Improving last-mile parcel delivery through shared consolidation and portering: A case study in London

Carlos Lamas-Fernandez[a], Antonio Martinez-Sykora[a], Fraser McLeod[b], Tolga Bektaş[c], Tom Cherrett[b] and Julian Allen[d]

[a]Centre for Operational Research, Management Science and Information Systems, Southampton Business School, University of Southampton, Southampton, Highfield, SO17 1BJ, United Kingdom;
[b]Faculty of Engineering and Physical Sciences, University of Southampton, Southampton, Highfield SO17 1BJ, United Kingdom;
[c]University of Liverpool Management School, University of Liverpool, Chatham Street, Liverpool, L69 7ZH, United Kingdom;
[d]Faculty of Architecture and the Built Environment, University of Westminster, London, NW1 5LS, United Kingdom

ARTICLE HISTORY
Compiled June 16, 2023

Abstract
Based on the insights derived from practical trials conducted with carriers operating in London, this paper proposes a last-mile delivery model that involves both driving and walking. A particularly innovative aspect of the model is the shared use of a third-party portering service between the carriers as an enabler to collaborative provision, and shared micro-consolidation points for temporary storage of goods. The model gives rise to a complex routing problem, for which a tabu search algorithm is developed. The algorithm is able to suggest vehicle routes that consolidate parcels in shared locations and porter paths that deliver items from various companies, reducing the overall delivery cost. We applied this algorithm to a manifest dataset obtained from two carriers operating in London. The results from the model suggested that sharing a third party portering service for the last-mile delivery element between the rival carriers in the case study reduced overall transport costs by up to 14%.

KEYWORDS
last-mile; distribution; micro-consolidation; routing; shared portering

1. Introduction

Demand for delivered goods and services in urban areas has been increasing due to urban population growth, increasing levels of e-commerce and customer expectations. The parcels market in the UK increased by 8%, to 2.8 billion delivered items, in 2019–20 (Ofcom, 2020). Next-day, same-day or even 'instant' (e.g., within two hours) delivery options are becoming increasingly common (Allen et al., 2018b). Last-mile carriers work under increasing pressure from customers to provide cheaper and faster services. Their operating environment is difficult to work in due to competition for road and

CONTACT C. Lamas-Fernandez. Email: C.Lamas-Fernandez@soton.ac.uk

kerbside space from other road users, and pressure from regulations that restrict access, parking or the type of vehicle used, for example, where a low emissions zone is in force.

Goods transport gives rise to air and noise pollution, road congestion, accidents and damage to infrastructure. A number of last-mile logistics innovations that aim to ameliorate such negative externalities include the use of innovative vehicles, collaborative working, vehicle scheduling and routing optimisation, regulatory measures, provision of infrastructure (e.g. electric vehicle charging), and, of interest to our paper, micro-consolidation points, and temporary storage locations (Ranieri et al., 2018). Micro-consolidation refers to bundling of goods close to final delivery addresses, and may be an attractive proposition for carriers as it becomes more difficult and unreliable to deliver to city centres. Micro-consolidation points tend to be smaller and more centrally located than urban consolidation centres. Furthermore, there is scope to rethink how parcels could be moved over the last mile, where drivers might rendezvous with walking or cycling couriers to perform the delivery transaction (McLeod et al., 2020), or to rethink how a team of workers move, enabling them to either walk, drive or share vehicles depending on what is more efficient in each situation (Coindreau et al., 2019; Fikar and Hirsch, 2015). These new transport modes align well with environmental policies being pursued by big cities. In London, for example, pedestrians, cyclists and buses tend to be prioritised by local authorities, leading to reduced road network capacity for private motorised vehicles and reduced speeds, with average speeds of 8mph in central London (Transport for London, 2017). The Mayor of London has pledged to further reduce car dependency, setting a target for 80% of all passenger transport trips in London to be made on foot, by cycle or using public transport by 2041[1].

The sharing of resources to transport more than one carrier's goods (often referred to as 'horizontal collaboration') can help improve the efficiency and reduce the delivery costs by improving vehicle fill and reducing vehicle distance travelled. In case of last-mile deliveries, a multi-stop vehicle operation, vehicles are often stationary at the kerbside while drivers walk to delivery points (Allen et al., 2018c). Horizontal collaboration also offers the potential to combine several parcels in a single walking round, hence reducing vehicle parking time and walking time per parcel delivered, which can lead to reductions in vehicle traffic, transport fossil fuel consumption, carbon emissions and local air pollution, and demand for kerbside parking by delivery vehicles (Allen et al., 2018c; Park et al., 2016).

Whilst last-mile distribution models that include micro-consolidation points (e.g. Arrieta-Prieto et al., 2022) and those that combine routing and walking have been studied (e.g. Allen et al., 2018a; Martinez-Sykora et al., 2020; Nguyen et al., 2019), and the idea of collaboration in freight distribution well-researched (see Gansterer and Hartl, 2018, for a thorough survey) with its benefits shown for standard problems such as vehicle routing (Krajewska et al., 2008; Muñoz-Villamizar et al., 2015), there is need to exploit the value of collaboration in more complex last-mile distribution settings, such as the combined walking and driving model presented here.

This paper describes two innovative distribution models that make use of a shared use of micro-consolidation points and a third-party shared portering resource. We make three main contributions:

(1) Based on an analysis of parcel courier rounds in London, we describe an operational delivery model that uses walking porters and micro-consolidation points.
(2) We develop a Tabu Search algorithm to jointly produce driving and walking

[1] `https://www.london.gov.uk/what-we-do/transport/our-vision-transport/mayors-transport-strategy-2018`

routes designed to evaluate the potential costs and benefits of switching to a shared portering model using micro-consolidation points.

(3) Using a significant dataset from carriers operating in London, we show, through the use of game-theoretical tools, the benefits of collaborative working and apportioning of costs in a fair way.

The rest of the paper is structured as follows. Section 2 presents a detailed description of the new business models we propose in this paper. In Section 3, we describe an algorithm that can be used to produce collaborative plans for implementation in practice. Numerical results on real data obtained from carriers operating in London are presented in Section 4. Conclusions are provided in Section 5.

