

International Journal of Logistics Systems and Management

ISSN online: 1742-7975 - ISSN print: 1742-7967
<https://www.inderscience.com/ijlsm>

Development and application of a cost-driven decision model for store replenishment logistics in the fast-food sector

Alejandro Vigo Camargo, Yavuz A. Bozer

DOI: [10.1504/IJLSM.2022.10057812](https://doi.org/10.1504/IJLSM.2022.10057812)

Article History:

Received:	27 December 2021
Accepted:	23 June 2022
Published online:	13 January 2025

Development and application of a cost-driven decision model for store replenishment logistics in the fast-food sector

Alejandro Vigo Camargo* and Yavuz A. Bozer

Department of Industrial and Operations Engineering,

University of Michigan,

1205 Beal Avenue, Ann Arbor, MI 48109, USA

Email: avigo@umich.edu

Email: yabozer@umich.edu

*Corresponding author

Abstract: In this paper, we address the store replenishment problem (SRP) in the fast-food sector, which is concerned with minimising the logistics costs associated with replenishing stores in a network over a fixed time horizon. We develop a cost function that captures transportation, labour, truck availability, and route-time overage costs and compare its performance against traditional, one-dimensional objectives, such as minimising only the travel distances or the fleet size. We formulate the SRP as a mixed-integer program and present a modified nearest-neighbour-based clustering heuristic to pre-generate a set of potential delivery routes. The model concurrently determines the fleet size, the delivery routes, and chooses between single-driver and team routes. Using real-world data from our industry collaborator we show that the proposed heuristic yields solutions with lower costs than the industry baseline. Furthermore, we show that superior results are obtained from an objective with multiple cost components compared to traditional one-dimensional objectives.

Keywords: store replenishment; clustering heuristic; supply chain logistics; fast-food supply chain.

Reference to this paper should be made as follows: Vigo Camargo, A. and Bozer, Y.A. (2025) 'Development and application of a cost-driven decision model for store replenishment logistics in the fast-food sector', *Int. J. Logistics Systems and Management*, Vol. 50, No. 1, pp.69–93.

Biographical notes: Alejandro Vigo Camargo is a PhD candidate in the Industrial and Operations Engineering (IOE) Department at the University of Michigan, with a focus in Supply Chain Logistics. He received his BS in Industrial Engineering from the University of Puerto Rico – Mayaguez (2015), and an MSE (2017) in IOE from the University of Michigan. His current research is focused in using operations research tools to optimise the outbound logistics of the food supply chain. His other areas of interest include inventory management and transportation.

Yavuz A. Bozer received his BSIE from the Middle East Technical University, Turkey and MS and PhD in Industrial and Systems Engineering from Georgia Institute of Technology. Currently, he is an IOE Department Professor at the University of Michigan. His research interests focus on the development of quantitative models to design and improve material handling systems. He is a leading researcher in facility logistics and design, and the co-author of

Facilities Planning (Wiley). His recent work includes device dispatching in material handling systems, patient transport in large hospitals, order picking systems in online retailing, and multi-node food supply networks.

1 Introduction

Although consumer spending at fast-food restaurants reached \$310 billion in 2017 (USDA, 2018), there is a scarcity of academic research on the supply chain logistics of this sector. A significant portion of food supply related literature is focused on production, with applications in the agriculture industry, and on reducing food waste; see, for example Lemma et al. (2014). While production and waste reduction are relevant, the highly-competitive fast-food sector relies on the timely and efficient replenishment of their stores, which is fertile ground for academic research.

This study is focused on the fast-food sector, where customer loyalty, driven by price and quality of the brand, is key for success (Shokri et al. 2014). The store networks for fast-food companies often span large regions or countries, and may involve thousands of locations. For example, as of 2018, Domino's Pizza reports over 5,400 stores in the US, with plans to expand to about 8,000 locations (Klein, 2018). Similarly, Yum! Brands operates over 43,000 locations globally, including 4,062 KFC stores, 7,447 Pizza Hut stores, and 6,611 Taco Bell stores in the US alone (Yum! Brands, nd). Given the large number of stores, their geographical spread, and the need to replenish them frequently, the store network is often divided into regions, with a distribution centre (DC) serving each region. Usually the store replenishment operations are coordinated through a company's central planning group, using their own DCs and trucks/drivers. Even when such operations are outsourced, using a central DC for each region is fairly common.

Similar settings are also found in the grocery-food sector. However, grocery stores carry a much larger number and variety of stock keeping units (SKUs), and in many cases (such as wine, soft drinks, and dairy products) they use a *hybrid model* where some goods are delivered through a DC while others are delivered by the vendors through a direct store delivery (DSD) model, which represents about 25% of their sales (GMA, 2008). Some of the major grocery-food suppliers/producers are moving away from the DSD model, opting instead to ship their goods to their customers' DCs (Staff, 2019). While such a trend brings the grocery-food sector closer to our study, given the above differences in scale and variety, and the fact that DSD is still common for groceries, we limit our attention to the fast-food sector.

Store replenishment is a primary source of cost (Akkerman et al., 2010). The travel costs between the DC and the stores is a significant contributor to the overall cost, which is perhaps why a majority of the publications reviewed in a survey of routing papers by Braekers et al. (2016) have distance-dependent objective functions. However, other factors, such as the labour cost associated with delivering inventory, the fleet cost, and whether routes are driven by a single driver or a team of two, can also play a major role in determining the total replenishment cost. Balancing the above costs while making resource allocation decisions and meeting multiple time/distance constraints is often a complex challenge, which perhaps persuaded some to use simpler, one-dimensional objectives. While distance minimisation alone may be suitable for some routing problems, for the fast-food store replenishment problem (SRP), using a one-dimensional

objective neglects key trade-offs that may lead to inferior solutions. In this paper we address the above challenge by developing a cost-based model that incorporates distance-based costs as well as cost of delivery labour and other cost elements into the objective function to minimise the overall cost of replenishment logistics. Using real-world data from our industry collaborator, we show how such an approach leads to better solutions and provides key managerial insights to improve competitiveness.

1.1 Objective and contribution

We present a cost-driven model to minimise the logistics costs associated with the SRP, which we define as determining the number of trucks, the delivery routes, and which routes will be single- or two-driver (team) routes to support a given delivery plan (i.e., set of orders for each store in a region) over a fixed time horizon, with a known DC location and known distances. In addition to fleet-sizing and vehicle routing, the model considers single vs team routes and the unloading labour. The objective is to minimise the total replenishment cost to satisfy store orders while accounting for route-time and truck-capacity constraints.

Given a delivery plan, the SRP is structurally similar to a vehicle routing problem (VRP), since the main concern is the construction of low-cost routes to deliver the necessary inventory to the stores. However, the SRP is also concerned with minimising the fleet cost and the labour cost associated with unloading inventory. Additionally, the SRP considers the added cost incurred by routes that extend beyond their time limits (overage routes). Unlike traditional vehicle routing models, the proposed model considers single-driver and team routes, which are a common practice (Goel et al., 2019). Although the same truck might be used for a single- or team-route, the limits on route length and usable truck capacity are different since team routes can complete longer routes and deliver larger quantities. Accordingly, single- and team-routes have a specific cost structure and constraint set associated with each. Capturing the interplay between the above costs is an essential component of the SRP.

