]

SAE International defines different levels of driving automation, ranging from L0 to L5, corresponding to no-driving automation to full-driving automation.[3] The current focus is on L4 technology (high automation), which aims at delivering automated trucks that can drive without any need for human intervention in specific domains, e.g., on highways. This level of automation fundamentally changes the role of drivers in the logistics system: It allows drivers to perform other tasks during long stretches of autonomous driving, or to take rests to comply with hours-of-service regulations, increasing productivity.[4] The driver may even be completely absent from the cab, resulting in substantial labor cost savings.

The trucking industry is actively involved in making L4 vehicles a reality. Daimler Trucks, one of the leading heavy-duty truck manufacturers in North America, is working with both Torc Robotics and Waymo, and will be testing the latest generation of L4 trucks in the Southwest in early 2021.[5] In 2020, truck and engine maker Navistar announced a strategic partnership with technology company TuSimple to develop L4 trucks, to go into production by 2024.[6] Truck manufacturers Volvo and Paccar have both announced partnerships with Aurora.[7] Other companies developing self-driving vehicles include Embark, Gatik, Kodiak, and Plus.[8]

Various companies are showing interest in autonomous transportation for their logis-

[1] Greene 2013.
[2] Short and Murray 2016; Slowik and Sharpe 2018.
[3] SAE International 2018.
[4] Short and Murray 2016.
[5] Engadget 2020.
[6] Transport Topics 2020.
[7] TechCrunch 2021.
[8] FleetOwner 2021; Forbes 2021; FreightWaves 2021.

tics and supply chain operations. This has resulted in collaborations with technology firms to develop L4 technology and test it through pilots.[9] These companies point out that driverless trucks allow for higher asset utilization, as there are no hours-of-service restrictions. Furthermore, automation would address the driver shortage, which is only expected to grow in the coming years.[10] Enhanced safety and reduced emissions are also cited as advantages.

The return on investment of L4 technologies will likely be an important driver for its adoption[11], and quantifying the impact of automation on operations is crucial to making the business case. This study contributes to the discussion by presenting an Autonomous Transfer Hub Network (ATHN) for freight transportation that combines autonomous trucks on highways with regular trucking operations for the first and last miles, as well as introducing optimization models to optimize the routing and dispatching of autonomous trucks. The proposed network is evaluated on an actual case study, Ryder's dedicated transportation business in the Southeast of the US, and compared with existing operations under various assumptions. The analyses indicate that the ATHN, in conjunction with optimization technology, may reduce costs in the range of 27% to 38% for a small ATHN and 29% to 40% for a more extensive network.

2 Autonomous Tranfer-Hub Networks

A study by Viscelli describes different scenarios for the adoption of autonomous trucks by the industry.[12] The most likely scenario, according to some of the major players, is the *transfer hub business model*.[13] Transfer Hub Networks (THN) make use of autonomous truck ports, or *transfer hubs*, to hand off trailers between human-driven trucks and driverless autonomous trucks (Figure 1). Autonomous trucks then carry out the transportation between the hubs, while conventional trucks serve the first and last miles.

Figure 2 presents an example of an autonomous network with transfer hubs. Orders

[9] National Public Radio 2019; Government Technology 2020; The Wall Street Journal 2020.
[10] Short and Murray 2016.
[11] Slowik and Sharpe 2018.
[12] Viscelli 2018.
[13] Roland Berger 2018; Viscelli 2018; Shahandasht, Pudasaini, and McCauley 2019.

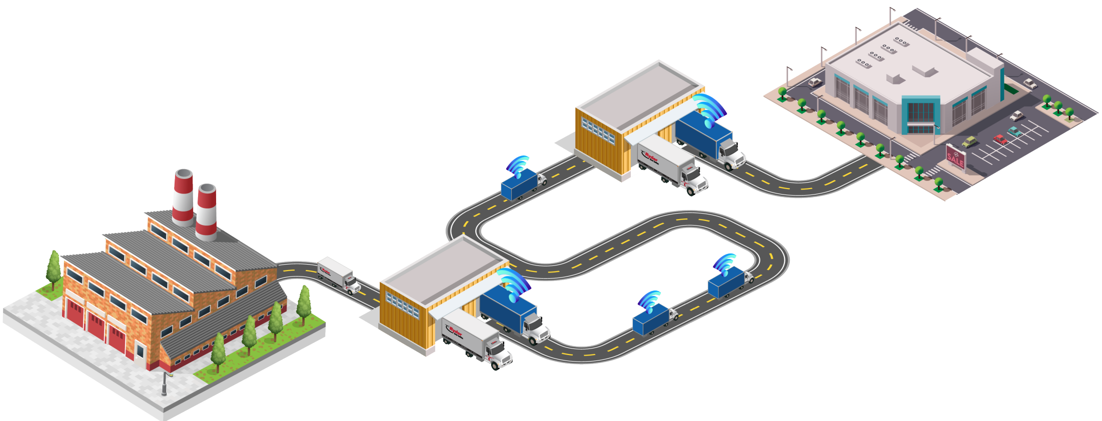

Figure 1: An Example of the Transfer Hub Business Model with Two Hubs.

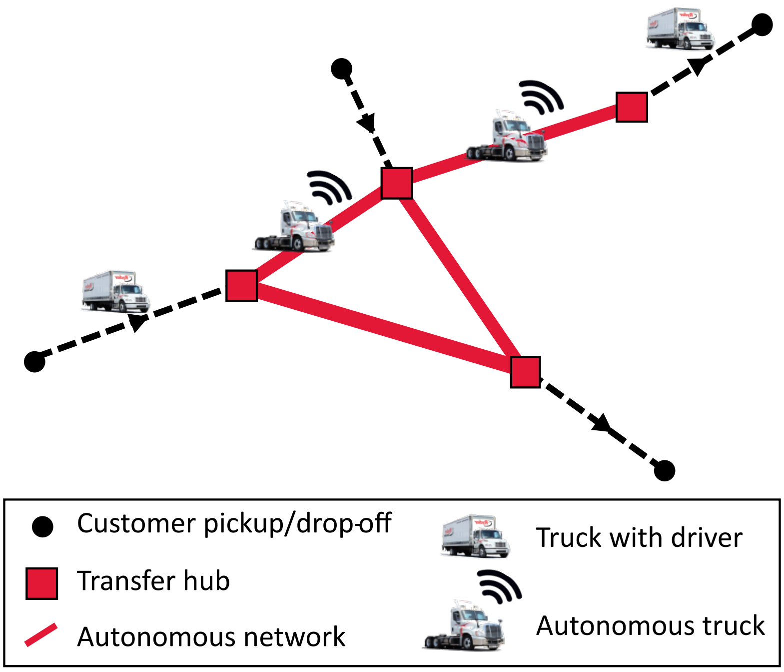

Figure 2: An Example of an Autonomous Transfer Hub Network.

are split into a first-mile leg, an autonomous leg, and a last-mile leg, each of which served by a different vehicle. A human-driven truck picks up the cargo at the customer location, and drops it off at the nearest transfer hub. A driverless autonomous truck moves the trailer to the transfer hub closest to the destination, and another human-driven truck performs the last leg.

ATHNs apply automation where it counts: monotonous highway driving is automated, while more complex local driving and customer contact is left to humans. Especially for long-haul transportation, the benefit of automation is expected to be high. Global consultancy firm Roland Berger estimates that operational cost savings may be between 22% and 40% in the transfer hub model, based on the cost difference between driverless trucks and conventional trucks.[14] These estimates are based on single trips and it is not clear that they can be realized in practice: in particular, they do not take into account the empty miles traveled by autonomous trucks to pick up their next orders. The impact on fleet productivity, flexibility, and driver retention are also key factors, which may be even more significant than labor cost differences.[15]

The flexibility offered by autonomy is especially important for dedicated carriers, for which empty return trips are often unavoidable. In an autonomous network, trucks can immediately start serving the next order, without returning back empty. Furthermore, autonomous trucks do not need to return to a domicile at the end of the day and can keep operating 24/7. This provides the opportunity to increase asset utilization by negotiating appointment times with the customers. In dedicated transportation, orders are typically placed multiple days in advance, which gives enough time to do so. An additional benefit in dedicated transportation is that training becomes more effective. Drivers are commonly trained to deal with specific customers, and automation allows them to spend less time on the highway, and more time putting this knowledge to use.

It is important to observe that, besides autonomy, the transfer hub model substantially differs from a traditional freight transportation network. For dedicated carriers, it is common that an order is carried out by a single driver and a single truck. The truck may even be dedicated to a specific customer. With the transfer hub model, every vehicle still transports at most one order at a time, but there is no one-to-one matching between a customer and a truck/driver. Drivers would exclusively serve first/last-mile legs for multiple orders on the same day. For example, a driver may bring freight from a transfer hub to a customer (last mile), drive empty to the next customer to pick up a new load (relocation), and then deliver to the hub (first mile), and so on. As a result, ATHNs will reshape the traditional business model in significant ways.

This main contribution of this paper is to provide, for the first time, a detailed quan-

[14] Roland Berger 2018.
[15] Greene 2013.

titative study of the benefits of ATHNs by considering a real case study where actual operations are modeled and optimized with high fidelity. It examines whether the predicted savings materialize when the network effects, e.g., empty miles for relocation, are taken into account.