2. Collaborative distribution modelling

The distribution model we envisage here assumes several carriers operating in an urban area, each operating from a single depot. Each carrier has a single van of sufficiently large capacity to perform deliveries which is realistic for the size of delivery area considered in this paper and supported by carrier data analysis and in portering trials (Clarke et al., 2018). We assume that portering, where final deliveries are undertaken on foot, is available as a mode of delivery service (Allen et al., 2018c). Porters use wheeled bags with a limited carrying capacity defined with respect to both weight and volume (e.g., up to 25kg or 250L). Parcels can be dropped off at micro-consolidation points by the drivers for pick up by porters. We assume that the drop-off occurs prior to the route delivery, and hence synchronising routes is not required. Each porter is subject to constraints on the maximum walking distance, and maximum carrying capacity measured by both weight and volume at any point in time. We assume that porters are always available to start work at a micro-consolidation point, and that they must also finish their journey at one in order to return their delivery bags. In practice, porters could either be hired or employed directly by carriers or via a third-party service provider. If the amount of deliveries made to a consignee is within the capacity limits, they are said to be *porterable*, otherwise their delivery needs to be made by the van.

As for the collaborative element, the carriers could either (i) only share the micro-consolidation points, but where each carrier is responsible for the final delivery to their consignees, which we denote by Shared Infrastructure (SI), or (ii) share both the infrastructure as well as a third-party portering service between the carriers, which we will name as Shared Portering Resource (SPR). Various operational practices can be employed within this model, including: (i) carriers dropping off pre-sorted loads in bags that are to be carried (or wheeled) by a porter, with no mixing of items between different carriers being undertaken; (ii) carriers dropping off loose packages that are sorted into bags at the micro-consolidation facility, with mixing of items between carriers allowed. The latter option is investigated in this paper on the basis that it provides more flexibility and should be the more efficient in terms of minimising the portering resource needed, although it is recognised that it likely entails higher infrastructure costs associated with the need for on-site sorting.

An illustrative example of how such a model can operate in practice is shown in Figure 1 assuming two carriers, namely Carrier 1 operating from the depot shown by the (pink) square on the bottom left, and Carrier 2 operating from the depot shown by the (teal) square on the top right. There are six micro-consolidation points for possible use each shown by the black diamonds, and a total of 14 consignees, each belonging to one of the two carriers, as indicated by the corresponding colours.

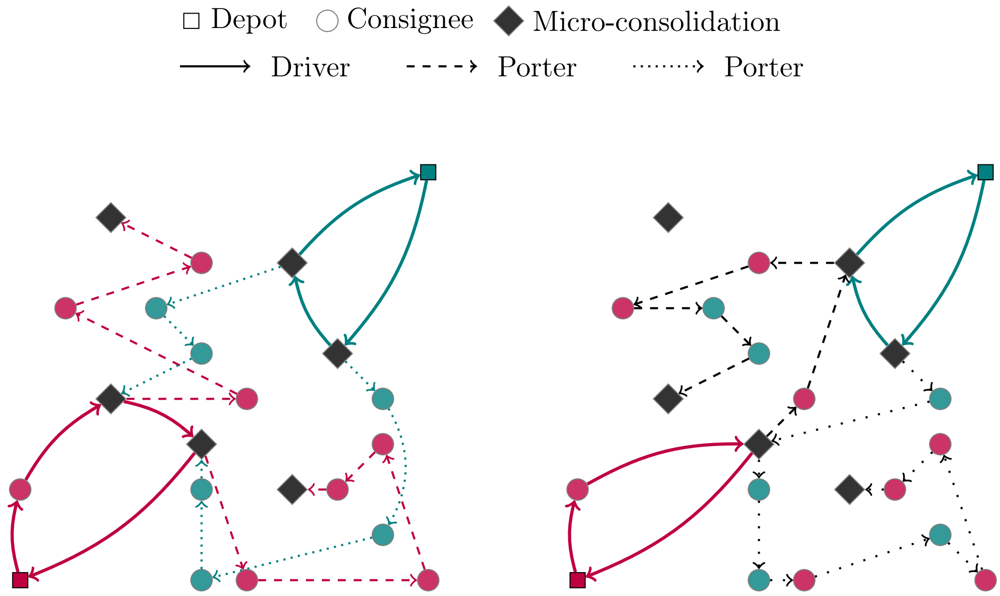

(a) A solution under an SI strategy

(b) A solution under an SPR strategy

Figure 1.: An illustrative example of the alternative models with two carriers (pink and teal).

Figure 1a shows a solution that employs the SI strategy where there are two vehicle routes one for each carrier. A consignee with items for delivery that exceeds the porter capacity, such as the one on the driver route of Carrier 1, is served by the van. Others are served by the porters, which are assumed otherwise to have unlimited capacity for this example. In this solution, there are two porters used by Carrier 1 and two porters used by Carrier 2, which follow the paths indicated by the dashed and dotted lines in Figure 1a, where micro-consolidation points are used as either start or end points of each path. In this solution, each carrier serves their own consignees. Compared to an SI model, SPR yields a more efficient solution as shown in Figure 1b. In particular, although the solution employs the same driver routes as in Figure 1a, it uses a total of two, as opposed to four porters altogether. Both porters stop to refill their bags at micro-consolidation points and then continue onward to deliver items for the remaining consignees on the route. In addition, this solution does not differentiate between the consignees of the two carriers, and as such a porter path might well include consignees from either of the carriers.

3. Optimisation Modelling and Algorithm Design

The two strategies described in the previous section, namely SI and SPR, each give rise to a complex routing problem that entails determining, for several collaborating carriers, the decisions concerning (i) the number and location of micro-consolidation points to be used, (ii) optimal driving routes for the drivers, and (iii) optimal walking routes for the porters. Our implementation aims to minimise the total cost of delivery, which is a function of the number of porters used in a solution, and the distance travelled by the porter(s) and the driver(s), which are amalgamated into a single objective function using cost as the unit of measure as will be explained in the following section. We assume that all porterable consignees are to be visited by porters, in an effort to reduce to the maximum vehicle time and kerbside usage.

In our models, we do not explicitly consider any time-sensitive parcels or any constraints on delivery time. This follows an analysis of a data set from a major carrier operating in London, covering the period 4–9 June 2018, which revealed that of about 13,000 consignments that were delivered, only 1.6% required delivery by 9am, 2.3% by 10am and 7.2% by 12pm, whereas the rest (88.9%) could be delivered anytime until 6pm on a given day. Given the relatively small proportion of time-sensitive parcels, we assume that these could be delivered separately to ensure the delivery deadlines are met, and focus our attention on parcels that do not have any time delivery constraints attached to them.

