


Journal Pre-proof

Improving the efficiency of patient diagnostic specimen collection with the aid of a multi-modal routing algorithm

Andy Oakey, Antonio Martinez-Sykora, Tom Cherrett



PII: S0305-0548(23)00129-6
DOI: <https://doi.org/10.1016/j.cor.2023.106265>
Reference: CAOR 106265

To appear in: *Computers and Operations Research*

Received date: 9 March 2022
Revised date: 24 April 2023
Accepted date: 25 April 2023

Please cite this article as: A. Oakey, A. Martinez-Sykora and T. Cherrett, Improving the efficiency of patient diagnostic specimen collection with the aid of a multi-modal routing algorithm. *Computers and Operations Research* (2023), doi: <https://doi.org/10.1016/j.cor.2023.106265>.

This is a PDF file of an article that has undergone enhancements after acceptance, such as the addition of a cover page and metadata, and formatting for readability, but it is not yet the definitive version of record. This version will undergo additional copyediting, typesetting and review before it is published in its final form, but we are providing this version to give early visibility of the article. Please note that, during the production process, errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

© 2023 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Improving the efficiency of patient diagnostic specimen collection with the aid of a multi-modal routing algorithm

Andy Oakey^a, Antonio Martinez-Sykora^b, Tom Cherrett^a

^a*Transportation Research Group, University of Southampton, Boldrewood Innovation Campus, Southampton, SO16 7QF, United Kingdom*

^b*Southampton Business School, University of Southampton, Highfield Campus, Southampton, SO17 1BJ, United Kingdom*

Abstract

The Sustainable Specimen Collection Problem (SSCP), in which diagnostic specimens are collected from GP surgeries (doctor's office/clinics) and subsequently transported to a hospital laboratory for analysis using more sustainable transport modes, is introduced in this paper. Using a weighted objective function, we solve a new multi-objective problem using cycle consolidation to limit driving time and the numbers of vans used whilst improving overall service quality, reducing costs and emissions. This particular heterogeneous vehicle routing problem is explored and applied to two real-world case studies in the UK, where 97 and 22 sites (respectively) are currently served, using a column generation based heuristic algorithm with some additional improvement heuristics. The results demonstrated a potential improvement in the system's maximum delivery time between 41% and 74% compared to business-as-usual activity using solely road vehicles. Road vehicle (van) fleets could be reduced by up to 40%, and the total driving time across the fleet by between 41% and 65%. Operational costs were estimated to increase by up to 38%, though additional workloads for gig-economy cycle couriers and improvement in specimen quality and service reliability may make this trade-off worthwhile. Tailpipe CO₂ emissions were also reduced by up to 43%. The proposed algorithm was effective, reducing computational time by up to 99% whilst achieving an average of 5% deviation from optimality.

Keywords: multimodal, diagnostic specimens, routing, pathology, specimen collection problem, mixed-mode, multi-objective, SSCP

1. Introduction

This paper aims to improve the logistics of a local healthcare diagnostic specimen collection system by combining driving and cycling. The objective is to carry samples from community healthcare facilities to an analysis laboratory at a nearby hospital; reducing the time samples spend in transit, whilst minimising the use of fossil fuelled vehicles as much as possible.

The study focuses on two case studies based in the south of the UK in which collections are made from (i) up to 97 general practitioner (GP) clinics/surgeries (often known as doctor's offices in other countries, hereafter referred to as surgeries) using a fleet of 10 vehicles; and (ii) up to 22 surgeries served by a fleet of 3 vehicles, each day. In the current business-as-usual (BAU) operation, only road vehicles (vans) are used, and no cycling logistics is currently considered. The vehicles are based at the hospital and visit the majority of surgeries on regular rounds, whilst other surgeries are served by collections on an ad-hoc basis outside of this arrangement. In principle, this

problem somewhat aligns with a vehicle routing problem (VRP) where the objective is to reduce the longest route, the total number of vehicles used, and the total driving time across the fleet (assuming no capacity constraints are applied to the vehicles). All specimens must be collected daily for analysis and are currently based on a timetabled appointment schedule such that samples are not ‘bled’ from patients after the final collection.

The aim of using cycle couriers is to consolidate the specimens from certain surgeries so that the primary collection vans do not have to visit so many collection points, reducing time, mileage, and tailpipe emissions. Specific surgeries are chosen to act as consolidation points with cycle routes serving the localised surgeries around those consolidation points. To enable cyclists to be used on an ad-hoc basis in this, this problem considers the use of gig-economy cyclists, who can typically be recruited on-demand via various providers to maintain flexibility over payment and workforce structures, particularly in city areas where collection densities are higher. Outside of peak meal-time hours, work for these riders is often sparse (Lord et al., 2020), meaning that they are an asset that can be put to use in other services such as medical deliveries. Hence, one research question explored in this paper is to identify what additional work can be generated for gig-economy workers outside of peak hours to enhance their potential income. This business model is used with the aim of: i) reducing the time samples spend in transit and the environmental impact of deliveries, and ii) increasing off-peak work opportunities for gig-economy cyclists who are often not paid for unoccupied time (Lord et al., 2020), typically outside of specific peak times of day (Bernal, 2020; On-Demand Workers Australia, 2018).

The number of vehicles used in this problem also dictates the performance of the system with regards to delivery times, congestion and emissions. Additional vehicles enable greater reductions in delivery times, as fewer sites need to be served, whilst fewer vehicles cut congestion and emissions through reductions in stem mileage. The problem presented in this paper captures this trade-off, where we formalise a new multi-objective problem using a weighted objective function, and propose a new algorithm capable of efficiently producing solutions. The key contributions of this paper are: (i) quantifying the extent to which the delivery times of diagnostic specimens can be reduced by a cycle courier supported delivery network; (ii) understanding the potential increased workload for the gig-economy under such a system; and (iii) identifying the approximate costs and emissions the delivery network would produce. The analysis also exposes some of the wider logistics challenges faced by specimen collection systems using real-world data from the UK.

1.1. Problem Description

Diagnostic specimens (commonly referred to as ‘pathology’ or ‘laboratory’ ‘samples’ or ‘specimens’) are routinely taken by primary care clinicians across the world to aid in the diagnosis of patient ailments, with roughly 1 in 3 (29%) visits requiring a diagnostic test (Ngo et al., 2017). After being taken, samples require transportation to a nearby laboratory, often at a hospital, for analysis so that patients can be correctly diagnosed and effectively treated (Cherrett and Moore, 2020). In the UK, specimen transport is traditionally carried out by Light Goods Vehicles (LGVs), with samples being taken from local surgeries to hospital laboratories using set vehicle rounds (NHS and Sedman, 2020; NHS and Nixon, 2015). The routing problem could be simplified to one with collections from a set of known nodes (surgeries/clinics) which are then delivered to a single node (hospital).

In the proposed problem, we identify a new approach to serving these sites using multiple-modes, whereby a combination of vans and cargo cycles are used to collect samples within the given time constraints (detailed below), in an arrangement similar to that of a two-echelon VRP. By allowing on-demand cyclists to consolidate loads from multiple sites, and LGVs to subsequently collect either via these consolidation sites or directly, the problem presents a novel approach to the overarching collection logistics challenge.

Samples typically have a fairly short time frame in which they must be analysed, generally within the day they are taken, and have specific requirements for storage and transportation (NHS and Sedman, 2020). As a result, samples should be delivered to the hospital promptly to enable swift diagnosis and maximise the effective use of laboratory staff undertaking the analyses. The Royal College of Nursing in Wales (2020) identified that community COVID-19 testing was significantly slower than in-hospital testing, with less than a third of community test results being turned around the same day, as opposed to 80% in hospital. Part of the problem stems from sample transportation, with many collection systems limiting final sample submission to mid-afternoon (e.g. 3:30 PM (Exeter Laboratory, 2017; Godfrey, 2020)) due to long end-to-end round durations. Samples are not typically permitted to be taken after this point, and any which are will have to travel by taxi or other ad-hoc means (NHS and Sedman, 2020; Wessex Academic Health Science Network, 2020).

As suggested by previous studies and after anecdotal discussions with hospital staff (McDonald, 1972; Allan, 2019), there may be delaying factors within the surgeries and hospital (e.g. administrative, staff scheduling, etc.), but these are beyond the control of the logistics carrier and are ignored in this research. The period of greatest importance is the time spent travelling in the vehicle as controlled conditions cannot be guaranteed, unlike at the origin surgeries and destination laboratory (Anaya-Arenas et al., 2016); thus, minimising the maximum delivery time to the hospital across all surgeries served was defined as the first objective. Reducing the durations of individual rounds, the time to deliver all samples, and the time samples spend in transit may offer the potential for more flexible collection scheduling and later final collections, as well as reducing sample degradation rates and the number of patients requiring repeat diagnostic tests as a result.

It should be noted that this research does not explore the scheduling side of this application because it is understood that collections can be managed as discrete events and current scheduling constraints are not known in sufficient detail due to commercial sensitivity. It is envisaged that the system could be used multiple times in a day, running the optimisation, shortly before performing the collections, for only those surgeries with loads available to collect at that time. The number of daily collections would remain a contractual issue, but those surgeries which produce more samples would likely experience more frequent and later collections. This matter is somewhat decision-maker dependent and could be a further development from the discretised collection approach.

There is an increasing need to reduce congestion and emissions in urban areas which contribute to poor air quality, slow transit times, and anthropogenic climate change, (European Commission, 2011) with policy makers often stating a parallel aim to move to alternative, more sustainable transport modes (European Commission, 2013). Health care providers are responsible for around 5% of the total national carbon dioxide footprint in developed nations (Pichler et al., 2019). Of this, 62% of this contribution can be attributed to medicines, medical equipment, and other supply chain sources (NHS, 2020a). The National Health Service (NHS) in the UK has set a goal to be

net-zero by 2040, and improving the efficiency of logistics operations will be key to achieving this (NHS, 2020a). To support this target, changes to logistics systems are being explored such as mode-shift and adopting different supply-chain management strategies (NHS, 2020a). Effective logistics is key to successful patient care and supply chain operations in healthcare systems and any changes made to existing transport systems to reduce environmental impacts should not affect the overall level of service (Landry and Philippe, 2004; Buntak et al., 2019).

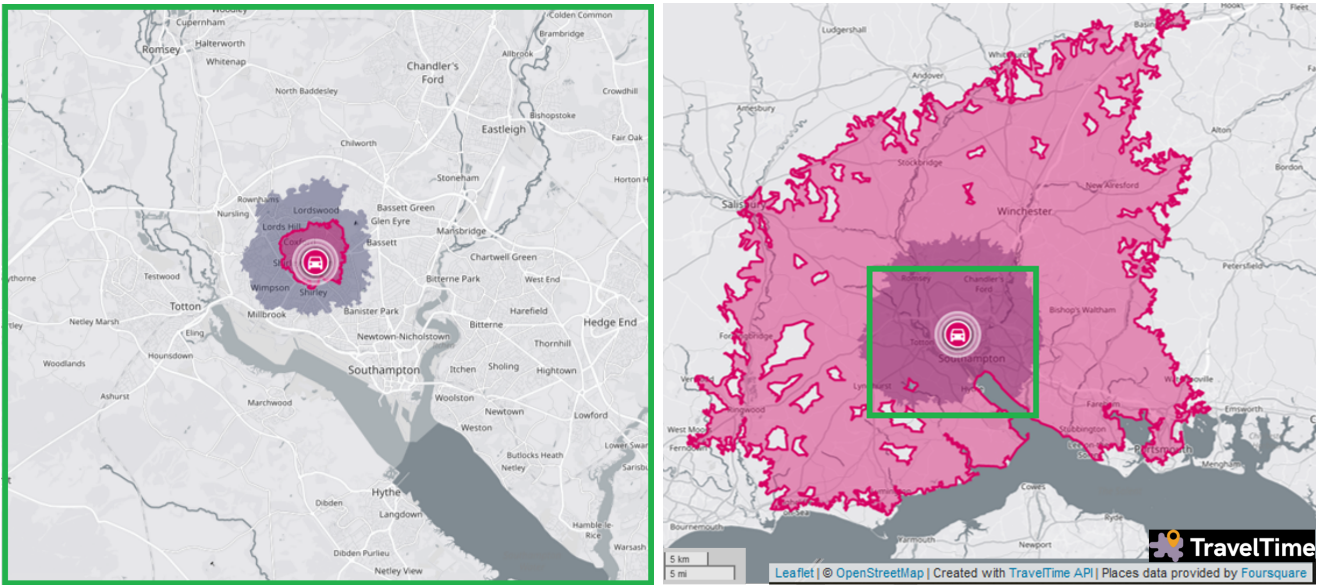
The three-aim objective was therefore to: (i) reduce the maximum time to serve all surgeries, (ii) reduce the number of vehicles needed to undertake the collections, and (iii) reduce the total round duration across the fleet. This is subsequently combined into a weighted-sum single objective to enable decision-makers (e.g., NHS procurement teams and laboratory management) to easily configure the problem according to their specific needs and aims, without the need for additional parameters to define and solve a lexicographic or other multi-objective optimisation (MOO). In this paper, the formulated problem is referred to as the Sustainable Specimen Collection Problem (SSCP).