We formulate the SRP as a mixed-integer programming (MIP) model, with an objective that incorporates multiple cost elements. Due to complex routing decisions and constraints, a set of pre-generated potential routes are used as input. To generate the routes, we utilise a simple clustering heuristic. The proposed model, including the objective function, was developed with extensive input from the supply chain team of a national fast-food company who was our collaborator for the study.

The main contributions of the paper are twofold. First, we developed a cost-based objective that incorporates multiple, relevant cost elements associated with fast-food store replenishment logistics. To the best of our knowledge, this is the first store replenishment study that incorporates single/team-driver routes, route overage costs, and fleet sizing decisions concurrently. We show that our cost function outperforms one-dimensional objective functions, frequently adopted in routing problems. Second, through our computational results, using real-world data and a series of sensitivity analyses, we obtain insights that can be critical for management in the fast-food sector. Furthermore, the proposed approach allows large-scale problems to be solved with a standard solver such as CPLEX.

The remainder of the paper is organised as follows. In Section 2, we review the relevant literature. The MIP formulation and the cost function are presented in Section 3. The route-generating clustering heuristic is described in Section 4. In Section 5, we

discuss the computational results and present benchmark comparisons. The sensitivity analyses we performed and the results are presented in Section 6. Lastly, Section 7 presents our conclusions and possible future research directions.

2 Literature review

Most problems that involve routing are modelled as a VRP, which is an extensively-researched, NP-hard problem (Lenstra and Kan, 1981), with books and survey articles devoted to the topic; see, for example Sharda et al. (2008), Laporte (2009), Toth and Vigo (2014), and Braekers et al. (2016). Food distribution has been one of the many applications for the VRP. For example, Tarantilis and Kiranoudis (2001) focused on perishable foods and developed a metaheuristic to minimise the distance-based cost for milk distribution. In a related paper, Tarantilis and Kiranoudis (2002) developed an algorithm to minimise the total distance travelled in a meat distribution application.

As VRP research expanded, more variants emerged. A recent ‘concise review’ of emerging VRP variants cites over 300 papers (Vidal et al., 2019). Two extensions of the VRP that are relevant to the SRP are the Period(ic) VRP (PVRP) and the multi-period VRP (MVRP). The PVRP focuses on selecting a delivery schedule from a predefined set of schedules for each node in a network, and building routes to support the selected delivery schedule (Francis et al., 2008). In the classical PVRP, the customers must be visited with certain frequencies during the time horizon. While both the SRP and PVRP include routing decisions over multiple periods, as explained below, there are specific differences between them.

The PVRP was introduced by Beltrami and Bodin (1974), and its first formulation was presented in Christofides and Beasley (1984) as an integer program (IP). Building upon the IP formulation of the PVRP, Cordeau et al. (1997) proposed a Tabu search heuristic to solve a mixed-integer formulation of the PVRP. Similarly, in Chao et al. (1995), which was later expanded by Gulczynski et al. (2011), the authors propose a two-step heuristic that first solves an IP to develop an initial delivery schedule, and then uses an improvement heuristic to find an ideal routing and schedule. Additionally, other heuristics developed originally for the VRP were extended to the PVRP by Alegre et al. (2007) and Hemmelmayr et al. (2009).

Francis et al. (2006) presented an approximating method to solve the PVRP with ‘service choice’, which allows the model to choose the frequency of visit for each customer, and change the structure of the delivery schedule accordingly. However, as is the case with most of the PVRP literature, the delivery frequencies are based on predetermined patterns (such as daily deliveries, a delivery every other day, etc.), narrowing the possibilities. Unlike the PVRP, the SRP assumes that the delivery schedule over the time horizon has been fixed and it does not necessarily follow a predetermined pattern.

Virtually all the PVRP papers assume the number of trucks is given, and they use it to limit the number of routes completed in each period. In contrast, the SRP treats the number of trucks as a decision variable. One exception is Gaudioso and Paletta (1992), who propose a heuristic to solve the PRVP with the main objective of minimising the number of trucks required. Their model is similar to the SRP in that the number of trucks is a decision variable and the maximum travel time for each route is limited. However, the SRP does not focus only on the number of trucks; it is concerned with minimising the

overall cost, which includes travel and labour cost in addition to the truck cost. Another key difference is that, although we also limit the maximum route-time, the SRP allows a route to exceed the limit by imposing an overage cost.

The MVRP, on the other hand, which was formally defined by Archetti et al. (2015), shares many similarities with the PVRP, the main one being its goal to build routes for a set of customers over multiple time periods. While the MVRP also selects a delivery schedule, it assumes the customers have delivery due dates, and in contrast to the PVRP, there is no fixed frequency for the deliveries. Both Archetti et al. (2015) and Larrain et al. (2019) formulated the MVRP with the objective of minimising the transportation cost, considering inventory holding and a late-delivery penalty. In López-Santana et al. (2018), the authors present a three-phase approach for the MVRP to minimise the total travel time. Wen et al. (2010) studied a dynamic extension of the MVRP, with customer demands updated each day. The authors present a three-phase heuristic to solve a multi-objective formulation of the problem that tries to minimise travel distance and customer wait time, while balancing work. Although the MVRP has a broader decision scope than the PVRP, the differences we described between the SRP and PVRP apply to the MVRP as well, particularly the fixed schedule, the number of trucks being a decision variable, and our multi-component cost function.

To the best of our knowledge, none of the published PVRP or MVRP studies explicitly consider single- vs. team-driving routes. Our formulation explicitly considers them, which bears some similarity to the heterogeneous VRP (HVRP). Introduced by Golden et al. (1984), the HVRP focuses on route optimisation with a fleet of vehicles of varying parameters, such as vehicle capacity. While the SRP assumes a homogeneous fleet of vehicles, the cost, the (effective) capacity, and the length limit of each route depends on whether it is a single-driver or team route. This is comparable to choosing between two types of vehicles in a non-homogeneous fleet. However, most HVRP publications do not consider multi-period time horizons, and they only focus on fleet dimensioning. Further details on the HVRP are presented in Koç et al. (2016).

An advantage of our formulation of the SRP lies in the objective function, which incorporates multiple cost elements that account for miles travelled, labour, truck cost, excess route-time costs, and single- vs team-routes (see Section 3). Instead of being one-dimensional, our objective function captures the inherent trade-offs between the above cost elements, which not only leads to superior solutions but also reveals managerial insights as we show later in the paper. The proposed SRP heuristic serves as a useful tool for planning and decision-making in the fast-food supply chain and similar settings.

3 Formulating the SRP

We present our assumptions, the cost model, and the resulting MIP model.

3.1 Modelling assumptions

The time horizon is modelled as a finite set of discrete time periods. For our application, each day represents a time period. Store demand is represented by the orders a store places (at most one order per store per day), which are fulfilled by the DC. Based on the store orders, a delivery plan is developed for the region; it specifies the delivery quantity

and the time period for each delivery at each store over the time horizon. Given the delivery plan, the following assumptions are made for the MIP model:

- 1 Each store must receive the planned delivery quantity in the specified time period; no early or late deliveries are allowed.
- 2 The stores are responsible for managing their own inventories and they each own the store inventory. Therefore, inventory carrying cost at the stores is not considered. We assume sufficient stock to support the delivery plan is available at the DC. Expiration dates and shelf life are considered by the DC in filling the store orders.
- 3 Routes are classified into single-driver and team routes (two drivers).
- 4 The trucks are identical. However, the usable capacity of a truck depends on the type of route (single-driver or team).
- 5 A truck can complete at most one route per time period.
- 6 The trucks and drivers are available at the beginning of each time period.
- 7 A store is visited by at most one truck per time period (no split deliveries).
- 8 Potential routes are predetermined using a route-generation heuristic.
- 9 Store demand (measured in pounds) is for a single, aggregate SKU.