3 Problem Description and Methodology

To quantify the impact of autonomous trucking, this paper introduces a data-driven methodology to design and operate an ATHN. Section 5 applies the methodology to the case study of Ryder's dedicated transportation business, and compares current operations and those of the ATHN. This section first reviews the input data before addressing two challenges:

- *designing the network*: selecting transfer hub locations and roads for autonomous travel (tactical decisions), and
- *operating the system*: planning the exact movements of every load, vehicle, and driver to carry out the orders (operational decisions).

Efficiently operating the system requires an optimization model, which is discussed in more detail in Section 4.

3.1 Input Data

The first input is the road system data which was obtained from OpenStreetMap.[16] Moreover, route distance and mileage were calculated with the GraphHopper library. The second input is a list of customer orders which were formatted as in Table 1. Every order has a unique *OrderNumber*, and every row corresponds to a stop for a particular order. Stops have a unique identifier *StopNumber*, and the *Stop* column indicates the sequence within the order. The columns *StopArrivalDate* and *StopDepartureDate* indicate the scheduled arrival and departure times, and *City* and *ZipCode* identify the location of the

[16] OpenStreetMap 2020.

StopNumber	OrderNumber	StopArrivalDate	StopDepartureDate	Stop	City	ZipCode	Status	Event
68315760	7366366	2-10-2019 09:01	2-10-2019 09:02	1	Atlanta	30303	LD	HPL
68315761	7366366	2-10-2019 16:29	2-10-2019 18:33	2	Tennessee	37774	LD	LUL
68315762	7366366	3-10-2019 11:00	3-10-2019 11:30	3	Atlanta	30303	MT	DMT
...	...	...	...	...	...	...	...	...
46798427	5207334	7-10-2019 02:35	7-10-2019 02:50	1	Alpharetta	30009	NaN	LLD
46798428	5207334	7-10-2019 08:10	7-10-2019 08:49	2	Macon	31201	LD	LUL
46798429	5207334	7-10-2019 15:16	7-10-2019 15:45	3	Alpharetta	30009	LD	LUL

Table 1: An Example of the Order Data.

stop. The *Status* column gives a code for the status of the vehicle on arrival, and the *Event* column indicates what happens at the stop.

The example data in Table 1 displays two orders. The first order is a trip from Atlanta to Tennessee and back. Based on the status code, the truck arrived in Tennessee loaded (LD) and returned to Atlanta empty (MT). The event codes show that a preloaded trailer was hooked in Atlanta (HPL), followed by a live unload (LUL) in Tennessee, after which the truck dropped the empty trailer (DMT) in Atlanta. The exact codes are not important for the purpose of this paper. What is important, is the ability to derive the parts of the trip when the truck is moving freight, and the parts when the truck is driving empty.

3.2 Designing the Network

The design of ATHNs needs to decide the locations of the transfer hubs, which are the gateways to the autonomous parts of the network. A natural choice is to locate the hubs in areas where many trucks currently enter or exit the highway system. Historical order data is used to identify these common highway access points. For a given order, the truck is routed through the existing road network, and the highway segments, and their access points, can easily be identified, as shown by Figure 3. The transfer hubs are then placed in areas with many access points. The case study will consider two different sets of hubs: a small network where only the cities with the most activity are selected, and a large network where less important locations are also included.

Once the hubs are selected, the road network that links them needs to be identified.

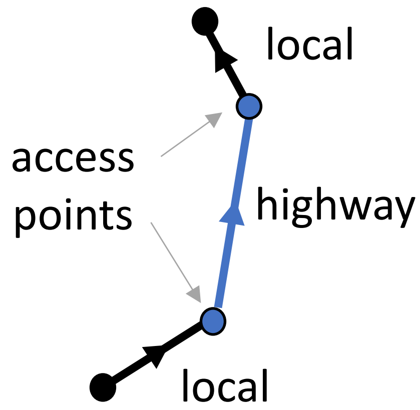

Figure 3: Identifying Highway Access Points.

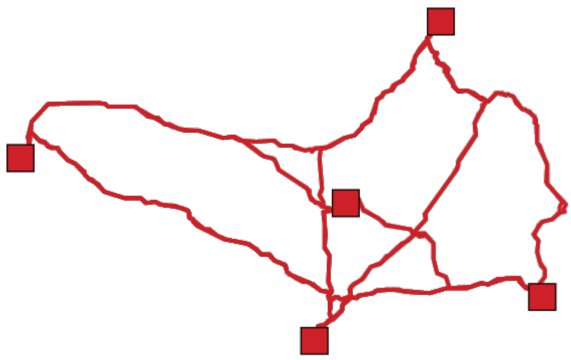

Figure 4: An Example of an Autonomous Network.

There is a link between two hubs if an autonomous vehicle can travel between them directly. Figure 4 depicts such a network: in this example, all the transfer hubs are connected by L4-supported roads, and the shortest route between two hubs may pass by another hub.

3.3 Operating the System

This section discusses how to produce operational plans for an ATHN. The operational plan considers orders within a given time horizon (e.g., one week ahead) and details the movements of every load, vehicle, and driver.

3.3.1 Order Selection

The first step in designing the operational plan consists of identifying the orders that may benefit from the autonomous network, and those that are better served by human-driven trucks. For example, a direct trip with a human-driven truck may be preferred

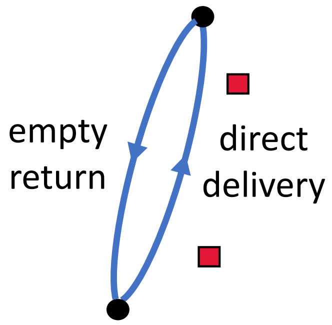

(a) Direct Delivery.

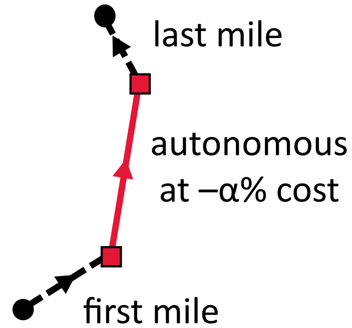

(b) Delivery through the ATHN.

Figure 5: Comparison Between Direct Delivery and Delivery through the ATHN.

for trips that are short, or for trips between locations that are far from any hub. This paper focuses on orders using autonomous vehicles; the remaining ones can be planned separately. The relevant orders are obtained by comparing two options, as visualized by Figure 5. The first option serves the order with a human-driven truck (Figure 5a). The case study focuses on orders that necessitate an empty return leg. Hence, the total mileage for the overall trip represents the cost of this option. The second option uses the autonomous network (Figure 5b), which amounts to driving to the nearest hub with a human-driven truck (first mile), shipping the load over the autonomous network to the hub closest to the destination, and then serving the last mile with another human-driven truck. Again, the total mileage is used to represent the cost, but the mileage over the autonomous leg is reduced by a fraction α to account for the benefits of autonomy. The value of α is a parameter, and the case study uses values ranging from $\alpha = 25\%$ to $\alpha = 40\%$.

For every order, the two options are compared and the cheapest one is selected. The operational plan to be presented next only considers the orders using the autonomous network.

3.3.2 Operational Plan

To build an operational plan for the selected orders, the idea is to split their routes into three legs: the first-mile leg, the autonomous leg, and the last-mile leg. The optimization model in Section 4 creates a schedule for the first/last-mile operations of each of

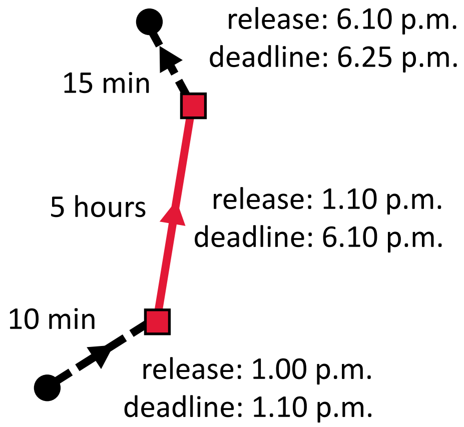

Figure 6: Example Release Times and Deadlines for an Order Picked Up at 1 p.m.

the transfer hubs, and a schedule for the autonomous operations between the hubs. It is important to emphasize that these schedules must be synchronized, such that freight arrives in time to be picked up by the next vehicle. Every leg is assigned a release time (earliest pickup time) and a deadline (latest delivery time), where the release time for the next leg is later than the deadline for the previous leg. In this case study, the original appointment time is used as the first-leg release time and the subsequent deadlines and release times are set to the earliest arrival time. Figure 6 demonstrates this synchronization with an example. After setting time windows for each of the legs, the first/last-mile operations and the autonomous operations are scheduled independently. The next section introduces an optimization model to create these schedules. Together, the schedules make up a complete operational plan that details the movements of every load, vehicle, and driver to serve all the orders.

The timing of the orders is a key factor for efficient asset utilization. For the case study, the timing is based on the original appointments. In the future, these appointment times may be negotiated with the customers to allow for a more efficient operation. To take this aspect into account, the model introduces *appointment flexibility*, which assumes that customers are willing to shift their appointments by a certain amount of time.