Whilst emissions are not directly addressed in the objectives of this problem, unlike other green VRPs (Demir et al., 2014; Kramer et al., 2015), the proposed model introduces cycle couriers to help reduce both the driven mileage and time, and the number of vehicles required, using riders to consolidate loads at selected surgeries prior to collections by road vehicles. The introduction of the two-echelon style system inherently improves the sustainability of deliveries, whilst taking advantage of the time benefits offered by road vehicles over longer distances (Gruber and Narayanan, 2019; Conway et al., 2017; Anderluh et al., 2017). Figure 1 demonstrates this concept with respect to one of the case studies used in the present study, with road transport (vans) travelling faster than bicycles, on longer-distance journeys beyond the immediate urban environment. Equally, over shorter distances in the urban environment, bicycles can travel faster, avoiding the effects of traffic and congestion (TravelTime, 2021).

The method by which cycles are introduced in this healthcare-specific logistics problem is of key interest where cycle routes take preference over vans wherever possible, with the goal of minimising driving time. They are therefore not contributing to the objective function as in the two-echelon vehicle routing problem, but are instead helping to minimise van driving time. The problem also differs from other similar routing problems addressed in the literature due to constraints on the length of routes, and the potential for waiting to be incurred if cargo cycles do not precede or synchronise with van collections at consolidation sites.

Chained cycle routes (i.e. one cycle route delivering to another to move samples closer to the destination in further legs) have not been considered in this system due to reasons of practicality. Limiting the system to two echelons reduces any risk with respect to handling, security, and chain of custody (Nybo et al., 2019).

Vehicle costs may vary depending on pricing structures, and the balance of fixed (e.g. insurance) to variable costs (e.g. fuel), among other factors; and are therefore not directly addressed in the problem objectives to limit the effect of uncertainties (FTA, 2020; Grasas et al., 2014). To minimise the number of objectives and simplify the problem, only the delivery time, number of vehicles, and driving time are directly considered. Where savings may be made with respect to sample delivery time, knock-on cost savings may be seen through better human resource and equipment use in the receiving laboratory. Equally, quality of care may be improved, though this is somewhat difficult



(a) 10 minutes travel time by cycle and car from the centroid - cycle travels further in urban environment.

(b) 45 minutes travel time by cycle and car from the centroid - road vehicle travels further once outside of the urban environment.

Figure 1: Travel time isochrones (distance travelled in fixed time) from the same fixed point; Southampton General Hospital (Case Study 1). Red = road vehicle, indigo = bicycle. Green box indicates geographic position of (a) within (b). (TravelTime, 2021) (Base Map ©OpenStreetMap contributors)

to define in terms of cost.

1.2. Related Works - Medical Sample Logistics & Routing

In England, the handling of diagnostic specimen collections is managed by a network of 29 pathology networks (NHS, 2017). Each of those networks operates collaboratively, transporting samples from the community surgeries to their assigned ‘spoke’ hospital for analysis, or onward to the ‘hub’ hospital if the sample requires more specialist analysis. Surgeries are allocated to a single hospital laboratory and, at present, areas do not overlap for contractual reasons (NHS, 2021); i.e., a surgery which is allocated to Southampton’s laboratory cannot also be served by another one.

The Southampton GP network has previously been investigated with respect to its potential for drone served deliveries (Cherrett and Moore, 2020), although the core routing problem relating to ground logistics was not addressed. The problem explored in this paper, the SSCP, is not entirely new, with past research exploring the challenge of collecting diagnostic specimens, but with different objectives such as minimising cost, the number of vehicles used, or the laboratory’s workload (McDonald, 1972; Grasas et al., 2014; Smith et al., 2015). Table 1 summarises and compares the objectives and approaches of the previous studies in this area. In this table, it is evident that problems and solutions can vary significantly; these problems and their differences are discussed below. The present study is also featured in the table for comparison, and is notably different in terms of the dynamics of the problem.

The problem was first discussed by McDonald (1972), who proposed several possible objectives, including the level of service to the patient, the time specimens spent in transit, and the cost of the operation. Cost was selected, though total vehicle travel time of the fleet was used as an approximation as it was more easily quantifiable. Constraints were also applied, most notably on the maximum transit time to limit calculation time and satisfy clinician’s requirements. McDonald (1972) also suggested a procedure that formalised the problem and suggested that time and cost constraints should be applied to maintain sample quality and limit expenditure.

Another study models the “Blood Sample Collection Problem” (BSCP) as a Capacitated Time-Constrained Open Vehicle Routing Problem (CTCOVRP), with constraints relating to capacity, and the maximum route duration after the initial collection (Grasas et al., 2014). Only one objective was set, to minimise the number of vehicles used; a pseudo-objective for cost. The BSCP differs from the SSCP in that road vehicle capacity is assumed to be ‘unlimited’ in the SSCP, given vehicles are only lightly loaded in existing rounds (NHS, 2020b; Quadir et al., 2019). The SSCP also identifies additional objectives relating to reducing driving time, and the delivery time of samples. Results from the BSCP study were produced using a genetic-based algorithm with a heuristic method, and achieved improvements over the BAU with run-times of less than 30 seconds. Grasas et al. (2014) found that the cost of existing operations could be reduced by 20-30% annually, under a 2-hour route time limit. Where a single objective has been used in the BSCP, decision-makers are provided with a solution which only addresses cost; however, it may be more beneficial to offer greater flexibility of choice in this setting, given the difficulty in defining quality-of-care in a cost form (The NHS Confederation et al., 2016).

Smith et al. (2015) proposed the concept of specimen transport as a new problem, addressing the same overall challenge as the SSCP and BSCP, with a variation in the objectives of the proposed problem; aiming to minimise transport costs, balance workloads between laboratories, balance workloads between vehicles, and minimise the number of vehicles used. In the SSCP, the need to balance workloads between laboratories is not needed, as it is assumed that the current

diagnostic networks have been created such that the demands of feeder surgeries can be met by their associated laboratory. With respect to sample lifespan, Smith et al.'s proposed routing problem featured collection/delivery time windows from each surgery to limit sample travel time. This is similar to the SSCP, where a schedule of collections is assumed and travel time is limited. Smith et al. (2015) did not test the model with experimental or case study data, noting that heuristics and metaheuristics would be required to reach good solutions.

Building on this concept, Elalouf et al. (2018) optimised their 'Blood Sample Supply Chain' towards sample expiry and cost through minimising the number of samples that were received outside of a time window after the point of production, in addition to the cost of the operation. Testing multiple approaches (heuristic, tabu search, etc.), it was found that a heuristic was the most effective in terms of the trade off between solution quality and computational time.

A further study investigated the potential to balance workloads in the diagnostic laboratory to maximise the total number of daily processed samples, and minimise laboratory idle and processing time as a primary objective (Yücel et al., 2013a). In a hierarchical optimisation, a secondary objective, to minimise transport costs, was also applied; initially as a single-vehicle optimisation problem (Yücel et al., 2013a), and, subsequently, as a multi-vehicle optimisation problem (Yücel et al., 2013b). Compared to the SSCP, there are significantly more decision variables and constraints which, while providing solutions that address the broader problem more holistically, take impractical lengths of time to solve (c. 4 hrs). Where demands from individual surgeries vary by day, loads are not always predictable, hence more flexibility through shorter run times to change routes may be required. As suggested by McDonald (1972), delays within hospitals and clinics are beyond the scope of the logistics provider, and are ignored by the present study. Furthermore, demand from within the hospital can also be quite changeable and can be difficult to predict (Allan, 2019).

Similar to this research, the Vehicle Routing and Scheduling Algorithm (VeRSA) for specimen deliveries proposed by Zabinsky et al. (2020) sought to reduce the delivery time of samples by optimising the final delivery time across all routes after an initial point when samples were ready to collect. Zabinsky et al.'s study also addressed the matter of vehicle re-use through a scheduling element of the problem. Using an effective branch and bound approach, the solutions found were often optimal, though in some larger instances, the approach took a considerable length of time.

The Biomedical Sample Transportation Problem (BSTP) proposed by Anaya-Arenas et al. (2016) sought to minimise travel distance whilst serving the demands of clinics, within given time windows, under the general assumption that at the origin surgery or destination lab, the sample can be held in controlled conditions. This meant that whilst time was not an immediate constraint, samples should not be waiting for more than a few hours, and their time in transit should be minimised. A similar assumption is also used in the SSCP, based on discussions with hospital staff (Allan, 2019), and reinforced by Wilson (1996) who suggested transport should be developed to limit damage to samples through prompt delivery and controlled intermediate storage. Constraints are applied to limit the maximum transportation time (180 mins) and maximum driver working time (Anaya-Arenas et al., 2016). In a series of experiments, solutions were found, though the study highlighted the computational demands of large numbers of instances limiting the success of the system (within a 1-hour time limit).

The SSCP most closely aligns with the BSTP study by Anaya-Arenas et al. (2016), with the assumptions around temporary sample holding, and the importance of transit time being fundamental to solving the problem. It should be noted that the SSCP simplifies the scheduling element of this problem where computational time is important to enable recalculation based on the demand experienced at the time of use. Naji-Azimi et al. (2016) further developed part of the work by Anaya-Arenas et al. (2016) by testing alternative approaches and introducing a constraint to prevent overloading the laboratory with samples at any point.

A slight variation on the core medical specimen problem was explored by Kergosien et al. (2014), with samples being taken at patients' homes by nurses. The objective of the study was to optimise the routing of nurses who take the samples in terms of minimising cost and travel time, whilst maximising the number of samples processed. Like many of the other studies, time constraints were applied, though, in this instance, they were only given to critical samples. Whilst not a direct comparison to the research in this paper, this does highlight the need for limiting the time samples spend in transit and out of controlled conditions.

Other medical collection-delivery optimisation problems have focused on blood stocks and humanitarian disaster relief, where the core objective has been to maximise the quantities delivered whilst reducing costs (Lodree et al., 2016), often involving constraints on transit time. Doerner and Hartl (2008) discussed the movement of blood donations for processing and delivery to hospitals. They modelled the problem of collections as one of multiple interdependent time windows, with the main constraint being donations requiring urgent processing throughout the day, with the sole objective of minimising cost (Doerner and Hartl, 2008).

Other medical-related studies have also sought to improve on the environmental effects of logistics, such as Ettazi et al. (2021), who considered fuel consumption in their routing problem for at-home care. The problem required synchronisation and precedence and used a meta-heuristic to solve instances. The main limitation of the approach was that feasible solutions were not returned when solving large instances. In a similar vein, Liu et al. (2013) investigated at-home health care deliveries, optimising for the cost of the overall operation using a genetic algorithm to solve the problem, with a consideration of precedence when coordinating patients and medicines.