As the first assumption suggests, the SRP is concerned with minimising cost by choosing the appropriate routes and fleet size needed to support the delivery plan. This assumption is aligned with the business model of our industry collaborator, where each store has a predetermined delivery plan. Furthermore, although the stores are charged a nominal delivery fee, all the transportation and delivery costs are borne by the company, which gives them a strong incentive to minimise the cost of supporting the delivery plan. While the delivery plan itself is also of interest, the development/optimisation of the delivery plan is beyond the scope of our paper.

The second assumption helps us focus on supporting the delivery plan instead of inventory management at the DC. Per our industry collaborator, stock shortages at the DC are rare, and if it occurs, they have systems in place to minimise the impact on the stores. Unless there is a major supply disruption, shortages at the DC are not common and they are rectified quickly.

The third assumption stems from an industry practice of using multiple drivers on a route to allow longer routes and larger delivery quantities, without violating federal hours-of-service regulations (MacMillan, 2018). Assumption 4 addresses single-driver vs. team routes, as each type of route has specific parameters. Assumptions 5, 6, and 7 are made primarily to simplify the model.

The eighth assumption, i.e., using a set partitioning approach for our formulation, is advantageous when capturing complex cost functions and intra-route constraints (Toth and Vigo, 2014). The last assumption, i.e., using an aggregate SKU, allows us to limit the total weight on a truck without tracking each SKU individually. Our industry collaborator also expresses their delivery data in total pounds using an aggregate SKU.

3.2 The cost model

The cost model and the MIP formulation (Section 3.3) are general-purpose and they capture the inherent trade-offs in the SRP. While the primary computational results in our study are based on specific parameter values furnished by our industry collaborator, whose input was instrumental in the development and validation of the model, the cost model and the MIP formulation have broad applicability, meaning they can be used in other SRP applications as long as the appropriate parameter values are specified.

Using a one-dimensional objective, such as just minimising distances, is common in the literature. However, a one-dimensional objective is at best a surrogate objective that captures only a portion of the problem, and it may fall short of adequately capturing key trade-offs inherent to the SRP. For example, in some cases, one may be able to reduce overall cost by using a team route instead of a single-driver route, even if the total distance travelled increases.

Working with our industry collaborator, we defined the costs associated with the SRP based on

- 1 the miles travelled
- 2 the labour associated with unloading inventory delivered (in pounds)
- 3 the number of trucks needed.

Another important aspect we gathered from our industry collaborator and incorporated into the model is that, although US federal regulations limit the amount of time a driver can work continuously to 14 hours (FMCSA, 2017), some routes require more time due to the size of the delivery region. In such cases, the driver is given an allowance to rest overnight, and the route becomes a two-day route. Team routes may also be used in such cases, since two drivers can alternate their working periods. Even with team routes, however, some routes might be long enough to qualify as a two-day route. Since each time period is modelled as a day, a fourth cost component, namely, an ‘overage cost’, was introduced (for both single and team routes) to capture the additional cost of those routes that exceed the one-day time limit. The above four cost components are used in the objective function of the MIP model, which is described in the next section.

3.3 MIP formulation

The MIP model for the SRP uses the following parameters:

N = set of stores in the region ($i = 1, \dots, N$)

R = set of routes ($r = 1, \dots, R$)

T = planning horizon ($t = 1, \dots, T$)

m_r = distance in miles of route r

f_{ir} = indicator parameter; equals 1 if store i is visited in route r , 0 otherwise

d_i^t = demand in pounds (*lbs*) of store i in time period t .

A distinctive aspect of the proposed formulation is that both single and team routes are considered. Single-driver routes have more restrictive limits on truck capacity and route

time length than team routes as well as different cost parameters. To capture the above differences, additional parameters are defined as follows:

$U_V^1(U_V^2)$ = upper limit of truck capacity, in pounds, for single (team) routes

$U_T^1(U_T^2)$ = upper limit of route time length, in minutes, for single (team) routes

$l_r = 1$ if route r is a team route; 0 if a single route

B_{lb} = unloading time, in minutes/pound unloaded

B_n = stop time, in minutes/stop (equal for each stop)

B_m = travel time, in minutes/mile travelled

C_V = daily cost/truck required

$C_m^1(C_m^2)$ = cost/mile travelled in a single (team) route

$C_{lb}^1(C_{lb}^2)$ = cost/pound delivered in a single (team) route

C_U = fixed overage cost for each route that exceeds the route-time upper limit.

The upper limits on truck capacity are determined based on the space and weight limitations of the truck as well as what managers consider to be a reasonable amount of inventory a driver can unload on a route. The upper limits on route times, on the other hand, are defined subject to federal hours of service regulation depending on single vs. team routes. Route times are calculated by considering the driving time, a fixed stopping time/store, and the unloading time/store.

The decision variables are defined as follows:

$a_r^t = 1$ if route r is used in time period t ; 0 otherwise

q_{ir}^t = number of pounds delivered to store i on route r in time period t

$x_r^t = 1$ if length of route r in mins in period t goes over the upper limit; 0 otherwise

p_r^t = amount of time in minutes that route r in period t exceeds the upper limit

K = fleet size.

The objective function is based on the four cost elements described in Section 3.2. The first term reflects the distance cost for each single or team route used:

$$\sum_t \sum_r [C_m^2 l_r + C_m^1 (1 - l_r)] m_r a_r^t \quad (1)$$

Similarly, the second term reflects the labour cost incurred due to unloading inventory, based on single or team route:

$$\sum_t \sum_r \sum_i [C_{lb}^2 l_r + C_{lb}^1 (1 - l_r)] q_{ir}^t \quad (2)$$

The third term calculates the overage cost for each route that exceeds the route time limit:

$$\sum_t \sum_r C_U x_r^t \quad (3)$$

Lastly, the final term in the objective function reflects the cost of providing K trucks over the time horizon:

$$C_V T K \quad (4)$$

The sum of the four components constitutes the objective function shown below. Given a delivery plan, and a set of potential routes, the proposed formulation determines the number of trucks required and the routes to be used in each period to minimise the total cost over the time horizon. The MIP formulation is presented as follows:

$$\begin{aligned} \min & \sum_t \sum_r [C_m^2 l_r + C_m^1 (1-l_r)] m_r a_r^t + \sum_t \sum_r \sum_i [C_{lb}^2 l_r + C_{lb}^1 (1-l_r)] \\ & + \sum_t \sum_r C_U x_r^t + C_V T K \end{aligned} \quad (5)$$

$$s.t. \sum_r q_{ir}^t = d_i^t \quad \forall i \in N, t \in T \quad (6)$$

$$\sum_i q_{ir}^t \leq a_r^t [U_{\tilde{r}}^2 l_r + U_V^1 (1-l_r)] \quad \forall r \in R, t \in T \quad (7)$$

$$\sum_r a_r^t f_{ir} \leq 1 \quad \forall i \in N, t \in T \quad (8)$$

$$\sum_r a_r^t \leq K \quad (9)$$

$$\begin{aligned} & \sum_i (B_{lb} q_{ir}^t) + \sum_{i \in \{i | d_i^t > 0\}} \left(B_n \frac{q_{ir}^t}{d_i^t} \right) + B_m m_r a_r^t \quad \forall r \in R, t \in T \\ & \leq [U_{\tilde{r}}^2 l_r + U_T^1 (1-l_r)] + p_r^t \end{aligned} \quad (10)$$