3.1. Formal problem description

In this section we formally describe the routing problem that arises under an SPR strategy (simply referred to as SPR), of which SI is a special case. In particular, the solution to an instance of SI can be obtained through merging the solution of a SPR solved individually for each carrier, where the sum of the costs of the individual SPRs yield the overall cost for the SI.

An instance of SPR, denoted by $\mathcal{I}$, is described by a set V of vertices, which is partitioned into four subsets, namely a set V_m of nodes as candidates for the micro-consolidation points, a set V_d of depots the locations of which are given, a set V_c of consignees that should be visited by porters, and a set V_h as the consignees that should be visited by the driver either because they are a high-volume/weight delivery, or a collection. Each consignee $v \in V$ is to be delivered a parcel that weighs w_v units and occupies a volume of l_v units. There exist several carriers, each referred to by the index set $s \in \{1, 2, \ldots |V_d|\}$, where $|V_d| \geq 2$. Each carrier s operates one vehicle of sufficient capacity out of a depot d_s, and its consignees are shown by the set $V^s = V_c^s \cup V_h^s$, where $V_c^s \subset V_c$ and $V_h^s \subset V_h$.

Porters are used to perform deliveries on foot to consignees in set V_c by following a path subject to the following constraints:

- *Maximum distance:* To avoid unrealistically long routes, the length of each porter path is restricted to a maximum distance d_{max}.
- *Item pick up:* In order to deliver an item, a porter must have first visited a micro-consolidation point whence the item would be picked up, implying that a porter path always starts at a micro-consolidation node.
- *Route end point:* Porter paths are required to end at a micro-consolidation node.
- *Bag capacity:* Porter bags are limited to a maximum weight W and a maximum volume L, both of which must not be exceeded at any point along a path.
- *Deliveries only:* Porters only deliver (as opposed to collect) goods.
- *Single item limit:* Limits on the maximum weight $\bar{W}$ and volume $\bar{L}$ apply to individual items carrier by a porter. This is a practical constraint to prevent porters from carrying items that would take a substantial amount of their bag. Any item exceeding either of these limits is delivered by a driver, hence these limits precisely define the sets V_c and V_h

There are no restrictions on the number of times that a porter can visit a micro-consolidation points on their path. The total weight and volume of the items delivered by a porter is therefore not necessarily bounded by the bag capacity, as bags might be refilled several times en-route. Drivers are not allowed to deliver porterable items, so in each instance there is a clear distinction between the sets V_c and V_h, where V_h contains only collections and consignees v such that either $w_v \geq \bar{W}$ or $l_v \geq \bar{L}$.

A feasible solution to an SPR instance is described by the tuple (P, Q), in which $P = \{P_1, \ldots, P_{\overline{r}}\}$ is the set of paths assigned to the porters, $\overline{r}$ is the total number of porters used in the solution and $P_r = \{1, \ldots, n_r\}$ is the sequence of n_r nodes in the path traversed by a porter $r \in \{1, \ldots, \overline{r}\}$, $Q = \{Q_1, \ldots, Q_{|V_d|}\}$ is the set of routes undertaken by the drivers, and $Q_s = \{1, \ldots, n_s\}$ is the route performed by the driver of carrier s, expressed as a sequence of nodes. We denote by $P_r(i)$ the node visited in the i^{th} position of path P_r and by $Q_s(i)$ the node visited by the driver of carrier s in the i^{th} position. $Q_s(1) = Q_s(n_s)$ corresponds to the depot of carrier s.

The use of a porter incurs a fixed cost c_1, whereas driving and walking incur per distance unit costs shown by c_2 and c_3, respectively. In particular, if $\mathcal{D}_d(Q_s)$ denotes the total length of the driver route Q_s and $\mathcal{D}_w(P_r)$ is the total length of the porter path P_r, then the cost of a feasible solution can be calculated as follows:

$$C(P, Q) = c_1|P| + c_2 \sum_{Q_i \in Q} \mathcal{D}_d(Q_i) + c_3 \sum_{P_i \in P} \mathcal{D}_w(P_i). \tag{1}$$

The *objective* of the problem is to generate routes for drivers and paths for porters, so that all consignees are visited and their parcels delivered at minimum total cost.

3.2. Tabu Search

Tabu Search (Glover, 1977) has been successfully applied to solve distribution problems, which we also use here to solve the problem described in the previous section. Our Tabu Search implementation is shown in Algorithm1.

The algorithm generates an initial solution using a two-phase constructive algorithm (`constructive_algorithm`, see Section 3.2.1). Then, it improves the solution iteratively using a `local_search` which makes use of a tabu list. The details of the local search are given in Section 3.2.2. Every n_{2opt} iterations, the algorithm tries to further improve porter paths applying an adapted version of the 2OPT algorithm (`adapted_2opt` described in Section 3.2.3). The resulting paths which are short enough to be combined together are further merged and re-optimised with 2OPT (lines 13-20). As the last step of the algorithms the driver routes are re-optimised by solving a Travelling Salesperson Problem (TSP) on their route points, which in our implementation is done using OR-Tools[2].

3.2.1. Constructive algorithm

The constructive algorithm is used to produce an initial feasible solution and operates in two phases. The first phase constructs the initial routes for the drivers and paths for the porters, but without any micro-consolidation points, which is obtained by solving a TSP instance for each carrier s over the set $\{d_s\} \cup V_h^s$ of nodes, solved using OR-Tools. To construct a porter path, say P_1, we first randomly insert a consignee $\overline{u} \in V_c$ into P_1. We then choose a consignee $u^* \neq \overline{u}$ that is closest in distance to $\overline{u}$ and insert it into P_1 if it is feasible to do so with respect to the weight, volume and distance constraints. We continue to insert other consignees in the best position that results in the minimum