Other relevant problems, such as do C. Martins et al. (2021) discussed the speedy optimisation of two-echelon VRPs in a medical/aid context using a heuristic approach to solve large instances, where rapid deployment was important, such as in humanitarian disaster relief. In a small set of test instances, the problem was solved within a few seconds, whilst larger test cases took a few minutes. The study used only aerial drones as the mode, making the decision process less challenging, unlike the SSCP, where multiple modes are considered. Nonetheless, the need for fast computation is still relevant to the SSCP if the use of discrete collections is to be adopted.

Osaba et al. (2019) modelled the collection of pharmacological waste as a clustered VRP with numerous complexities, solving the problem using a Bat algorithm. A key feature of their study was the use of a cost constraint to limit the maximum cost of a route and the clustering of sites, similar to the maximum delivery time concept and cycle consolidation approach applied in the SSCP. The presented solution approach generated effective solutions, though the absence of multiple modes makes this method less applicable to the present study.

Table 1: Comparison of previous investigations of the specimen collection problem (and similar). MOO = Multi-objective optimisation. *= for maximum number of instances tested.

Author(s)	Problem Title	Key Dynamics	Objective(s)	Formulation, Approach	Calculation Runtime*
McDonald (1972)	VRP Case Study - Specimen Collection	Vans Only, Time Constrained	Total Travel Time	Single Objective, Heuristic	Not Stated
Grasas et al. (2014)	Blood Sample Collection Problem (BSCP)	Vans Only, Capacitated, Time Constrained	No. of Vehs.	Single Objective, Genetic Algorithm	<30 secs
Smith et al. (2015)	Pathology Laboratory Service Delivery	Vans Only, Capacitated, Sample/Lab Type Constrained	Cost, Lab Workload, Veh. Workload, No. of Vehs.	Hierarchical MOO, Theorised	Theorised Only
Yücel et al. (2013a)	Vehicle Collection for Processing Problem (CfPP)	Vans Only, Single Vehicle, Time Constrained, Lab Capacity Limit	Lab Performance, Transport Costs	Hierarchical MOO, Heuristic	c. 4 hrs
Yücel et al. (2013b)	Multiple Vehicle Collection for Processing Problem (mCfPP)	Vans Only, Time Constrained, Lab Capacity Limit	Lab Performance, Transport Costs	Hierarchical MOO, Theorised	Theorised Only
Anaya-Arenas et al. (2016)	Biomedical Sample Transportation Problem (BSTP)	Vans Only, Several Time Constraints, Uncapacitated	Travelled Distance	Single Objective, Heuristic	1 hr (limit)
Liu et al. (2013)	Home Healthcare Problem (HHC), variant of VRP with simultaneous pickup and delivery and time windows (VRPSDPTW)	Vans Only, Time Constrained, Precedence	Cost	Single Objective, Genetic Algorithm with Tabu Search	72 hr limit, though likely less
Zabinsky et al. (2020)	Vehicle Routing and Scheduling Algorithm (VerSA)	Vans Only, Time Constrained, Vehicles Re-Used	Total Duration from Goods Ready to Delivery	Single Objective, Branch and Bound	2 hr (limit)
Naji-Azimi et al. (2016)	Vehicle Routing Problem with Desynchronized Arrivals (VRPDA)	Vans Only, Time Constrained, Uncapacitated, Synchronised Laboratory Arrival Penalty	Travelled Distance, Sum of Travel Times, Number Deliveries within Any Time Period	Weighted (Multi-Term) Objective, Heuristic	1 hr (limit)
Elalouf et al. (2018)	Blood Sample Supply Chain	Vans Only, Time Constrained, Late Laboratory Arrival Penalty	Cost and Number of Samples Delivered On-Time	Multi-Objective, Heuristic/Tabu/Bisection Search/Advanced Heuristic	Up to 30-40 minutes
Oakey et al. (2022, present study)	Sustainable Specimen Collection Problem	Vans and gig-economy cycles, Time Constrained, Multi-Echelon	Longest Collection Round Duration, Number of Vehs., Total Travel Time	Weighted (Multi-Term) Objective, Column Generation with Improvement Heuristics	c. 30 secs

1.3. Related Works - Heterogeneous Logistics & Routing

In this study, multi-modal, or heterogeneous logistics systems relate to the use of multiple modes (e.g., van and cycle), to transport goods from one point to another. Cycle logistics systems have been seen to offer significant environmental and cost benefits in urban delivery networks (Sibilski and Targa, 2020; ECF, 2012; Marujo et al., 2018). In late 2019, a Newcastle based cargo cycle company trialled a service with the NHS to deliver light goods, including diagnostic specimens (NHS, 2020b). A single cargo cycle covered 25 stops across an 8-mile route over a three month period, reducing carbon output by 212kg and costs by £6,250 (an approximate reduction of 848kg and £25,000 per annum, respectively).

Time benefits can also be seen when using cycle couriers in urban logistics. Conway et al. (2017) found that cargo cycles in New York were capable of achieving competitive travel times in some cases, though this did depend on the road layout, route, and cycle type. A further study by Gruber and Narayanan (2019), based on trips in Germany indicated similar trends, with cycles being favourable during working-day hours (up to 7 pm), and vehicles being favourable over longer distances (>a few km) or over hilly terrain. Replacing large vehicles with two-wheeled or small cargo cycles on busy routes offered the best savings in terms of time and emissions (Conway et al., 2017; Gruber and Narayanan, 2019).

When combining cargo cycles with traditional road vehicles, there are often fixed or temporary locations used for trans-shipment activities (Marujo et al., 2018). Verlinde et al. (2014) and Marujo et al. (2018) investigated the potential for mobile depots and de-consolidation style delivery systems in the urban environment which could be likened to a reversal of the proposed concept to solve the SSCP. Both studies identified significant emissions savings from using such setups, whilst operational costs associated with the mobile depot systems were generally higher than traditional setups. To reduce this cost impact, situating depots in areas of higher consignee density resulted in improved performance (Marujo et al., 2018). It was also noted that there was a slight operational delay from additional loading/unloading activities, but consignees did not notice any drop in performance (Verlinde et al., 2014).

One of the challenges when using cargo cycles for consolidation, in combination with vehicles is that this requires some type of synchronisation or pairwise temporal precedence. In (Bredström and Rönnqvist, 2008) the authors proposed a general mathematical programming model which expanded the standard (homogeneous) vehicle routing problem to include the scheduling aspect with time windows and temporal constraints between routes. In (Anderluh et al., 2021) the authors developed a large neighborhood search algorithm to solve the two-echelon vehicle routing problem with vehicle synchronization, where the two routes should coincide simultaneously. The problem addressed has multiple objectives, considering terms capturing the economic cost as well as the social and environmental benefits by including greenhouse gas emissions and disturbance.

In another study, walking porters and cyclists were combined with traditional road vehicles in a more dynamic setup, varying routes depending on demand, parcel sizes/weight, and delivery drop density (McLeod et al., 2020). Emissions reductions of 45% and cost savings of up to 39% over BAU were made if 50% of parcels were served by cyclists/porters, enabling a reduction in the number of road vehicles required. McLeod et al. (2020) also proposed that casual workers from the gig-economy could be used to support a multi-modal system when workers are not occupied with gig-economy deliveries due to varying demand (e.g. takeaway riders could deliver parcels in

the off-peak periods). The concept could enable more sustained and reliable employment for gig workers where they would undertake medical deliveries alongside traditional food deliveries, and is adopted in this research, where the costs for using cycle couriers has been taken from a gig-economy provider. Additionally, range limits are applied to cycle couriers to prevent the overuse of cyclists and excessive cycling durations. In many cases, it is likely that the resultant reduction in van use will improve the sustainability of specimen collection rounds, however, care must be taken to ensure delivery times are not significantly delayed by the introduction of consolidation points.

The problem addressed in this research has some similarities with the min-max Generalised Vehicle Routing Problem (GVRP), which is an extension of the well known Capacitated Vehicle Routing Problem (CVRP). The GVRP was first proposed by Ghiani and Improta (2000) and since then there have been several works proposing new integer linear formulations and exact algorithms (Bektaş et al., 2011; Pop et al., 2012), as well as metaheuristic algorithms (Hà et al., 2014; Biesinger et al., 2018). A common feature considered in the GVRP is that only one node from each cluster can be visited and the clusters do not overlap. In the case of the SSCP, the clusters are not pre-defined and are determined by the algorithm during the optimisation process. Each surgery has a defined catchment area, based on the surgeries within a given cycling radius. The specimens from surgeries in the same catchment area can eventually be collected by a vehicle, visiting only one surgery in the area, if the samples have previously been collected from the other surgeries in the same catchment.

The two-echelon vehicle routing problem considers two different levels (echelons) in which different vehicle types can be used and solved using various mathematical models, exact algorithms or math-based heuristics (Perboli et al., 2011; Baldacci et al., 2013; Cuda et al., 2015). It is worth noting that if the surgeries are considered as the potential intermediate facilities where the samples of other surgeries might be kept, then the problem can be viewed as a two-echelon problem. However, both (i) the objective addressed by the SSCP; and (ii) the possibility of driving directly to the surgeries, instead of visiting them using cycle routes defines a new problem that is not easily translated into a traditional two-echelon problem. The objective of the SSCP seeks to improve delivery times by minimising the largest time across the network.

Heterogeneous two-echelon problems with multiple objectives are also not widely explored. An early example of this was by Eitzen et al. (2017), in a theorised problem that looked to appease multiple stakeholders in an urban delivery system, with respect to costs, the number of vehicles, and emissions. It was found that improvements could be achieved across all objectives, despite the complication of the two-echelon and heterogeneity.

2. Real-World Case Studies

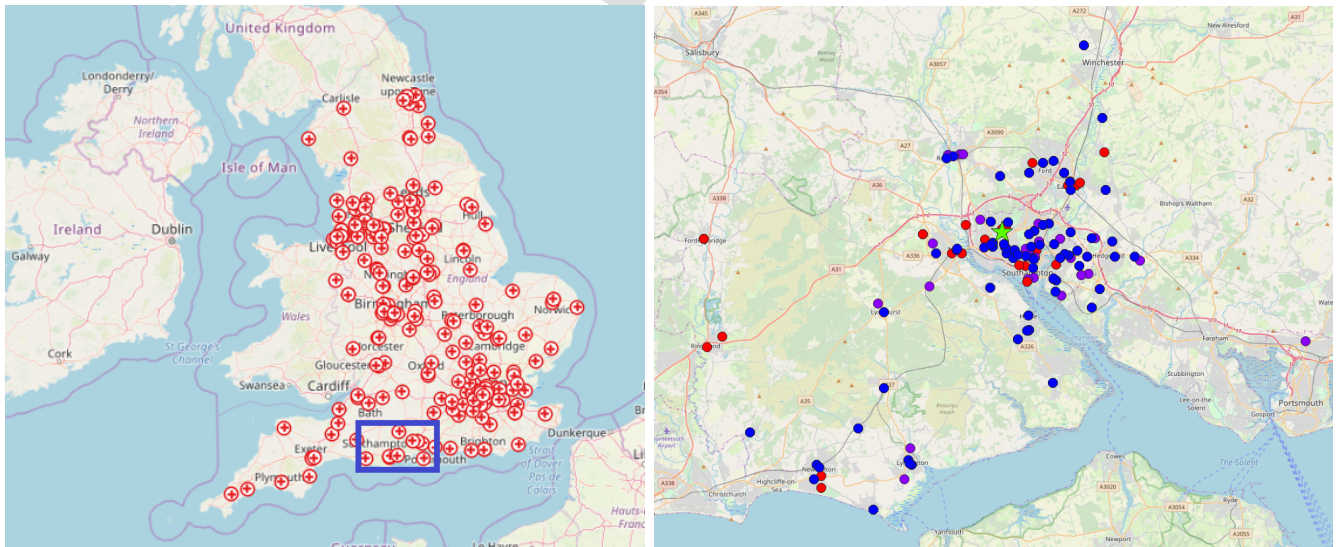
Two case studies involving current patient specimen delivery services were used in this research; Southampton, UK, and the Isle of Wight (IOW), UK. Both cases featured specimen delivery in which samples were routinely collected from a set of known surgeries and delivered to a single delivery point at the major local hospital, within the ‘South 6 pathology network’ in southern England (Figure 2a). The findings from the Southampton study were used to support the formulation of the problem, and the development and original testing of the algorithm. The IOW study

was used in further testing of the algorithm to understand its effectiveness in other collection areas of differing scale and geography.