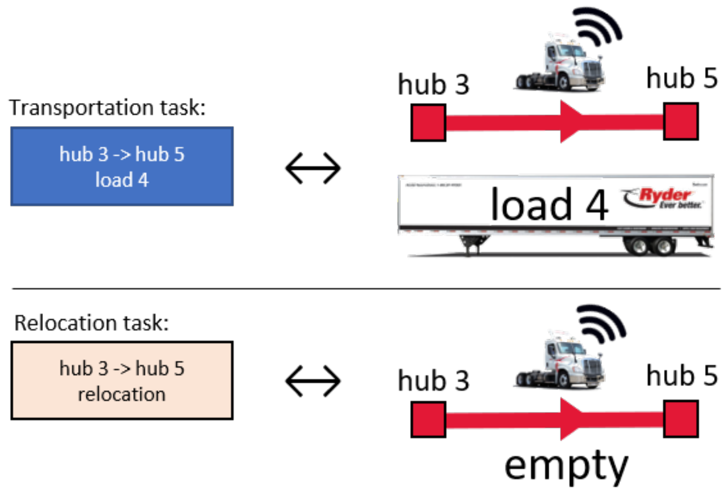

Figure 7: An Example of Transportation and Relocation Tasks.

4 Optimization Model

This section presents an optimization model to create a schedule to serve a set of legs with a fleet of vehicles. The same model is applied both to scheduling the autonomous trucks, and to scheduling the first/last-mile operations for the regular trucks at each of the hubs. Appendix A provides a mathematical description of the general model. For simplicity, this section presents the model in the context of autonomous operations between the transfer hubs, and mentions at the end how this translates to the first/last-mile operations.

The two most important inputs to the model are the *autonomous network*, and a list of *transportation tasks*. The autonomous network consists of a list of hubs and the L4-supported roads between them. A transportation task corresponds to serving a single leg on the autonomous network. It consists of an origin hub and a destination hub, a release time for pickup at the origin and a deadline for delivery at the destination, and an id number that indicates which load is transported. Figure 7 shows an example of a transportation task. In this case, the task *hub 3* $\rightarrow$ *hub 5, load 4* indicates that the autonomous truck picks up the load with id number 4 at hub 3, drives to hub 5, and drops off the load. In subsequent examples, the width of the task box indicates the duration of the task, which consists of loading, driving, and unloading. The loading and unloading times are parameters to the model.

In addition to transportation tasks, the model includes *relocation tasks* that corre-

Figure 8: An Example Schedule with Five Loads and Two Trucks.

spond to autonomous trucks moving empty from one location to another. Figure 7 gives an example for the task *hub 3* $\rightarrow$ *hub 5, relocation* for which the truck drives from hub 3 to hub 5 without any freight. Relocation tasks are necessary to reposition trucks when no transportation tasks are available at the current location. Similar to transportation tasks, relocation tasks are associated with a duration and a distance.

The autonomous network operations can be modeled as a *schedule* of transportation and relocation tasks. Figure 8 presents an example with five loads and two trucks. In the schedule, every row corresponds to a truck, and the horizontal axis represents time. Tasks are assigned to vehicles and are given a starting time, reflected by their positions. In the example, *Truck 1* starts at hub 3, where it is waiting for its first task. After some time, the truck picks up load 4 and brings it to hub 5. There it waits again, before driving empty to hub 1 to pick up load 5, and transport it to hub 6. After a short amount of time, a new task becomes available at its current location, and the truck moves load 3 from hub 6 to hub 2.

The objective of the optimization model is to create a feasible schedule that serves all the transportation tasks and minimizes the total relocation distance. The distance driven for the transportation tasks is the same in every schedule, and therefore does not need to be considered in the optimization.

Before stating the complete model, two additional parameters are introduced: K for the number of vehicles, and Δ for the *appointment flexibility* in minutes. If $\Delta = 60$ for example, then loads may be picked up one hour before the release time and may be delivered one hour after the deadline. This is visualized by Figure 9, where a_t is the release time of transportation task t, and b_t is the deadline. For $\Delta = 0$, only Example 2 would

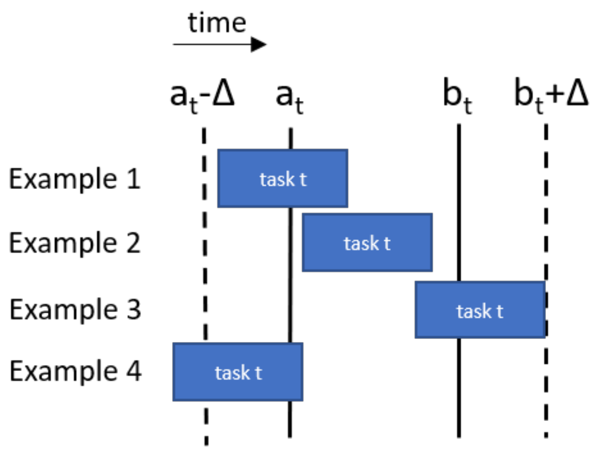

Figure 9: Visualization of the Flexibility Constraints for Transportation Task t.

be allowed, as the pickup is after the release time and the dropoff is before the deadline. With the additional flexibility, Example 1 and 3 are also feasible, while Example 4 is still not allowed. The amount of flexibility depends on the willingness of customers to move their appointments, and the case study investigates how flexibility affects efficiency.

Figure 10 now presents the complete model, which essentially describes what a schedule is supposed to look like. It is worth mentioning that every transportation task is assigned exactly once (Constraint 1), while the relocation tasks may appear multiple times, even for the same truck. Constraints 4 and 5 only affect transportation tasks, as there are no restrictions to when relocations are performed.

The model in Figure 10 is solved by *Constraint Programming* (CP), which is an optimization technique that is known for efficiently finding high quality solutions to scheduling problems and other problems. After translating the model to a modeling language, the *CPLEX CP Optimizer* version 12.8 by IBM is used to solve the problem.[17]

While the model has been presented for the autonomous operations, the same idea can be applied to the first/last-mile operations. Instead of moving between transfer hubs, the trucks move between a single hub and customer locations, but everything else remains the same. One remark, though, is that autonomous trucks do not have idling costs, while this may be different for human drivers. To include waiting costs, the model may be extended by including *waiting tasks*, but this extension is not considered in this paper.

[17] Van Hentenryck 1999; Laborie et al. 2018.

- Objective: Minimize the total relocation distance.
- Constraints:
  1. Every transportation task is assigned to exactly one of K trucks.
  2. Tasks cannot overlap in time.
  3. The origin of the next task is equal to the destination of the previous task.
  4. The start time of a transportation task is at or later than the release time subtracted by Δ.
  5. The end time of a transportation task is at or before the deadline plus Δ.
- Decisions:
  1. Which tasks are assigned to each vehicle?
  2. What is the starting time for each assigned task?

Figure 10: Optimization Model for Operating the Autonomous Network.

5 Case Study

To quantify the impact of autonomous trucking on a real transportation network, a case study is presented for the dedicated transportation business of Ryder System, Inc., commonly referred to as *Ryder*. Ryder is one of the largest transportation and logistics companies in North America, and provides fleet management, supply chain, and dedicated transportation services to over 50,000 customers. Its dedicated business, *Ryder Dedicated Transportation Solutions*, offers supply-chain solutions in which Ryder provides both drivers and trucks, and handles all other aspects of managing the fleet. Ryder's order data is used to design an ATHN, and to create a detailed plan for how it would operate. This allows for a realistic evaluation of the benefits of autonomous trucking.

5.1 Data and Assumptions

Ryder prepared a representative dataset for its dedicated transportation business in the Southeast of the US, reducing the scope to orders that were strong candidates for automation. The dataset consists of trips that start in the first week of October 2019, and stay completely within the following states: AL, FL, GA, MS, NC, SC, and TN. It contains

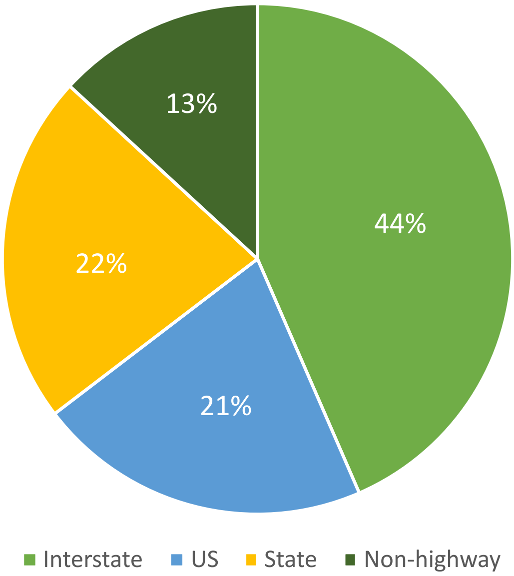

Figure 11: Mileage per Road Type.

11,264 rows, which corresponds to 2,090 orders, formatted as in Table 1. Every order has a unique *OrderNumber*, and every row corresponds to a stop for a particular order. Stops have a unique identifier *StopNumber*, and the *Stop* column indicates the sequence within the order. The columns *StopArrivalDate* and *StopDepartureDate* indicate the scheduled arrival and departure times, and *City* and *ZipCode* identify the location of the stop. The *Status* column gives a code for the status of the vehicle on arrival, and the *Event* column indicates what happens at the stop.

The provided orders are long-haul trips, with an average trip length of 431 miles. Most of this distance is driven on highways. Figure 11 breaks down the mileage per road type, and shows that 65% of the distance is driven on interstates and US highways. This number goes up further to 87% if state highways are included. Significant highway usage is typical for long-haul transportation, and indicates that a significant part of each trip can potentially be automated.