[2]`https://developers.google.com/optimization/routing/`

Algorithm 1 Tabu Search Algorithm

$\mathcal{I} \leftarrow$ Problem instance
$counter \leftarrow 0$
$tabu\_list \leftarrow \emptyset$
$(P, Q) =$ `constructive_algorithm`$(\mathcal{I})$
while time limit not reached **do**
  if $counter > n_{2opt}$ **then**
    $(P, Q), tabu\_list =$ `local_search`$((P, Q),\ tabu\_list,\ \mathcal{I})$
    $counter = counter + 1$
  else
    for $P_i \in P$ **do**
      $P_i =$ `adapted_2opt`$(P_i, \mathcal{I})$
    end for
    for $P_i, P_j \in P$ **do**
      Concatenate P_i and P_j into P_{ij} (removing the last visit to the micro-consolidation of P_i)
      $P_{ij} =$ `adapted_2opt`$(P_{ij},\ \mathcal{I})$
      if $\mathcal{D}_w(P_{ij}) \leq d_{max}$ **then**
        $P = P \setminus \{P_i, P_j\}$
        $P = P \cup P_{ij}$
      end if
    end for
  end if
end while
for $Q_c \in Q$ **do**
  Solve the TSP defined by the set Q_c and update Q
end for

possible increase in the distance of path P_1. In doing so, we allow for sufficient distance to remain for the porter path to start and finish at a micro-consolidation point. The process is repeated until no more consignees can be added to P_1, and any unassigned consignees at this stage are added to a new porter path, say P_2, constructed by the same procedure, and continue to do so until all consignees are assigned to a porter path.

In the second phase, we insert micro-consolidation points into the solution so that each porter path starts at a micro-consolidation point. As each micro-consolidation point is at an intersection of a porter path and a driver route, this will require extending the path as well as the routes yielding a distance increase in both. As the starting point of a given porter path, we therefore choose a micro-consolidation point that minimises the sum of the increased driver and porter distances. Once this is done, each porter path is extended to end a micro-consolidation point by choosing one that is closest to the final consignee visited on the path. As the first consignee of each porter path is selected randomly, each run of the algorithm is likely to result in a different set of paths. A full pseudo-code of the constructive algorithm is given in the Online Appendix.

3.2.2. Local Search

Local search is used to iteratively refine an initial feasible solution by exploring a neighbourhood of solutions generated by moves that reassign a consignee from one porter path to another. This, in turn, changes the number and path of the porters and the micro-consolidation points used, and consequently the routes of the drivers. Each time a move is performed, it is added to a Tabu List which prescribes the number of iterations for which the move cannot be used again.

Whilst the neighbourhood defined above is simple and efficient, there are problem-specific cases that need to be considered. To illustrate, consider a move that reassigns consignee $v \in P_r$ to another path. The cost of the move is evaluated by computing three parts: the reduction in cost derived from removing v from P_r, the cost of inserting v into all other paths in $P \setminus \{P_r\}$ or into a new path, and any increase or reduction in cost due to changing any of the driver routes.

The following conditions are used in computing the cost of removing v from P_r:

- if $P_r \setminus \{v\} \cap V_c = \emptyset$ then porter r is no longer needed, in which case is removed from the solution and its cost deducted from the total cost.
- let $i^* \in P_r \cap V_m$ be the micro-consolidation node where parcels to be delivered to consignee v are picked up. If i^* is the only micro-consolidation node visited by porter r, then it is from P_r, reducing further the cost of P_r. In addition, if the driver delivering the parcels of consignee v to i^* is not delivering any other parcel there, then i^* is removed from the route of the driver too.

After removal, the consignee v is either inserted into an existing porter path, or a new path is created exclusively for this consignee. As for the former, we only consider insertions that would result in feasible paths that respect the capacity constraints. The cost of insertion is computed by considering the following cases.

- *The removed consignee v is inserted into an existing path $P_{r'}$*, where $r' \neq r$.
  Let $q^* \in \{2, \ldots, q-1\}$ denote the position in $P_{r'}$ into which consignee v will be inserted. Let v belong to a carrier s. If porter r' is already visiting a micro-consolidation point that is also visited by the vehicle of carrier s, and the porter has enough capacity left to accommodate the parcel, then v can be inserted into

position q^*. Otherwise, the insertion is performed by choosing the cheapest of the following options:

(1) Inserting one of the previous micro-consolidation nodes visited by path r' into the vehicle route of carrier s, provided that there is sufficient capacity left.
(2) Inserting a new micro-consolidation node to the path of porter r' before position q^*. In this case, we consider the insertion of all micro-consolidation nodes, including those already visited by porter r' into all feasible positions. Due to weight and volume restrictions, not all positions may be feasible (e.g., if the micro-consolidation point is visited early on in the path, the porter might not have capacity to keep the parcel in the bag whilst delivering or picking up others.
(3) Inserting a new micro-consolidation point into both the route of vehicle s and the path of porter r'. In this case, we perform the best feasible insertion of the micro-consolidation node for both the porter and the driver.

- *Consignee v is assigned to a new porter.*
  In this case a new path will be created with three nodes, namely two micro-consolidation nodes, one for the start and one for the end of the path, and the consignee v. The end node is chosen as the closest micro-consolidation point to v. As for the start node, we consider all the micro-consolidation points available, and choose the one that yields the lowest increase in cost, where both the walking distance from v and the extra cost arising from the need of the driver to visit that micro-consolidation node are taken into account. This is done by checking the cost of the cheapest insertion of micro-consolidation node in the driver route. Note that there might be some micro-consolidation nodes that will not require such cost because the driver of the carrier might be already visiting them.

The neighbourhood described above does not explicitly consider the move of micro-consolidation points, but these are still explored during the course of the Tabu Search within the special cases described above. Once a move on a consignee v is applied, it is moved into a Tabu List, the size t_n of which specifies the number of iterations for which v cannot be moved again.

3.2.3. Adapted 2OPT heuristic

In any feasible porter path $P_r = \{1, \ldots, m_i, v_i, \ldots, v_j, m_j, \ldots, n_r\}$, any permutation of consignees located between two consecutive micro-consolidation points m_i and m_j is also feasible. To obtain the best possible sequence of consignees $\{v_i, \ldots, v_j\}$, one can solve a TSP instance on these nodes, which we perform using the well known 2OPT heuristic. In order to adapt 2OPT to work with full routes, we need to make sure that permutations of parts of the path containing micro-consolidation nodes result in a porter path that is still feasible. The pseudocode for this procedure is presented in Algorithm 2. In a nutshell, for each pair of nodes i and j with the exception of the first and the last micro-consolidation nodes we analyse whether the classical 2OPT movement, i.e., reversing the visiting order of the nodes between i and j (see line 6 of Algorithm 2), is feasible and can reduce the distance travelled by the current porter.