$$p_r^t \leq [U_{\tilde{r}}^2 l_r + U_T^1 (1-l_r)] x_r^t \quad \forall r \in R, t \in T \quad (11)$$

$$x_r^t \leq 1 \quad \forall r \in R, t \in T \quad (12)$$

$$K \geq 0 \quad \text{integer} \quad (13)$$

$$a_r^t \in \{0,1\} \quad \forall r \in R, t \in T \quad (14)$$

$$q_{ir}^t \geq 0 \quad \forall i \in N, r \in R, t \in T \quad (15)$$

$$x_r^t \in \{0,1\} \quad \forall r \in R, t \in T \quad (16)$$

$$p_r^t \geq 0 \quad \forall r \in R, t \in T \quad (17)$$

Constraints (6) ensure that store demand is met in each period, while constraints (7) enforce the truck capacity. Constraint set (8) prevents split deliveries. Constraint set (9) ensures that the number of routes used in any period does not exceed the number of

trucks. Set (10) calculates the estimated route time and determines if it exceeds the time limit. The left-hand side of the constraint contains three terms that calculate the amount of time needed for unloading the inventory, the fixed time associated with each stop on the route, and the drive time, respectively. On the right-hand side, the variable p_r^t captures the excess route time that goes over the upper limit. Set (11) identifies the number of routes that exceed the time limit to assign the appropriate overage cost in the objective function. While constraints (12) ensure that the route time for overage routes is bounded. Constraints (13) to (17) define the variables and the non-negativity constraints.

4 Clustering heuristic for route generation

One challenge associated with the set partitioning-based approach we adopted is generating potential routes *a priori*. For example, with just 50 stores, if every route with up to 7 stores is feasible, it would yield over 118 million routes. To address this challenge, we use a simple clustering heuristic to pre-generate a manageable set of potential routes. It is adapted from the nearest-neighbour-based clustering and routing heuristic (nCAR) presented by Sarkar et al. (2018). Our goal is not to develop an optimum-seeking algorithm for the SRP, but rather to use a simple heuristic to obtain solutions in reasonable times that allow us to investigate the impact of the various cost elements in the objective function on the solution structure and at the same time glean managerial insights made possible by a cost model with multiple elements. Given the groundwork we laid with the proposed model, subsequent studies can explore other solution techniques, such as column generation/pricing schemes, to obtain better solutions.

The proposed clustering heuristic considers each period sequentially, and it allows for more than one set of parameters to be considered when building the routes, which lets the heuristic accept different parameters between single and team routes when determining the composition and size of the routes. Usable truck capacity is a constraining factor when determining which store will be added next to a cluster, while route time limits are used to constrain the cluster size.

For each period, given a set of stores that must be visited, the heuristic creates multiple clusters of increasing size (measured by the number of stores). For each cluster, a route is generated by finding the shortest path through the stores in the cluster, beginning and ending at the DC. The parameters for truck capacity and route time for both single and team routes are used as input, along with a distance matrix, store demand for each period, and the maximum number of stops allowed on a route, say, Z .

Given the above parameters, the following heuristic is employed, where z serves as an iteration counter, and s reflects the current cluster size. The steps are executed for each time period t , considering first the parameters for single-driver routes and then team routes:

- 1 Let N_t represent the (sub)set of stores that must be visited in period t .
- 2 Set $z = 1$. Select store $i \in N_t$ as the seed store of cluster i_z and set $s = 1$. Initialise the cluster demand to the demand of store i , i.e., set $d_{i_z} = d_i^t$, and the cluster time length to the driving time from the DC to store i plus the time associated with visiting store i , which is equal to the unloading time, $d_i^t B_{lb}$, and the fixed time per stop, B_n .

- 3 If $s = z$, or the cluster time length exceeds the maximum route length in time (U_T^1, U_T^2) , go to Step 5. Otherwise, go to Step 4.
- 4 Determine the next store k to add to the cluster as follows:
 - 4a For each store already in the cluster, identify its nearest neighbour in distance. (Ties are broken by the smallest store number.) Verify that adding the nearest neighbour to the cluster would not violate the truck capacity. If it does, then consider the next closest neighbour. Continue in this fashion until a feasible neighbour is found, or there are no neighbours that can be added. The above is repeated for each store already in the cluster, leading to a set of potential stores.
 - 4b If the set of potential stores from Step 4a is empty, then set $z = Z$ and go to Step 7. Otherwise continue to Step 4c.
 - 4c Given the set of potential stores identified in Step 4a, evaluate the increase in travel distance by inserting each potential store, one at a time, before and after its nearest neighbour in the cluster. Choose the position that yields the smallest increase. Record the increase in distance for each potential store.
 - 4d Set the store that leads to the smallest increase in distance as store k and add it to the cluster. Update the cluster size (i.e., set $s = s + 1$), the cluster demand, (i.e., set $d_{i_z} = d_{i_z} + d_k^i$), and the cluster time length to include store k . Go to Step 3.
- 5 Compare the new cluster with previously saved clusters to avoid duplicates.
- 6 Determine the shortest path to visit all the stores in the cluster from the DC. Save the shortest path as a route and add it to the set of routes R , identifying it as a single or team route.
- 7 If $z = Z$, then remove store i from N_t and go to Step 2. Otherwise, set $z = z + 1$, and go to Step 3.

The shortest path in Step 6 was obtained via Google's OR-Tools Python library (Google, 2019). Although the routes are generated by time period, the model is not restricted to use any route for a specific time period; rather, it considers the entire pool of routes when deciding which routes to use. The above heuristic allows some flexibility in the order in which the routes are built since the final set of routes does not change with the order in which the seed stores are selected, or whether single or team routes are generated first.

5 Computational results

In this section we discuss two sets of results. For the first set, we solved a series of test instances of varying sizes with randomly generated data to evaluate the size limitations of the SRP heuristic. We also compare the results obtained from the SRP heuristic with baseline results obtained from alternative, one-dimensional objective functions. For the second set, using real-world data provided by our industry collaborator, we developed multiple problem instances to discuss the solutions obtained from the SRP heuristic and to compare them to the industry baseline as well as the other one-dimensional baselines. The SRP heuristic was programmed with 64-bit Python 3.7.4, while the MIP model was

solved with CPLEX 12.7 on a Windows computer with a 64-bit 2.50 GHz Intel Core i7 and 8 GB of RAM.

5.1 Randomly generated test instances

We randomly generated a series of test instances of varying sizes to evaluate how the proposed heuristic performs with different data scenarios as shown in Table 3. The two main factors we varied on each instance is the problem size, defined by the number of stores, and the delivery frequency, defined by how many times during the time horizon each store needs replenishment.

The largest problem size tested was 350 stores which is a representative average upper limit on the number of stores in a region for our industry collaborator. Store locations were randomly scattered following a uniform distribution within a 600×600 grid, with the DC placed at the centre (plots of the locations on each instance are included in the Appendix). The distance matrix was calculated for each instance assuming Euclidean distances. The frequency of delivery for each store was randomly generated following a normal distribution, with the average changing according to the level shown in Table 1. Once a frequency was determined for each store, a uniform distribution was used to randomly assign each demand to a period within the time horizon. The demand for each store was also generated using a normal distribution based on the parameters shown in Table 3. Data from our industry collaborators were used to define the values for all other parameters in the model (truck capacity, unit costs, etc.), these values are discussed in Section 5.2.