5.2 Network Design

Ryder's order data is used to design two different autonomous transfer hub networks. Figure 12a presents the *small network* that was created by locating 17 transfer hubs in the areas where Ryder trucks are most frequently accessing the highway system. The *large*

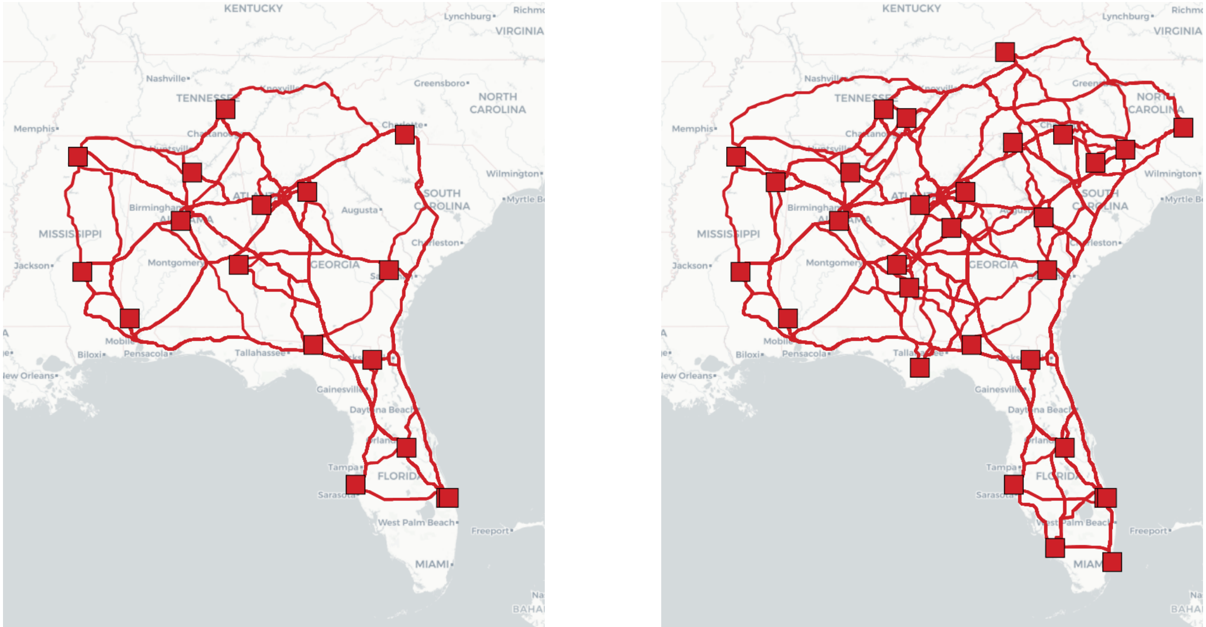

(a) Small Network (17 hubs).

(b) Large Network (30 hubs).

Figure 12: Autonomous Network Designs for the Southeast.

network in Figure 12b was designed by adding 13 additional hubs to the small network, in locations with fewer highway access points. The exact hub locations are masked, but the figures are accurate within a 50 mile range. The additional transfer hubs extend the network further northeast into North-Carolina, and further south into Florida. It also makes the network more dense in the center of the region.

5.3 Parameters and Settings

The operational plans for the autonomous THN considers the 494 most challenging orders. These orders consist of a single delivery, followed by an empty return trip. Because they require empty travel for 50% of the trip, these challenging orders are an excellent target for cost savings. They also make up 24% of the dataset, and account for 53% of the empty mileage.

To prevent empty mileage, Ryder actively searches among internal customer demand and on the spot markets for backhauls: orders that can be completed on the return trip. The provided dataset includes these backhauls, which means that the challenging orders are orders for which driving empty could not be avoided, mainly due to the asymmetry in the demand. The case study focuses on optimizing these trips that are otherwise difficult to deal with. Depending on the revenue from the backhauls, it may be possible to obtain additional savings by scheduling orders on the ATHN, instead of performing a backhaul, but this is not considered in this paper.

The following scenario is defined as the *base case* for the upcoming experiments. The base case uses the small network, presented by Figure 12a, and driving on autonomous roads is assumed to be $\alpha = 25\%$ cheaper than driving a conventional truck. This results in a conservative estimate of the benefits of autonomous trucking, as driverless transportation is predicted to be 29% to 45% cheaper per mile.[18] The time for loading/unloading or swapping trailers is estimated at 30 minutes, and the appointment time flexibility is set to one hour ($\Delta = 60$). The number of autonomous trucks is set to $K = 50$. To study the effect of, e.g., time flexibility, the parameters are varied and the results are compared to the base case.

5.4 Base Case Results

Out of the 494 challenging orders, 437 (88%) are found to potentially benefit from the autonomous network. These orders are split into separate autonomous and first/last-mile legs, while the direct orders are considered separately. Figure 13 visualizes these legs and the direct trips. It can be seen that many first/last-mile trips are very local, although there are exceptions, especially for orders going to North-Carolina.

The optimization model is used to schedule the autonomous trucks, and the first/last-mile operations at each of the hubs. Scheduling the autonomous trucks is the most challenging: the model has close to 110,000 decision variables and more than 110,000 constraints. The model is given an hour of CPU time for the base case and returns the best found solution, which is visualized in Figure 14, and which only requires 50 autonomous trucks. This figure shows both the transportation tasks and the relocation tasks for the first week of October.

Figure 14 shows that the transportation tasks are close together, and only a relatively

[18] Engholm, Pernestål, and Kristoffersson 2020.

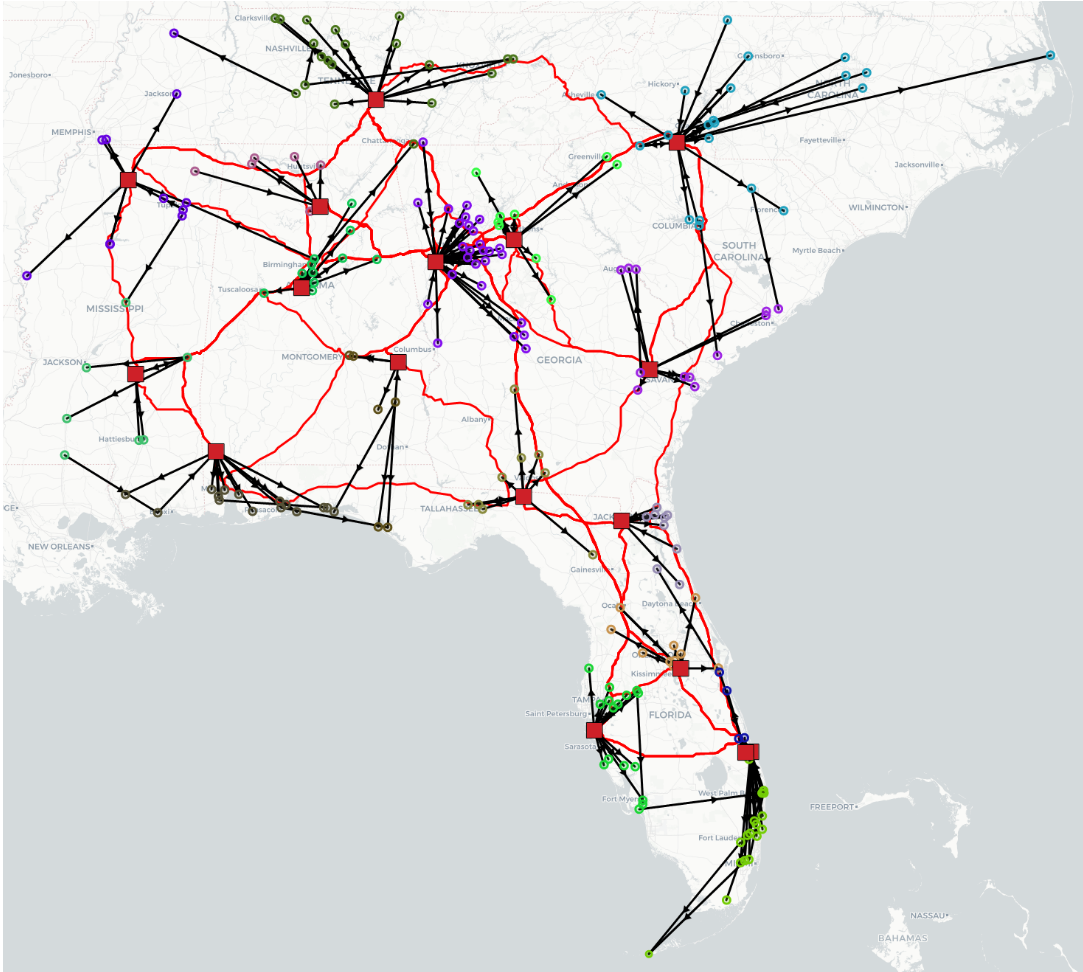

Figure 13: Autonomous Network (red lines) and First/Last-mile Trips (black arrows).

small amount of relocation is necessary. It is interesting to see that only a small number of autonomous trucks is driving during the weekend. This is because the appointment times are still based on the current agreements with the customers, and having drivers work during the weekends is typically avoided. For autonomous trucks, this would not be a problem, which again underlines that making full use of autonomous transportation requires adapting the business model and current practices. In Section 5.7, the importance of time flexibility is considered in more detail.