It is possible that the routes include micro-consolidation points in consecutive positions, visited by the same driver. Such redundant visits are removed from a route when detected.

Algorithm 2 Adapted 2OPT

Require: $P \subset V_m \cup V_c$

Compute *current_cost* of P
$improvements = TRUE$
while *improvements* **do**
  for $i, j \in \{2, \ldots, |P_i| - 1\}$, $(i < j)$ **do**
    $\overline{P} = P$
    $\overline{P} = flip(\overline{P}(i : j + 1), 0))$
    Compute *new_cost* of $\overline{P}$
    if Order micro-consolidation points in $\overline{P}$ is not valid **then**
      continue (check next iteration of i, j)
    end if
    if Weight (Volume) on each nodes $i \in \overline{P}$ exceeds W (L) **then**
      continue (check next iteration of i, j)
    end if
    if *new_cost*$<$ *current_cost* **then**
      $P = \overline{P}$
      break (end for loop)
    end if
  end for
end while

4. Computational Experiments

In this section, we present computational results obtained by applying the Tabu Search on a real data set. We start by describing the input data used in the experiments and subsequent sections present results on the value of sharing a portering resource and an exposition on how the benefits could be shared using game-theoretical tools. The fine-tuning and exact values of parameters used in the algorithm are reported in the Online Appendix. The algorithm has been implemented in Python and the results obtained in a single 2.6 GHz core with 16 GB of memory, part of the IRIDIS HPC facility.

4.1. Input data

An instance of our problem contains the information of the daily operation of two carrier companies, including a list of consignees, each of them with a requirement for delivery or collection of one or more items (with known weight, volume and the carrier that serves them) a list of carrier depots, a list of land assets that can be used as micro-consolidation points, and data on workforce and travel costs. In the case where multiple items were to be delivered to one location by the same company, these were merged into a single (larger) item in a pre-processing step.

Parcel data describing numbers of items being delivered to or collected from addresses within the EC3 postcode area of London were obtained from two major carriers, denoted C1 and C2. The data were collected during the first week of June 2018 and each instance corresponds to a weekday (Monday to Friday). As these data did not include any reliable parcel size or weight information, we measured and weighed 489 packages (boxes, bags) at one of the operator's depots and randomly sampled volume/weight paired data obtained from this survey. Table 1 summarises the key information on the instances, where AW and AV are, respectively, the average weight (kg) and volume (L) per item, while AWL and AVL represent the average weight (kg) and volume (L) to be delivered or collected from one single location. As the table shows, carrier C1 deals with much larger volumes as compared to C2, delivering twice or three times the amount of parcels to a larger number of locations.

Table 1.: Details of the number of items, weights and volumes of the items of the estimated carrier operation for one week in central London.

Day	**Carrier**	**Items**	**Locations**	**AW**	**AV**	**AWL**	**AVL**
Monday	C1	300	163	1.12	12.10	2.06	22.27
	C2	102	91	1.04	11.67	1.16	13.09
Tuesday	C1	220	133	1.09	14.22	1.81	23.53
	C2	117	102	1.29	13.93	1.48	15.98
Wednesday	C1	314	188	1.29	13.52	2.16	22.58
	C2	108	94	0.88	11.52	1.01	13.24
Thursday	C1	272	154	1.19	13.36	2.11	23.60
	C2	114	101	1.45	14.21	1.63	16.04
Friday	C1	285	178	1.16	12.60	1.86	20.17
	C2	113	101	0.96	9.38	1.08	10.49

As part of their wider policy to reduce vehicle impacts in the heart of the capital, Transport for London have expressed an interest in how their various land assets might be used for logistics activities to reduce negative externalities associated with delivery. To this end, 13 potential sites for micro-consolidation points were extracted from a list of over 60 covering the EC3 postal area provided by Transport for London, as shown in Figure 2. These ranged from underground stations to buildings and areas of land where a vehicle or fixed asset could be sited.

To estimate walking and driver costs, we first computed the pairwise distances between all consignees, micro-consolidation points and depots (origin-destination matrices) and then assigned a cost per distance unit. Both the walking and driving distances used in the experiments have been computed at street level, using OpenStreetMap maps along with routing software Open Source Routing Machine (OSRM) (Luxen and Vetter, 2011). The cost of the different delivery workers has been derived based on a previous portering trial with one of the companies, and are £25 (driver) and £12.5 (porter) per hour worked, and converted into meters assuming driving and walking speeds of 3 m/s and 1.2 m/s, respectively. We also assume a fixed cost of £20 per porter used.

We work under the assumption that both companies can use a single vehicle to perform the operation of one day over the EC3 area, since this reflects the business-as-usual operations of the two carriers working on this part of London. This vehicle is always necessary, therefore it has no fixed assigned. In contrast, the unit cost of the porters plays an important role as, depending on the allocation to routes it might be possible to perform the same amount of deliveries with a variable number of porters and total travelled distances.

4.2. Quantifying the value of sharing a portering resource

In this section we compare the cost of operation of the two proposed models, SI and SPR, using the original data set with two carriers described in Section 4.1. To quantify

