Title

Balancing of Truck Parking Demand by a Centralized Incentives/Pricing System

Authors

Vital, Filipe
Ioannou, Petros

Publication Date

2022-03-01

DOI

10.7922/G2NG4NZZ

Balancing of Truck Parking Demand by a Centralized Incentives/Pricing System

March 2022

A Research Report from the National Center for Sustainable Transportation

Filipe Vital, University of Southern California

Petros Ioannou, University of Southern California

TECHNICAL REPORT DOCUMENTATION PAGE

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
NCST-USC-RR-22-10	N/A	N/A

4. Title and Subtitle Balancing of Truck Parking Demand by a Centralized Incentives/Pricing System	**5. Report Date** March 2022 **6. Performing Organization Code** N/A
7. Author(s) Filipe de Almeida Araujo Vital, https://orcid.org/0000-0001-5987-5993 Petros Ioannou, Ph.D., https://orcid.org/0000-0001-6981-0704	**8. Performing Organization Report No.** N/A
9. Performing Organization Name and Address University of Southern California METRANS Transportation Consortium University Park Campus, VKC 367 MC:0626 Los Angeles, California 90089-0626	**10. Work Unit No.** N/A **11. Contract or Grant No.** USDOT Grant 69A3551747114
12. Sponsoring Agency Name and Address U.S. Department of Transportation Office of the Assistant Secretary for Research and Technology 1200 New Jersey Avenue, SE, Washington, DC 20590	**13. Type of Report and Period Covered** August 2020 – December 2021 **14. Sponsoring Agency Code** USDOT OST-R

15. Supplementary Notes

DOI: https://doi.org/10.7922/G2NG4NZZ

16. Abstract

Due to hours-of-service (HOS) regulations, commercial drivers are required to stop and rest regularly, thus reducing fatigue-related crashes. Nevertheless, if the parking infrastructure cannot cope with the demand generated by these required stops, new issues arise. In particular, this is the case for long-haul trucking, which is the focus of this work. Drivers often have difficulty finding appropriate parking, leading to illegal parking, safety risks, and increased pollution and costs. In this project, the researchers consider the issue of coordinating the parking decisions of a large number of long-haul trucks. More specifically, how to model the behavior of a region's driver population and how it could be influenced. Understanding how truck parking demand is affected by the interaction of individual drivers' selfish planning behaviors (in the sense that they minimize their own costs, not the overall system cost) and how parking prices affect optimal schedules are important steps in developing a system able to balance demand. The study presents a formulation that uses a modified TDSP (Truck Driver Scheduling Problem) mixed-integer programming model which tracks parking usage by dividing time into time-slots and charging drivers per time slot used. Results show that if truck drivers are following optimal schedules, then parking prices would be effective in changing which locations and time slots would be chosen by each driver. However, price adjustments can cause demand to shift in unexpected and not always beneficial ways, likely due to HOS regulations and client constraints limiting the possible alternative schedules. Therefore, further study is required to better understand the system's properties and how to avoid or dampen these oscillations. Furthermore, due to HOS rules and client constraints, it might be impossible to divert demand from specific time slots and locations sufficiently. Nevertheless, this model could still aid in identifying these spots and contribute to the evaluation of infrastructure investment needs.

17. Key Words	18. Distribution Statement
Parking demand estimation, truck driver scheduling problem, hours of service regulation, demand balancing	No restrictions.

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price
Unclassified	Unclassified	43	N/A

About the National Center for Sustainable Transportation

The National Center for Sustainable Transportation is a consortium of leading universities committed to advancing an environmentally sustainable transportation system through cutting-edge research, direct policy engagement, and education of our future leaders. Consortium members include: University of California, Davis; University of California, Riverside; University of Southern California; California State University, Long Beach; Georgia Institute of Technology; and University of Vermont. More information can be found at: ncst.ucdavis.edu.

Disclaimer

Acknowledgments

This study was funded, partially or entirely, by a grant from the National Center for Sustainable Transportation (NCST), supported by the U.S. Department of Transportation (USDOT) through the University Transportation Centers program. The authors would like to thank the NCST and the USDOT for their support of university-based research in transportation, and especially for the funding provided in support of this project.

NCST

Balancing of Truck Parking Demand by a Centralized Incentives/Pricing System

A National Center for Sustainable Transportation Research Report

March 2022

Filipe Vital, Department of Electrical and Computer Engineering, University of Southern California

Petros Ioannou, Department of Electrical and Computer Engineering, University of Southern California

NCST

[page intentionally left blank]

TABLE OF CONTENTS

List of Tables

List of Figures

Balancing of Truck Parking Demand by a Centralized Incentives/Pricing System

EXECUTIVE SUMMARY

Due to hours-of-service (HOS) regulations, commercial drivers are required to stop and rest regularly, thus reducing fatigue-related crashes. Nevertheless, if the parking infrastructure cannot cope with the demand generated by these required stops, new issues arise. In particular, this is the case for long-haul trucking, which is the focus of this work. Drivers often have difficulty finding appropriate parking, leading to illegal parking, safety risks, and increased pollution and costs. In previous projects, we focused on efficiently planning a single truck's long-haul trips while accounting for HOS regulations and parking availability information, without considering how each driver's parking decisions would affect the entire system. However, although single-vehicle planning methods can help us understand the impact of parking conditions and HOS regulations on trip planning, they cannot be applied on a large scale. If adopted by a large enough number of drivers, the parking availability information used would become invalid, and parking demand may turn unbalanced again. As single-vehicle methods are not aware of the impact of each driver's decisions on the overall parking availability, they could send a large number of drivers to the same location and time slot regardless of the location's parking capacity. In this project, we consider the issue of coordinating the parking decisions of a large number of long-haul trucks. More specifically, we study how to model the behavior of a region's driver population and how it could be influenced. Understanding how truck parking demand is affected by the interaction of individual drivers' selfish planning behaviors (in the sense that they minimize their own costs, not the overall system cost) and how parking prices affect optimal schedules are important steps in developing a system able to balance demand.

We present a formulation that uses a modified TDSP (Truck Driver Scheduling Problem) mixed-integer programming model which tracks parking usage by dividing time into time-slots and charging drivers per time slot used. The parking rates depend on both the time and location being considered. Each driver or trip is defined by an origin, a destination, a departure time constraint, a delivery time constraint, a set of initial conditions for the values restricted by working hours regulations (e.g., driving or elapsed time since last off-duty period), and an hourly operational cost. This heterogeneity means that different planners measure the advantage of using a particular parking location and time differently. Due to the complexity of solving the problem for a large number of drivers and modeling how to manage parking availability and drivers interactions within a single optimization problem, we calculate each drivers' schedule separately using a common price matrix. We assume that penalties regarding parking difficulties faced in some locations and times could be included in the parking prices. If the pricing strategy used can find a price matrix that avoids overcrowding, then the solution generated by this model will also be a solution to the problem when all drivers and their interactions a modeled as a single optimization problem. Assuming we have information on the usual conditions of truck drivers operating in a region (or a large number of drivers/companies

willing to provide their trip information or keep their planning system continuously connected to the pricing coordinator), the system could use a sample population to simulate the effect of price changes before actually implementing the new prices.

Simulations were used to study how drivers react to price changes. Results show that the scheduling model is very sensitive to even small changes in parking prices, which is conducive to using parking prices as a means to influence demand. However, under the price update rule tested, the system oscillates significantly before reaching a valid solution, and the initial iterations might see an increase in peak demand instead of the intended demand redistribution. In summary, if we consider that truck drivers are following optimal schedules, then parking prices would be effective in changing which locations and time slots would be chosen by each driver. However, price adjustments can cause demand to shift in unexpected and not always beneficial ways, likely due to HOS regulations and client constraints limiting the possible alternative schedules. Therefore, further study is required in order to better understand the system's properties and how to avoid or dampen these oscillations. Furthermore, due to HOS rules and client constraints, it might be impossible to divert demand from certain time slots and locations sufficiently. Nevertheless, this model could still aid in identifying these spots and contribute to the evaluation of infrastructure investment needs.

Introduction

According to the American Transportation Research Institute surveys, truck parking is currently one of the trucking industry's main issues (American Transportation Research Institute, 2019). Due to hours-of-service (HOS) regulations, commercial drivers are required to stop and rest regularly, thus reducing fatigue-related crashes. Nevertheless, if the parking infrastructure is unable to cope with the demand generated by these required stops, new issues arise. In particular, this is case for long-haul trucking (de Almeida Araujo Vital et al., 2020), which is the focus of this work. First, not finding appropriate rest locations may cause drivers to work beyond the allowed time limit, increasing chances of fatigue-related crashes. Second, drivers may choose to park illegally (road shoulders, ramps, abandoned lots, etc.), which also poses safety risks. Third, the shortage of truck parking may adversely affect industry costs in multiple ways, such as increased fuel consumption due to idling or looking for parking, and higher accident-related costs and insurance premiums. Fourth, the increase in fuel consumption will also negatively affect the environment. Emissions generated by truck idling can cause substantial deterioration of the surrounding region's air quality.

Recent surveys found that most drivers use unauthorized parking locations at least once a week (Martin & Shaheen, 2013; Rodier et al., 2010; U.S. Department of Transportation, 2015). Furthermore, other studies report that many drivers often spend more than 30 minutes looking for parking (American Transportation Research Institute, 2018; NCDOT, 2017), or park one hour earlier than required in order to guarantee parking (Boris & Brewster, 2018), both of which reduce productivity. Nevertheless, less than 50% of truck stops reported working overcapacity, and the reported difficulties with parking usually refer to the period between 7PM and 5AM. Therefore, redistributing parking demand over time and space may help mitigate the effects of the parking shortage.

In previous projects, we focused on efficiently planning a single truck's long-haul trips while accounting for HOS regulations and parking availability information, without considering how each driver's parking decisions would affect the entire system (Vital & Ioannou, 2021, 2019). These solutions are appropriate when used only by a small number of vehicles that do not significantly affect the system. However, if adopted by a large enough number of drivers, the parking availability information used would become invalid, and parking demand may turn unbalanced again. This can easily happen if uncoordinated individual selections of the same parking rest area at the beginning of the trip lead to reaching overcapacity at arrival time, violating the initial assumption of parking availability at the time of arrival. Therefore, we approach the issue of uncoordinated parking selection by studying methods to coordinate the actions of multiple truck drivers/companies. Our aim is not to directly plan all trucks' trips, but to study ways to indirectly influence planners to better utilize the available parking capacity. We contemplate a central parking coordinator (CPARC) system that will have access to historical and real-time parking availability from a region's truck parking locations, as well as information on the characteristics of the region's usual truck trips, e.g., location and time interval when drivers expect to enter or leave the region, and drivers' remaining allowed driving time when entering the region. While such information flow may not be available today, the industry and infrastructure are expected to become more and more connected over time. The development of methodologies that take advantage of this expected connectivity in order to identify and quantify the benefits to industry and environment is very crucial in accelerating information technologies in the area. The CPARC system would be able to estimate parking availability/demand at a region's truck parking locations. Such a system could be used to implement pricing schemes to control parking demand and achieve a better balance between supply and demand.

In this project, we consider the issue of coordinating the decisions of a large number of long-haul trucks. More specifically, we study how to model the behavior of a region's driver population and how it could be influenced. The planning models developed in previous projects can be used to simulate how companies would revise their trip plans in response to changes in different factors, such as parking prices or parking availability estimates. This information can then be used to study how to control parking demand.

Related Work

The truck parking shortage is a serious issue in several American states as their current parking infrastructure cannot accommodate peak demand. This could be solved by increasing capacity or decreasing peak demand. We are interested in the latter. As the general objective is to manage demand, the possible approaches depend on the level of control we have over the system.

Demand-Side Management

If a model of how the demand reacts to certain parameters is known, or if demand can be rejected, it is possible to directly control demand. In this case, simply controlling demand is often not the objective, so demand control is used along with some allocation, scheduling or routing system to maximize an objective or satisfy particular constraints. For example, mobility-on-demand systems with electric vehicles where the distribution of clients' willingness to pay is known, allowing demand to be directly affected by the price. In this case, the number of clients being serviced at each location can be controlled directly by price, and the demand at charging stations is controlled directly by routing decisions. However, the range of feasible demands at charging stations will be indirectly affected by the price and demand changes occurring throughout the network. There is also work on dynamic pricing applied to smart grid in order to maximize profit or reduce peak demand.