Table 1 Parameters for the random test instances

<i>Variable</i>	<i>Level</i>	<i>Parameter value</i>
Problem size	3 levels	$N = 10$
		$N = 100$
		$N = 350$
Order frequency	Low	$\mu = 1, \sigma = 0.60$
	Average	$\mu = 2, \sigma = 0.60$
	High	$\mu = 4, \sigma = 0.60$
Store demand		$\mu = 2030, \sigma = 760.15$

For each combination 5 random test instances were generated following the procedure described above, for a total of 45 test instances. Using the SRP heuristic, we were able to solve all 45 instances within a time limit of 3 hours. Table 2, where low and high demand frequency are denoted by ‘L’ and ‘H’, respectively, shows the average total runtime for each combination as well as the portion of the total time dedicated to the CPLEX solver under the MIP Runtime columns.

As expected, the total runtime increases with the number of stores in the network. Within instances of the same size, the total runtime increases as the order frequency increases. The majority of the overall runtime is dedicated to the clustering heuristic and building the MIP model. These two processes account for 94% of the total runtime on average across all the test instances. On the other hand, the average solver runtime across all the instances was under a minute, which indicates that, with the proposed approach, we are able to find solutions for various scenarios in a practical time. A key factor to

consider is that, as the number of stores and the order frequency increase, the clustering heuristic generates a substantially larger number of routes, which leads to an increase in the size of the MIP since some variables and constraints are indexed by the number of routes and stores. This explains the significant increase in runtime for the instances with 350 stores which generate approximately 10,000 routes each.

Table 2 Average runtime (in secs.) per test instance

<i>Instances</i>	<i>Total runtime</i>		<i>MIP runtime</i>	
	<i>Avg.</i>	<i>Std. Dev.</i>	<i>Avg.</i>	<i>Std. Dev.</i>
N10-L	2.05	0.81	0.25	0.09
N10	3.71	1.13	0.56	0.18
N10-H	3.33	0.92	0.34	0.15
N100-L	34.67	3.40	0.95	0.15
N100	60.40	7.77	1.92	0.69
N100-H	76.99	7.48	3.18	0.93
N350-L	1,913.68	198.05	17.51	2.80
N350	5,072.02	724.94	56.72	1.05
N350-H	10,109.55	1,395.58	139.26	9.63

To evaluate the performance of our proposed multi-dimensional objective function, we solved the same set of instances with alternative/one-dimensional objectives for comparison as follows:

- 1 minimising only the truck cost (i.e., the fleet size)
- 2 minimising only the cost of miles travelled (i.e., using only a distance-based objective as most studies do).

For both of the alternative objectives, once a solution was found by the model, we calculated the total cost of the solution using our proposed cost model to have a direct comparison. In Table 3, we summarise the average percentage difference between our proposed multi-dimensional objective function (MIP Soln) and the two alternative one-dimensional ones using the same set of instances.

Table 3 Comparison of alternative objective functions


<i>Instances</i>	<i>MIP Soln.</i>	<i>Min dist.</i>	<i>% Gap</i>	<i>Min truck</i>	<i>% Gap</i>
N10-L	\$9,153.02	\$9,413.92	2.7%	\$9,329.00	1.8%
N10	\$13,009.37	\$13,205.86	1.6%	\$14,941.01	15.0%
N10-H	\$19,321.88	\$19,457.41	0.7%	\$21,266.34	9.9%
N100-L	\$39,917.25	\$39,932.94	0.0%	\$47,861.54	19.8%
N100	\$53,877.55	\$54,367.47	0.9%	\$65,996.03	22.5%
N100-H	\$86,304.88	\$86,713.60	0.5%	\$99,367.30	15.1%
N350-L	\$136,899.52	\$138,081.19	0.9%	\$161,656.75	19.4%
N350	\$166,462.99	\$167,316.36	0.5%	\$195,828.80	17.6%
N350-H	\$238,665.22	\$239,848.54	0.5%	\$273,002.85	15.5%

These results show that on average the multi-dimensional objective function outperforms both of the commonly used one-dimensional objectives across all the instances. Overall the solutions obtained when only truck cost is considered are 15% higher than the ones obtained with our objective function on average. The solutions obtained with the distance-based objective have a significantly smaller gap, being overall about 1% over our solutions. However, this shows how incorporating multiple cost factors into the decision-making, while increasing the complexity of the problem, can lead to better solutions within reasonable time.

5.2 Industry data

Two problem instances were developed using data from our industry collaborator. Each instance is based on orders received from 181 stores over 6 days. The two instances are representative of the operations of our collaborator, where plans are made on a weekly basis with no Sunday deliveries. Both instances involve the same 181 stores but use different delivery plans. Figure 1 shows the location of the stores with the map details omitted to maintain data confidentiality.

Figure 1 Map of store and DC location (see online version for colours)



Store demand and a distance matrix are the main inputs to our approach. Actual road distances for the distance matrix were obtained from the Google Maps API (Google, 2020). Additional input data includes the truck capacity and the maximum route time for both single and team routes, the average minutes/mile travelled, and the appropriate values for the cost parameters. Since the parameter values and cost estimates used for the test instances represent actual values used by our industry collaborator, they are not shown for confidentiality reasons. However, we can report their relative relationships. For example, the truck capacity for team routes is 35% higher than that of single-driver routes, as two drivers are allowed to handle higher quantities for delivery. Similarly, the maximum route-time is 30% higher for team routes. However, compared to single routes,

team routes are 20% higher in cost/mile travelled, and 35% higher in cost/pound delivered. The overage cost for exceeding the route time limit was set as the daily truck cost, to account for the incremental usage of the trucks on such routes.

The results obtained from the SRP heuristic are compared to the industry baseline cost, which was calculated employing the data obtained from our collaborator and their current routes, using the same cost model presented in Section 3. Our industry collaborator uses a commercial routing software to construct their routes and update them manually as new stores are opened and to accommodate weekly changes in demand.

5.3 Discussion of industry data results

The ‘MIP solution’ we get from solving the model with CPLEX is exact. However, since the MIP model considers only the routes generated by the clustering heuristic, the MIP solution is a heuristic solution. The overall runtime for the SRP heuristic is 11.37 (14.8) mins. for instance 1 (2), with the MIP solver taking only 6–9 secs. of this time. The bulk of the time is used to generate the routes (3,582 and 3,294 routes for instance 1 and 2, respectively) and build the MIP model. In addition to the industry baseline, we solved the same two instances with alternative/one-dimensional objectives for comparison as follows:

- 1 minimising only the truck cost (i.e., the fleet size)
- 2 minimising only the cost of miles travelled (i.e., using only a distance-based objective as most studies do)
- 3 using the same cost function but allowing only single-driver routes.