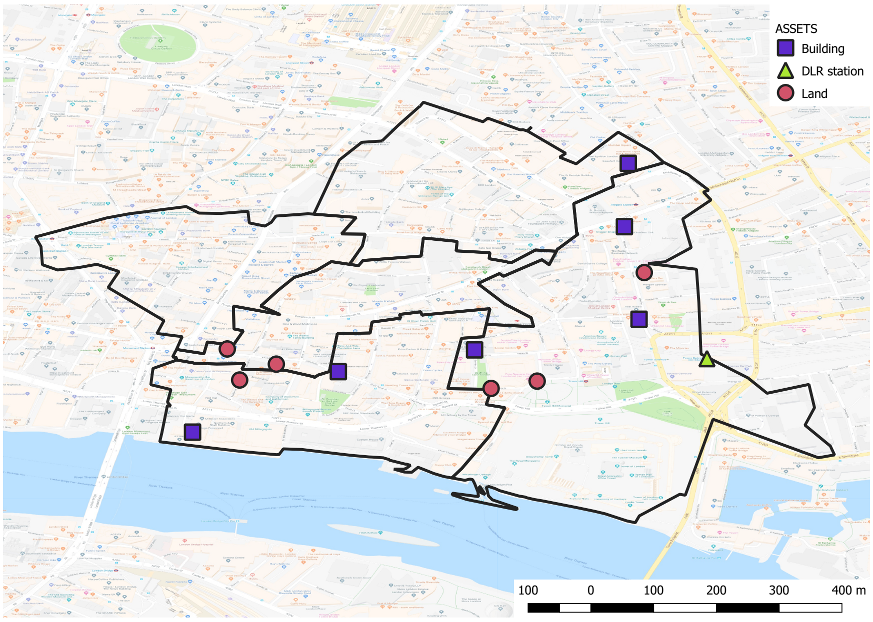

Figure 2.: Distribution of some locations that Transport for London (TfL) has identified in the central the EC3 postcode district with potential to be used as micro-consolidation points. These include buildings, Docklands Light Railway (DLR) stations and other land assets (*ie.* terrain owned by TfL). Postcode boundaries data source: Pope (2017). Attribution: Postal Boundaries ©GeoLytix copyright and database right 2012 Contains Ordnance Survey data ©Crown copyright and database right 2012 Contains Royal Mail data ©Royal Mail copyright and database right 2012 Contains National Statistics data ©Crown copyright and database right 2012. Map data ©2018 Google.

the effect of problem parameters on the overall cost, we explore different configurations that vary the maximum allowed size (weight and volume) of porterable items, the capacity of porter bags (weight and size) and the maximum distance that porters can walk on a given day. We denote each configuration by a triplet $(\gamma_1, \gamma_2, \gamma_3)$. The first setting γ_1 is the maximum item size specified by the weight and volume (kg, L), for which we use the settings A $= (5, 50)$, B $= (8, 80)$ and C $= (10, 100)$. The second setting γ_2 indicates the bag capacity in weight and volume (kg, L), for which we use A $= (20, 200)$ and C $= (25, 250)$. The last setting shows the maximum travel distance d_{max}, for which we consider the values $\gamma_3 = 8, 10, 12, 14$ km. For example, the triplet (C, A, 14) indicates that the algorithm is run where porters are allowed to carry items weighing up to 10kg and with a volume of up to 100L each, using bags with capacity of 20kg in weight and 200L in volume, and where each porter is allowed to walk a maximum distance of 14km. For each configuration, we run the algorithm ten times and calculate the average daily costs. Table 2 shows the resulting weekly cost averages for models SI and SPR. More detailed results, broken down by day of the week and company, are available on the Online Appendix of this article.

As the results indicate, there is always a positive impact on sharing the portering resource, which yields costs reductions that range from 4.69% to 14.23%, irrespective of the configuration used. While there does not seem to be a clear pattern that favours any particular configuration, we observe that the maximum size of an item that can be carried by a porter has an impact on the amount of driving. The immediate consequence

Table 2.: Weekly cost averages (£) for the two sharing models SI and SPR

Configuration	**SI**	**SPR**	**Cost reduction**
(A, A, 8)	270.92	244.37	9.80%
(A, C, 8)	259.60	228.42	12.01%
(B, A, 8)	281.60	255.75	9.18%
(B, C, 8)	267.16	237.67	11.04%
(C, A, 8)	281.44	254.84	9.45%
(C, C, 8)	266.84	237.48	11.00%
(A, A, 10)	254.82	224.04	12.08%
(A, C, 10)	246.58	211.50	14.23%
(B, A, 10)	261.48	233.84	10.57%
(B, C, 10)	249.80	218.11	12.69%
(C, A, 10)	262.25	231.42	11.76%
(C, C, 10)	248.83	218.20	12.31%
(A, A, 12)	224.66	208.13	7.36%
(A, C, 12)	214.95	199.75	7.07%
(B, A, 12)	241.90	216.85	10.36%
(B, C, 12)	221.93	202.29	8.85%
(C, A, 12)	245.30	220.53	10.10%
(C, C, 12)	223.02	201.03	9.86%
(A, A, 14)	216.22	203.54	5.87%
(A, C, 14)	212.14	189.64	10.60%
(B, A, 14)	219.00	208.73	4.69%
(B, C, 14)	210.69	196.38	6.79%
(C, A, 14)	219.88	206.93	5.89%
(C, C, 14)	207.64	196.33	5.45%
Min.	207.64	189.64	4.69%
Max.	281.60	255.75	14.23%
Avg.	242.03	218.57	9.54%

of varying this parameter is a change in the split between the items that are delivered by porters or drivers.

To understand the implications of the effect of porterable item sizes on the amount of driving needed, we present the distance driven for each day of the week in Table 3, for different item size limits as specified above. In particular, the table shows, for each of the two operational models SI and SPR, the amount of driving undertaken (in km) for each day of the week, along with the average in the last column.

Table 3.: Total driving distance (km) for the different delivery models and limits on porterable items.

Model	**Item size**	Monday	Tuesday	Wednesday	Thursday	Friday	Average
SI	A	28.4	31.2	29.5	30.3	30.0	29.9
	B	26.9	29.3	26.1	27.3	28.9	27.7
	C	24.2	28.6	25.4	26.2	26.4	26.2
SPR	A	28.5	31.2	29.6	30.2	30.1	29.9
	B	26.9	29.4	26.1	27.3	29.1	27.8
	C	24.2	28.7	25.3	26.2	26.5	26.2

The results in Table 3 suggest that, whilst there does not appear to be a significant difference between SI or SPR, increasing the weight and volume limits for single items has a noticeable effect in the driving in either case. In particular, in the SI case, the reduction in driving is 2.2km on average when the maximum item size limits increase from $(\bar{W}, \bar{L}) = (5, 50)$ to $(8, 80)$, and a further average reduction of 1.5km is achieved when items of sizes up to $(10, 100)$ are allowed, with these figures being very similar for the SPR case.

4.3. Splitting the benefits of sharing

The results presented in previous section suggest that a shared portering resource can result in significant savings in the overall cost. The cost of the SPR model is to be jointly borne by the collaborating carriers, although the exact amount that each carrier would have to contribute is not immediate. This raises the question of how the benefits of using a third party portering resource should be divided. The successful application of collaborative models in practice is closely linked to this question, as companies will naturally seek to receive a discount in relation to their contribution to the overall delivery operation. This question can be answered by resorting to concepts found in co-operative game theory, such as the core, the least core, the nucleolus or the Shapley value (see, e.g. Krajewska et al., 2008; Yang et al., 2020).