In (Moradipari & Alizadeh, 2020), Moradipari and Alizadeh address the problem of managing demand on public EV charging stations. Their system consists of a central operator that allocates resources according to users' value of time, charging demand and travel preferences. The system's objectives are to provide fair service with short wait times to customers while managing the effects of EVs on the grid. A set of service options is provided to users, each one with different prices and probabilities of being assigned to particular charging stations and expected waiting time. Incentive compatible pricing-routing policies that maximize either a measure of social welfare or the central operator profits are presented. They assumed that the users do not observe the exact wait times, the expected wait time is constant at equilibrium and given by a function of the arrival rates and routing probabilities generated by the chosen policy. They suggest the use of queueing models to define the expected waiting time functions, but this is not covered in this study. The value of time (VoT) used to model users utility and choice is considered a random variable with known distribution.

In (Turan et al., 2019), Turan et al. use reinforcement learning to generate a dynamic policy controlling ride prices and routing/charging decisions for an autonomous-mobility-on-demand (AMoD) fleet of autonomous EVs. The problem is modeled as a markov decision problem with states determined by electricity prices at each node, customers queue lengths for each origin-destination pair, and the number of vehicles at each node and their energy levels. The decision policy is defined by a deep neural network trained using Trust Region Policy Optimization. The system objective is to maximize operator profit. Although expected waiting time is unknown to customers and do not affect demand, the operator is penalized for waiting time and thus tries to reduce total waiting time.

In (Tucker et al., 2019), Tucker et al. design a pricing framework for online electric vehicle (EV) parking assignment and charge scheduling. Each user is defined by the requested time interval, acceptable locations, required energy and utility obtained from each location. Energy prices at each location vary with time. The system allocates how the energy received by each user varies with time. As long as users demands are met, each location can control charging behavior in order to optimize its own costs. Each location can generate a certain amount of solar energy at no cost, and can buy a limited amount of energy from the grid. After users send their information, the system generates a set of options defined by cable reservation and charging schedule, along with a price for each option. If a request is rejected, the user's utility is set to zero and it is assumed that the user parked at an auxiliary parking lot without charging capabilities. The offline method assumes that all requests for a certain time interval are known, and maximizes the social welfare, defined as the difference between the total user utility and operational costs. A mathematical formulation of the offline problem and its Fenchel dual are described. In the online method, heuristic pricing functions are used along with an auction mechanism. The system decides whether to accept or not a request depending on the current prices and the user's potential utility gain. Prices are updated after each user request is processed. The proposed mechanism is shown to be α-competitive.

In (Tian et al., 2018), Tian et al. address the problem of optimizing a parking operator's revenue using a parking reservation system. Parking requests are modeled by a Poisson process with an arrival rate dependent on the parking price. However, it is assumed that demand for a certain time interval depends only on the price for that same interval in that particular parking location. Solutions are proposed for exponential and linear demand functions. The authors justify the assumptions on price/demand relations as that being the point of view of the parking manager. The manager is not aware of the decision process of the users and cannot observe it, he/she can only analyze how the changes and price affect the arrival rate.

Direct Allocation

If there is some base demand that is uncontrollable, but we can directly control how that demand is distributed, we can look at it as a resource allocation problem or a routing with load balancing problem. Smart parking systems that have some flexibility in how to allocate parking reservation requests are examples of the resource allocation point of view. In the case of load balancing, we can think of the problem of routing multiple vehicles/packets over a transportation/communications network where the cost of an edge depends on the number of routes using it. Some routing studies use dynamic pricing schemes to incentivize user participation.

In (Capdevila et al., 2013), Capdevila et al. used a multi-agent system for the management of parking reservations among requesting trucks. When a vehicle enters the road network it sends its origin, destination and preferred parking to the system manager. If the rest area has available spots a temporary reservation will be made. If the rest area does not have available spots the negotiation protocol is initiated. Each driver receives a list of possible rest areas to be graded according to his/her preferences. Each driver's vote is weighted according to their maximum allowed driving time and current driving time. These weights give priority to drivers that are closer to reaching their legal driving limit. The scores for each driver are summed for each feasible solution and the solution with the largest score is selected. Note that this algorithm assumes that there is at least 1 feasible solution for a given problem. Following the selection all drivers and rest areas are notified of the new allocation. While the number of trucks requesting parking reservations for a given parking lot is smaller than the number of available spaces all of them are granted spots, but when there are more reservations than available spaces a negotiation protocol is used to choose the parking allocation. The negotiation involves a voting procedure that takes into account the preferences of each truck and its allowed stops. The system's robustness to changes in the available parking areas and the system's scalability were tested through simulation and the results were promising, showing a substantial reduction of the necessity for drivers to park in illegal areas. Similar resource allocation problems were also treated before in the context of urban parking. In (Geng & Cassandras, 2011) the resource allocation problem was defined as a sequence of Mixed Integer Linear Programming problems solved over time subjected to a set of fairness constraints. (Doulamis et al., 2013) uses interval scheduling algorithms to try to optimally allocate parking spaces.

In (Xie et al., 2015), Xie et al. approach the problem of assigning health care workers to home visits within certain time-windows. Practitioners have different skill sets, time constraints and client preferences. Similarly, visits have time-window constraints and skill requirements. The payment sought by each practitioner depends on their skill level and the costs incurred to provide the service. The system wants to minimize service costs while guaranteeing that all visits are covered by qualified practitioners. An iterative bidding framework is proposed where providers calculate feasible schedules' costs and bid for the schedule with the largest payoff. The health agency provisionally chooses schedules that satisfy its constraints while minimizing costs. Auction follows certain bidding rules regarding when bids can be changed and by how much, and terminates when no valid bids are updated.

In (Kordonis et al., 2020), Kordonis et al. propose mechanisms to coordinate truck drivers routing decisions to balance the traffic load and improve the overall traffic conditions and time delays experienced by both truck and passenger vehicle drivers. The mechanisms use monetary incentives and fees to steer individual drivers' decisions towards a system optimum without penalizing drivers compared to the user equilibrium. They propose fairness measures and 2 sets of constraints that would encourage drivers to participate either as a group (either everyone or no one) or individually (each driver see it as beneficial to participate regardless of other drivers' decisions). The effect of routing assignments on cost/travel time is assumed known. Passenger vehicle assignments are assumed known and fixed. Demand for each OD is a random variable with known distribution.

Indirect Allocation

If we can control only parameters that affect demand distribution indirectly, then the problem resembles work using dynamic pricing to indirectly influence agents' decisions in congestion pricing, and work on anticipatory route guidance, which uses traffic predictions information to influence drivers' routing choices. This type of system can also be viewed as a non-cooperative game where we a looking for a pricing policy/mechanism that leads to an equilibrium state with particular properties.

In (Hollander & Prashker, 2006), Hollander and Prashker present a survey of applications of non-cooperative game theory in transportation problems. The problems are categorized based on the players involved: travelers x demon, travelers x travelers, travelers x authorities, authority x authority. Problems such as congestion pricing would fall into the categories of games between travelers or between travelers and authorities, depending on how the pricing scheme is included in the problem. For example, the pricing scheme can be given as an input, and a game between travelers is then used to study the equilibrium results from different inputs. The game itself would not output a policy but could be used to study a policy's impacts. Another possibility is that the authority can be explicitly modeled as a player with its own objective and constraints. In this case, studies often use bi-level formulations where the upper level optimizes the authority's objective, whereas the lower level is the user-equilibrium problem that defines how travelers react to the authority's decisions. This game between an authority and a collective of travelers would output a policy recommendation.

In (Kaufman et al., 1991), Kaufman et al. study anticipatory route guidance that accounts for the behavior of anticipatory vehicles' impact on the system. If a prediction is given to drivers, they will change their behavior and invalidate the prediction. So, this paper calculates what traffic prediction should be given to drivers so that they behave as predicted, i.e., a self-fulfilling prediction. This is calculated iteratively by switching between solving dynamic assignment problems and time-dependent shortest path problems until the routing policy converges.

In (Chen et al., 2020), Chen et al. use a bi-level optimization approach to choose the location and capacity of EV charging stations such that construction costs and drivers' travel time and waiting time are minimized. The lower level calculates vehicle routing and charging behavior at equilibrium given a set of locations and capacities, whereas the upper level optimizes the decisions regarding location and capacity subject to service level constraints at each charging station. The effect of routing choices on travel time is a non-linear function taken from the literature. Waiting time at charging stations is modeled as a queueing system, but due it being computationally expensive to use, an approximate function is proposed based on Monte-Carlo simulation. The problem is then reformulated as a single-level mathematical programming with complementary constraints, which solved by using standard NLP solvers to solve a sequence of relaxed problems.

Significant differences to truck parking

Dynamic pricing and incentive schemes are often studied for issues in the energy (Black & Tyagi, 2010; Dutta & Mitra, 2017; Gabr et al., 2018; Song, 2012), transportation (Papadopoulos et al., 2019; Turan et al., 2019; Yang et al., 2020), and communications (Falowo et al., 2009) sectors. However, some of the assumptions made are not reasonable for the truck parking management problem.

In (Falowo et al., 2009), Falowo et al. applied dynamic pricing to solve a load balancing problem in wireless networks. In this case, one of their objectives was to achieve uniform load distribution, which may not be necessary or reasonable in our case. Our main objective is to avoid overloading any parking facility at any time, so the load should be balanced enough to avoid peaks that exceed capacity, but not necessarily uniform.

In (Turan et al., 2019), Turan et al. used reinforcement learning to manage (ride prices, routing and charging decisions) an electric autonomous mobility-on-demand system. Ride prices were used to control customers' arrival rate at each node. However, the authors considered that the base arrival rates are not time-dependent, and that the customers' willingness-to-pay distribution (how much each client is willing to pay for a ride) is the same at all locations and times. In the case of truck parking management, due to the several factors that influence the planners' costs, the willingness of each planner to pay to use a certain parking slot and the arrival rate may vary greatly with time and location. In addition, some dynamic pricing schemes treat demands at different locations and times as if they are independent and can simply be eliminated. When demands are reduced due to higher prices, it is assumed that customers gave up, it does not affect the demand at other times/locations. In our case, demand is usually not simply reduced, it is shifted. Although rerouting may cause the total parking time of a truck to be reduced (possibly turned into extra driving time), most of it is only moved to a different location or time slot. This issue was also raised in (Song, 2012) with regards to energy consumption, as some customers react to dynamic prices by changing their consumption time instead of only reducing consumption. Nevertheless, estimating clients' behavior in order to predict when and by how much demand will be shifted or reduced is very hard.

In the congestion pricing /coordinated routing problem, all drivers using a certain route/link are affected by the high usage rate of that route/link. However, in truck parking, only drivers that arrive after a parking lot is full would experience cost increases. The cost of drivers arriving early is not affected by the high occupancy, so they might be harder to influence.

In the coordinated routing problem all routes are assumed to be contained within a single time interval, so even though routes are composed of multiple links, the time dimension is ignored and all costs and effects affect all links at the same time. As trucks have large limitations on available routes, it is also somewhat reasonable to assume that the number of possible routes is small. Especially if we assume that these formulations focus on short-haul as they consider that trips are completed within a single time interval.

In truck parking, the time dimension is important as the time a driver occupies a certain rest area has a large impact on the times he/she is likely to stop next. And due to the focus on long-haul and the possibility of a large number of rest areas existing along a route, the number of possible combinations of rest areas used, stopping times and rest durations that form a schedule is very large and not straightforward to reduce.

Problem Description

Figure 1 shows a simplified diagram of how truck parking demand is generated. Each trucking company plans its trips according to some public information, such as regulations, traffic data and parking data when available, and its own private information regarding its own operational constraints and parameters, such as clients' requirements and drivers' remuneration. For simplicity, we refer to each truck driver/company that needs to plan a trip as a planner. Each planner acts in its own best interest, planning a route and schedule that minimizes its operational cost. The drivers will then follow their itinerary and try to park at the planned location and time. However, the planners do not possess information on how the decisions of other planners will affect the future state of the system. So, assuming planners have access to the same public information, if a certain parking lot is low cost compared to others and usually available at a certain time, all planners will assume that they can use it at that time. Unable to see the whole picture, all trucks may be routed to the same parking lots, causing an unbalanced parking capacity usage. Certain facilities may be working overcapacity, while others may have plenty of parking spaces left. Not being able to find parking at the expected location, some drivers may be forced to park illegally, which may pose significant safety and financial risks. For the drivers that could not find parking, the cost increase of choosing a sub-optimal route/schedule from the beginning may be lower than the cost increase caused by not finding parking at the planned route. However, in order to consciously decide against using the optimal route, the planner needs to know in advance whether others' decisions will turn its own decision infeasible.