Figure 2 Cost comparison, instance 1 (see online version for colours)




Figure 2 (3) shows the cost of the MIP solution versus the alternative baselines for instance 1 (2). For both instances, the MIP solution outperforms all the baselines. Relative to the industry baseline, our solution yields a reduction of about 4% and 13% in total cost for each instance. This is a significant cost reduction for VRPs where, considering the large transportation volumes typically associated with these problems, even a 2% reduction is considered substantial (Hasle and Kloster, 2007). Looking at the breakdown by cost component, the graphs show that the MIP solution does not incur in

the lowest cost by component. For example, the distance-only objective solution has a miles travelled cost that is 3% lower than that of the MIP solution, similarly the truck cost objective has a solution with an overage cost 11% lower than the MIP solution for

instance 1 (4% for instance 2). However, when we consider all the cost dimensions together, the MIP solution identifies trade-offs that result in a lower overall cost. Table 4 presents a detailed comparison of the MIP solution and the industry baseline.

Figure 3 Cost comparison, instance 2 (see online version for colours)





Table 4 Baseline versus MIP results

	<i>Baseline – 1</i>	<i>MIP Soln – 1</i>	<i>Baseline – 2</i>	<i>MIP Soln – 2</i>
<i>Total cost</i>	\$54,283.48	\$51,946.13	\$56,284.60	\$49,399.04
<i>Miles cost</i>	\$25,661.31	\$23,936.69	\$27,191.51	\$22,129.20
<i>Del cost</i>	\$15,585.85	\$13,936.15	\$16,353.05	\$13,789.10
<i>Overage cost</i>	\$4,147.92	\$4,296.06	\$2,962.80	\$3,703.50
<i>Truck cost</i>	\$8,888.40	\$9,777.24	\$9,777.24	\$9,777.24
<i>Total runtime (min)</i>	-	16.53	-	6.55
<i>MIP runtime (min)</i>	-	0.18	-	0.11
<i>Miles</i>	19,834.29	19,091.66	20,066.70	18,184.66
<i>Trucks</i>	10	11	11	11
<i>Visits</i>	408	408	395	395
<i>Miles/visit</i>	48.61	46.79	50.80	46.04
<i>Delivery trips</i>	55	60	58	56
<i>Routes over limit</i>	28	29	20	25
<i>% Over</i>	50.9%	48.33%	34.5%	45%
<i>No. of team routes</i>	22	2	38	4

The percentage differences observed in these comparisons may change based on the parameter values used and specific problems tested (recall that the proposed model can be implemented with any set of user-provided data). Nevertheless, the above results

support our main message that the SRP problem is a multi-dimensional problem with labour costs (driving as well as delivery) and equipment (truck) costs, and therefore, using traditional one-dimensional objectives may lead to inferior solutions. By considering multiple cost elements concurrently, our model is able to quantify trade-offs and ultimately identify superior solutions, which is important not only for reducing costs but also for providing managerial insights. Furthermore, some of the trade-offs may not be intuitive or obvious at first glance, such as achieving a lower total cost with longer routes or more delivery trips, which makes the proposed model an interesting and effective heuristic. The breakdown shown for the MIP solutions and the industry baseline in Table 4 highlights some of the key trade-offs which provides valuable managerial insights.

Figure 4 Truck contents and route length for the MIP solution for instance 1 (see online version for colours)




The MIP solution for the first instance shows a slight increase in the fleet size and the number of overage routes. However, the overall cost is reduced since 90% fewer team routes are deployed. Similarly, for the second instance, the heuristic solution reduces the number of team routes by 95%, which, combined with the reduced total miles travelled, yields a 13% reduction in total cost. Although our industry collaborator relies heavily on team routes, our solutions suggest that extending single routes by incurring overage costs is, in general, better than extending them through team driving (one would also need to

consider some of the non-tangible trade-offs between single and team driving before making such changes).

Further insights can be gained by examining the truck capacity utilisation and the route lengths (in time). When we compare the truck contents on each route with the route times, it becomes clear that the time limit is the tighter constraint. The graphs in Figure 4 (Figure 5) show the truck content in pounds for each truck used in the solution and the route time for each trip in the first instance under the MIP solution (industry baseline). The guidelines in the graphs show the upper limits for the truck capacity and the route times. The guideline for the truck capacity on team routes is not shown due to the scale chosen to best depict the data.

Figure 5 Truck contents and route length for the baseline for instance 1 (see online version for colours)



The graphs for the truck capacity utilisation show that, for instance 1, only 5 routes (10%) in the baseline use more than 90% of the truck capacity, while 15% of the routes in the heuristic solution are over the 90% threshold. Instance 2 shows similar results with 5% and 18% of the routes over the 90% threshold for the baseline and MIP solution, respectively (graphs for instance 2 are shown in the Appendix). However, in both the baseline and the MIP solution for the two instances, the *average* truck capacity utilisation ranges between 61%–65%, with the capacity utilisation of individual trucks ranging from 5% to 99%. Such a wide range reflects the challenging nature of the SRP. The routes with low truck utilisation are single-store routes, which are also known as *out-and-back* deliveries. For the baseline, with the exception of one, the routes with the lowest truck capacity utilisation have a route time close to the upper limit, which indicates that the store visited was far from the DC. In contrast, in the SRP heuristic solution, single-store routes do not exceed 50% of the upper time limit. This suggests that, although these stores are at a reasonable distance from the DC, the SRP heuristic replenishes them


individually, while more distant stores are part of multi-stop routes. Our results imply that further research on the merits of mixing out-and-back deliveries with multi-stop routes would be well-justified.

For route times, the baseline (MIP solution) shows that about 51% (48%) of the routes are over the one-day limit for the first instance, and 35% (45%) for the second instance. These results indicate that, as we stated earlier, there is a tendency to accept overage costs instead of increasing the number of routes or changing other elements in the solution. This result can be influenced by the overage cost; we study how it interacts with the solution structure in the sensitivity analysis discussed in Section 6.

5.4 Split deliveries

The above results disallow split deliveries (SDs), which is consistent with the operations of our industry collaborator. SDs are often disallowed in the literature as well, although the SDVRP has received more attention in recent years; see, for example, Archetti and Speranza (2008) for a review of SDVRP and its comparison with the traditional VRP. Even though SDs is not used by our collaborator, there are no systemic rules against SDs, and exploring its impact seems justified due to the low average truck capacity utilisation (61%–65%). With SDs, a higher truck utilisation can be achieved, which in turn may reduce the number of routes and total distance for a lower overall cost.

Figure 6 Cost impact of split deliveries, instance 1 (see online version for colours)



Removing constraint (4) in Section 3, and using the same two problem instances, we obtain the results shown in Figures 6 and 7, which indicate that, compared to the baseline, SDs reduce the total cost by 7% and 15% for the first and second instance, respectively. Although the optimality gap is very small, we note that the SD results are based on incumbent solutions obtained from the MIP model after it reached an upper limit of 3 hours without being able to confirm their optimality. Table 5 summarises the performance of the SRP heuristic with and without SDs.

For both instances, allowing SDs results in a further reduction of about 3% in total cost. This can motivate further study on the impact of incorporating SDs into store replenishment in the fast-food sector. Factors such as the model runtime and the additional disruption to store operations caused by SDs should also be considered.