2.1. Case Study 1: Southampton

The Southampton case study used the diagnostic specimen network delivering to the analysis laboratory at Southampton General Hospital, UK as its case study (Figure 2a). Two separate data sets were provided by Southampton General Hospital Pathology; one providing the details of specimens produced by 78 feeder surgeries during November 2018, and the other detailing the route schedules during September 2018. When combined, the BAU operations suggested that specimens were collected from 97 GP surgeries (postcodes) (Figure 2b), and a collection service may need to service all of these sites in a single day. It should be noted that postcodes were used as they geo-coded more reliably instead of site names, even though it is possible for more than one surgery to share a postcode, this was not common, and such sites were found to be close enough that they could transfer goods internally without causing delays. Other sites were also visited by some rounds for collection/delivery of other items such as internal mail and paper records.

In this study, only the sample collections were of interest due to the core objective of reducing the time from when the patient was bled to when the sample was received at the hospital. The ancillary delivery services covered by these rounds (internal mail, paper medical records) were understood to be less urgent, of low volume, or were due to be phased out. To enable a closer comparison between the BAU and computational results, a modified version of the BAU was used where any ancillary service stops were removed and stop times adjusted where sites were for specimen collection only. A comparison of the Key Performance Indicators (KPI) of the two versions is shown in Table 2. The values described below were from the modified schedule.



(a) NHS England Trust Locations (box contains South 6 network in the Solent area) (NHS, 2019)

(b) GP surgery locations in the Southampton area. Blue and purple points indicate surgeries from which specimens must be collected. Red surgeries are currently visited by the rounds for purposes other than specimen collection. The green star indicates SGH.

Figure 2: Solent area GP surgeries and NHS trust locations in the wider UK context (Base Map ©OpenStreetMap contributors)

Ten vehicles, primarily medium-sized LGVs (Vauxhall Vivaro or Ford Transit), served the

Southampton-based surgeries on weekdays, whilst only two were used on weekends (Quadir et al., 2019). On weekdays, vehicles covered an average of 113km per day, stopping 20 times, over a period of 4 hours 13 minutes. Each surgery was visited an average of 1.96 times per day.

The specifications of vehicles varied, though all were assumed to be “Vans of 3.5 tonnes GVW –diesel” defined under the Manager’s Guide to Distribution Costs for the purposes of analysis (FTA, 2020). Calculations used a per-mile rate (£0.464 /mi) for vehicle costs to cover fuel, tyres, maintenance, tax, insurance, depreciation, and overheads, a per hour rate (£10.78/hr) for driver wages, and a per-mile rate (0.45 kg/mi) for CO₂ tailpipe emissions. As a result, on each weekday each vehicle was emitting approximately 31.8kg of CO₂; a daily total of 318kg across the fleet (FTA, 2020). Collectively, the rounds cost approximately £782 each day.

The maximum time between departing Southampton General Hospital and returning (i.e. the delivery interval) was 285 minutes with the mean duration (per collection) being 135 minutes. If a direct driving route was taken, each surgery was an average of 12 minutes 15 seconds drive (one-way) from the hospital and a quicker and more effective transit option should have been possible (GraphHopper, 2020). Not every surgery produced the same volume of samples, however, all sites had to be visited to ensure a collection was made during the day. Figure 3 demonstrates this variation, with thicker lines indicating a greater number of samples.

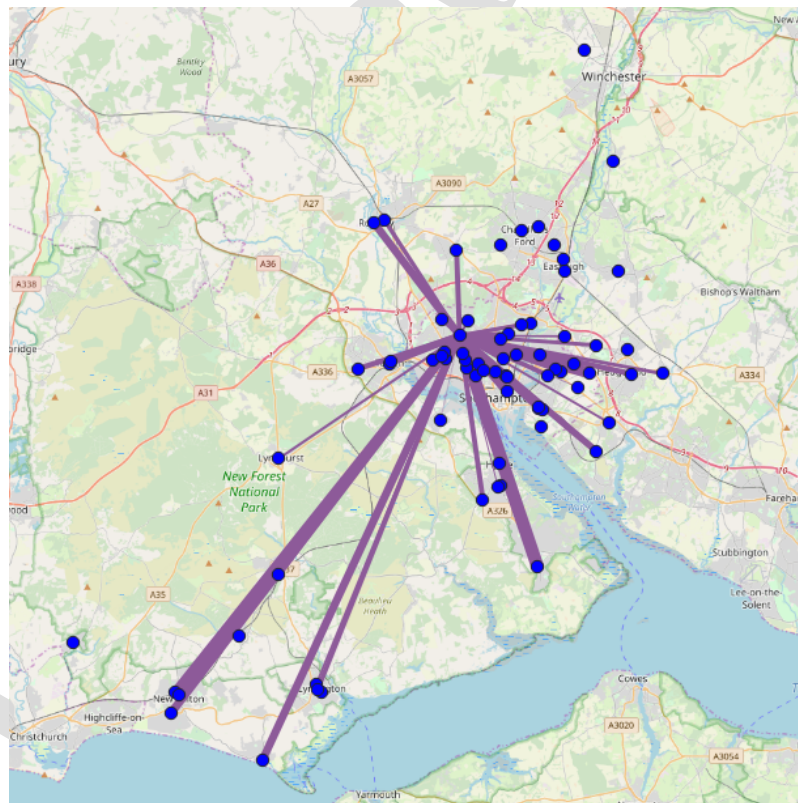



Figure 3: Flowline map of feeder surgeries. Surgeries with sample production information are connected to the hospital by a line with thickness relative to the load produced (some flows are significantly smaller so lines are not visible on this plot; all blue points produce specimens.)

Only sites from the sample production dataset are shown. (Base Map ©OpenStreetMap contributors).

With regards to the timing of collections, every surgery was typically served by a morning and afternoon collection which correlated with a surge in observed samples received just after midday, and again at around 17:00 (Figure 4a). There was minimal variation in the peaks of when samples were taken (Figure 4b), with an average of 79% of each day's samples taken by practitioners by midday. Only 12% of samples would have been checked-in to the hospital by the same time as they were still in transit. It should be noted that the tail on the plots, particularly the 'taken' curve, related to samples in the dataset that were marked as taken or received outside of working hours due to either (i) errors in the dataset; (ii) where samples had been taken at the hospital but assigned to the surgery; or (iii) where samples had been checked in late. The last scheduled delivery of samples from any surgery to the hospital was 18:35.

According to laboratory staff, check-in processes required at the hospital accounted for some of the delays seen in the 'receipt' peak, though unfavourable routing was responsible for the majority of this delay (Allan, 2019). It was assumed that the check-in procedure took a fixed amount of time, and in order to enhance the receipt times at the diagnostic lab, delivery times must be improved. Following discussions with local clinicians, there appeared to be limited scope to alter the timings of when samples were taken due to surgery opening hours, though there was often demand for later final collections, in addition to the existing timetabled collections, to enable a better patient service offering (Wessex Academic Health Science Network, 2020). This aligned with the suggestions made by McDonald (1972) and Anaya-Arenas et al. (2016), highlighting that reducing the transit time and the time to deliver all samples are critical considerations.



(a) Comparison of sample production time and receipt time. Most samples are produced in the morning but are received in the afternoon.

(b) Change in sample production across each day. A clear peak is seen in the mid-morning. Average production is shown in red.

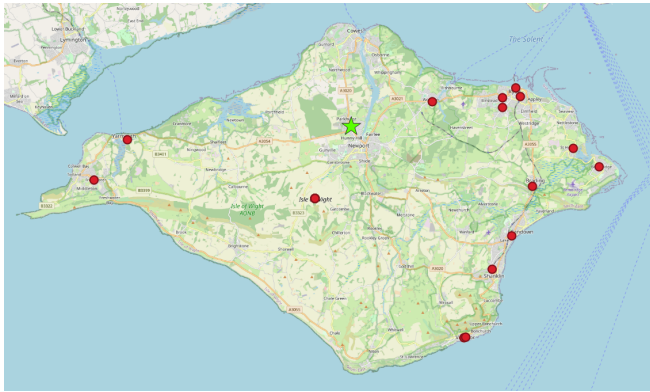
Figure 4: Southampton - Plots of sample production/receipt over time. Last delivery to the hospital is 18:35; thus, data beyond this time may be erroneous.

2.2. Case Study 2: Isle of Wight


The second case study in this research was used to test the algorithm's performance in a different setting. The case study was based on the Isle of Wight (IOW), also in the Solent region of the UK (Figure 2a). Surgeries on the island are more sparsely distributed than in Southampton, and the analysis laboratory is at St. Mary's Hospital near the centre of the island (Figure 5a). As in the Southampton case study, historic round and sample collection data were made available, enabling

a comparison of how the algorithm performed in other areas with different characteristics and stops.

At the time of the study, 22 sites (postcodes) were served by 3 vehicles on a daily basis with all surgeries receiving an average of 1.64 collections per day. The data covered a 5-day period in June 2020, where a total of 1637 samples were produced; an average of 327 across the island per day, or 15 per surgery per day. In BAU routing, three vehicles each travelled an average of 90km each day. Assuming the same vehicle type as in Southampton case study, this resulted in 75.7kg of CO₂ being generated across the fleet each day; 25.2kg per vehicle. Including service time (but excluding breaks), vehicle rounds took an average of 3 hours 17 minutes at a cost of £184 each day. The maximum time between departure and return to St. Mary's Hospital (i.e. the delivery interval) was 125 minutes with the mean duration being 112 minutes. As with the Southampton case study, a lag between taking and receiving the samples was also observed (Figure 5b).



(a) GP surgery locations on the IOW. The green star indicates St. Mary's Hospital, where samples are analysed. (Base Map ©OpenStreetMap contributors).



(b) Comparison of sample production time and receipt time. As with the main case study, most samples are produced in the morning but are received in the afternoon. Two clear delivery peaks are seen.

Figure 5: Isle of Wight Case Study

2.3. Business-as-usual Summary

The two case studies used in this paper are summarised in Table 2. It should be noted that the computational experiments produced results based on a single collection at all surgeries, with all vehicles departing simultaneously (as detailed in the problem description), whilst the case studies feature additional collections. In reality, the model would be used multiple times per day, generating different collection rounds for varying smaller instances. Table 2 reflects these multi-collections, with rounds being scaled to an equivalent of 1 collection at each surgery. Average van driving time was taken as the total driving time divided by the number of vans, whilst the average collection round time was the mean duration between laboratory drops per stop served.

A consultation with NHS diagnostics laboratory staff took place to discuss the challenges that were being faced during day-to-day operations. The key concern that was highlighted was a need to improve failed samples through an improvement in the arrival times. Pressures from senior management to manage costs and environmental impacts were also important. The objectives selected for optimisation (max. round time, qty. of vehicles, total fleet time) were chosen from the KPIs in Table 2 following the consultation, after the availability of reliable data was considered. Other objectives, such as the average stops per van are somewhat incidental from the optimised

KPIs; meanwhile, costs are noisy approximations based on a per mile and per hour basis and are only used as a rough indicator.