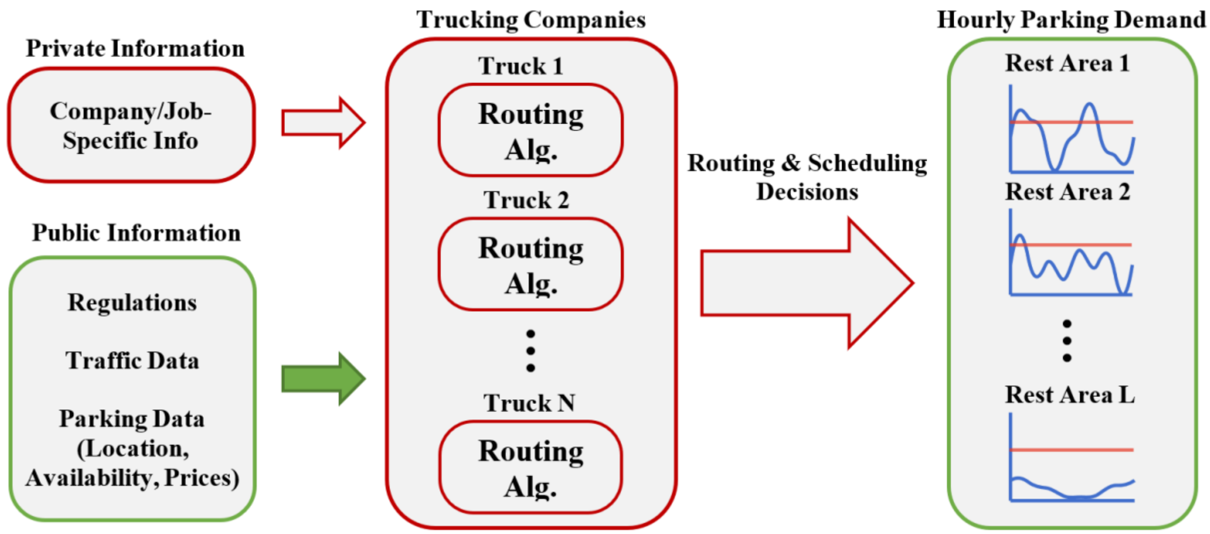

Figure 1. Simplified representation of how parking demand is generated. Green boxes represent information that is unique to the system, regardless of whether it can be measured. Red boxes represent information that is specific to each company. The red horizontal lines on the parking demand plots represent each facility's maximum capacity.

Planners' decisions may be influenced by factors such as traffic, fuel costs, parking availability, parking rates, and illegal parking penalties. Many cost-defining factors may be particular to a certain planner, such as company policy about illegal parking, driving speed limits, fuel consumption, and driver remuneration. The same is true for the set of feasible schedules, as it can depend on restrictions particular to each planner, such as the details of each trip (client location, time-windows, etc.) and drivers' remaining driving time. This heterogeneity means that different planners measure the advantage of using a certain parking location and time differently.

Our objective is to study how to influence this cost calculation so that planners having flexible schedules will opt to use low-demand parking slots instead of high-demand ones. Consider the effects of the following factors: hourly parking price; and illegal parking penalties. If a certain parking lot is expected to work overcapacity at a particular time, the parking rates for that time and place can be increased, motivating drivers to park at different times or locations. The price value would convey how good or bad (for the system) it is for the planner to use that resource at that time. As each planner is solving a different optimization problem (the jobs, vehicles, clients, and constraints may all be different), each truck has a different cost for changing itineraries. Therefore, as the prices increase, companies that have reasonable alternative route options would change their itineraries to minimize costs. Whereas companies that are in urgent need of that resource due to less flexible conditions, would keep their itinerary and accept the cost increase. Different from urban parking dynamic pricing schemes that often aim to maximize parking lot profits, in this case, we aim for a better utilization of the available parking capacity and a reduction of the cases of illegal parking. The illegal parking penalty could help model and control planners' unwillingness to switch routes, as well as measure the quality of alternative routes. If a planner considers more cost-effective to risk parking illegally than to switch routes, then it means that all alternatives are too expensive. In the future, the penalty values could also be used to measure the need and benefits of infrastructure investments in certain areas. However, as the penalty's effect depends on the probability of finding parking, which varies according to other planners' decisions, we do not use it as a control input. Furthermore, it would be hard to control the ``real'' penalty values as they depend on the fines imposed by the government, on how strictly highway officers are enforcing parking restrictions, and other factors completely out of our control, such as accident rates, and average litigation costs.

In Figure 2, we show the basic diagram of the CPARC including a parking pricing manager. The CPARC system would first provide the planners with initial parking rates, calculated according to historical data. Planners would then calculate their routes and communicate their desired parking slots to CPARC. Assuming full participation, CPARC's Demand Estimator would be able to perfectly calculate the hourly demand for each parking slot, and verify which ones are overcapacity. The pricing manager would calculate new prices, which would then be sent to the planners. In the case of partial participation, CPARC could generate demand estimates by using the participants' demands along with historical parking availability data. Similarly, the prices would be adjusted according to the demand estimates and sent back to the planners. This cycle of planning, demand estimation, and price update would continue until an acceptable solution is found, or until a time or iteration limit is reached.

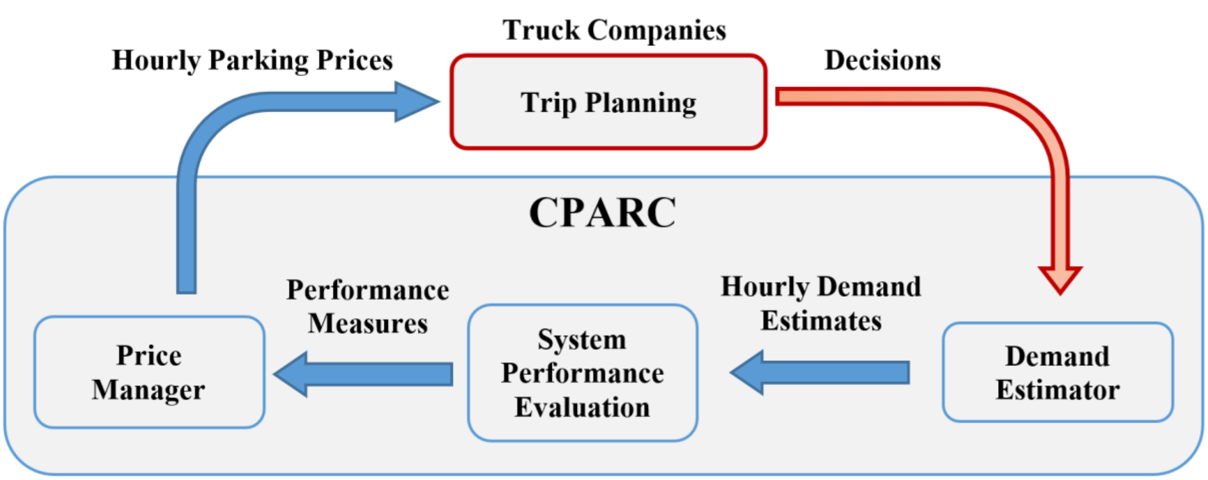

Figure 2. System Diagram. The blue boxes represent our system's components and output. The red box represents the distributed system encompassing all truck companies and their decisions.

One of the challenges in developing this kind of system is estimating planners' reactions to price changes so that appropriate prices can be determined. Therefore, we focus on how to use models for the Truck Driver Scheduling Problem to simulate planners' behavior. In our analysis, we assume that the objective of each planner is to minimize its overall cost, part of which is associated with the parking cost. By changing the parking cost at a particular location, the overall cost may no longer be optimum when compared with a lower cost parking, which may require the planner to modify the initial route. Parameters unique to each planner, such as client locations, delivery constraints and driver hourly wage, can be sampled from a given distribution, whereas parameters such as travel time, parking locations, parking price and HOS regulations are the same for all planners.

Performance measures to determine what is an acceptable solution and how distant we are from achieving it depend on each region's particular context, and the pricing strategy will depend on the measures used. For example, we could measure the parking shortage's severity by measuring the excess demand at each parking lot. Let $D_i(t)$ represent facility i's parking demand at time t, C_i is facility i's parking capacity. Then we can define the following performance measures:

$$E_i(t) = \max(0, D_i(t) - C_i) \tag{1}$$

$$\Delta_i = \int_0^{\infty} E_i(t)dt \tag{2}$$

$$\Delta = \sum_i \Delta_i \tag{3}$$

where $E_i(t)$ measures the excess demand of location i at time t in parking spaces, Δ_i is a measure of the parking shortage at location i in $parking\ spaces \cdot hours$, and Δ measures the parking shortage of the whole system in $parking\ spaces \cdot hours$. Ideally, an acceptable solution would have Δ equal to zero, so no parking facility has a demand larger greater than it can support. However, in order to properly adjust the prices, we also need to know which facilities need to have their demand decreased and at what time. This information is provided by $E_i(t)$ and Δ_i. It may also be necessary to track how the price changes are affecting the costs of each planner and of the whole system. If the system cost decreases, it may be possible to use the gains of some planners to compensate for the losses of others. This ideal situation would allow for the system to be maintained without extra investments besides infrastructure and management costs.

Preliminaries

Truck Driver Scheduling Problem

Consider the truck driver scheduling problem (TDSP) under parking availability constraints presented in (Vital & Ioannou, 2019). The problem consists of scheduling the rest stops for a single long-haul truck trip with a known route and a single client while taking into account the USA HOS regulations and estimated parking availability windows for all rest areas along the route. It is assumed that the rest areas are located on the route and require no detours to be accessed. The parking availability time-windows are assumed known. The route has $n + 1$ nodes, 2 of which are the origin, node 0, and destination of the truck, node n. The other $n - 1$ are rest areas located along the route. For each node $i \in \{0,1,\dots,n\}$ the variable $x_i = (x_{i,a}, x_{i,d})$ represents the arrival and departure times of the truck at that node. Each rest area i has T_i parking availability time-windows $[t_{i,\tau}^{min}, t_{i,\tau}^{max}]$, where $\tau \in \{1,2,\dots,T_i\}$ indicates the time-window's index. The time-windows restrict the arrival time at that node and are only in effect when the truck has to stop at that specific node, driving by it is not constrained by the time windows. For each location and time-window, a binary variable $y_{i,\tau}$ represents if that specific time window is being used (yes:1, no:0). Driving by without stopping is represented by the variable $y_{i,0}$ (drive by:1, stop:0). The travel time $d_{i,i+1}$ in between nodes is considered known and independent of time. The planning horizon is denoted by t_{hor}. The driver must reach its destination before the specified planning horizon.

The schedule must comply with the HOS regulations (see (Vital & Ioannou, 2019) for regulations considered in the model). R is defined as the set of different types of rest periods described in the regulation. For each $r \in R$, t_r defines the minimum duration of that type of rest period. C is the set of constraints imposed by the regulation. $C_1 \subseteq C$ is the set of constraints controlling the maximum elapsed time between off-duty periods. $C_2 \subseteq C$ is the set of constraints controlling the maximum accumulated driving time between off-duty periods. The HOS regulations also limit the on-duty time over the last 7 days, but we will consider a simplified version of the constraint which limits the on-duty time since the last 34h rest. This simplified constraint is more restrictive than the original and still guarantees regulation-compliance. For each constraint $c \in C$, t_c is the time limit imposed by the regulation and $R_c \subseteq R$ is the set of rest types that can reset this counter. The binary variable $z_{i,r}$ indicates whether a rest of type r is taken at location i (yes:1, no:0). The driver cannot take more than 1 type of rest at the same location. If no type of rest is schedule for a rest area, the driver cannot stop there. The departure time from the origin must be within the interval $\left[t_0, t_{dep}\right]$. It is assumed that the driver has been off-duty for long enough before the departure time, so that all constraints' counters are reset before departure.