The routes that are driven by the autonomous trucks consist of serving autonomous legs of different orders, with relocations in between. Figure 15 shows a representative single truck route from the schedule. Blue arrows indicate that freight is being moved, and red arrows indicate that the truck is driving empty. The truck starts at the west coast of Florida, where it picks up a load that has been delivered to the hub by a first/last-

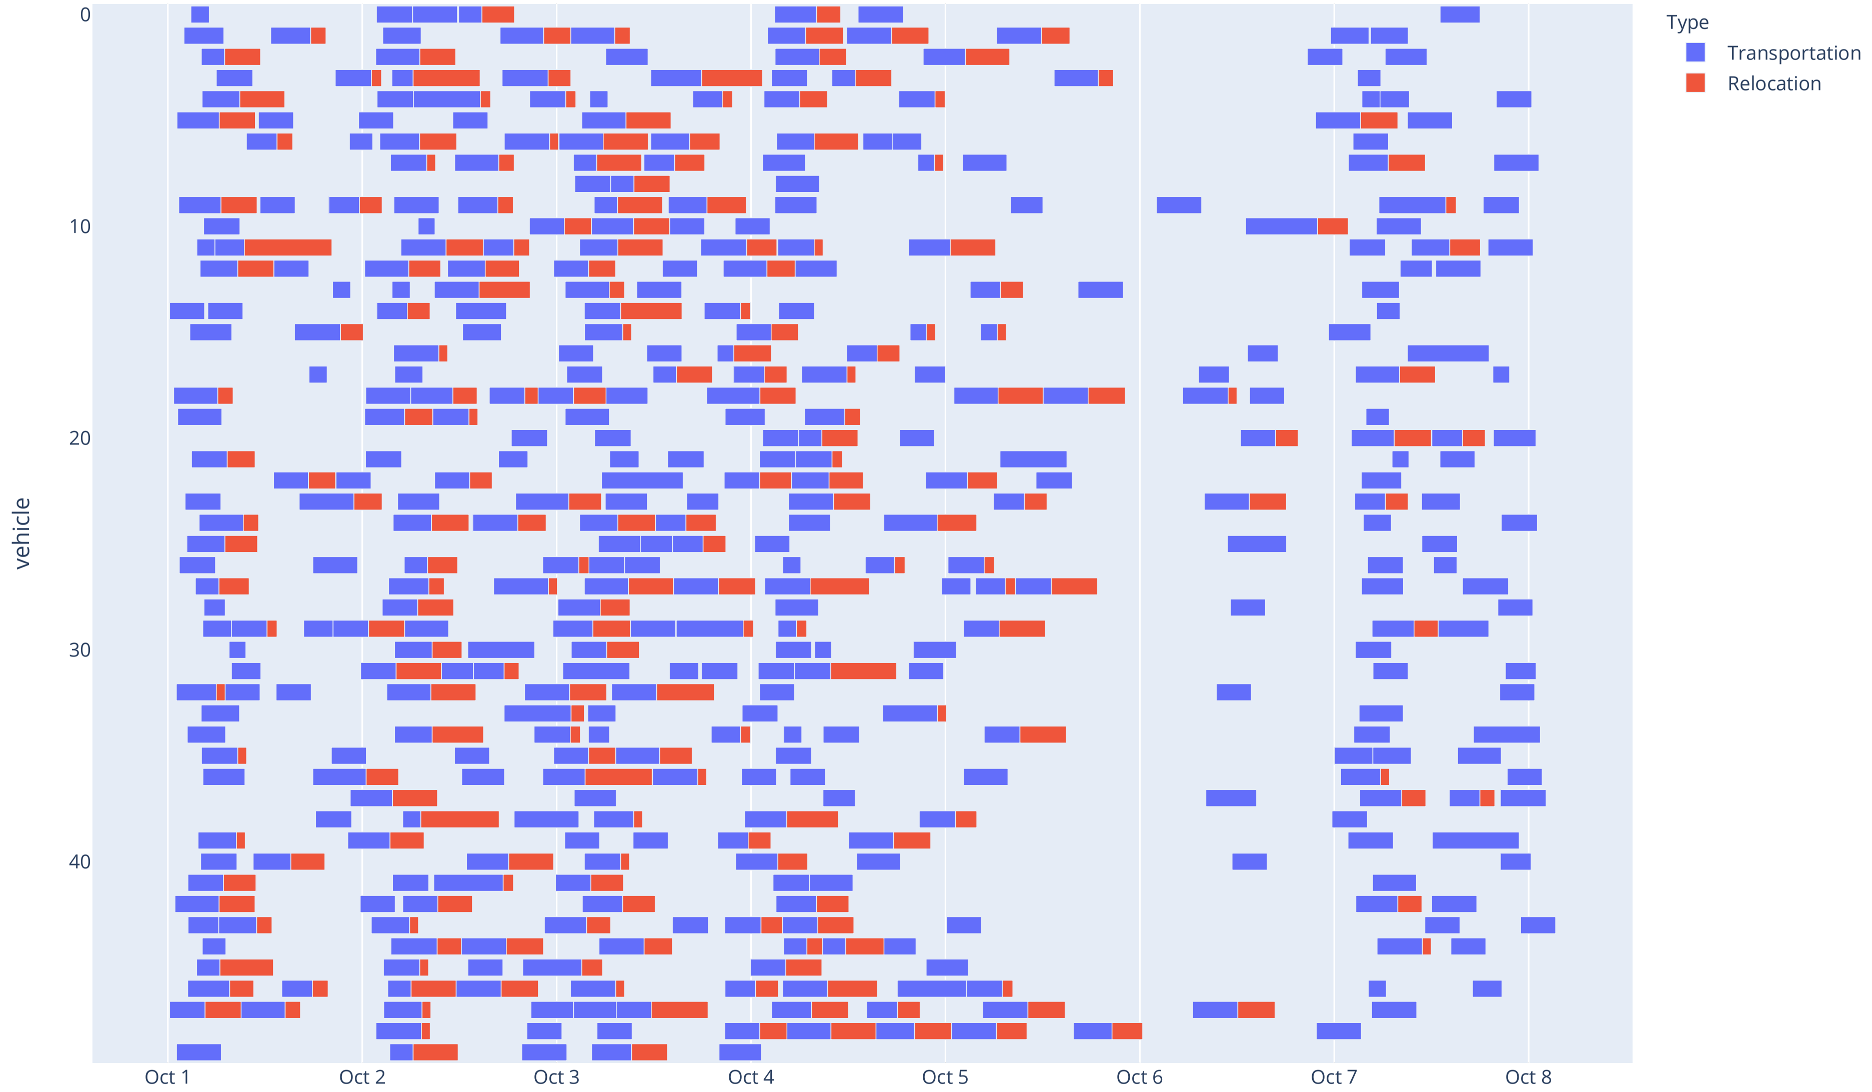

Figure 14: Autonomous Truck Schedule for the Base Case.

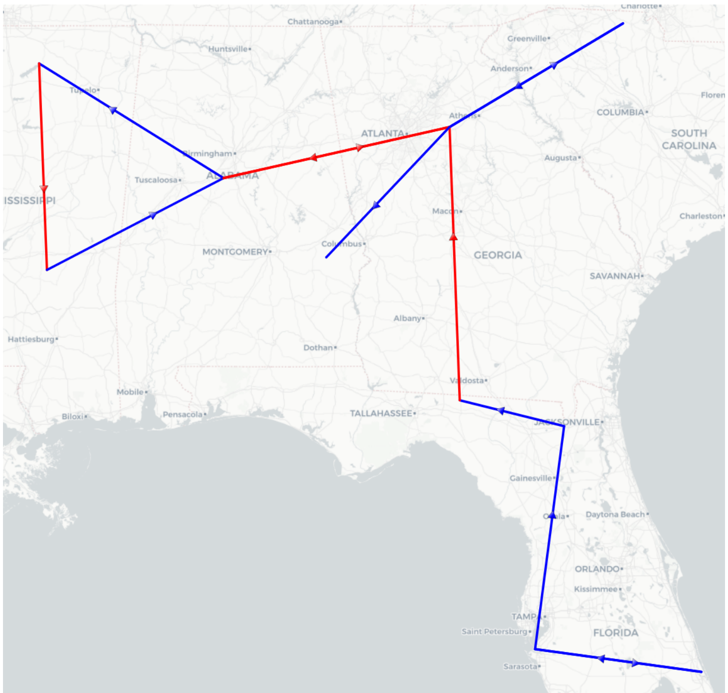

Figure 15: Route of an Autonomous Truck in the Base Case (blue is loaded, red is empty).

mile truck driver. The freight is then transported to the east coast, where it is unloaded so that a regular truck with driver can complete the last-mile. The autonomous truck immediately starts serving the autonomous leg of the next order, returning to the Tampa area. After that, legs are served that are going north. The first time a relocation is needed, is when the truck makes a delivery near Valdosta, close to the Georgia and Florida border. No freight is immediately available, and the vehicle drives empty to the next location to pick up a load there. The routes are clearly complex, which emphasizes the power of optimization: It is unlikely that this solution can be found by manual planners, but it is possible with optimization techniques.

Compared to the autonomous trucks, the total distance driven by regular trucks is relatively short. The optimization model was used to schedule the first and last-mile operations at selected hubs, and it was found that the amount of work is often insufficient to keep drivers occupied throughout the week. To prevent driver idle time, it may be beneficial to outsource these legs, or to consolidate them with other operations in the area. In terms of mileage, the percentage of empty miles at the hubs is typically under 25%, which is used as an estimate for the first/last-mile efficiency in the remainder.