Whilst it is not our intention to show the application of these concepts here as they have been well-studied elsewhere, we suffice to provide an example to illustrate how the costs might be allocated on a collaboration where three carriers are involved. Using our data set, we construct instances where three carriers have the option to collaborate. To generate carrier data, we treat the five days of C1 as five large carriers, labelled L1–L5, and the data for each of the five days of C2 as five small carriers, labelled S1–S5. Using these ten carriers, we generate a total of 120 instances, considering combinations of triplets combining large and small carriers as follows: (Large, Large, Large), (Large, Large, Small), (Large, Small, Small) and (Small, Small, Small).

For each instance, we executed the Tabu Search for five times, each under a 10 minute time limit, for each possible coalition. In particular, under the SI model each carrier operates individually, whereas under the SPR model there are coalitions formed by two of the three carriers, and a grand coalition of three carriers altogether. For the purposes of this analysis and in order to quantify the value of sharing as much as possible, we use the best value obtained out of the five runs for each coalition.

With these values in hand, we look at two aspects of splitting the benefit: is it worth collaborating and, if it is, how should the benefits be split? To answer the first question, we looked at the concept of the core of a cooperative game. In particular, a point in the core of a cooperative games shows how the total cost can be apportioned between collaborating players so that none would be better off by leaving the coalition or by choosing a smaller coalition (Chalkiadakis et al., 2011). More formally, for a set $C = \{\text{C1, C2, C3}\}$ of collaborating carriers, let $\Omega(C)$ denote the total operating cost of a coalition C. If the vector $\omega = \{\omega_1, \omega_2, \omega_3\}$ shows the amount that each carrier should contribute, the core is formed by the set of payouts ω for which the following two conditions hold:

$$\sum_{i \in C} \omega_i = \Omega(C) \tag{2}$$

$$\sum_{i \in S} \omega_i \geq \Omega(S) \qquad \forall S \subset C. \tag{3}$$

The experiments indicated that, out of the 120 instances analysed, the core was never empty, indicating that there is a financial incentive for all carriers to adopt the SPR model as opposed to using SI. Nevertheless, the core is a set and not all the allocations contained in it are necessarily fair in terms of contribution, *i.e.* the couriers that contribute most to the coalition might not get the highest utility. To evaluate this further, we computed the Shapley Value (Chalkiadakis et al., 2011), which finds a fair allocation of costs for carriers involved in the coalition, and the nucleolus (Schmeidler, 1969), which seeks to minimise dissatisfaction across all possible coalitions. The nucleolus was calculated using the algorithm and implementation from Benedek et al. (2021).

Table 4 shows the results of these experiments, where the first column shows the size of the collaborating partners, the second column shows the number of instances tested, the third column shows the absolute savings in monetary terms (£) for the set (on average) and the fourth column shows the average savings achieved by the SPR model over the SI model, relative to the total cost of the operation, which are the same regardless of the cost allocation. The following four columns show the range of savings produced when allocating costs with the Shapley Value and the nucleolus. Please refer to the Online Appendix of this article for the individual results of each of the 120 instances.

Table 4.: Number of instances generated of each type, based on 5 large (L1 - L5) and 5 small (S1 - S5) instances, and their savings in absolute (£) and relative costs when using the SPR for the delivery and either the Shapley Value or the nucleolus to allocate the costs.

Type	**Instances**	**Abs. savings**	**Rel. savings**	**Shapley Value Range**		**Nucleolus Range**	
		Avg	Avg	Min	Max	Min	Max
(L, L, L)	10	85.93	18.35%	10.07%	23.83%	5.17 %	24.19 %
(L, L, S)	50	76.03	19.06%	11.73%	30.15%	11.4 %	32.08 %
(L, S, S)	50	73.33	22.98%	7.02%	37.11%	2.36 %	45.12 %
(S, S, S)	10	78.24	27.39%	19.62%	35.65%	14.37 %	43.68 %

The collaboration between smaller companies seems more profitable in relative terms, saving on average more than a quarter of the total operation. However, the saving is similar to those of the other coalitions in absolute terms. Nonetheless, all collaborations resulted in savings, with the nucleolus allocations having a wider range of savings (2.36%, 45.12%), compared to the Shapley Value (7.02%, 37.11%) The distribution of the savings under the different allocations can be seen in Figure 3, that also shows that the median payment of large companies tends to be lower with the nucleolus allocation.

Finally, in our experiments the Shapley Value lied outside of the core in 10 out of 120 the instances tested, while the nucleolus, by definition, always belongs to the core when it is non-empty. This suggest that in this case a nucleolus allocation might be preferable to ensure the stability of the SPR business model.

5. Conclusions

This study investigated two innovative business models for collaboration on last-mile delivery in an urban area. We proposed an optimisation tool based on a tabu search algorithm that can be used by decision managers for three main purposes. Firstly, it can provide a low-risk approach to evaluate the potential benefit of switching to a portering model using shared micro-consolidation nodes in the city. Secondly, it shows the benefit of collaboration between competitors by using a shared portering resource

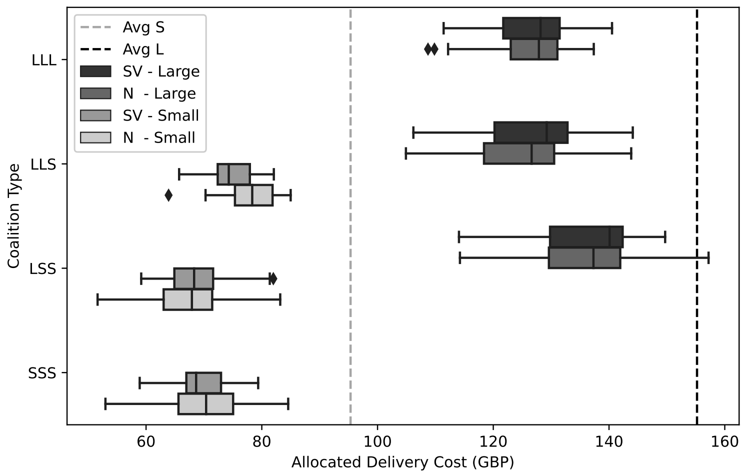

Figure 3.: Distribution of the allocation of delivery costs using either the Shapley Value (S) or the Nucleolus (N), for different coalition types involving small (S) and large (L) companies. The vertical lines represent the average delivery cost of small and large companies when they do not collaborate.