$$\text{Minimize} \quad \text{Total travel time} = x_{n,a} - x_{0,d} \tag{4}$$

s.t.:

$$x_{i,d} + d_{i,i+1} = x_{i+1,a} \forall 0 \leq i \leq n-1 \tag{5}$$

$$x_{i,a} + \sum_{r \in R} t_r \, z_{i,r} \leq x_{i,d}, \forall 1 \leq i \leq n \tag{6}$$

$$x_{i,d} \leq x_{i,a} + \left(1 - y_{0,\tau}\right) t_{hor}, \forall 1 \leq i \leq n \tag{7}$$

$$y_{i,0} + \sum_{\tau=1}^{T_i} y_{i,\tau} = 1, \forall 1 \leq i \leq n \tag{8}$$

$$\sum_{\tau=1}^{T_i} y_{i,\tau} = \sum_{r \in R} z_{i,r}\,, \forall 1 \leq i \leq n-1 \tag{9}$$

$$\sum_{\tau=1}^{T_i} y_{i,\tau} \, t_{i,\tau}^{min} \leq x_{i,a}, \forall 1 \leq i \leq n \tag{10}$$

$$x_{i,a} \leq t_{hor} - \sum_{\tau=1}^{T_i} \left[y_{i,\tau} \left(t_{hor} - t_{i,\tau}^{max} \right) \right] \forall 1 \leq i \leq n \tag{11}$$

$$x_{k,a} - x_{i,d} \leq t_c + t_{hor} \sum_{j=i+1}^{k-1} \sum_{r \in R_c} z_{j,r}, \forall 0 \leq i < k \leq n, c \in C_1 \tag{12}$$

$$\sum_{j=i}^{k-1} d_{j,j+1} \leq t_c + t_{hor} \sum_{j=i+1}^{k-1} \sum_{r \in R_c} z_{j,r}, \forall 0 \leq i \leq k \leq n, c \in C_2 \tag{13}$$

$$x_i \in [0, t_{hor}]^2, y_i \in \{0,1\}^{T_i+1}, z_i \in \{0,1\}^{|R|}, \forall 1 \leq I \leq n \tag{14}$$

$$x_{0,d} \in \left[0, t_{dep}\right], y_{n,0} = 0 \tag{15}$$

The objective function (4) is set to minimize the total trip duration. Constraint (5) guarantees that the arrival time equals the departure time of the previous location plus the driving time. Constraint (6) states that the vehicle must not depart before the arrival time plus the minimum rest time decided for that location. Constraint (7) controls what happens when the driver does not stop at a certain location. If the vehicle does not stop at location i, the arrival time equals the departure time. This constraint works with constraints ((6), (8), (9)) to assure this. Equality will hold when $y_{i,0} = 1$. If $y_{i,0} = 0$, then constraint (7) is always true as t_{hor} is large. Constraint (8) states that at any location, either exactly 1 time-window is used or the vehicle does not stop. Constraint (9) states that the driver only stops if an off-duty period is scheduled. Constraints (10) and (11) check the time-windows. Arrival must happen after the beginning and before the end of the chosen time window. Constraint (12) checks that the time elapsed since the last rest in $R_c, c \in C_1$ is less than t_c. Constraint (13) checks if the accumulated driving time between rest periods in $R_c, c \in C_2$ is less than t_c. Constraint (14) sets the variables' domains, and (15) guarantees that the departure time from the origin is within the required period and that the vehicle will stop at the destination.

Cooperative Game

In cooperative games, the focus is on the overall system cost, not each agent's cost. We can look at it as if a single entity controls all vehicles, and that entity is willing to let some vehicles operate at higher costs as long as the overall cost is reduced. Let V be the set of different vehicles or trips usually found in the region of interest. Each $v \in V$ contains all the parameters necessary to describe the trip (e.g., origin, destination, time-windows, HOS constraints initial condition, etc). Let $TDSP(v)$ be the set of feasible schedules for vehicle/trip v, without considering parking availability constraints, we will also refer to $TDSP(v)$ as the set of pure strategies for vehicle v. Although the TDSP models described in section Truck Driver Scheduling Problem consider time-related variables as continuous, here we assume that the decision space is discretized. Let $A_v \in TDSP(v)$ represent a pure strategy for vehicle v, and a pure strategy profile $A = (A_v)_{v \in V}$ be the vector of pure strategies assigned to each vehicle. The set of all possible pure strategy profiles is defined as $S = \prod_{v \in V} TDSP(v)$. The function $\alpha_v(s): TDSP(v) \mapsto R$ is a mixed strategy for vehicle v, defined as the probability of vehicle v

choosing to use schedule $s \in TDSP(v)$. The mixed strategy profile $\alpha = \{\alpha_v(\cdot)\}_{v \in V}$ is the vector of functions $\alpha_v(\cdot)$ indicating the mixed strategy adopted by each vehicle. Given a system-wide objective function $\Theta(\alpha)$, the system optimum can be defined as the solution to the following optimization problem:

$$\min_{\alpha} \Theta\,(\alpha) \tag{16}$$

s.t.:

$$\sum_{s \in TDSP(v)} \alpha_v(s) = 1, \forall v \in V \tag{17}$$

$$\alpha_v(s) \geq 0, \forall v \in V, s \in TDSP(v) \tag{18}$$

Parking capacity constraints are not included as we consider that $\Theta(\alpha)$ already accounts for all relevant costs, such as costs due to overcapacity parking facilities, trip delays, etc. If feasible, overcrowding would be avoided by setting high associated costs sufficiently high. Nevertheless, this formulation assumes a cooperative relationship where all vehicles are controlled by a single agent, and it is acceptable to increase the costs of certain vehicles without limits as long as the overall system cost is reduced.

Non-cooperative Game

The trucking industry is a competitive market composed by a large number of agents, with large and small companies commanding varying fractions of the vehicles in operation and a large number of truck-owners operating independently. Although it is reasonable to assume cooperation within large companies, that is not the case for the overall system. The way agents behave in a competitive market is better modeled as a non-cooperative game, where each agent is trying to optimize its own objective depending on how it expects competitors will behave. This scenario is usually analyzed in the literature by assuming that agents have some information on each other’s' intentions and can plan accordingly. If the agents are able to reach a stable solution, i.e., no one could improve their objective by changing behavior given that all other agents' behaviors remain unchanged, that solution is called a Nash Equilibrium (NE). A game may have zero, one or multiple equilibria, which are usually less efficient than the system optimum. Therefore, the system manager is interested in either pushing the system towards the NE that best suits the system's objective, or implementing policies that generate NEs with better system-wide performance. The ratio between the worst-case NE's cost and the system optimum is usually referred to as the Price of Anarchy (Christodoulou, 2008), it is often used in the literature to study the quality of a game's equilibria.

Let $F_v(A)$ be the cost of vehicle v under a pure strategy profile $A \in S$ and let (A_{-v}, s_v) represent the strategy profile obtained when, starting from a pure strategy profile A, vehicle v switches to strategy s_v. For convenience, $F_v(\alpha)$ is also used to represent the expected cost of vehicle v under a mixed strategy profile α, and, in this case, it is defined as:

$$F_v(\alpha) = \sum_{A \in S} p(A) F_v(A) \tag{19}$$

$$p(A) = \prod_{v \in V} p_v(A_v) \quad (20)$$

where $p(A)$ is the probability that pure strategy profile A is used, and $p_v(A_v)$ is the probability of vehicle v using pure strategy A_v. A mixed strategy profile α constitutes a NE if and only if it satisfies the following constraints:

$$F_v(\alpha_{-v}, s_v) \geq F_v(\alpha), \forall v \in V, s_v \in TDSP(v) \quad (21)$$

$$\sum_{s \in TDSP(v)} \alpha_v(s) = 1, \forall v \in V \quad (22)$$

$$\alpha_v(s) \geq 0, \forall v \in V, s \in TDSP(v) \quad (23)$$

If the system manager were able to recommend schedules to all users, but those users would only follow the recommendations if it optimizes their objective given their expected behavior of the other agents, the best recommendation could be calculated by finding $\min_{\alpha} \Theta(\alpha)$ such that (21)-(23) are satisfied. For example, a planning software used by a large number of vehicles would be able to both recommend schedules and show planners the expected cost of other options. The cost estimate is generated assuming that the other users will follow the recommended schedule. Although unable to directly control the vehicles and force the adoption of a system optimum strategy, the system manager would be able to at least steer the system towards the most beneficial NE. Anticipatory routing systems (Dong, 2008; Kaufman et al., 1991) interfere with the information provided to drivers, instead of directly charging for certain routes or imposing decisions. $F_v(\alpha)$ depends on predicting the strategy α_{-v} used by other agents and how the different strategies from a strategy profile α interact. Therefore, it can be affected by the information provided to drivers regarding the expected effect of others on the system. However, if the provided information does not match what drivers actually experience, they would lose trust in the system and stop using it. So, the manager is restricted to using policies that make drivers behave as predicted. If we consider that the importance drivers give to parking conditions is already adequate and they simply lack the information to act on it, we could consider a system similar to anticipatory routing. In this case, the system would give drivers a parking availability prediction such that they realize that prediction.

Now, consider the case when the system manager is unable to recommend schedules, or all NE have undesirable costs. In this case, we want to change the system such that the new system's NEs have acceptable costs. As described in (21), the NEs depend on the agents' perception of cost. So in order to shift the NEs towards more desirable solutions, we need to somehow influence the agents' costs. Good examples are the pricing mechanisms in (Kordonis et al., 2020; Moradipari & Alizadeh, 2020). In pricing mechanisms the manager can directly charge or give monetary incentive for the usage of each resource, so it would be equivalent to adding an extra cost $\gamma_v(\alpha)$ to the cost function $F_v(\alpha)$ used by agents to optimize their decisions. Studies using this type of approach often explore concerns regarding user participation, fairness of the prices/incentives used, whether the manager or users are making or losing money, and whether users are truthful when providing information to the system. This approach can be described as solving the following problem:

$$\min_{\gamma} \Theta\,(\alpha, \gamma) \tag{24}$$

s.t.: $$F_v(\alpha_{-v}, s_v) + \gamma_v(\alpha_{-v}, s_v) \geq F_v(\alpha) + \gamma_v(\alpha), \forall v \in V, s_v \in TDSP(v) \tag{25}$$

$$\sum_{s \in TDSP(v)} \alpha_v(s) = 1, \forall v \in V \tag{26}$$

$$\alpha_v(s) \geq 0, \forall v \in V, s \in TDSP(v) \tag{27}$$

where $\gamma = (\gamma_v)_{v \in V}$, and $\Theta(\alpha, \gamma)$ is a modified objective function to account for any impact the pricing policy γ might have on the system cost, e.g., the objective function might include the balance of incentives given and fees collected by the system.

Formulation

Agent interaction

In the case of truck parking, the interaction of agents will be based on the demand of truck parking locations. This effect could be modeled by the recourse function used in (Vital & Ioannou, 2021) to estimate the cost of not finding parking or maybe a penalty for scheduling a stop at a full rest area. One issue is whether we assume that drivers account for this cost when planning their trips. In theory, drivers do not know other drivers' decisions and how they will affect parking availability. We could argue that freight movement within a region follows certain patterns and that drivers would adapt to these patterns over time and reach a NE as if every driver knew other drivers' decisions. With this assumption we can calculate the impact of parking shortage penalties both on individual drivers and on the whole system. One option is to assume that the parking shortage has no impact on the drivers that arrive before rest areas reaching capacity. However, it might be easier to consider that all drivers are affected equally, similar to congestion pricing problems. It might also be interesting and reasonable to consider that parked drivers can be affected to a smaller extent too as overcrowding might cause problems within rest areas and drivers can be affected by accidents caused in the rest area's surroundings. Let $\rho_i(t_{arrival}, t_{dep}, o_i)$ represent penalties due to parking conditions perceived by drivers parking from t_0 to t_1 at a location i with parking demand profile o_i. This function defines how much drivers care about parking conditions, and whether every driver using an overcrowded parking facility incurs penalties or only those that arrived after it reached capacity.