Table 2: Summary of KPIs seen in the Case Studies. Scaling is used to make rounds representative of a single collection at all surgeries. Soton = Southampton

Scenario	Maximum Sample Coll. Round Time (mins)	Number of Vehicles (qty.)	Total Fleet Round Time (mins)	Avg. Van Driving Time (mins)	Avg. Collection Round Time (mins)	Avg. Stops Per Van	Total Costs (driver+ vehicle)	CO ₂ (kg)
Soton Original	N/A	10	4380	438	N/A	28.9	£1,188	389
Soton Spec. Only	285	10	2530	253	135	19.6	£782	318
Soton Scaled Spec. Only (1.96x)	285	10	1291	129	135	10	£399	162
IOW Original	125	3	592	197	112	10.7	£184	75.6
IOW Scaled (1.64x)	125	3	361	120	112	6.5	£112	46.1

Scaled = factoring each parameter by the average number of visits at each site to make it the equivalent of 1 collection. Southampton average collections/day = 1.96, IOW = 1.64. Spec. only refers to a modified BAU case that removes stops that are not for specimen collections.

3. Mathematical Formulation

In order to enhance the time of arrival of samples at the diagnostics lab, a new approach is explored in this paper. This uses a combination of road vehicles (vans/LGVs), and pedal cycles. The pedal cycles start at a surgery or the hospital and collect from surgeries local to them (satellite surgeries) before returning to their origin surgery (the consolidation point). The road vehicles complete longer distance rounds starting from the hospital, serving the consolidation points and any others which fall outside of the range of consolidation rounds, taking advantage of the faster travel speeds and lower carbon impact of cycling in short distance urban environments whilst maintaining the benefit of faster trunk mileage from road vehicles. The introduction of cycles in this way has not been proposed previously, though cargo cycles have been used in specimen collections in the past (NHS, 2020b).

3.1. Master Problem

The BAU activity presents a problem in which a set of known locations/nodes produce samples that need transporting to a single location/node as fast as possible without incurring excessive cost, congestion, or environmental impact.

Let S be the set of surgeries that require a collection, and H as the Target Hospital to which samples are delivered and vans are based. An individual surgery is denoted by $s \in S$. The set of all nodes, including the target hospital is defined as $S' = S \cup \{H\}$.

Two modes are available in this problem, V , which represents a van; and C , which represents a cycle. Subsequently, the set of van routes is defined as R^V , and the set of cycle routes is defined as R^C . We define T as the service time required at each stop in a route, and the time to travel between a pair of surgeries (i, j) as $t_{i,j}^V$ for a van, and $t_{i,j}^C$ for a cycle.

Let $r_v = (H, s_1, \dots, s_{n_v}, H) \in R^V$ be a van route, where n_v denotes the number of surgeries being visited, $s_i \in S, \forall i \in \{1, \dots, n_v\}$. Note that all vans are based at the Hospital H and must return to H . Additionally, n_v^s denotes the number of stops from the hospital to surgery s in r_v . Each van route has an associated time, denoted by t_{r_v} . The sum of all durations between surgeries visited and the sum of all service times, t_{r_v} , is calculated as $t_{r_v} = t_{H,s_1}^V + \sum_{i=1}^{n_v-1} t_{i,i+1}^V + t_{s_{n_v},H}^V + n_v T$. It should be highlighted that t_{r_v} excludes any delays caused by cycle routes, and T is embedded in the value.

Similarly, let $r_c^s = (s, s_1, \dots, s_{n_c}, s) \in R^C$ be a cycle route based on surgery s in which all surgeries $\{s_1, \dots, s_{n_c}\}$ are being served by cycle, and the samples are delivered to surgery s . We denote the number of surgeries being visited as n_c . Cyclists are capacity constrained such that they cannot carry more than three surgeries' worth of samples in one round, based on a typical gig-economy backpack (Deliveroo, 2020); $n_c \leq 3$. It is worth highlighting that the cycle route should start and end at the same surgery $s \in S$.

Each cycle route also has an associated time, denoted by $t_{r_c^s}$. The sum of all durations between surgeries visited and the sum of all service times, $t_{r_c^s}$, is calculated as $t_{r_c^s} = t_{s,s_1}^C + \sum_{i=1}^{n_c-1} t_{i,i+1}^C + t_{s_{n_c},s}^C + n_c T$. It should be noted that T is embedded in $t_{r_c^s}$. Additionally, cycle routes are subject to a time constraint of 25 minutes or less to ensure the cycle elements can be managed as discrete gig-economy tasks (Allen et al., 2021); $t_{r_c^s} \leq 25$ minutes.

We define a **collection round**, $\bar{r} = (r_v, R_v^C)$, as the combination of a single van route $r_v \in R^V$ with a subset of cycle routes $R_v^C \subseteq R^C$ such that for any given cycle route $r_c^s \in R_v^C$ based at surgery s , it is satisfied that $s \in r_v$, i.e., any cycle route in R_v^C is based in a surgery that is being visited by a van route r_v . Furthermore, other than surgery s where each r_c^s begins/ends, the cycle routes in R_v^C do not share any other surgeries; i.e. no surgery is served by multiple cycles. We define the set of all the collection rounds as \bar{R} , and each collection round is denoted by $\bar{r} = \{(r_v, R_v^C) \mid r_v \in R^V, R_v^C \subseteq R^C\}$, $\bar{r} \in \bar{R}$. Note that, since R_v^C could be empty, then it can be satisfied that $R^V \subseteq \bar{R}$. The set of surgeries served by all of the constituent routes of \bar{r} is denoted by $S_{\bar{r}}$.

Cycle routes that start and end at the hospital are permitted in order to serve the catchment area of the hospital directly. To account for this in the model formulation, a dummy van route, $r_{v_0} \in R^V$, is created. Starting and ending at the hospital, with no intermediate stops ($r_{v_0} = (H, H)$) and a travel time of zero ($t_{r_{v_0}} = 0$), r_{v_0} enables a collection round where surgeries are cycle served

only, \bar{r}_0 .

In a collection round, we define the waiting time for the van at a given surgery, $w(\bar{r}, s)$, as the difference between the durations of the cycle routes based on s , $t_{r_c^s}$, and the duration of the van route (including any previous waiting time) up to surgery s , $t_{\bar{r}_v, s}$. The waiting time, $w(\bar{r}, s)$, is only considered when it is positive (i.e. the cycle takes longer than the van). The sum of the travel times, service times, and waiting times from the start of the van route in \bar{r} up to surgery s is defined as $t_{\bar{r}_v, s}$, whilst $t_{\bar{r}}$ denotes the total collection round duration and is equal to $t_{\bar{r}_v, H}$; i.e. the time for the vehicle to return to the hospital after first departure, including any waiting time incurred.

$$w(\bar{r}, s) = \max\{t_{r_c^s} - t_{\bar{r}_v, s} \quad \forall r_c^s \in R_v^C \quad s \in r_v, 0\}$$

$$t_{\bar{r}_v, s} = t_{H, s_1}^V + \sum_{i=1}^s t_{i, i+1}^V + n_v^s T + \sum_{i=0}^s w(\bar{r}, s)$$

Collection rounds are subject to a time constraint of 90 minutes or less (McDonald, 1972), to ensure timely delivery of samples; $t_{\bar{r}} \leq 90$ minutes. It should be noted that whilst this constraint is not met in business-as-usual rounds, staff have expressed a desire to meet this criteria. Vans are assumed to depart at the same time; thus, the durations all commence at this point. If a waiting time is incurred due to a longer cycle route, this will slow the van's progress and increase the collection round duration. If there is no waiting time incurred, then the collection round duration is equal to the van route duration. Some example collection round scenarios are given in Table 3.

Table 3: Example collection round scenarios, illustrating potential delays. Delays only demonstrated at the first stop, other delays may occur. All cycle routes deliver to first stop in these examples unless otherwise stated. CR = Collection Round

Scenario	Van Route in CR	Cycle Routes in CR, $t_{r_c^s}$	Time to First Stop	Delay at First Stop	Van Route Duration	CR Duration
Van + 1 cycle, van delayed	(H, s_1, s_2, s_3, H)	(s_1, s_4, s_5, s_1) , 25 mins	20 mins	5 mins	80 mins	85 mins
Van + 2 cycles, no van delay	(H, s_1, s_2, s_3, H)	(s_1, s_5, s_1) , 15 mins; (s_1, s_4, s_1) , 10 mins	20 mins	0 mins	80 mins	80 mins
Van + no cycles, no van delay	(H, s_6, s_7, s_8, H)		30 mins	0 mins	85 mins	85 mins
Cycle direct to hospital	H, H	(H, s_9, s_{10}, H) , 23 mins	0 mins	23 mins	0 mins	23 mins

A binary decision variable, $x_{\bar{r}}$ is introduced to select collection rounds:

$$x_{\bar{r}} = \begin{cases} 1 & \text{if the collection round is used in the solution} \\ 0 & \text{otherwise;} \end{cases}$$

The master problem is a multi-objective problem in which a balance between (i) the latest delivery times of samples (first term in (1)); (ii) the number of collection rounds (vans) used (second term in (1)); and (iii) the total driving and servicing time (last term in (1)) must be found in any solution. These three objectives are combined by weighting with coefficients α_1 , α_2 , and α_3 . In Section 3.2 we explain how these weights are set by taking into consideration the BAU, and with that approach we aim to solve the problem using a single objective function (Section 4).

To capture the maximum collection round duration across all rounds/surgeries, a continuous variable u is introduced (Equation 2). The latest delivery time/interval to the Hospital for the samples of any surgery is denoted by u and is based on the $t_{\bar{r}}$ values for all rounds.

Full collection round durations were used, as opposed to only the duration after the first pickup, due to cycling being handled by a 3PL. It has been assumed that the 3PL have the flexibility to complete cycle collections at any time after the collection rounds are permitted to start. For example, if rounds (all van routes and cycle routes) are permitted to start at 09:00, samples cannot be collected any earlier; however, the 3PL flexibility means cycles can complete their deliveries any time before the vans **depart** the target consolidation site.

Using the full collection round duration enables a conservative upper bound to be captured. Should contractual arrangements allow full control over the departure times at all sites for both vans and cycles, the term t_{H,s_1}^V could be removed from the calculation of $t_{\bar{r},v,s}$ for a more accurate calculation of the time samples spend in transit. Where waiting time is included in the calculation of $t_{\bar{r},v,s}$, the duration of the cycle rounds contributes to $t_{\bar{r}}$ if waiting times are incurred. This concept is demonstrated in Figure 6.

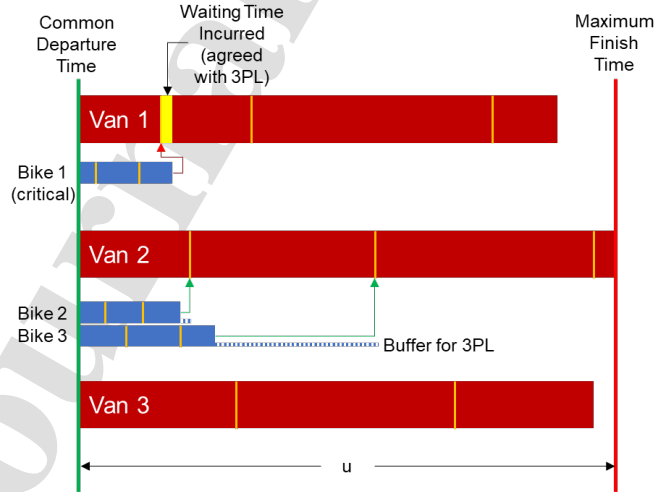


Figure 6: Demonstration of the critical path that governs the maximum duration, u .