Figure 7 Cost impact of split deliveries, instance 2 (see online version for colours)




Table 5 Performance of the SRP heuristic, split deliveries

	<i>MIP soln. No SD Instance 1</i>	<i>MIP soln. SD Instance 1</i>	<i>MIP soln. No SD Instance 2</i>	<i>MIP soln. SD Instance 2</i>
<i>Total cost reduction</i>	4.3%	7.1%	13.4%	15.2%
<i>Optimality gap</i>	0%	0.59%	0%	0.50%
<i>Total runtime (min)</i>	16.53	204.84	10.65	193.90
<i>MIP runtime (min)</i>	0.18	180.73	0.15	181.12

6 Sensitivity analyses

In this section, we perform sensitivity analyses to better understand how each cost component impacts the solution. The analysis was performed by re-solving the first industry data instance with the SRP heuristic (with no SDs) while changing individual cost parameter values from zero to three times their original value, and keeping the other parameters fixed. As expected, when a cost parameter is changed, the total cost changes; however, we show the *magnitude* of the change relative to the change in the parameter value as well as possible changes in the *structure* of the solution. While one may also vary the cost parameters two or more at a time, our approach is needed as a starting point before trying more complicated changes.


Figure 8 Sensitivity analysis graphs (see online version for colours)

Figure 8 includes four graphs that show how each cost component and the overall cost change as a particular cost parameter was increased. Since the distance cost consistently

represents the largest portion of the total cost, it is not surprising that increasing the cost/mile travelled in both single and team routes has the biggest impact in the total cost. While the lines for the other cost components do not change in the graph, changes in the unit distance cost affects the total miles travelled every time it is increased. More specifically, when the unit costs are increased from 0 to $3x$, the total miles travelled decreases by 14%. This change also led to more than doubling the number of overage routes, and a reduction from 21% to 9% in the proportion of team routes used.

On the other hand, the overage cost appears to have the most impact in the *structure* of the solution, while causing the least increase in total cost. As the overage cost increases, the other components of the model adjust to keep the total cost from increasing significantly. As the overage cost increases from 0 to $3x$, the number of overage routes decreases by almost 60%, which leads to a small increase (3%) in the total miles travelled, and a significant change in single versus team routes, with the proportion of team routes going from 0% to 32%. This shows that as overage cost increases, team routes become advantageous since they allow larger delivery quantities and longer route times. Changes to the overage cost are the only scenario where we observed a clear and significant change in the solution structure and the percent contribution of the individual cost components. Although the change in the solution structure is significant, remarkably, the change in the total cost is relatively small as shown in the graph.

Changes in the unit delivery cost caused an anticipated reduction in the total number of team routes used. However, the biggest reduction happens when the unit costs are increased from 0 to $1x$, with the proportion of team routes decreasing by 80%. Further increases in this cost only lead to slight changes in the solution structure. The line graph shows how the total cost increases with the unit delivery costs, while the other cost components show almost no change. Similarly, examining the impact of changes in the truck cost, we can conclude that, within the range of values we examined, the solution structure is not sensitive to changes in the truck cost. As this cost parameter increased, the total number of trucks used did not change, and as the graph shows, it did not cause any significant changes in other cost components.

7 Conclusions

We introduced a MIP formulation for the SRP and developed a heuristic to minimise the overall logistics costs associated with replenishing stores in a specified region. The proposed approach derives its strength from its simplicity and effectiveness. It is simple in that we are able to obtain results in a reasonable time on a modest computer, using a commercial solver like CPLEX. It is effective in that using real-world data, our model identifies better solutions with a lower total cost than the baseline cost established from our industry collaborator's operations.

The two central contributions of this work are the cost-driven multi-dimensional objective function we developed, and the key managerial insights we obtain from the solutions. The proposed SRP heuristic uses a cost model (developed jointly with our industry collaborator) which captures trade-offs that are inherent to the SRP, and it allows us to capture distance-based and labour-based costs as well as equipment and time overage costs. Our numeric results based on industry data show that using traditional one-dimensional objectives, such as minimising only the total distance or the fleet size, may lead to inferior solutions. Furthermore, our cost model considers single-driver and

team routes. To our knowledge, this study may be the only one to consider team routes, overage cost, and fleet-sizing decisions simultaneously. Additionally, the general structure of the cost model, and the user-defined parameter values, allows for the proposed SRP heuristic to be generalised and applied to other distribution logistics problems with multiple nodes.

The results from this study also provide key managerial insights for decision-makers in the fast-food industry. Our results show that total cost can be reduced with extended single-driver routes as opposed to using a larger number of team routes. Through a series of sensitivity analyses, we established that among the cost components, the distance cost and the overage cost for extended routes show the highest impact on the overall cost and the solution structure, respectively. This can motivate further study into how the overage cost is defined and what other factors may influence extending the routes beyond their current one-day limit.

Although the results from the SRP heuristic showed a reduction in total cost, there are potential areas of study that can further improve the solutions, as shown for example with the SD results. Other future directions to study include a non-homogeneous fleet of trucks (which would allow for lighter or shorter routes, while simplifying urban-area deliveries) and incorporating multiple groups of SKUs.

References

- Akkerman, R., Farahani, P. and Grunow, M. (2010) 'Quality, safety and sustainability in food distribution: a review of quantitative operations management approaches and challenges', *OR Spectrum*, doi: 10.1007/s00291-010-0223-2.
- Alegre, J., Laguna, M. and Pacheco, J. (2007) 'Optimizing the periodic pick-up of raw materials for a manufacturer of auto parts', *European Journal of Operational Research*, Vol. 179, No. 3, pp.736–746, doi: 10.1016/j.ejor.2005.03.063.
- Archetti, C. and Speranza, M.G. (2008) 'The split delivery vehicle routing problem: a survey', in Golden, B., Raghavan, S. and Wasil, E. (Eds.): *The Vehicle Routing Problem: Latest Advances and New Challenges*, Springer, Boston, MA, pp.103–122, doi: https://doi-org.proxy.lib.umich.edu/10.1007/978-0-387-77778-8_5.
- Archetti, C., Jabali, O. and Speranza, M.G. (2015) 'Multi-period vehicle routing problem with due dates', *Computers and Operations Research*, 61, pp. 122–134, Elsevier, doi: 10.1016/j.cor.2015.03.014.
- Beltrami, E.J. and Bodin, L.D. (1974) 'Networks and vehicle routing for municipal waste collection', *Networks*, Vol. 4, No. 1, pp.65–94, doi: 10.1002/net.3230040106.
- Braekers, K., Ramaekers, K. and Van Nieuwenhuyse, I. (2016) 'The vehicle routing problem: State of the art classification and review', *Computers and Industrial Engineering*, Vol. 99, pp.300–313, Elsevier Ltd, doi: 10.1016/j.cie.2015.12.007.
- Chao, I-M, Golden, B.L. and Wasil, E. (1995) 'An improved heuristic for the period vehicle routing problem', *Networks*, Vol. 26, No. 1, pp.25–44, doi: 10.1002/net.3230260104.
- Christofides, N. and Beasley, J.E. (1984) 'The period routing problem', *Networks*, Vol. 14, No. 2, pp.237–256, doi: 10.1002/net.3230140205.
- Cordeau, J.F., Gendreau, M. and Laporte, G. (1997) 'A tabu search heuristic for periodic and multi-depot vehicle routing problems', *Networks*, Vol. 30, pp.105–119, doi: 10.1002/(SICI)1097-0037(199709)30:2%3C105::AID-NET5%3E3.0.CO;2-G.
- FMCSA (2017) *Summary of Hours of Service Regulations*, Federal Moto Carrier Safety Administration [online] <https://www.fmcsa.dot.gov/regulations/hours-service/summary-hours-service-regulations> (accessed 21 January 2020).