Individual Behavior

We assume that each driver seeks to minimize his/her own costs, and that they are aware of the penalty costs that will be incurred after equilibrium is reached. For simplicity, we consider that all drivers at a parking facility are penalized equally when the facility is overcapacity and this penalty will be considered as part of the parking fees imposed by the system. As in section Truck Driver Scheduling Problem, each vehicle is subject to HOS constraints and client delivery time-windows. However, we are now interested in using this model to estimate parking demand, so we do not include parking availability time-window constraints. Furthermore, we need to account for parking costs and initial conditions and to track the parking demand generated by each vehicle. Therefore, we consider the following modified TDSP formulation:

$$\min \beta^v \left(x^v_{e^v,a} - x^v_{s^v,d}\right) + \sum_{i=s^v}^{e^v} p^v_i + \sum_{i=s^v}^{e^v} \rho^v_i(\{y_i\}, o_i) \tag{28}$$

s.t.:

$$x^v_{i,d} + d_i = x^v_{i+1,a}, \forall s^v \le i \le e^v - 1 \tag{29}$$

$$x^v_{i,a} + \sum_{r \in R} t_r \, z^v_{i,r} \le x^v_{i,d}, \forall s^v \le i \le e^v - 1 \tag{30}$$

$$x^v_{i,d} \le x^v_{i,a} + \left(1 - y^v_{i,0}\right) t_{hor}, \forall s^v \le i \le e^v - 1 \tag{31}$$

$$y^v_{i,0} + \sum_{\tau=1}^{T^v_i} y^v_{i,\tau} = 1, \forall s^v \le i \le e^v \tag{32}$$

$$y^v_{d,i,0} + \sum_{\tau=1}^{T^v_i} y^v_{d,i,\tau} = 1, \forall s^v \le i \le e^v - 1 \tag{33}$$

$$y^v_{d,i,0} = y^v_{i,0}, \forall s^v \le i \le e^v - 1 \tag{34}$$

$$1 - y^v_{i,0} = \sum_{r \in R} z^v_{i,r}, \forall s^v \le i \le e^v - 1 \tag{35}$$

$$\sum_{\tau=1}^{T^v_i} \left(y^v_{i,\tau} t^{min}_{i,\tau}\right) \le x^v_{i,a}, \forall s^v \le i \le e^v \tag{36}$$

$$x^v_{i,a} \le t_{hor} - \sum_{\tau=1}^{T^v_i} \left[y^v_{i,\tau}\left(t_{hor} - t^{max}_{i,\tau}\right)\right], \forall s^v \le i \le e^v \tag{37}$$

$$\sum_{\tau=1}^{T^v_i} \left(y^v_{d,i,\tau} t^{min}_{i,\tau}\right) \le x^v_{i,d}, \forall s^v \le i \le e^v - 1 \tag{38}$$

$$x^v_{i,d} \le t_{hor} - \sum_{\tau=1}^{T^v_i} \left[y^v_{d,i,\tau}\left(t_{hor} - t^{max}_{i,\tau}\right)\right], \forall s^v \le i \le e^v - 1 \tag{39}$$

$$x^v_{k,a} - x^v_{i,d} \le t_c + \sum_{j=i+1}^{k-1} \sum_{r \in R_c} z^v_{j,r}(t_c + t_r), \forall s^v \le i < k \le e^v, c \in C_1 \tag{40}$$

$$\sum_{j=i}^{k-1} d_j \le t_c + t_c \sum_{j=i+1}^{k-1} \sum_{r \in R_c} z^v_{j,r}, \forall s^v \le i < k \le e^v, c \in C_2 \tag{41}$$

$$x_{k,a}^{v} - x_{s^v,a}^{v} + h_c^v \leq t_c + \sum_{j=s^v}^{k-1} \sum_{r \in R_c} z_{j,r}^{v}(t_c + t_r), \forall s^v < k \leq e^v, c \in C_1 \quad (42)$$

$$h_c^v + \sum_{j=s^v}^{k-1} d_j \leq t_c + t_c \sum_{j=s^v}^{k-1} \sum_{r \in R_c} z_{j,r}^{v}, \forall s^v < k \leq e^v, c \in C_2 \quad (43)$$

$$p_i^v \geq \sum_{k=1}^{T_i^v} CP_{i,k}\left(y_{d,i,k}^{v} - y_{i,k+1}^{v}\right), \forall s^v \leq I \leq e^v - 1 \quad 44)$$

$$x_i^v \in [0, t_{hor}]^2, y_i^v \in \{0,1\}^{T_i^v+1}, \forall s^v \leq I \leq e^v \quad (45)$$

$$y_{d,\,i}^{v} \in \{0,1\}^{T_i+1}, z_i^v \in \{0,1\}^{|R|}, \forall s^v \leq I \leq e^v - 1 \quad (46)$$

$$x_{s^v,d}^{v} \in \left[\underline{t}_{dep}^{v}, \bar{t}_{dep}^{v}\right] \quad (47)$$

$$o_{i,\tau} = \sum_{v \in V}\left(\sum_{k=1}^{\tau} y_{i,k}^{v} - \sum_{k=1}^{\tau-1} y_{d,i,k}^{v}\right), \forall i, \tau \quad (48)$$

Table 1 presents a description of the variables and parameters used. Most variables are defined as in section Truck Driver Scheduling Problem, with the superscript v representing to which vehicle that variable belongs to. Time is divided in time slots forming a partition of the interval $[0,\ t_{hor}]$ and we included an extra set of variables $\{y_d\}$ that track vehicles' departure time slots at each location. The time slots on rest areas are not used to restrict arrival time, but to track parking demand and calculate parking costs. Each time slot at each parking location is assigned a certain cost, which is reflected in the variables $\{CP_{i,k}\}$. A consequence of calculating parking costs like this is that we automatically discretize the set of pure strategies available for each vehicle. The variables $\{x_i^v\}$ representing the arrival and departure times at each location are continuous, so, originally, the TDSP can have an uncountable number of feasible solutions. However, as the different vehicles only interact through parking demand and its effect on parking costs, varying arrival or departure times within a given time slot will not change how the schedule affects parking demand. Therefore, the pure strategies are defined only be the time slots used, not by the detailed schedule described by the arrival and departure times.

Ideally, (28) should be affected by o_i through the term $\sum_{i=s^v}^{e^v} \rho_i^v(\{y\},\{o\})$ as described in section Agent interaction, and the problem should be solved for all vehicles simultaneously, including the equilibrium conditions described in section Non-cooperative Game, in order to find the generated NE. However, this would be intractable. Therefore, we chose to solve the problem iteratively, considering that the drivers and system manager would consider data from the previous iteration when estimating penalties or deciding parking rates. In addition, we assume that any overcrowding penalty can be included in the parking rates and thus stop using the term $\sum_{i=s^v}^{e^v} \rho_i^v(\{y\},\{o\})$. This way, we can solve (28)-(47) for each vehicle in parallel, then use (48) to calculate the parking demands needed to adjust parking rates. When necessary, we will refer to the simplified model that removes the ρ term and solves the optimizes each vehicle independently as the *decoupled model*.

For the purpose of estimating parking demand, we can see each car as a system that receives the vehicle's parameters (initial conditions, start/end locations, departure/delivery constraints and hourly operational cost) and all rest areas' parking rates, and returns the trip duration, parking cost and parking demand for each location and time slot. If the parameters' distribution is known for vehicles of a certain region, it can be used to estimate the costs and parking demand those vehicles generate within the region.

Table 1. Variables and Parameters

Variables

Symbol	Description	Unit
$\left(x_{i,a}^{v}, x_{i,d}^{v}\right)$	Vehicle v's arrival/departure times from location i	h
$\left(y_{i,\tau}^{v}, y_{d,i,\tau}^{v}\right), \tau > 0$	Vehicle v stopped at/departed from location i within time slot τ. (True:1, False:0)	
$y_{i,0}^{v} = y_{d,i,0}^{v}$	Vehicle v does not stop at location i. (True:1, False:0)	
$z_{i,r}^{v}$	Vehicle v takes a rest of type r at location i. (True:1, False:0)	
p_{i}^{v}	Parking cost incurred by vehicle v at location i	$
$o_{i,\tau}$	Location i's parking demand at time slot τ.	

Parameters

Symbol	Description	Unit
d_i	Travel between locations $(i, i+1)$.	h
s^{v}, e^{v}	Start/end locations of vehicle v.	
β^{v}	Vehicle v's hourly operational cost.	$/h
$\left[\underline{t}_{dep}^{v}, \overline{t}_{dep}^{v}\right]$	Vehicle v's departure time-window.	h
t_{hor}	Planning time horizon	h
T_{i}^{v}	Number of time slots considered by vehicle v at location i.	
$\left[t_{i,\tau}^{v,min}, t_{i,\tau}^{v,max}\right]$	Limits considered by vehicle v for time slot of index τ at location i.	h
$CP_{i,k}$	Location i's cumulative parking cost at time slot k. (Sum of the cost of all time slots from 0 to k)	$
R	Set of rest types defined in the HOS regulation.	
C	Set of constraints defined in the HOS regulation.	
$C_1 \subset C$	Set of constraints restricting elapsed time.	
$C_2 \subset C$	Set of constraints restricting accumulated driving time.	
t_c	Time limit related to constraint $c \in C$.	h
$R_c \subset R$	Set of rest types that can reset constraint $c \in C$.	
t_r	Minimum duration for rest of type $r \in R$.	h
h_{c}^{v}	Vehicle v's initial condition relative to constraint c.	h
ρ_{i}^{v}	Function describing the penalty incurred by vehicle v at location i given the time slots used ($\{y_i\}$) and the demand profile o_i.	

Equilibrium

A game where each player's set of strategies is described as the set of feasible time slot usage configurations subject to (29)-(47) constitutes a finite game, and thus it has a Nash equilibrium (Nash, 2020). Furthermore, if we consider the set of penalty functions ρ for which no penalty is applied at parking locations operating at or under capacity, we can say that any solution to the *decoupled model* that keeps all parking facilities working at or under capacity is a NE. Therefore, if we are able to find parking prices for which the resulting *decoupled model* satisfies parking capacity constraints, this solution will be a NE. This happens because all drivers will already be following their optimum route, and the resulting parking demand will not generate penalties that might make drivers switch strategies.

Demand Estimation

Let v be a random vector with known distribution representing the parameters of vehicles within a region of interest. We estimate parking demand by sampling v and solving a TDSP for each sample. The results indicate how this vehicle population tend to behave under the current parking rates. The results can be scaled to match the region's actual traffic volume.

When simulating a region's vehicles we need to define each vehicle's start and destination nodes, the time they ``enter'' de region, any departure or delivery time constraints, and the initial status of drivers' HOS counters. We assume that the origin-destination matrices representing the intent of drivers starting trips at different times/locations is known and can be used to generate reasonable samples of the trips start/end nodes and the time drivers are available to start the trip. In addition, we need to generate the initial state of the drivers' HOS constraints counters. As these counters are not independent, we need to consider their coupling when defining the probability distributions for a region's vehicles. Consider the following counters:

- $\eta^b \in [0, t_{eb}]$: Accumulated driving time since last *break*
- $\eta^r \in [0, t_{ar}]$: Elapsed time since last *daily rest*
- $\psi^r \in [0, t_{er}]$: Accumulated driving time since last *daily rest*
- $\psi^w \in [0, t_{aw}]$: Accumulated on-duty time since last *weekly rest*

where t_{eb}, t_{ar}, t_{er}, and t_{aw} represent the time limits defined by the regulation for the associated constraints. *break*, *daily rest* and *weekly rest* refer to off-duty periods of at least 30min, 10h and 34h, respectively. Note that the regulation regarding 30min *breaks* restricted the elapsed time since the last *break* when (Vital & Ioannou, 2019) was written, but it was recently changed to limit the accumulated driving time since the last *break*. After accounting for how different activities affect each counter, we have that the following inequalities must hold:

$$\eta^b \leq \psi^r \leq \eta^r, \qquad \psi^r \leq \psi^w \tag{49}$$

The convex polyhedron defined by these inequalities is the set of possible initial conditions for HOS counters and can be used to generate valid initial conditions. One possible approach is to use rejection sampling to uniformly sample the polyhedron. We can also generate samples by sampling one resource at a time, using the sampled value for the smaller variables to define the domain of the larger ones, i.e.:

$$\eta^b \in [0, t_{eb}], \qquad \psi^r \in [\eta^b, t_{ar}] \tag{50}$$

$$\eta^r \in [\psi^r, t_{er}], \qquad \psi^w \in [\psi^r, t_{aw}] \tag{51}$$

This method does not sample the polyhedron uniformly, but we do not need to check the validity of generated samples.

Sensitivity to HOS conditions and uniform time slot prices

The following figures were generated by independently simulating 100 vehicles with the same origin, destination, departure/delivery time constraints. The trip has 23h of driving time and evenly spaced rest areas with 1h of travel time between any two adjacent rest areas. Each set of initial conditions was generated by sampling random integers from the intervals described until values satisfying (49) were generated. Figure 3 presents the parking demand generated without parking charges and with HOS initial conditions restricted to $\{0,1\}$. Even though there is a small number of valid initial conditions and no influence from parking prices we already see the impact that even small changes in the initial conditions have on parking demand.

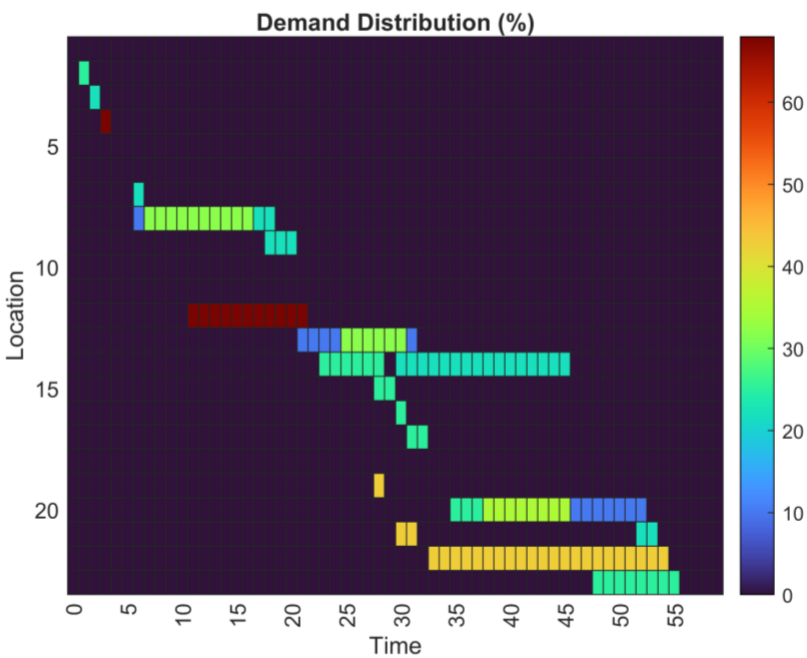

Figure 3. Distribution of the parking demand generated by 23h trips without parking charges and with HOS initial condition restricted to {0,1}.