			Mileage	% of total	Cost without auton. trucks	Cost adj.	Cost
Current network		Loaded	96,669	50%	\$ 193,338	1.00	\$ 193,338
		Empty	96,698	50%	\$ 193,396	1.00	\$ 193,396
		Total	193,367	100%	\$ 386,734	1.00	\$ 386,734
Autonomous transfer hub network	Autonomous	Loaded	91,618	67%	\$ 183,235	0.75	\$ 137,426
		Empty	44,217	33%	\$ 88,433	0.75	\$ 66,325
		Total	135,834	100%	\$ 271,668	0.75	\$ 203,751
	First/last mile	Loaded	29,286	75%	\$ 58,573	1.00	\$ 58,573
		Empty*	9,762	25%	\$ 19,524	1.00	\$ 19,524
		Total	39,049	100%	\$ 78,097	1.00	\$ 78,097
	Total		174,883		\$ 349,766		\$ 281,848
	Savings		18,484		\$ 36,969		\$ 104,886
	Savings (%)		10%		10%		27%

* estimated

Table 2: Cost Table for the Base Case (437 orders).

Table 2 quantifies the impact of autonomous trucking on the operating costs for the 437 selected orders. Infrastructure investments are discussed separately in Section 6. The *Mileage* column indicates the total miles driven for both the current network and for the ATHN. The numbers are separated based on whether the distance was driven while loaded or empty, and percentages are shown in the *% of total* column. The *Cost without autonomous trucks* converts the miles into dollars, using \$2 per mile as an approximation for the cost of human-driven trucks. Recall that driving autonomously is assumed to be $\alpha = 25\%$ cheaper, which is reflected by the *Cost adjustment* column. The *Cost* column presents the cost when autonomous trucks are available, and is obtained by multiplying the cost without autonomous trucks by the cost adjustment factor.

Compared to the current network, the ATHN allows for significant savings for the selected orders: Table 2 shows that the total cost goes down by 27%. At a cost of \$2 per mile, this corresponds to \$104,886 per week, or \$5.5M per year. The *Mileage* column shows that almost 80% of the mileage in the ATHN can be automated, which partly explains the large savings. What is very interesting to observe is that the total mileage for the ATHN is actually *less* than the total mileage for the direct trips in the current network. In the transfer hub network, there is no need to return back empty after a delivery, and there is no need to limit working hours or to return to a domicile at the end of the day. As a result, only 33% of the automated distance is driven empty, compared to 50% for the current system. This means that even if autonomous trucks would be as expensive as trucks with drivers, costs would still go down by 10% due to the additional flexibility that automation brings.

The cost calculations above estimate the cost of the first/last-mile operations based on the mileage. However, the operations at the hubs may not be viable as a separate business, and outsourcing or consolidating with local demand may change the economics. For example, the outsourcing contracts could include a fixed fee to compensate for the fact that most trips are short. This paper does not consider the first/last-mile economics in detail, as the case study results show that the autonomous part of the network is by far the most important, but studying the business models for the local trips is an interesting direction for future research.

5.5 Impact of the Size of the Network

A larger autonomous network results in shorter first/last-mile trips, and may have a larger area of coverage. To evaluate the impact of the size of the network, the calculations for the base case are repeated using the large network (Figure 12b) with 30 hubs, instead of the small network (Figure 12a) with 17 hubs. For the large network, 468 of the 494 orders (95%) may benefit from the autonomous network, compared to only 88% for the base case. This immediately implies that there is more potential for savings. It also means a higher utilization of the autonomous trucks, as more legs are served by the same 50 vehicles. Figure 16 shows the new schedule, which is clearly more dense than the schedule for the base case.

Table 3 shows that the relative cost savings for the large network (29%) are similar to those for the small network (27%). This means that the average benefit of automation is similar for both designs, *for the orders that are automated*. However, the large network allows more trips to benefit from automation, which is why the cost savings of $ 116,582 are 11% higher than the savings for the small network ($ 104,886). The average benefit

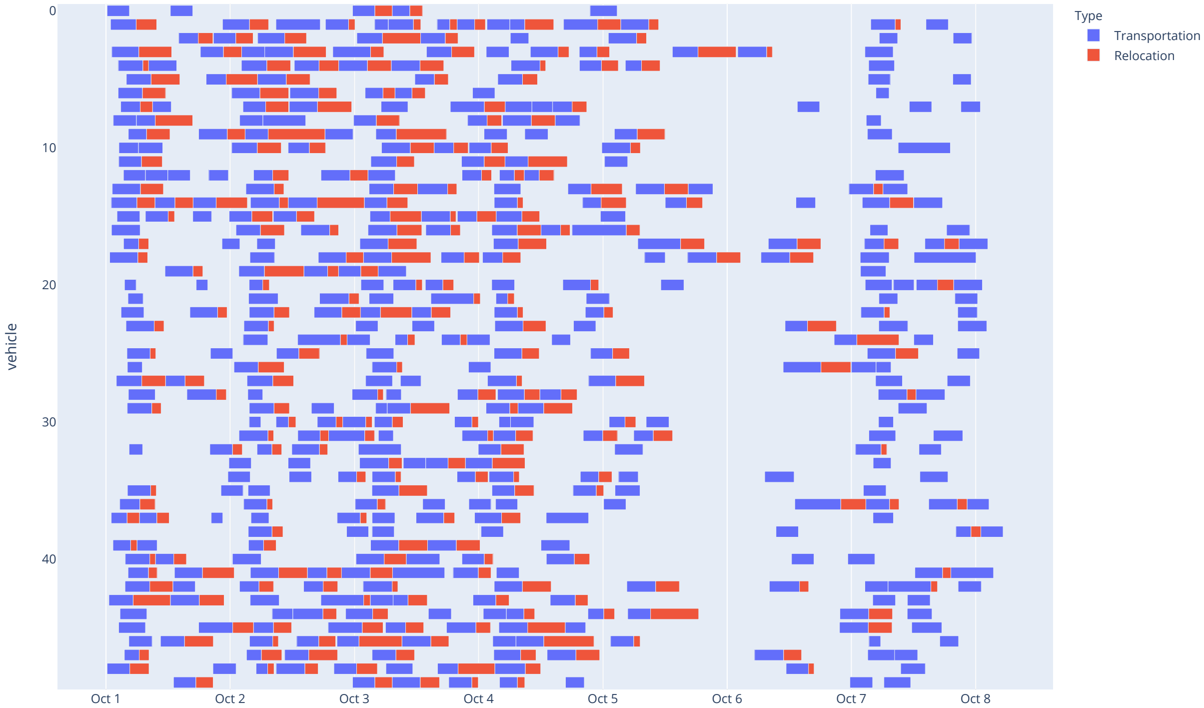

Figure 16: Autonomous Truck Schedule for the Large Network.

of automation is similar for the two designs due to two effects that cancel out. First, the same autonomous trucks have to serve more orders on the large network. This increases the utilization of the vehicles, but also increases the percentage of empty miles from 33% to 35%. The reason for this increase is that there is less time available to wait around at a hub for the next order, as the trucks are needed to perform other orders in the meantime. On the other hand, the first and last-mile trips are shorter due to the additional hubs, which saves costs.

5.6 Impact of the Cost of Autonomous Trucking

For the base case, it was assumed that autonomous trucks are $\alpha = 25\%$ cheaper per mile than trucks with a driver. However, this number is yet far from certain, and higher cost reductions have also been reported in the literature. To investigate the impact of the cost of autonomous trucking, Table 4 presents results for α ranging from 25% to 40%, for both the small and the large network. The column *Autom. orders* gives the number of

			Mileage	% of total	Cost without auton. trucks	Cost adj.	Cost
Current network		Loaded	101,213	50%	\$ 202,425	1.00	\$ 202,425
		Empty	96,698	50%	\$ 193,396	1.00	\$ 193,396
		Total	202,476	100%	\$ 404,953	1.00	\$ 404,953
Autonomous transfer hub network	Autonomous	Loaded	97,326	65%	\$ 194,653	0.75	\$ 145,990
		Empty	53,247	35%	\$ 106,493	0.75	\$ 79,870
		Total	150,573	100%	\$ 301,146	0.75	\$ 225,860
	First/last mile	Loaded	23,442	75%	\$ 46,883	1.00	\$ 46,883
		Empty *	7,814	25%	\$ 15,628	1.00	\$ 15,628
		Total	31,256	100%	\$ 62,511	1.00	\$ 62,511
	Total		181,829		\$ 363,657		\$ 288,371
	Savings		20,648		\$ 41,296		\$ 116,582
	Savings (%)		10%		10%		29%

* estimated

Table 3: Cost Table for the Large Network (468 orders).

orders that may benefit from automation, and are considered in the ATHN. The relative cost savings (*Rel. savings*) state the cost reduction compared to serving these orders with conventional trucks. The *Cost savings* column gives the absolute cost savings in dollars. The final column compares the absolute savings to the savings obtained for the baseline (small network, $\alpha = 25\%$).