over sharing the infrastructure only, which in our case study was shown to reduce up to 14% of the weekly transportation costs, as well as providing a sound basis for investigating how this benefit should be shared between carriers. Should companies be willing to adopt these models, the presented optimisation algorithms could be used as the basis of specialised software that would enable the allocation and routing of parcels in practice. A challenge of this implementation would be the availability of data (eg. parcel weights/dimensions), which might not to be readily available for some couriers. This is also true for micro-consolidation points, which in practice would have limitations in capacity and access that have not been modelled in this study. Finally, we provide an indication on how savings might be shared and demonstrate that, for our case study, the size of the operation did not have an impact on the profitability of the collaboration. A limitation of this analysis is that it is based purely on financial incentives, and it does not take into account market dynamics. It is possible that some companies would refuse to enter a coalition that benefits a direct competitor, even if it also benefits them. However, city regulations motivated by environmental targets can incentivise this kind of collaboration, even in a competitive environment.

Acknowledgments

The authors gratefully acknowledge the EPSRC for funding this work through its financial support of Freight Traffic Control 2050 (`www.ftc2050.com`), EPSRC Grant Reference: EP/N02222X/1. Responsibility for the contents of the paper rests with the authors.

References

Allen, J., Bektaş, T., Cherrett, T., Bates, O., Friday, A., McLeod, F., Piecyk, M., Piotrowska, M., Nguyen, T., and Wise, S. (2018a). The scope for pavement porters: addressing the challenges of last-mile parcel delivery in London. *Transportation Research Record: Journal of the Transportation Research Board*, 2672(9):184–193.

Allen, J., Piecyk, M., and Piotrowska, M. (2018b). Same-day delivery market and operations in the UK. Technical report, University of Westminster.

Allen, J., Piecyk, M., Piotrowska, M., McLeod, F., Cherrett, T., Ghali, K., Nguyen, T., Bektaş, T., Bates, O., Friday, A., Wise, S., and Austwick, M. (2018c). Understanding the impact of e-commerce on last-mile light goods vehicle activity in urban areas: The case of London. *Transportation Research Part D: Transport and Environment*, 61:325–338.

Arrieta-Prieto, M., Ismael, A., Rivera-Gonzalez, C., and Mitchell, J. E. (2022). Location of urban micro-consolidation centers to reduce the social cost of last-mile deliveries of cargo: A heuristic approach. *Networks*, 79(3):292–313.

Benedek, M., Fliege, J., and Nguyen, T.-D. (2021). Finding and verifying the nucleolus of cooperative games. *Mathematical Programming*, 190(1):135–170.

Chalkiadakis, G., Elkind, E., and Wooldridge, M. (2011). Computational aspects of cooperative game theory. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 5(6):1–168.

Clarke, S., Allen, J., Cherrett, T., McLeod, F., and Oakey, A. (2018). Report on the portering trial, TfL consolidation demonstrator project.

Coindreau, M.-A., Gallay, O., and Zufferey, N. (2019). Vehicle routing with transportable resources: Using carpooling and walking for on-site services. *European Journal of Operational Research*, 279(3):996–1010.

Fikar, C. and Hirsch, P. (2015). A matheuristic for routing real-world home service transport systems facilitating walking. *Journal of Cleaner Production*, 105:300–310.

Gansterer, M. and Hartl, R. F. (2018). Collaborative vehicle routing: a survey. *European Journal of Operational Research*, 268(1):1–12.

Glover, F. (1977). Heuristics for integer programming using surrogate constraints. *Decision Sciences*, 8(1):156–166.

Krajewska, M. A., Kopfer, H., Laporte, G., Ropke, S., and Zaccour, G. (2008). Horizontal cooperation among freight carriers: request allocation and profit sharing. *Journal of the Operational Research Society*, 59(11):1483–1491.

Luxen, D. and Vetter, C. (2011). Real-time routing with OpenStreetMap data. In *Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems - GIS '11*, page 513, New York, USA. ACM Press.

Martinez-Sykora, A., McLeod, F., Lamas-Fernandez, C., Bektaş, T., Cherrett, T., and Allen, J. (2020). Optimised solutions to the last-mile delivery problem in London using a combination of walking and driving. *Annals of Operations Research*, 295(2):645–693.

McLeod, F., Cherrett, T., Bektas, T., Allen, J., Martinez-Sykora, A., Lamas-Fernandez, C., Bates, O., Cheliotis, K., Friday, A., Piecyk, M., et al. (2020). Quantifying environmental and financial benefits of using porters and cycle couriers for last-mile parcel delivery. *Transportation Research Part D: Transport and Environment*, 82:102311.

Muñoz-Villamizar, A., Montoya-Torres, J. R., and Vega-Mejía, C. A. (2015). Non-collaborative versus collaborative last-mile delivery in urban systems with stochastic demands. *Procedia CIRP*, 30:263–268.

Nguyen, T., Bektaş, T., Cherrett, T., McLeod, F., Allen, J., Bates, O., Piotrowska, M., Piecyk, M., Friday, A., and Wise, S. (2019). Optimising parcel deliveries in London using dual-mode routing. *Journal of the Operational Research Society*, 70(6):998–1010.

Ofcom (2020). Annual monitoring update on the postal market. Financial year 2018-19. Technical Report.

Park, H., Park, D., and Jeong, I.-J. (2016). An effects analysis of logistics collaboration in last-mile networks for cep delivery services. *Transport Policy*, 50:115–125.

Pope, A. (2017). GB Postcode Area, Sector, District, [Dataset].

Ranieri, L., Digiesi, S., Silvestri, B., and Roccotelli, M. (2018). A review of last mile logistics innovations in an externalities cost reduction vision. *Sustainability*, 10(3):782.

Schmeidler, D. (1969). The nucleolus of a characteristic function game. *SIAM Journal on Applied Mathematics*, 17(6):1163–1170.

Transport for London (2017). TLRN performance report quarter 2, 2017/18.

Yang, F., Dai, Y., and Ma, Z.-J. (2020). A cooperative rich vehicle routing problem in the last-mile logistics industry in rural areas. *Transportation Research Part E: Logistics and Transportation Review*, 141:102024.