Figure 4 presents the parking demand generated with hourly operational costs (in $/h) are integers sampled from $[60,80]$, and each HOS initial condition is an integer between 0 and its regulation limit minus 1h, e.g., $\eta^b \in [0, t_{eb} - 1]$. Figure 4a does not charge for parking and Figure 4b uses parking charges of $5 per time slot. Despite charging the same value at all times and locations, as the parking costs are calculated per time slot, setting a non-zero parking rate already serves to discourage drivers from taking short stops. Consequently, Figure 4b has less narrow demand peaks compared to Figure 4a.

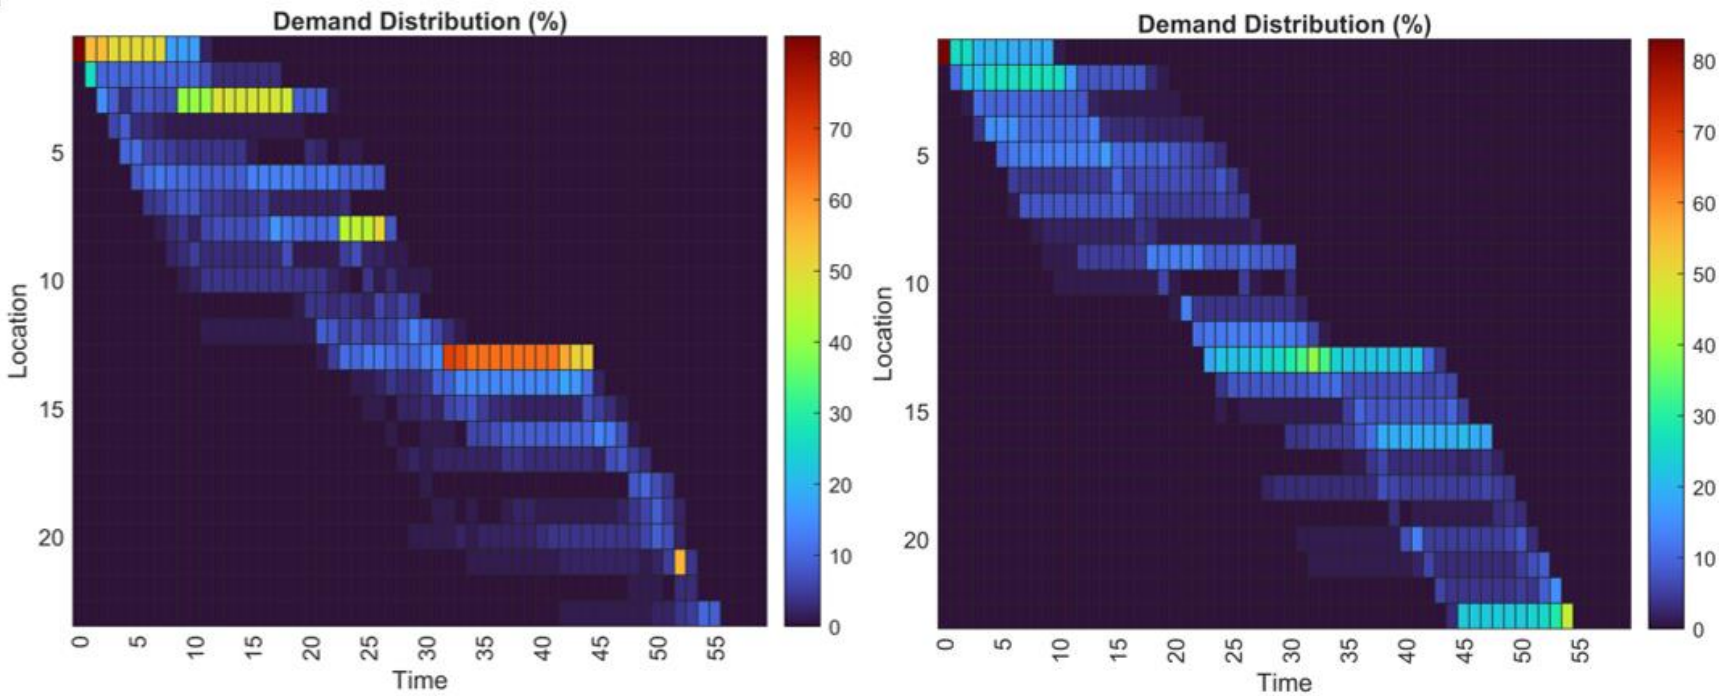

(a) Parking rate: $0

(b)Parking rate: $5 per time slot

Figure 4. Distribution of the parking demand generated by 23h trips with hourly operational cost in the interval $[60, 80]$ and with each HOS initial condition varying between 0 and its regulation limit minus 1h.

Response to price changes

Price Update

We consider keeping truck stops demand below capacity as the main objective. Consider a simple pricing strategy of increasing parking prices at locations/times where demand is above an upper threshold and decreasing it when demand is below a lower threshold. To avoid parking rates diverging, prices will be restricted to a certain range. Let O and T be matrices representing, respectively, the estimated demand and target demand for each location and time slot and let $[\mathrm{l}, u]$ be an interval used to define the range of demand values that do not required direct intervention (deadband). We define the error matrix E as follows:

$$E_{i,j} = \begin{cases} 0, if\ \ell \le \dfrac{O_{i,j}}{T_{i,j}} \le u \\ O_{i,j} - T_{i,j}, o.w. \end{cases} \tag{52}$$

The error matrix is used to feed an integral controller. As prices are restricted to a given range, a back calculation anti-windup method is used to correct the integrator state.

Experiment

Consider a route with 23 locations where the travel time between any pair of adjacent locations is 1h. We define 3 vehicle populations with the parameters listed in Table 2.

Table 2. Population Parameters

	All	**P1**	**P2**	**P3**
Marginal Hourly Cost ($/h)	[50, 70]			
Arrival Time (h)	[12, 18]			
Max Wait (h)	12			
Start Location		0	7	5
End Location		[5, 15]	[14, 20]	[17, 23]
Resource 1 (h)		[3, 7]	[0, 7]	[0, 7]
Resource 2 (h)		[3, 13]	[0, 13]	[0, 13]
Resource 3 (h)		[3, 10]	[0, 10]	[0, 10]
Resource 4 (h)		[3, 30]	[0, 30]	[0, 30]
Size		100	100	30

Resources 1, 2, 3 and 4 refer to, respectively, the counters for the HOS rules' 8h (30min break) driving limit, 14h limit, 11h driving limit and 60h on-duty time limit. Parameters defined by an interval are sampled uniformly from that interval. All vehicles consider daily delivery time-windows of $[8,18]$ at the end node.

Consider that each location represents a region with a finite truck parking capacity equal to a certain percentage of the total truck population being considered. In this example we set the capacity to 12% of the truck population, i.e., no location can accommodate more than 27 trucks at any given time. Parking demand is not considered at the end node, but it is counted at the start node. The deadband was set to 10-100% of the parking capacity at each location. Parking rates are restricted to the interval $[-20,20]$. Parking demand converged to values below capacity at every location after 14 iterations. Plots of the demand and parking price at each location and time slot for some time steps are presented in Figure 5. Table 3 presents the average time cost, parking cost, total cost and the maximum parking demand at each step of the experiment.

Table 3. Experiment Results

Step	Time Cost	Parking Cost	Total Cost	Max Demand
0	1571.18	0	1571.18	0.182609
1	1571.18	-0.0340742	1571.15	0.23913
2	1571.18	-0.329091	1570.86	0.321739
3	1571.18	-0.744367	1570.44	0.334783
4	1571.18	-0.751305	1570.43	0.326087
5	1571.18	-0.391361	1570.79	0.265217
6	1571.18	-0.103988	1571.08	0.217391
7	1571.18	0.00304427	1571.19	0.191304
8	1571.18	0.130596	1571.32	0.191304
9	1571.18	0.198699	1571.38	0.195652
10	1571.18	0.240001	1571.42	0.16087
11	1571.18	0.281829	1571.47	0.134783
12	1571.18	0.291691	1571.48	0.134783
13	1571.18	0.302905	1571.49	0.126087
14	1571.18	0.302278	1571.49	0.117391

Driver Sensitivity

In this example, the scheduling model showed very high sensitivity to parking prices. The scenario considers marginal hourly costs in the range of $50-$70, but the highest parking rate used to balance demand was less than $1/h. Table 3 shows that the average parking cost is less than 0.02% of the trip's average total cost. Nevertheless, it is important to note that required parking rates depend on the parking capacity of the different locations, the number of vehicles and their parameters. As this example considers only 3 different vehicle populations, a large part of the parking spaces are not utilized, making it easier for demand to be redistributed.

Transient Behavior

The controller succeeds in adjusting the demand to within the target level. However, parking demand oscillates significantly in the beginning, reaching levels higher than in the original (uncontrolled) setting. Table 3 shows that the maximum demand increases from 18% to 33% of the truck population in steps 0 to 3, before starting to decrease. It is only at step 10 that we see an improvement relative to the original condition. If this strategy were to be applied directly to a region, it could disrupt the system for a while before it starts showing benefits. Therefore, if we can use this kind of simulation to estimate the behavior of a region's truck population, we can mitigate these transient effects.