Table 4 shows that, as autonomous trucking gets cheaper, and as more hubs are added to the network, the savings compared to the current system go up. Additionally, more orders start using the ATHN, which increases the absolute savings further. Even though the autonomous trucks only perform the transportation between the hubs, the relative cost savings for the complete system often exceed the mileage cost reduction for autonomous trucks (α). This again shows that the benefit of autonomous trucks is not only the lower cost per mile, but also the additional flexibility. Compared to the base case, cheaper autonomous transportation results in significantly larger savings. Similar as in the previous section, increasing the size of the network does not strongly impact the average cost benefit per order, but does increase the total amount of orders that can be automated, which leads to more profits. In the best case (large network, $\alpha = 40\%$), the total benefit of the ATHN is \$ 161,762 per week for the challenging orders, which corresponds to \$8.4M savings per year.

Network	α	Autom. orders	Rel. savings	Cost savings	Additional savings comp. to base case
Small	25%	437	27%	$ 104,886	+0%
	30%	439	32%	$ 122,396	+17%
	35%	443	35%	$ 135,572	+29%
	40%	443	38%	$ 148,984	+42%
Large	25%	468	29%	$ 116,582	+11%
	30%	469	33%	$ 131,798	+26%
	35%	472	37%	$ 149,586	+43%
	40%	472	40%	$ 161,762	+54%

Table 4: Overview of Cost Savings under Different Assumptions.

Three scenarios are considered to estimate the impact of automation on the full dataset. In the pessimistic scenario, automation does not yield any savings for the orders other than the challenging orders. In the middle scenario, the relative savings for the other orders is half that of the challenging orders. And in the optimistic scenario, the same relative savings can be obtained as for the challenging orders. For the base case (small network, $\alpha = 25\%$) this results in 6%, 16%, and 27% savings for the pessimistic, middle, and optimistic scenario, respectively. For the large network with a low cost for autonomous transportation ($\alpha = 40\%$), the estimates range from 9% to 40%, with 24% for the middle scenario.

5.7 Impact of Appointment Flexibility

For the base case, the appointment flexibility was assumed to be $\Delta = 60$ minutes. Deviating from a previously agreed appointment must be negotiated with the customer, but if there are significant benefits in terms of efficiency, this may be worth the effort. To determine the impact of appointment flexibility, Table 5 presents results for the base case (small network, $\Delta = 60$), in which the value of Δ is varied. The model is given four hours of CPU time for each setting. The columns are similar to the previous table, and show the number of automated orders, the relative and absolute savings, and the additional savings compared to the base case. Note that the appointment flexibility does not affect the amount of orders that may benefit from automation, which is 437 for all four experiments.

Network	Δ	Autom. Orders	Rel. savings	Cost savings	Additional savings comp. to base case
Small	30	437	29%	\$ 110,887	-0.8%
	60	437	29%	\$ 111,802	0.0%
	90	437	30%	\$ 117,354	5.0%
	120	437	30%	\$ 116,865	4.5%

Table 5: Overview of Cost Savings under Different Values of Δ.

Table 5 reveals that, if the appointment flexibility is already limited to one hour, limiting it further to 30 minutes to increase the service level is relatively inexpensive: the cost savings would only go down by 0.8%. Increasing the flexibility by 30 minutes, on the other hand, goes a long way. Using $\Delta = 90$ instead of $\Delta = 60$ results in 5% additional savings. This indicates that the impact of appointment flexibility can be substantial. Also note that the additional benefit is almost half that of the additional benefit for extending the network (+11%).

It is surprising to see that increasing Δ from 90 to 120 actually leads to a schedule that is less efficient, while more flexibility is available. This can be explained by how the optimization model operates. Finding the best schedule is a very challenging task, and increasing the flexibility increases the number of possible schedules, which makes this task even more challenging. As the optimization model only has limited time to come up with a good solution, additional flexibility may make the search more difficult, and result in a less efficient outcome. Further investigation is needed to determine the actual additional savings that can be realized for $\Delta = 120$. The results do suggest that no schedule could easily be found that was significantly better than the schedule for $\Delta = 90$, which hints that the advantage of additional flexibility is leveling off after $\Delta = 90$.

5.8 Impact of Loading/Unloading Time

The time for loading/unloading or swapping trailers is assumed to be 30 minutes in the base case. However, because autonomous trucks operate 24/7, there may be a need to extend this time to allow for refueling or recharging, maintenance, or safety checks. To investigate the impact of longer loading and unloading times on operations, an operational plan is made for the base case in which both loading and unloading time are extended by one hour (i.e., every transportation task is extended by two hours). The current number of autonomous trucks is not sufficient to create a feasible plan, and five additional trucks are added ($K = 55$). The optimization model is given one hour to optimize the schedule.

Figure 17 shows the resulting schedule for the base case with extended loading and unloading times, which is reflected by the longer transportation tasks. The operational plan is much more dense than for the base case, and there is little waiting during the weekdays. The obtained cost savings are \$ 104,619 (27%), which is almost identical to the base case (\$ 104,886). It is concluded that significantly longer loading and unloading times can be incorporated in the operational plan, at the cost of adding a small number of extra vehicles. For the case study, scheduling more time for loading and unloading requires an investment into five additional autonomous trucks, but does not affect the operational cost.

Figure 17: Autonomous Truck Schedule for the Small Network with 90 Minutes Loading/Unloading.

6 Discussion

Autonomous trucking will impact many aspects of the freight industry, from the *business model* to the *use of IT*, and the necessary *investments*. The ATHN model is the most likely future scenario, and this paper showed that it can produce substantial cost savings. Besides the cost advantages, ATHNs can also offer different types of services. Traditionally, orders are picked up in the morning on weekdays, but autonomous trucking allows for more flexibility. For example, next-morning delivery is possible by picking up the freight on day one, transporting it with an autonomous truck over night, and making the delivery the next morning. This would not only benefit customers, but also increase vehicle utilization. The job of a truck driver would also change, and the study by Viscelli discusses the impacts in more detail.[19] Human drivers will only drive the first and last miles, and may be involved in the handling of the freight at the hubs. It is also possible that drivers will be more involved in administrative or monitoring tasks.

The switch to ATHNs is likely to come with a large consolidation in the industry. Roland Berger predicts that consolidation is inevitable, based on the investment cost for autonomous technology.[20] This paper brought forward a new argument that supports the same conclusion: *The possibility to immediately serve the next order on the autonomous network, without an empty return trip, accounts for a large part of the cost reduction.* To benefit from this flexibility, the network has to be sufficiently large, and there have to be sufficiently many orders, leading to consolidation. Additionally, it was shown that a larger network allows more orders to benefit from automation, and further improves efficiency. In the future, autonomous network operators may be large logistics companies, such as Ryder, or autonomous technology firms that offer the technology on a subscription basis. The case study showed that the first/last-mile operations at the hubs may not be viable as a separate business (they need to be complemented by a local demand), which will make outsourcing and consolidation even more likely.

For ATHNs, IT will become an integral part of the business, and not only because

[19] Viscelli 2018.
[20] Roland Berger 2018.

of L4 technology. Operating a transfer hub model is more complex: it requires a system to track the orders and make sure the freight is available at the right place at the right time. The role of optimization will be paramount to realize the potential savings. The case study demonstrated that the routes driven by the autonomous trucks are complex, and it is unlikely that human planners will be able to schedule the orders with the same efficiency. Finally, optimization may lead to new market mechanisms, including dynamic pricing, since certain orders will create network imbalances and hence empty miles.

Different investments are necessary to operate an ATHN, including investments in L4 technology and in autonomous trucks. The composition of the fleet is likely to change: the case study highlighted that a relatively small amount of autonomous vehicles is needed to operate the system. This is due to the flexibility of autonomous trucks, and their ability to operate 24 hours per day. Investments are also necessary to set up the transfer hubs. This cost has not yet been quantified precisely and will highly depend on what the hub itself will look like, which could range from an empty parking lot to a full cross-docking facility. The choice for a particular type of hub depends on its intended utilization, and there may be differences throughout the network. An interesting question is then to decide how many transfer hubs to build, and of what type. The case study showed that a larger network leads to additional benefits, but it has yet to be investigated how the investment costs affect the network structure. Optimization models can assist in making these decisions.

7 Conclusion

Autonomous freight transportation is expected to completely transform the industry, and the technology is advancing rapidly, with different players developing and testing high automation L4 trucks. A crucial factor for the adoption of autonomous trucks is its return on investment, which is still uncertain. This study contributes to the discussion by quantifying the benefits of the autonomous transfer hub model, which is one of the most likely future scenarios. The benefits are estimated based on a real transportation network, taking into account the detailed operations of the autonomous network.

This paper presented data-driven methods to design and operate an Autonomous Transfer Hub Network (ATHN), and an optimization model was introduced to create the operational plan. These methods were applied to real order data provided by Ryder Dedicated Transportation Solutions. For some of the most challenging orders in the Southeast (orders that make a single delivery and return empty), operational cost may be reduced by 27% to 38% for a small ATHN and 29% to 40% for a more extensive network, which could save an estimated \$5.5M to \$8.4M per year on these orders. The savings are mainly attributed to a reduction in labor cost, but the increased flexibility of autonomous trucks also plays a significant role: Even if autonomous trucks would have the same cost per mile as human-driven trucks, cost savings would still be possible. The operation of the autonomous network is more complicated than the current operations, and it was found that optimization techniques are crucial to obtain these savings.