- Francis, P., Smilowitz, K. and Tzur, M. (2006) 'The period vehicle routing problem with service choice', *Transportation Science*, Vol. 40, No. 4, pp.439–454, doi: 10.1287/trsc.1050.0140.
- Francis, P., Smilowitz, K. and Tzur, M. (2008) 'The period vehicle routing problem and its extensions', in Golden, B., Raghavan, S. and Wasil, E. (Eds.): *The Vehicle Routing Problem: Latest Advances and New Challenges*, pp.73–102, New York.
- Gaudioso, M. and Paletta, G. (1992) 'A heuristic for the periodic vehicle routing problem', *Transportation Science*, Vol. 26, No. 2, pp.86–92, doi: 10.1287/trsc.26.2.86.
- GMA (2008) *Powering Growth through Direct Store Delivery* [online] <https://www.yumpu.com/en/document/read/6749704/direct-store-delivery-grocery-manufacturers-association> (accessed 16 November 2020).
- Goel, A., Vidal, T. and Kok, A.L. (2019) *To Team Up or Not – Single versus Team Driving in European Road Freight Transport*, Rio de Janeiro, Brasil [online] <https://wl.cirrelt.ca/~vidalt/papers/Team-vs-Single.pdf> (accessed 20 May 2020).
- Golden, B. et al. (1984) 'The fleet size and mix vehicle routing problem with time windows', *Computers & Operations Research*, Vol. 11, No. 1, pp.49–66, [https://doi-org.proxy.lib.umich.edu/10.1016/0305-0548\(84\)90007-8](https://doi-org.proxy.lib.umich.edu/10.1016/0305-0548(84)90007-8).
- Google (2019) *Google OR-Tools*, Google [online] <https://developers.google.com/optimization> (accessed 21 January 2020).
- Google (2020) *Google Maps Platform*, Google [online] <https://cloud.google.com/maps-platform/> (accessed 21 May 2020).
- Gulczynski, D., Golden, B. and Wasil, E. (2011) 'The period vehicle routing problem: New heuristics and real-world variants', *Transportation Research Part E: Logistics and Transportation Review*, Vol. 47, No. 5, pp.648–668, Elsevier Ltd, doi: 10.1016/j.tre.2011.02.002.
- Hasle, G. and Kloster, O. (2007) 'Industrial vehicle routing', in Hasle, G., Lie, K., and Quak, E. (Eds.): *Geometric Modelling, Numerical Simulation, and Optimization*, Springer Berlin Heidelberg, Heidelberg, Berlin, pp.397–435, doi: https://doi-org.proxy.lib.umich.edu/10.1007/978-3-540-68783-2_12.
- Hemmelmayer, V.C., Doerner, K.F. and Hartl, R.F. (2009) 'A variable neighborhood search heuristic for periodic routing problems', *European Journal of Operational Research*, Vol. 195, No. 3, pp.791–802, Elsevier B.V., doi: 10.1016/j.ejor.2007.08.048.
- Klein, D. (2018) *Red-Hot Domino's and the Road to 30,000 Restaurants*, QSR [online] <https://www.qsrmagazine.com/pizza/red-hot-dominos-and-road-30000-restaurants> (accessed 21 January 2020).
- Koç, Ç. et al. (2016) 'Thirty years of heterogeneous vehicle routing', *European Journal of Operational Research*, Vol. 249, No. 1, pp.1–21, doi: 10.1016/j.ejor.2015.07.020.
- Laporte, G. (2009) 'Fifty years of vehicle routing', *Transportation Science*, Vol. 43, No. 4, pp.408–416, doi: 10.1287/trsc.1090.0301.
- Larrain, H. et al. (2019) 'Exact solution methods for the multi-period vehicle routing problem with due dates', *Computers and Operations Research*, Vol. 110, pp.148–158, Elsevier Ltd, doi: 10.1016/j.cor.2019.05.026.
- Lemma, Y., Kitaw, D. and Gatew, G. (2014) 'Loss in perishable food supply chain: an optimization approach literature review', *International Journal of Scientific & Engineering Research*, Vol. 5, No. 5, pp.302–311.
- Lenstra, J.K. and Kan, A.H.G.R. (1981) 'Complexity of vehicle routing and scheduling problems', *Networks*, Vol. 11, No. 2, pp.221–227, doi: 10.1002/net.3230110211.
- López-Santana, E., Rodríguez-Vásquez, W.C. and Méndez-Giraldo, G. (2018) 'A hybrid expert system, clustering and ant colony optimization approach for scheduling and routing problem in courier services', *International Journal of Industrial Engineering Computations*, Vol. 9, No. 3, pp.369–396, doi: 10.5267/j.ijiec.2017.8.001.
- MacMillan, C. (2018) *What You Need to Know About Team Truck Driving Jobs*, Smart Trucking [online] <https://www.smart-trucking.com/team-truck-driving/> (accessed 25 April 2020).

- Sarkar, C., Paul, H.S. and Pal, A. (2018) 'A scalable multi-robot task allocation algorithm', in *Proceedings – IEEE International Conference on Robotics and Automation*, IEEE, pp.5022–5027, doi: 10.1109/ICRA.2018.8460886.
- Sharda, R. et al. (2008) *The Vehicle Routing Problem: Latest Advances and New Challenges*, Springer Science+Business Media, LLC (Operations Research/Computer Science Interfaces, 1387-666X;43), Boston, MA [online] <http://dx.doi.org/10.1007/978-0-387-77778-8>.
- Shokri, A., Oglethorpe, D. and Nabhani, F. (2014) 'Evaluating sustainability in the UK fast food supply chain: Review of dimensions, awareness and practice', *Journal of Manufacturing Technology Management*, Vol. 25, No. 8, pp.1224–1244, doi: 10.1108/JMTM-04-2013-0031.
- Staff, E. (2019) 'Nestle joins kellogg, hostess and others before it to exit DSD to reduce logistics costs', *Supply Chain Digest*, June [online] <http://www.scdigest.com/ontarget/19-06-04-1.php?cid=15536> (accessed 16 November 2020).
- Tarantilis, C.D. and Kiranoudis, C.T. (2001) 'A meta-heuristic algorithm for the efficient distribution of perishable foods', *Journal of Food Engineering*, Vol. 50, No. 1, pp.1–9.
- Tarantilis, C.D. and Kiranoudis, C.T. (2002) 'Distribution of fresh meat', *Journal of Food Engineering*, Vol. 51, No. 1, pp.85–91.
- Toth, P. and Vigo, D. (Eds.) (2014) *Vehicle Routing: Problems, Methods, and Applications*, Society for Industrial and Applied Mathematics : Mathematical Optimization Society (MOS-SIAM series on optimization), Philadelphia [online] <http://epubs.siam.org/doi/book/10.1137/1.9781611973594>.
- USDA (2018) *Expenditure for Food at Home and Away from Home*, USDA [online] <https://www.ers.usda.gov/topics/food-markets-prices/food-service-industry/market-segments/> (accessed 21 January 2020).
- Vidal, T., Laporte, G. and Matl, P. (2019) 'A concise guide to existing and emerging vehicle routing problem variants', *European Journal of Operational Research*, doi: 10.1016/j.ejor.2019.10.010.
- Wen, M. et al. (2010) 'The dynamic multi-period vehicle routing problem', *Computers and Operations Research*, Vol. 37, No. 9, pp.1615–1623, doi: 10.1016/j.cor.2009.12.002.
- Yum! Brands (no date) Yum.com [online] <https://www.yum.com/wps/portal/yumbrands/Yumbrands> (accessed 24 April 2020).