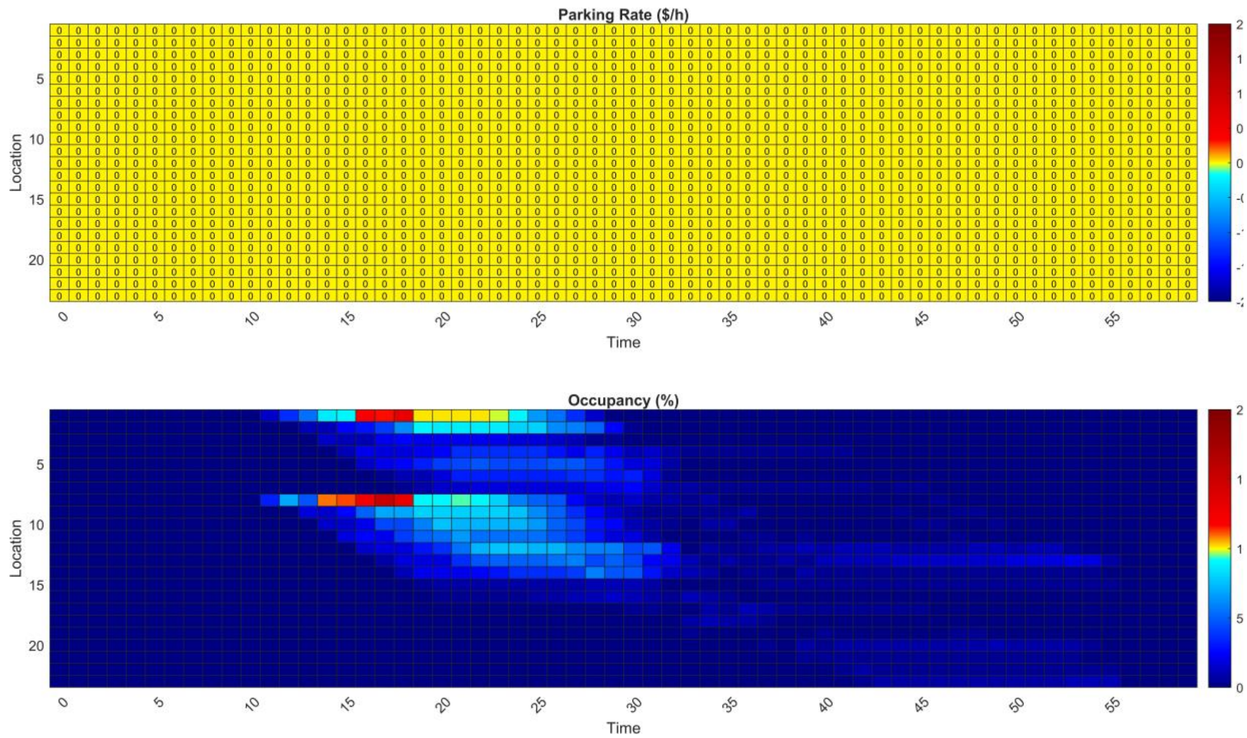

t=0

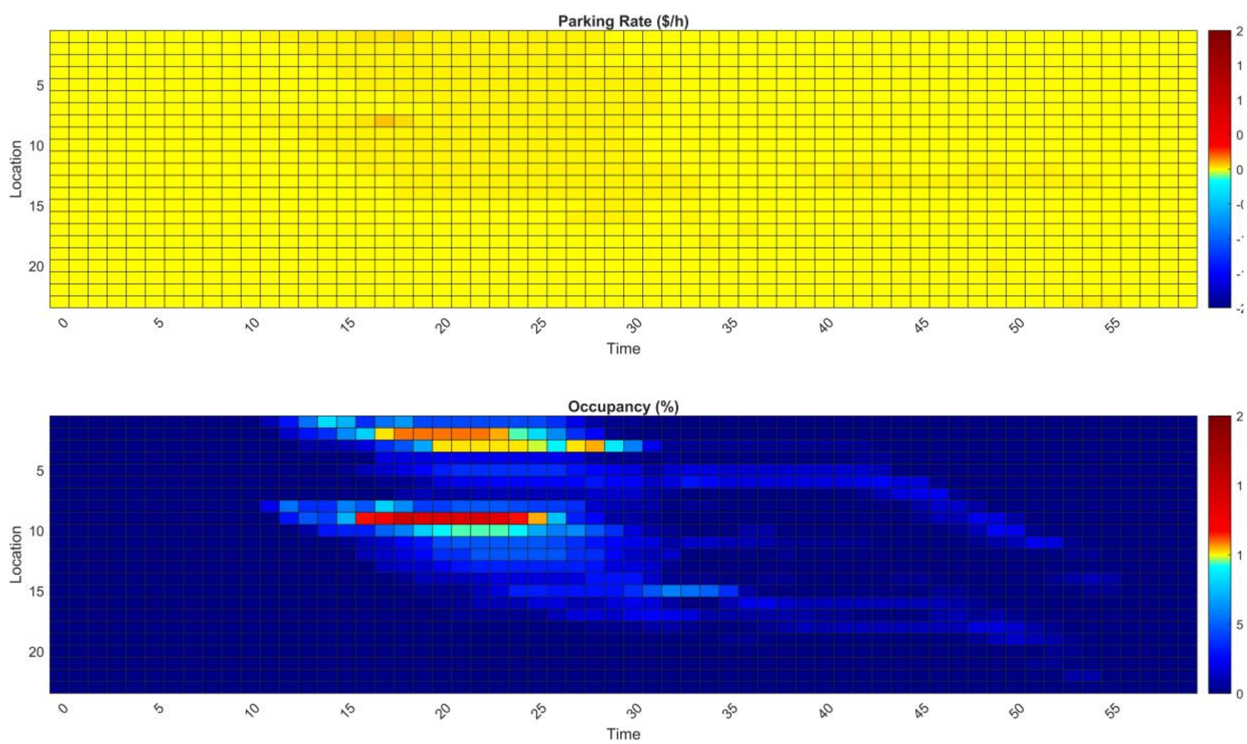

t=1

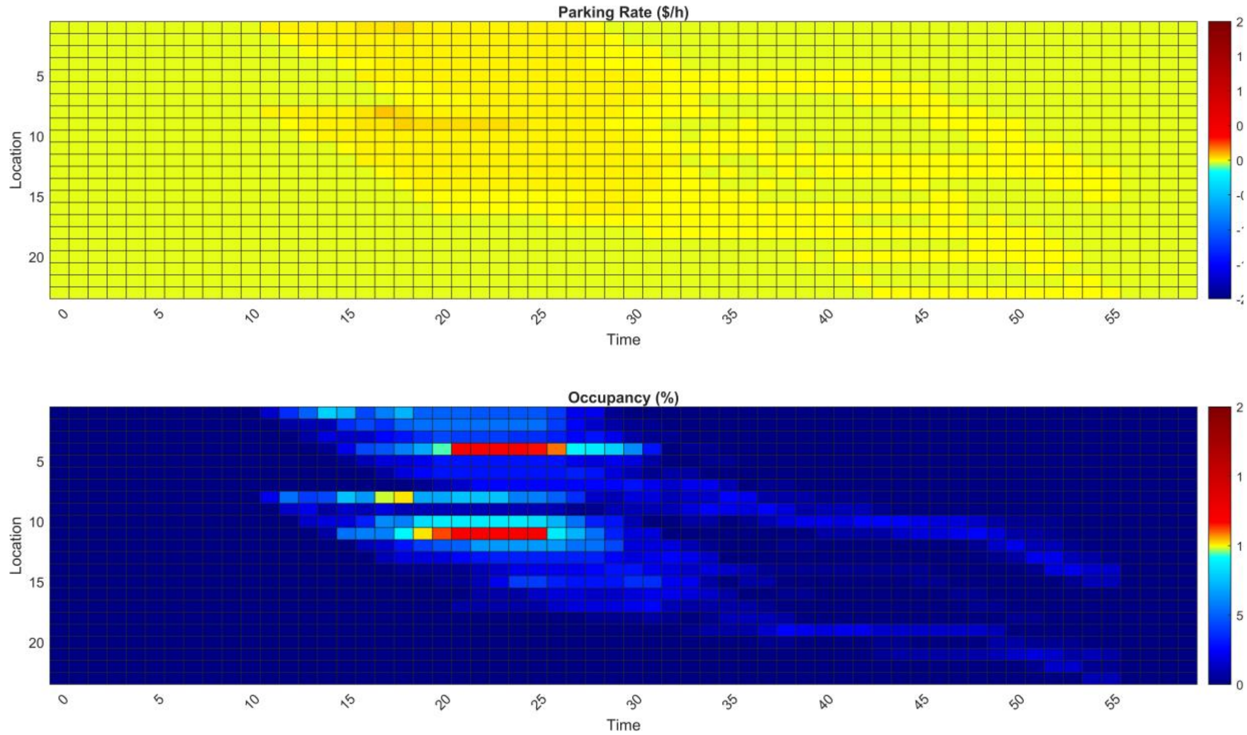

Parking Rate ($/h)
Location
Time
Occupancy (%)
Location
Time

t=2

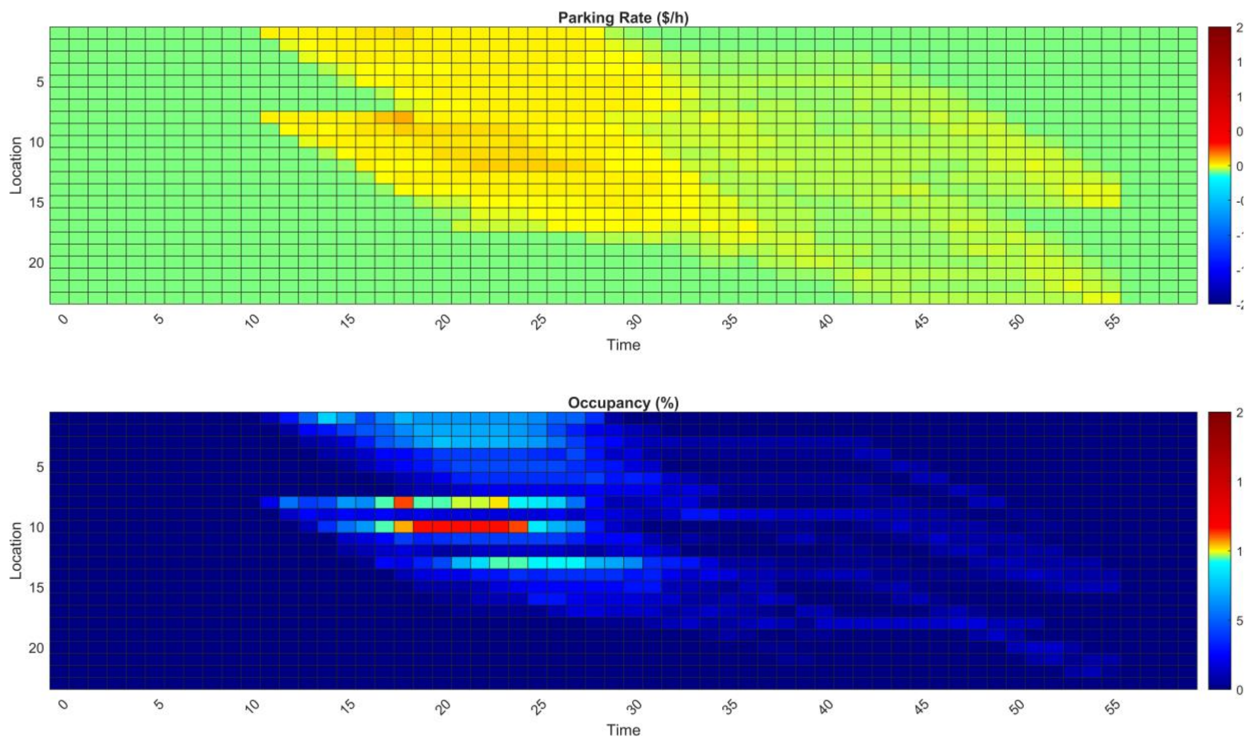

Parking Rate ($/h)
Location
Time
Occupancy (%)
Location
Time

t=7

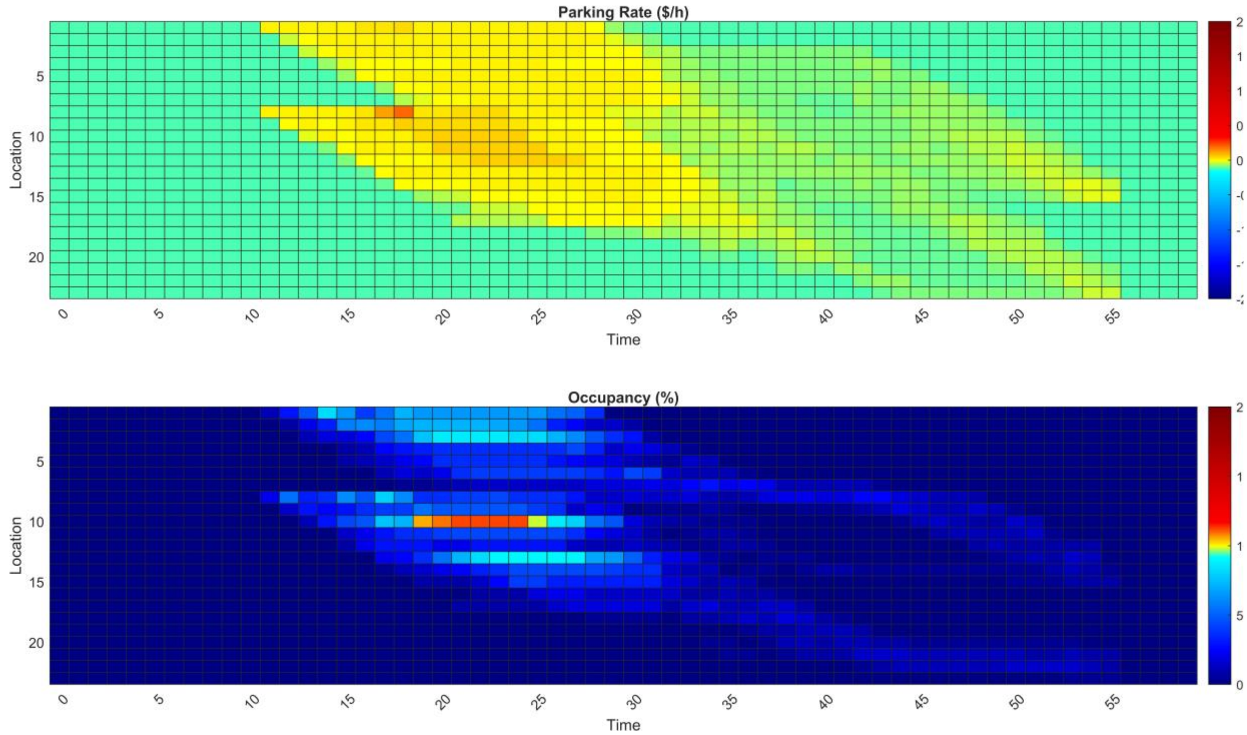

t=10

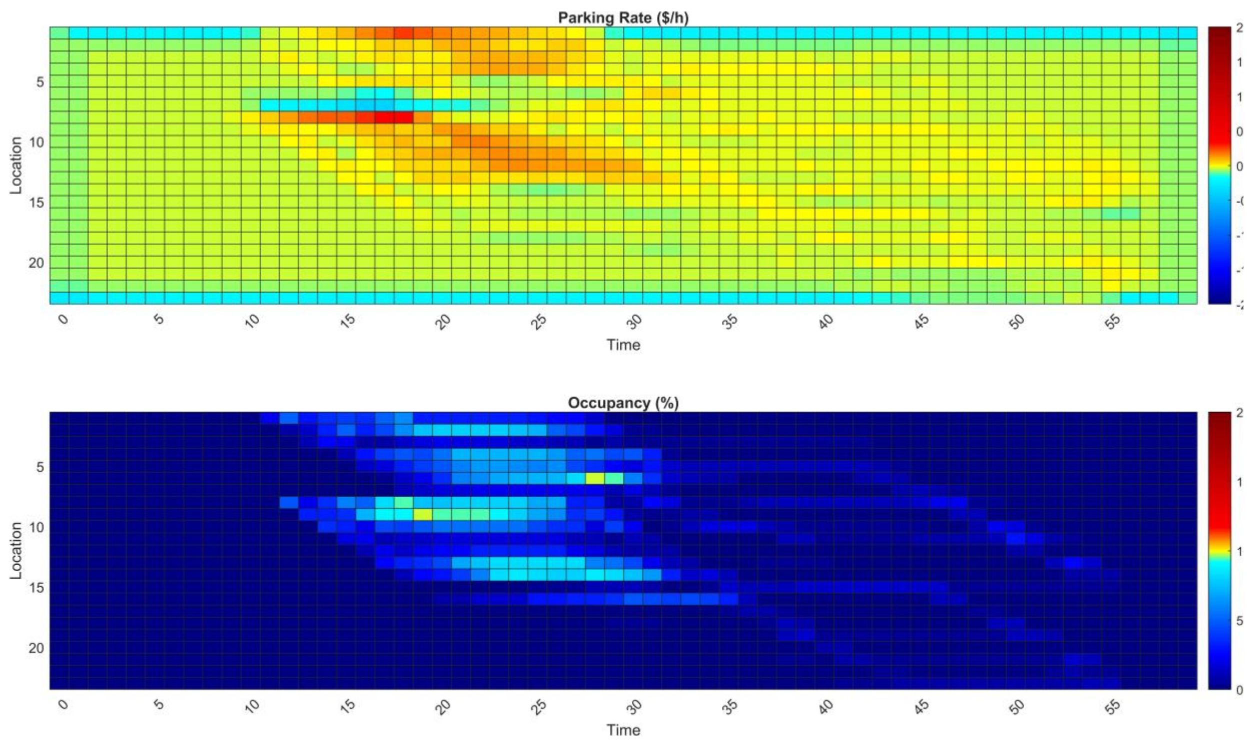

t=14

Figure 5. Example results: Demand and prices for time steps 0, 1, 2, 7, 10 and 14.

Conclusion

In this project, we consider the issue of coordinating the parking decisions of a large number of trucks. More specifically, we study how to model the behavior of a region's driver population and how it could be influenced. Understanding how truck parking demand is affected by the interaction of individual drivers' selfish planning behaviors (in the sense that they minimize their own costs, not the overall system cost) and how parking prices affect optimal schedules are important steps in developing a system able to balance demand.

We present a formulation that uses a modified TDSP (Truck Driver Scheduling Problem) mixed-integer programming model which tracks parking usage by dividing time into time-slots and charging drivers per time slot used. The parking rates depend on both the time and location being considered. Due to the complexity of solving the problem for a large number of drivers and modeling how to manage parking availability and drivers interactions within a single optimization problem, we calculate each drivers' schedule separately using a common price matrix. We assume that penalties regarding parking difficulties faced in some locations and times could be included in the parking prices. If the pricing strategy used can find a price matrix that avoids overcrowding, then the solution generated by this model will also be a solution to the problem when all drivers and their interactions a modeled as a single optimization problem. Assuming we have information on the usual conditions of truck drivers operating in a region (or a large number of drivers/companies willing to keep their planning system continuously connected to the pricing coordinator), the system would use a sample population to simulate the effect of price changes before actually implementing the new prices.

Simulations were used to study how drivers react to price changes. Results show that the scheduling model is very sensitive to even small changes in parking prices, which is conducive to using parking prices as a means to influence demand. However, under the price update rule tested, the system's demand distribution oscillates significantly before reaching a valid solution and the initial iterations might see an increased maximum demand instead of the intended demand redistribution. In summary, if we consider that truck drivers are following optimal schedules, then parking prices would be effective in changing which locations and time slots would be chosen by each driver. However, price adjustments can cause demand to shift in unexpected and not always beneficial ways, likely due to HOS regulations and client constraints limiting the possible alternative schedules. Therefore, further study is required in order to better understand the system's properties and how to avoid or dampen these oscillations. Furthermore, due to HOS rules and client constraints, it might be impossible to divert demand from certain time slots and locations sufficiently. Nevertheless, this model could still aid in identifying these spots and contribute to the evaluation of infrastructure investment needs. It is important to note that we assume that drivers would be using a planning tool to optimize their itineraries and accept the routes generated. Drivers might have their own preferences regarding itineraries and do not necessarily follow optimal schedules.

References

American Transportation Research Institute. (2018). *MAASTO Truck Parking Survey Analysis* (Issue May).

American Transportation Research Institute. (2019). *Critical Issues in the Trucking Industry – 2019* (Issue October).

Black, J. W., & Tyagi, R. (2010). Potential problems with large scale differential pricing programs. *IEEE PES T&D 2010*, 1–5. https://doi.org/10.1109/TDC.2010.5484309

Boris, C., & Brewster, R. (2018). A Comparative Analysis of Truck Parking Travel Diary Data. *Transportation Research Record: Journal of the Transportation Research Board*, *2672*(9), 242–248. https://doi.org/10.1177/0361198118775869

Capdevila, M., Tomas, V. R., Garca, L. A., & Farron, M. P. (2013). Dynamic management of parking spaces in road rest areas with automatic negotiation. *Proceedings - 2013 IEEE International Conference on Systems, Man, and Cybernetics, SMC 2013*, 3609–3614. https://doi.org/10.1109/SMC.2013.615

Chen, R., Qian, X., Miao, L., & Ukkusuri, S. V. (2020). Optimal charging facility location and capacity for electric vehicles considering route choice and charging time equilibrium. *Computers and Operations Research*. https://doi.org/10.1016/j.cor.2019.104776

Christodoulou, G. (2008). Price of Anarchy. In M.-Y. Kao (Ed.), *Encyclopedia of Algorithms* (pp. 665–667). Springer US. https://doi.org/10.1007/978-0-387-30162-4_299

de Almeida Araujo Vital, F., Ioannou, P., & Gupta, A. (2020). Survey on Intelligent Truck Parking: Issues and Approaches. *IEEE Intelligent Transportation Systems Magazine*. https://doi.org/10.1109/MITS.2019.2926259

Dong, J. (2008). Anticipatory traveler information and pricing for real -time traffic management. *ProQuest Dissertations and Theses*, *June*.