To ensure all sites are served, a further constraint is also added (Equation 3), whilst $x_{\bar{r}}$ in Equation 4 must be binary.

$$\min \quad \alpha_1 u + \alpha_2 \sum_{\bar{r}, \bar{r} \neq \bar{r}_0} x_{\bar{r}} + \alpha_3 \sum_{\bar{r}, \bar{r} \neq \bar{r}_0} t_{\bar{r}} x_{\bar{r}} \quad (1)$$

$$u \geq t_{\bar{r}} x_{\bar{r}}, \quad \forall \bar{r} \in \bar{R}. \quad (2)$$

$$\sum_{\bar{r}; i \in S_{\bar{r}}} x_{\bar{r}} \geq 1 \quad \forall i \in S \quad (3)$$

$$x_{\bar{r}} \in \{0, 1\}, \quad \forall \bar{r} \in \bar{R} \quad (4)$$

3.2. Objective Function Calibration

The generalised multi-objective problem is modelled with a weighted multi-term objective function (Equation 1), enabling a Pareto front of solutions to be found by varying the coefficients of each term. The values of the business-as-usual case were used to define and normalise the terms α_1 , α_2 , and α_3 , whilst a multiplier varied the relative importance of each term:

$$\alpha_1 = \frac{w_1}{\hat{u}} \quad \alpha_2 = \frac{w_2}{\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} \hat{x}_{\bar{r}}} \quad \alpha_3 = \frac{w_3}{\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} \widehat{t_{\bar{r}} x_{\bar{r}}}}$$

where w_1 , w_2 , and w_3 are the multipliers, and \hat{u} , $\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} \hat{x}_{\bar{r}}$, and $\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} \widehat{t_{\bar{r}} x_{\bar{r}}}$ are the values of the objective function terms under the business-as-usual case.

In a real-world application of the model, the multipliers could be user-defined to allow decision-makers to weight the relative importance of each objective and achieve a balanced outcome which satisfy their requirements.

This step is seen in Figure 9, noted by a superscript 2. The calculated α values define the objective weights.

A series of objective function weights ratios were tested to understand the sensitivity of the business-as-usual inputs and identify a Pareto front of solutions, whilst also demonstrating the trade-offs decision-makers could make.

4. Column Generation Based Heuristic Model

The master problem presented by the SSCP has been solved in this study using a column generation based heuristic, whereby all the surgeries being visited by the same vehicle in one single column; i.e., one column is one collection round. These surgeries can either be visited directly by the road vehicle or indirectly by using any combination of cycle couriers and the road vehicle. By including this information in columns, efficient computation and processing was achieved, and the problem could be solved quickly using an initial set of heuristically generated routes which were iteratively improved using further heuristics.

In the implementation explored in this study, van routes are constructed such that they can only stop once in each cycle-able catchment area; i.e., during construction, if a van stops at surgery s_1 , no cycle route can exist to the van's next stop (s_2) as it is out of cycling range (Figure 7). This assumption was made to (a) enable van routes to be solved independently of cycle routes (i.e., not as combined collection rounds); and (b) reduce the number of van routes in the solution space and improve calculation times whilst still aligning with the objective function. The approach

reduces the number of required van stops in each route; thus, the driving time of each route is also decreased in the solution.

Both case studies are explored in this paper, using a ‘worst case’ scenario, in which it is assumed that all surgeries require visiting, even though this may vary slightly day-to-day. Under BAU activity, road vehicles are already in use, however cycle couriers are not. As a result, it is envisaged that an existing 3PL provider would be responsible for the cycle logistics, enabling them to increase the off-peak (i.e. mid-morning and mid-afternoon) job offerings for their workers (current peaks - lunchtime, dinnertime), which have often been cited as insufficient for the number of workers signed-up (Lord et al., 2020; Bernal, 2020). The model explores single collections which will need to be scheduled throughout the day; most likely at the times which are low-demand for gig-workers. Modelling assumes cost constraints based on existing operations of the gig-economy courier company, “Stuart” (Stuart, 2020).

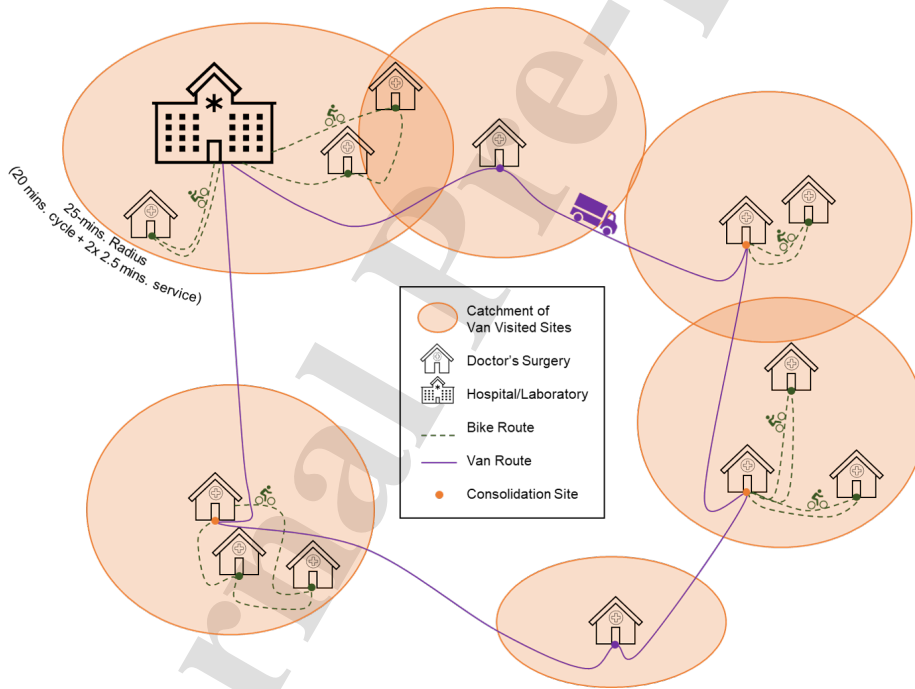


Figure 7: Example of collection round structures used in the modelled approach to the SSCP.

4.1. Initial Route Generation Heuristic

In the modelling process, two initial sets of routes are generated prior to solving. One set contains the cycle route options whilst the other contains the road vehicle route options. In the case of the Southampton study, the problem’s parameters result in a set of 3170 cycle routes, so all routes can be fully enumerated without issue. As part of the cycle route generation, a list of surgeries that fall in the ‘catchment’ of each surgery is defined using the maximum cycle time and service time, enabling the model to be solved for vans routes, before post-processing for exact cycle routes.

Following the calculation of cycle catchments, the vehicle routes are generated using a greedy-style heuristic under the constraints within the master problem (Equations 2-4), in addition to the following constraints:

- An initial maximum number of stops, or shorter (maximum changes during heuristic processes);
- The surgery catchment lists for each surgery (produced in the cycle route generation stage);
- A maximum number of surgeries to shortlist for each next stop during generation.

For the first stop from the hospital, all surgeries are tested to ensure a route exists to serve every surgery (Algorithm 1, Line 3). For subsequent stops, up to a maximum number of stops (n_{max}), new routes are made (Algorithm 1, Line 6) with additional surgeries within range (T_{max}) but not in the cycle catchment already served (Algorithm 1, Line 5). This is handled by Algorithm 2, where a recursive function is called until the maximum stops criteria is reached, or there are no more sites in range. The additional sites are selected (Algorithm 3) from two subsets: (i) the closest N ($N \leq n_{max}$) sites; (ii) those not within the N closest. From the closest, L sites are randomly selected, and from the others, O sites are selected. The position of this stage of the algorithm is seen in Figure 9, noted by a superscript 1.

It should be noted that the route generation in this study allows vehicle routes to be created where vans arrive at consolidation surgeries before some cycle routes are completed, meaning a delay is incurred in the vehicle route. This could be prevented in the vehicle route generation, however the cycle route(s) which cause the delay may not be selected by the algorithm. After generating the van routes, the waiting times for each route are calculated using the cycle routes associated for each stop on the van route. This approach is loosely related to the Covering Vehicle Routing Problem (Buluc et al., 2022; Semet and Taillard, 1993), but the overall problem is different due to the second echelon affecting the main vehicle routes.

On completion of the route construction and waiting time calculations, the routes are input into a Gurobi optimisation environment and solved using a branch and bound optimisation approach.

Algorithm 1 Initial Van route construction

Input: S : Set of surgeries to serve, H : Target Hospital; T_{max} : Maximum route duration; n_{max} : Maximum route stops; M : O-D time matrices for $\{S, H\}$; L : shortlist sites; N : shortlist sites selected; O : other sites selected;

```

1:  $R \leftarrow \{\}$ 
2: for Surgery  $s \in S$  do
3:    $r \leftarrow (H, s, H)$  ▷ Create out-and-back route
4:    $R = R \cup \{r\}$ 
5:    $nextSites \leftarrow toTestFilter(r, S, M, L, N, O)$  ▷ Identify potential next sites to test,
   Algorithm 3.
6:    $newRoutes(r, R, nextSites, T_{max}, n_{max}, M, L, N, O)$  ▷ Algorithm 2.
7: end for
8: return  $R$ 

```

Algorithm 2 Recursive function: newRoutes. Van route construction - subsequent site additions.

Input: r : previous route; R : route list; \bar{S} : set of possible surgeries; T_{max} : Maximum route duration; n_{max} : Maximum route stops; M : O-D time matrices for $\{S, H\}$; L : shortlist sites; N : shortlist sites selected; O : other sites selected;

```

1: if r.length <  $n_{max}$  and  $\bar{S} \neq \emptyset$  then                                ▷ Maximum stops not been reached
2:   for  $a \in \bar{S}$  do
3:      $q \leftarrow (H, \dots, a, H)$                                        ▷ Create route with a inserted into route  $r$ 
4:      $R = R \cup \{q\}$ 
5:      $\bar{S} \leftarrow toTestFilter(q, \bar{S}, T_{max}, M, L, N, O)$            ▷ Identify potential next sites to test,
   Algorithm 3
6:      $newRoutes(q, R, \bar{S}, n_{max}, M)$ 
7:   end for
8: else
9:   return  $R$ 
10: end if

```

Algorithm 3 Next stops options function: toTestFilter. Van route construction - identify potential next stops.

Input: r : previous route; \bar{S} : set of surgeries; T_{max} : Maximum route duration; M : O-D time matrices for $\{S, H\}$; L : shortlist size; N : shortlist sites selected; O : other sites selected;

```

1:  $\hat{S} = \bar{S} \setminus \{r\}$                                        ▷ List of sites not already covered by the existing route
2: Sort  $\hat{S}$  by duration required to add to route  $r$ 
3: Remove from  $\hat{S}$  all the sites outside of the range ( $> T_{max}$ ) of the current route  $r$ 
4:  $shortList \leftarrow \hat{S}[0 : L]$                                ▷ Copying first  $L$  elements of  $\hat{S}$  to  $shortList$ 
5:  $otherList \leftarrow \hat{S} \setminus shortList$ 
6:  $toTest = \{\}$ 
7: for  $i=0$  To  $N$  do
8:   Select  $s$  randomly from  $shortList$ 
9:    $toTest = toTest \cup \{s\}$ 
10:   $shortList = shortList \setminus \{s\}$ 
11: end for
12: for  $j = 0$  To  $O$  do
13:   Select  $s$  randomly from  $otherList$ 
14:    $toTest = toTest \cup \{s\}$ 
15:    $otherList = otherList \setminus \{s\}$ 
16: end for
17: return  $toTest$ 

```

4.2. Improvement Heuristics

After the first iteration of solving is completed (Figure 9, noted by a superscript 2), the van route list is ‘cleaned’ to eliminate all routes from the previous iteration that are longer than the longest route/shorter than the shortest route/fewer stops than the shortest route from the potential options. All cycle routes are kept for all iterations. Subsequently, two heuristic methods are used to improve routing options available for solutions; the longest route redistribution heuristic (LRRH), and the shortest route redistribution heuristic (SRRH). Both create new route variants by moving served surgeries between van routes and positioning them in the target route at the location where there is the smallest increase in route duration (Figure 8). One surgery is added per route variant per iteration. Other variants are made by rearranging the position of surgeries within routes to allow for where a later road vehicle arrival may mean u is reduced due to better timing with the arrival of cycle routes at consolidation surgeries.