It was also explored how different assumptions impact the ATHN. Increasing the size of the autonomous network mainly increases the number of orders that can benefit from automation, while the average benefit per automated order remains similar. The impact of the cost per mile for autonomous trucking was also studied. As autonomous trucks become cheaper, it is cost-efficient to automate more orders, and existing trips become cheaper as well. Due to the additional flexibility, it was found that the system benefit of automation often exceeds the benefit of the lower cost per mile. It was analyzed how appointment flexibility impacts the efficiency, and it was found that allowing deviations to be even 30 minutes larger can go a long way. Finally, the impact of loading and unloading time was studied, and it was concluded that significantly longer loading and unloading can be incorporated in the operational plan at the cost of adding a small number of extra vehicles.

This paper quantified the impact of autonomous trucking on a real transportation network, and substantial benefits were found in terms of labor costs and flexibility. These results strengthen the business case for autonomous trucking, and major opportunities may arise in the coming years. To seize these opportunities, transport operators will have to update their business models, and use optimization technology to operate the more complex systems.

Acknowledgments

This research was funded through a gift from Ryder. Special thanks to the Ryder team for their invaluable support, expertise, and insights.

References

Engadget (2020). *Waymo and Daimler team up to develop self-driving trucks*. https://www.engadget.com/waymo-daimler-self-driving-truck-partnership-142519288.html.

Engholm, Albin, Anna Pernestål, and Ida Kristoffersson (2020). "Cost Analysis of Driverless Truck Operations". In: *Transportation Research Record* 2674.9, pp. 511–524.

FleetOwner (2021). *TuSimple among autonomous truck companies to join Self-Driving Coalition*. https://www.fleetowner.com/technology/autonomous-vehicles/article/21152006/tusimple-among-autonomous-truck-companies-to-join-selfdriving-coalition.

Forbes (2021). *Plus Partners With IVECO To Develop Automated Trucks For Global Deployment*. https://www.forbes.com/sites/richardbishop1/2021/04/12/plus-partners-with-iveco-to-develop-automated-trucks-for-global-deployment.

FreightWaves (2021). *Gatik, Isuzu to partner on autonomous truck platform*. https://www.freightwaves.com/news/gatik-isuzu-to-partner-on-autonomous-truck-platform.

Government Technology (2020). *Arizona Self-Driving Truck Pilot Expands UPS Partnership*. https://www.govtech.com/transportation/Arizona-Self-Driving-Truck-Pilot-Expands-UPS-Partnership.html.

Greene, William (2013). "Autonomous Freight Vehicles: They're Heeeeere!" In: *Autonomous Cars – Self-Driving the New Auto Industry Paradigm*. Ed. by Ravi Shanker et al. Morgan Stanley & Co. LLC, pp. 85–89.

Laborie, Philippe et al. (2018). "IBM ILOG CP optimizer for scheduling". In: *Constraints* 23.2, pp. 210–250.

National Public Radio (2019). *U.S. Postal Service Tests Self-Driving Trucks*. https://www.npr.org/2019/05/21/725524334/u-s-postal-service-is-testing-self-driving-trucks.

OpenStreetMap (2020). *Planet dump retrieved from https://planet.osm.org*. https://www.openstreetmap.org.

Roland Berger (2018). *Shifting up a gear – Automation, electrification and digitalization in the trucking industry*. https://www.rolandberger.com/publications/publication_pdf/roland_berger_trucking_industry.pdf.

SAE International (2018). *Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles*.

Shahandasht, Mohsen, Binaya Pudasaini, and Sean Logan McCauley (2019). *Autonomous Vehicles and Freight Transportation Analysis*. Tech. rep. The University of Texas at Arlington.

Short, Jeffrey and Dan Murray (2016). *Identifying Autonomous Vehicle Technology Impacts on the Trucking Industry*. American Transportation Research Institute, http://atri-online.org/2016/11/15/identifying-autonomous-vehicle-technology-impacts-on-the-trucking-industry/.

Slowik, Peter and Ben Sharpe (2018). "Automation in the long haul: Challenges and opportunities of autonomous heavy-duty trucking in the United States". In: *The International Council on Clean Transportation*.

TechCrunch (2021). *Aurora and Volvo partner to bring autonomous long-haul trucks to North America*. https://techcrunch.com/2021/03/30/aurora-and-volvo-partner-to-bring-autonomous-long-haul-trucks-to-north-america/.

The Wall Street Journal (2020). *Robot Trucks Are Seeking Inroads Into Freight Business*. https://www.wsj.com/articles/robot-trucks-are-seeking-inroads-into-freight-business-11598954400.

Transport Topics (2020). *Navistar, TuSimple Partner to Launch Self-Driving Trucks in 2024*. https://www.ttnews.com/articles/navistar-tusimple-partner-launch-self-driving-trucks-2024.

Van Hentenryck, Pascal (1999). *The OPL optimization programming language*. MIT press.

Viscelli, Steve (2018). *Driverless? Autonomous Trucks and the Future of the American Trucker*. Center for Labor Research and Education, University of California, Berkeley, and Working Partnerships USA.

A Mathematical Description of the Optimization Model

```
1  range Trucks = ...;
2  range Tasks = ...;
3  range Sites = ...;
4  range Horizon = ...;
5  range Types = Sites union { shipType };
6  int or[Tasks] = ...;
7  int de[Tasks] = ...;
8  int pickupTime[Tasks] = ...;
9  int loadTime = ...;
10 int flexibility = ...;
11 int travelTime[Types,Types] = ...;
12 int travelCost[Types,Types] = ...;
13
14 dvar interval task[t in Tasks] in Horizon
15               size travelTime[or[t],de[t]] + 2*loadTime;
16 dvar interval ttask[k in Trucks,t in Tasks] optional in Horizon
17               size travelTime[or[t],de[t]] + 2*loadTime;
18 dvar interval load[Trucks,Tasks] optional in Horizon size loadTime;
19 dvar interval ship[k in Trucks,t in Tasks] optional in Horizon
20               size travelTime[ort],de[t]];
21 dvar interval unload[Trucks,Tasks] optional in Horizon size loadTime;
22 dvar sequence truckSeq[k in Trucks]
23   in append(all(t in Tasks)load[k,t],all(t in Tasks)ship[k,t],all(t in Tasks)unload[k,t])
24   types append(all(t in Tasks)or[t],all(t in Tasks)shipType,all(t in Tasks)de[t]);
25 dvar int emptyMilesCost[Trucks,Tasks];
26 dvar int truckEmptyMilesCost[Trucks];
27
28 minimize sum(k in Trucks) truckEmptyMilesCost[k];
29
30 constraints {
31
32    forall(t in Tasks)
33       startOf(task[t]) >= pickupTime[t] - flexibility;
34       startOf(task[t]) <= pickupTime[t] + flexibility;
35
36    forall(k in Trucks,t in Tasks)
37       span(ttask[k,t],[load[k,t],ship[k,t],unload[k,t]]);
38       startOf(ship[k,t]) == endOf(load[k,t])
39       startOf(unload[k,t]) == endOf(ship[k,t])
40
41    forall(k in Trucks)
42       alternative(task[t],all(k in Trucks) ttask[k,t])
43
44    forall(k in Trucks,t in Tasks)
45       emptyMilesCost[k,t] = travelCost[destination[t],typeOfNext(truckSeq[k],ttask[k,t],
             destination[t],destination[t])];
46
47    forall(k in Trucks)
48       truckEmptyMilesCost[k] = sum(t in Tasks) emptyMilesCost[k,t];
49
50    forall(k in Trucks)
51       noOverlap(truckSeq,travelTime);
52
53 }
```

Figure 18: Formulation for Scheduling Freight Operations on an ATHN.

The model is depicted in Figure 18 using OPL syntax. The data of the model is given in lines 1–12. It consists of a number of ranges (line 1–5), information about the tasks (lines 6–8) that include their origins, destinations, and pickup times, the time to load/unload a trailer (line 9), the flexibility around the pickup times (line 10), and the matrices of travel times and travel costs. These matrices are defined between the sites but also include a dummy location `shipType` for reasons that will become clear shortly.

The main decision variables are the interval variables `task[t]` that specify the start and end times of task `t` when processed by the autonomous network, and the optional interval variables `ttask[k,t]` that are present if task `t` is transported by truck `k`. These optional variables consist of three subtasks that are captured by the interval variables `load[k,t]` for loading, `ship[k,t]` for transportation, and `unload[k,t]` for unloading. The other key decision variables are the sequence variables `truckSeq[k]` associated with every truck: these variables represent the sequence of tasks performed by every truck. They contain the loading, shipping, and unloading interval variables associated with the trucks, and their types. The type of a loading interval variable is the origin of the task, the type of an unloading interval variable is the destination of the task, and the type of the shipping interval variable is the specific type `shipType` that is used to represent the fact that there is no transition cost and transition time between the load and shipping subtasks, and the shipping and destination subtasks. The model also contains two auxiliary decision variables to capture the empty mile cost between a task and its successor, and the empty mile cost of the truck sequence.

The objective function (line 28) minimizes the total costs of empty miles. The constraints in lines 32–34 specify the potential start times of the tasks, and are defined in terms of the pickup times and the flexibility parameter. The SPAN constraints (line 37) link the task variables and their subtasks, while the constraints in lines 38–39 link the subtasks together. The ALTERNATIVE constraints on line 42 specify that each task is processed by a single truck. The empty mile costs between a task and its subsequent task (if it exists) is computed by the constraints in line 45: they use the TYPEOFNEXT expression on the sequence variables. The total empty mile cost for a truck is computed in line 48. The NOOVERLAP constraints in line 51 impose the disjunctive constraints between the tasks and the transition times.