Doulamis, N., Protopapadakis, E., & Lambrinos, L. (2013). Improving service quality for parking lot users using intelligent parking reservation policies. *Proceedings - 27th International Conference on Advanced Information Networking and Applications Workshops, WAINA 2013*, 1392–1397. https://doi.org/10.1109/WAINA.2013.219

Dutta, G., & Mitra, K. (2017). A literature review on dynamic pricing of electricity. *Journal of the Operational Research Society*, *68*(10), 1131–1145. https://doi.org/10.1057/s41274-016-0149-4

Falowo, O. E., Zeadally, S., & Chan, H. A. (2009). Dynamic pricing for load-balancing in user-centric joint call admission control of next-generation wireless networks. *International Journal of Communication Systems*, *23*(3). https://doi.org/10.1002/dac.1062

Gabr, A. Z., Helal, A. A., & Abbasy, N. H. (2018). Dynamic pricing; different schemes, related research survey and evaluation. *2018 9th International Renewable Energy Congress (IREC)*, *Irec*, 1–7. https://doi.org/10.1109/IREC.2018.8362562

Geng, Y., & Cassandras, C. G. (2011). Dynamic resource allocation in urban settings: A “smart parking” approach. *Proceedings of the IEEE International Symposium on Computer-Aided Control System Design*, 1–6. https://doi.org/10.1109/CACSD.2011.6044566

Hollander, Y., & Prashker, J. N. (2006). The applicability of non-cooperative game theory in transport analysis. *Transportation*, *33*(5), 481–496. https://doi.org/10.1007/s11116-006-0009-1

Kaufman, D. E., Smith, R. L., & Wunderlich, K. E. (1991). An iterative routing/assignment method for anticipatory real-time route guidance. *Vehicle Navigation and Information Systems Conference, 1991*, 693–700. https://doi.org/10.1109/VNIS.1991.205814

Kordonis, I., Dessouky, M. M., & Ioannou, P. A. (2020). Mechanisms for Cooperative Freight Routing: Incentivizing Individual Participation. *IEEE Transactions on Intelligent Transportation Systems*, *21*(5), 2155–2166. https://doi.org/10.1109/TITS.2019.2915549

Martin, E. W., & Shaheen, S. A. (2013). Truck Parking and Traffic on I-5 in California: Analysis of a Clipboard Survey and Annual Average Daily Traffic Data. *TRB 92nd Annual Meeting Compendium of Papers*.

Moradipari, A., & Alizadeh, M. (2020). Pricing and Routing Mechanisms for Differentiated Services in an Electric Vehicle Public Charging Station Network. *IEEE Transactions on Smart Grid*, *11*(2), 1489–1499. https://doi.org/10.1109/TSG.2019.2938960

Nash, J. (2020). NON-COOPERATIVE GAMES. In *Classics in Game Theory* (pp. 14–26). Princeton University Press. https://doi.org/10.2307/j.ctv173f1fh.8

NCDOT. (2017). *North Carolina Statewide Multimodal Freight Plan: Truck Parking Study* (Issue January).

Papadopoulos, A.-A., Kordonis, I., Dessouky, M., & Ioannou, P. (2019). Coordinated Freight Routing With Individual Incentives for Participation. *IEEE Transactions on Intelligent Transportation Systems*, *20*(9), 3397–3408. https://doi.org/10.1109/TITS.2018.2876326

Rodier, C. J., Shaheen, S. A., Allen, D. M., & Dix, B. (2010). *Commercial Vehicle Parking in California : Exploratory Evaluation of the Problem and Solutions* (Issue March).

Song, R. (2012). Pre-shifting customer behavior in response to dynamic pricing events. *2012 IEEE Power and Energy Society General Meeting*, 1–3. https://doi.org/10.1109/PESGM.2012.6345401

Tian, Q., Yang, L., Wang, C., & Huang, H. J. (2018). Dynamic pricing for reservation-based parking system: A revenue management method. *Transport Policy*, *71*(August), 36–44. https://doi.org/10.1016/j.tranpol.2018.07.007

Tucker, N., Ferguson, B., & Alizadeh, M. (2019). An online pricing mechanism for electric vehicle parking assignment and charge scheduling. *Proceedings of the American Control Conference*, *2019-July*, 5755–5760. https://doi.org/10.23919/acc.2019.8814409

Turan, B., Pedarsani, R., & Alizadeh, M. (2019). *Dynamic Pricing and Management for Electric Autonomous Mobility on Demand Systems Using Reinforcement Learning*. 1–14. http://arxiv.org/abs/1909.06962

U.S. Department of Transportation. (2015). *Jason's Law Truck Parking Survey Results and Comparative Analysis* (Issue August).

Vital, F., & Ioannou, P. (2021). Scheduling and shortest path for trucks with working hours and parking availability constraints. *Transportation Research Part B: Methodological*, *148*, 1–37. https://doi.org/10.1016/j.trb.2021.04.002

Vital, F., & Ioannou, P. (2019). Long-Haul Truck Scheduling with Driving Hours and Parking Availability Constraints. *2019 IEEE Intelligent Vehicles Symposium (IV)*, *June*, 620–625. https://doi.org/10.1109/IVS.2019.8814011

Xie, Z., Sharath, N., & Wang, C. (2015). A game theory based resource scheduling model for cost reduction in home health care. *2015 IEEE International Conference on Industrial Engineering and Engineering Management (IEEM)*, *2016-Janua*, 1800–1804. https://doi.org/10.1109/IEEM.2015.7385958

Yang, H., Shao, C., Wang, H., & Ye, J. (2020). Integrated reward scheme and surge pricing in a ridesourcing market. *Transportation Research Part B: Methodological*, *134*, 126–142. https://doi.org/10.1016/j.trb.2020.01.008

Data Summary

Products of Research

The data generated are simulation data presented in plots and tables in the final report.

Data Format and Content

The data is presented as tables and plots in the final report, and consists of occupancy rates, and average trip durations and costs.

Data Access and Sharing

The data is included in the final report.

Reuse and Redistribution

The data is published as part of the final report.