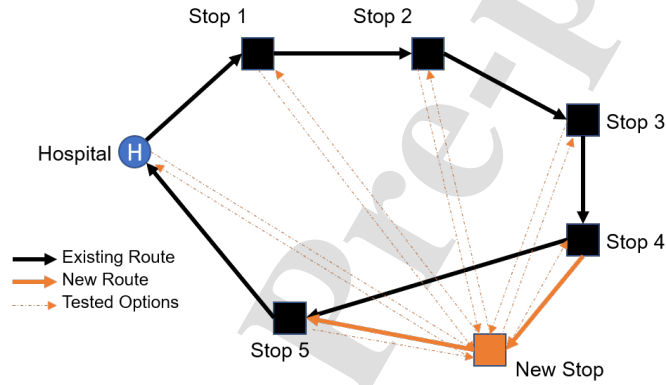


Figure 8: Representation of the positioning of a redistributed surgery stop. Dotted lines indicate possible paths to the additional surgery. Solid orange lines indicate the new path of the lengthened route.

4.2.1. Longest Route Redistribution Heuristic (LRRH)

Decreasing the length of the longest vehicle routes generally reduces the time from the surgery to the hospital (Objective 1). The ‘‘Longest Route Redistribution’’ heuristic (LRRH) aims to spread the surgeries served by the longest route(s) to the shorter routes, thus eliminating the longest route(s). This heuristic can be applied to the surgeries **directly** served by the longest route(s) (faster; fewer surgeries), or the surgeries within the catchment of the longest route’s stops (i.e. the surgeries served by the longest collection round). Surgeries that are redistributed are positioned in each route at the point which causes the least increase in route duration (Figure 8). The position of this stage of the algorithm is seen in Figure 9, noted by a superscript 3. In the case of all solutions presented in this study, this algorithm was applied to those routes which were longer than the average route selected in the previous iteration.

4.2.2. Shortest Route Redistribution Heuristic (SRRH)

Increasing the length of the shortest vehicle routes increases the average length of the selected routes, potentially eliminating a vehicle from the solution (Objective 3), decreasing the total vehicle time (Objective 2), and spreading loads between drivers more evenly. The ‘‘Shortest Route Redistribution’’ heuristic (SRRH) aims to spread the surgeries served by the shortest route(s) to the longer routes, thus eliminating the shortest route(s). This heuristic can be applied to the surgeries directly served by the shortest route(s) (faster; fewer surgeries), or the surgeries within

the catchment of the shortest route's stops (i.e. the surgeries served by the shortest collection round). Surgeries that are redistributed are positioned in each route at the point which causes the least increase in route duration. The position of this stage of the algorithm is seen in Figure 9, noted by a superscript 4. In the case of all solutions presented in this study, this algorithm was applied to those routes which were shorter than the average route selected in the previous iteration.

4.3. Post-Processing Cycle Selection

For each surgery, the algorithm assumes that the least favourable feasible set of cycle routes are used (with respect to delivery time) based on the same surgery when calculating $t_{\bar{r}}$; a conservative step that gives an upper bound of the u term (and total driving, if delays occur) in the objective, captured through complete enumeration. This assumption was made for because (i) the choice of cycle routing is controlled by the 3PL provider and as such may not be optimal with respect to delivery times; and (ii) routes which incur delays are less likely to be selected, which is favourable with respect to the objective function. Where no delay is incurred, $t_{\bar{r}} = t_{r_v}$; thus, knowledge of the exact cycle selection is less important.

The assumption that all sites within a cycle-able distance of van stops are served may result in duplication of cycle effort if a site occurs in multiple collection rounds, potentially at the detriment to the core objective. To prevent duplicate service, post-processing is used to select the final cycle routing (within the delays identified in the core vehicle selection) and minimise the number of riders used, as a 3PL might. Post-processing also enables more accurate cost estimates to be produced. The position of this stage of the algorithm is seen in Figure 9, noted by a superscript 5.

4.4. Heuristic Process Flow

The method followed by the algorithm (Figure 9) starts with the initial route generation. An iterative approach using the SRRH and LRRH then occurs in tandem with the cleaning process to eliminate ineffective routes, until the solution stagnates and is no longer dominated by new solutions. The selected vehicle routes are then input to the cycle route post-processing model. Only one iteration is required at this point due to the cycle routes being fully enumerated from the point of generation.

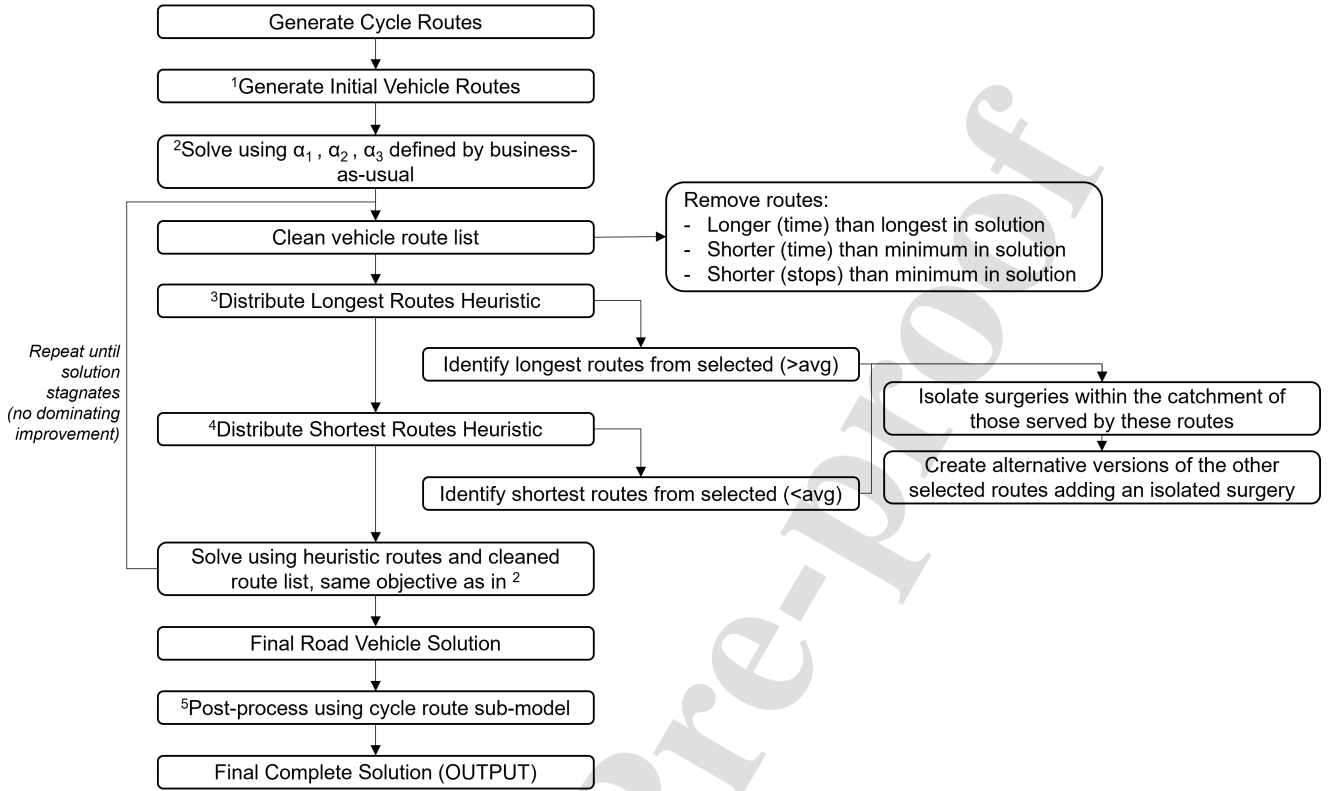


Figure 9: Process flow of the optimisation algorithm. Superscript values indicate the position of the algorithm stages described throughout this chapter.

5. Computational Experiments

A series of computational experiments were undertaken to optimise the model and identify the parameters which provided the best solutions. All tests were conducted on a system running an AMD Ryzen 5 3500U 2.1 GHz processor with 16GB RAM (Windows 10, 64-bit). Tests were coded and run in Java v13 in an Eclipse Development Environment, and were solved using Gurobi v9.5.1.

In the sensitivity testing of different objective function coefficients, the ratios between w_1 , w_2 , and w_3 were varied. In the experiments to test the heuristic performance, the number of routes initially generated prior to solving was varied by changing the maximum number of stops per route, or the number of next stop options tested. Other parameters and weightings remained fixed. The majority of tests were carried out on the more complex case study, Southampton. Further tests varying the weighting of the objective function were undertaken using the IOW case study to understand the algorithm's application to other areas and data sets.

5.1. Parameters

To allow for practicality constraints and limit computational time, several parameters were set (Table 4). Transit times and distances were specified by a locally hosted GraphHopper Routing engine which provided an asymmetric O-D time matrix (GraphHopper, 2020). The asymmetry accounted for turn restrictions and elevation changes. The 'car' and 'racing-bike' profiles were used to represent the two modes; 'racing-bike' was chosen over 'bike' to ensure routing prefers roads over tracks (GraphHopper, 2020; Reid, 2018). It was assumed riders do not have pedal assistance

which would further enhance their speed profile. Traffic was assumed to be ‘free-flowing’. The service time was assumed to be 2.5 minutes per stop, approximately based on investigations into freight driver dwell times carried out under the Freight Traffic Control 2050 (FTC2050) project (Allen et al., 2018).

The next-stop shortlist size (L in Algorithm 3) of 5 remains constant throughout all tests, in order to maintain route building speed. This parameter is unlikely to cause much impact on final results as the selection of the next stops from the shortlist is dictated by the ‘Stops from shortlist’ value (N in Algorithm 3). The ‘Other stops in range’ value (O in Algorithm 3) also prevents the shortlist length affecting results. Both of these supporting values vary between tests of the next stop heuristic parameters but remain at 2 and 1, respectively, for all other tests.

It is envisaged that the current provider for the NHS vehicles would be used to complete the vehicle rounds. For the purposes of this study, an hourly pay rate, mileage rate, and emissions rate have been assumed based on the FTA Distribution Manager’s cost guide (£10.78, £0.464/mi, 0.45 kg CO₂/mi) (FTA, 2020). Round durations have been rounded up to the nearest 15 minutes to emulate realistic payment practicalities. Vehicle costs have also been assumed based on the same guidance. Gig cycle couriers are often paid on a per-job basis (Lord et al., 2020); a structure that is assumed for modelling cost estimation purposes (i.e. one route equates to one job, £6.75).

Table 4: Run parameters used in tests.

Parameter	Value
Initial maximum surgeries per round - weight tests (vehicle)*	5
Initial maximum surgeries per round - algorithm tests (vehicle)	4
Next nearest stop shortlist	5
Stops from shortlist tested*	2
Other stops in range tested*	1
Vehicle Cost incurred per mile	£0.464 ¹
Wage Cost incurred per driving hour	£10.78 ¹
Cost incurred per cycle job	£6.75 ²
CO ₂ production rate per vehicle mile	0.45 kg CO ₂ /mile ¹
Service Time per stop	2.5 minutes ³

*Parameter used during weight value tests, varied during heuristic parameter tests

¹ (FTA, 2020), ² (Stuart, 2020), ³ (Allen et al., 2018)

5.2. Solution Limits

Where the weighting of the objective function is varied during experiments, the limits of the solutions are governed by each of the three objective terms. Under the given constraints, the maximum value for u (the maximum time to delivery) is 90 minutes. The minimum is limited by the travel time by road vehicle to the surgery which is furthest (wrt. travel time) from the hospital, with no cycle consolidation options. In the case of the Southampton case study, this is the Milford Medical Centre, with a return journey time of 70.24 minutes. The service time also must be considered, hence the minimum u falls at 75.24 minutes (70.24 mins + 2 × (2.5 minutes) = 75.24 minutes).

The number of collection rounds, $\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} x_{\bar{r}}$, has an upper limit equal to the number of surgeries in S . There is no clear lower bound whilst other constraints limit the maximum journey time

and stop limit. Should these limits be removed, the lower limit is 1 round, with the solution being the travelling salesman problem.

The total collection round time, $\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} t_{\bar{r}} x_{\bar{r}}$, has an upper limit dictated by the sum of the return journey times and service times when one vehicle is used per surgery. There is no clear lower bound whilst other constraints limit the maximum journey time and stop limit. Should these limits be removed, the lower limit is the travel time and service time of the single-vehicle travelling salesman problem.

In the subsequent tests using the Isle of Wight data, a lower limit for u of 43.98 minutes applies (return travel time + service time = 38.98 mins + $2 \times (2.5 \text{ mins}) = 43.98 \text{ mins}$). This time is given by the Grove House Surgery, Ventnor, or the Ventnor Medical Centre, Ventnor. Both surgeries reside on the same, one-way street; thus, have the same transit time to the hospital.

The minimum achievable number of collection rounds (i.e. vehicles), objective 2, and total driving time, objective 3, are not easily determined. This is due to the problem being a generalisation of the Capacitated Vehicle Routing Problem (CVRP), and the Travelling Salesman Problem (TSP) with added complexities from the introduction of cycle consolidation, meaning the vehicle stops are not initially known.

Table 5: Upper and lower limits of objective function outcomes. Soton = Southampton

Objective	Lower Limit (Soton)	Upper Limit (Soton)	Lower Limit (IOW)	Upper Limit (IOW)
u	75.24 mins	115 mins	43.98 mins	115 mins
$\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} x_{\bar{r}}$	Indeterminate	97 rounds	Indeterminate	23 rounds
$\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} t_{\bar{r}} x_{\bar{r}}$	Indeterminate	2854 mins	Indeterminate	382 mins

5.3. Objective Input Sensitivity Analysis

Where the master problem presents multiple objectives, many different solutions can exist. To access these and offer decision-makers the choice of trade-offs towards each, the weights of the multi-objective function (Equation 1), w_1 , w_2 , and w_3 , can be modified to favour the different terms. Weight sensitivity tests were carried out in which weights were independently varied, with w_1 and w_2 varying from 0 to 10, and w_3 from 1 to 10. This allowed testing of the algorithm's robustness and ensured that it could be used in any combination of weighting. Testing w_1 with weights set to 0 also allowed understanding of the relationship between the constraining u and objectifying it, whilst setting w_2 to 0 allowed understanding of the relationship between the second and third objectives, which are closely linked.

5.3.1. Southampton

Using the same initially generated set of routes, one full run (including heuristic stages) was completed for each weight using the Southampton area data. The results of the even weight tests (e.g. 2-2-2, 2-2-4, etc.) are tabulated in Table A.10 (Appendix Appendix A).

The objective weightings that produced dominating solutions (i.e. better than all other solutions in one or more objective terms) were identified to produce a Pareto front of solutions (Table 6, Figures 10a, c, e). They could be offered to decision-makers to assist in choosing a weighting


that best suits their needs. The configurations which gave dominating solutions were isolated and displayed in Figures 10b, d, and f. Where some configurations give better solutions in one or more terms of the objective function, their benefits in other potential performance indicators (e.g. cost and emissions) may be more limited.

Interestingly, the dominant solutions in the Southampton dataset all featured a modelled objective function with a zero-weighted term. Whilst in these tests the weighting was set to 0, in practice they could be set to a negligible value when compared to the other objectives to achieve the same result; thus, these solutions are still relevant to this problem. The selection of the 0-5-5 weighted objective solution achieves a minimum in the second and third objectives, which somewhat align; though, the 90-minute constraint on u prevents excessive collection round durations. Furthermore, the other dominating solutions from objective weightings 4-0-2 and 10-0-1 will have occurred due to the second and third objectives aligning sufficiently in this use case. In use cases with a greater sensitivity to the second objective, the weighting w_2 would be more important.


Generally, increasing w_2 and w_3 relative to w_1 results in a higher u result, and a lower $\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} x_{\bar{r}}$ and $\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} t_{\bar{r}} x_{\bar{r}}$, as expected. In turn this reduces costs and CO₂, confirming McDonald's expectation (McDonald, 1972). In some cases the trade-off decision makers must make is marginal; for example, the dominating solution resulting from the 4-0-2 weighting makes only a 3.9% (3.33 mins) reduction in u compared to the cheapest solution weighting (0-5-5), whilst costing 5.6% (£27.35) more. Equally, u in the 4-0-2 solution is 10.4% (7.79 mins) longer than the solution with the fastest u value (10-1-1); however, this trade-off enables a reduction of 1 round and an 12.7% (64.87 mins) reduction in total driving time.

Table 6: Solutions from the weight testing results for the Southampton data set. The coefficients, w_1 , w_2 , and w_3 , are the ratios used to weight objectives in each test. Only dominating solution configurations are shown.


w_1	w_2	w_3	u (mins)	$\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} x_{\bar{r}}$ (qty)	$\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} t_{\bar{r}} x_{\bar{r}}$ (mins)	Run Cost (£)	Run CO ₂ (kg)
0	5	5	86.36	6	420.47	486.57	92.19
4	0	2	83.03	6	445.12	513.92	101.02
10	0	1	75.24	7	509.99	550.92	128.33




(a) Solutions with respect to objective function terms (all solutions). Line indicates Pareto front found in testing.




(b) Solutions with respect to objective function terms (dominating solutions only).




(c) Maximum time to delivery (u) vs. Total Solution Costs (all solutions).



(d) Maximum time to delivery (u) vs. Total Solution Costs (dominating solutions only).



(e) Total collection round time ($\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} t_{\bar{r}} x_{\bar{r}}$) vs. Total Solution Costs (all solutions).



(f) Total collection round time ($\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} t_{\bar{r}} x_{\bar{r}}$) vs. Total Solution Costs (dominating solutions only).

Figure 10: Southampton Objective Function Weight Tests. Marker shapes indicate the number of rounds used in the solution ($\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} x_{\bar{r}}$). Left = all solutions from tests, right = dominating solutions only.

5.3.2. Isle of Wight


To test the robustness of the algorithm and demonstrate its effectiveness in a variety of test environments, the investigation of the objective function weighting was repeated for the IOW data set. The dominating solutions are shown in Table 7. Table A.11 (Appendix Appendix A) displays the full results for the even weight values.

The IOW data produced a wider range of vehicle options, likely due to the low density and distribution of the surgeries covered by the delivery service. The algorithm determined a minimum number of rounds (3) under the timing constraints, whilst the lower limit of u is only reached when there are more rounds (9). This is due to the surgery distribution causing large quantities of stem mileage. As with the Southampton dataset, there are many cases of marginal trade-offs to be decided on by the decision-maker. No dominating solution with 6 or 8 rounds exists, as the surgery which limits u in these solutions can be served by solutions with 5 or 7 rounds, respectively.


As with the Southampton tests, there are some dominating solutions which were achieved when a zero-weighting was used on one of the objective terms. As previously alluded to, if the relative objective weights are similar, some solutions may be achieved with multiple weight combinations. This was the case with the IOW study, with many weight ratios producing the same route selection (only one example for each key solution is shown). The selection of the 0-6-5 weighted objective solution achieves a minimum in the largely aligning second and third objectives, whilst the solutions from the 1-2-6 and 3-3-1 weightings allow a greater balance to be struck between the objectives. The importance of the second objective is seen in the dominating solution from the 4-0-1 weighting, where the sensitivity to the number of vans used is demonstrated and a somewhat excessive total of 9 vans are used.

Table 7: Dominating solutions from the weight testing results for the IOW data set. Many weight ratios produced the same route selection; only one example for each key solution is shown.


w_1	w_2	w_3	u (mins)	$\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} x_{\bar{r}}$ (qty)	$\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} t_{\bar{r}} x_{\bar{r}}$ (mins)	Run Cost (£)	Run CO ₂ (kg)
0	6	5	74.08648	3	213.12	125.39	46.34
1	2	6	61.65555	4	215.41	129.87	48.19
3	3	1	61.18522	4	223.47	136.48	52.28
4	0	1	43.98037	9	354.79	194.21	87.16
7	0	3	54.04353	5	238.3	144.27	54.78
10	1	1	50.42818	6	290.98	167.05	70.09




(a) Objective Function (all solutions). Line indicates Pareto front found in testing.




(b) Objective function (dominating solutions only).




(c) Maximum time to delivery (u) vs. Total Solution Costs (all solutions).



(d) Maximum time to delivery (u) vs. Total Solution Costs (dominating solutions only).



(e) Total collection round time ($\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} t_{\bar{r}} x_{\bar{r}}$) vs. Total Solution Costs (all solutions).



(f) Total collection round time ($\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} t_{\bar{r}} x_{\bar{r}}$) vs. Total Solution Costs (dominating solutions only).

Figure 11: IOW Objective Function Weight Tests. Marker shapes indicate the number of rounds used in the solution ($\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} x_{\bar{r}}$). Left = all solutions from tests, right = dominating solutions only.

5.4. Construction Heuristic Parameterisation

Varying the number of routes produced in the initial generation stage of the algorithm enabled testing of the performance of the heuristic methods and identify the settings which provided the most effective balance of computational time and solution quality. The variation in the number of routes was achieved through (i) changing the initial maximum stop limit, or (ii) changing the number of stop options tested at each stop using the next stop shortlist selection and other stops in range parameters. Each set of parameters was run 10 times, and an average of the results was taken (Table 8, Figure 12). The Southampton dataset and objective function weights of $w_1 = 4$, $w_2 = 0$, $w_3 = 2$ were used, based on the sensitivity tests conducted, to give a reasonable balance between the objectives. The redistribution heuristics were applied to the longest and shortest routes (relative to the average route) in the existing solution.

Results indicated a positive correlation between the number of initial vehicle routes and the run time of the VRP algorithm (Figure 12a). As the next stop parameter values are increased, more van routes will be made using similar surgery combinations which serve almost the same cycle catchment area, resulting in a longer run time in the optimisation environment. As the maximum stop value increases, the routes are able to visit a wider range of surgeries beyond the current surgery catchment, therefore with fewer options for the same surgeries but more routes overall. Thus, increasing the initial number of routes using the maximum stop value generally results in a faster run-time per route than using the next stop parameters.

Final solutions contained a higher percentage of routes originating from the SRRH, suggesting that the SRRH was more effective than the LRRH in all cases (Figure 12b). The number of routes generated by the LRRH was also lower than the SRRH. In an isolated series of 5 test runs (initial next stop = 5, next stop parameters = 2, 2) an average of 4 routes and 13 routes were generated by the LRRH for the longest and second-longest routes, respectively, in the first iteration. In comparison, 77 routes and 59 routes were generated by the SRRH for the shortest and second shortest routes, respectively. New routes beyond the time length of the current maximum were not permitted, hence many LRRH routes were discarded, accounting for the disparity. Solutions from tests with fewer initial routes contained a higher percentage of routes that originated from the heuristic methods, whilst more initially generated routes led to lower usage of heuristic routes.

When varying the initial maximum number of stops, the best value for the first objective term, u , was achieved when the maximum was set to 2 stops or 3 stops (75.24 mins) (Figure 12c). The best values for the second and third objective terms, $\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} x_{\bar{r}}$ and $\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} t_{\bar{r}} x_{\bar{r}}$, were achieved when the maximum was set to a maximum of 6 stops (6 rounds/438.02 mins) (Figures 12e, 12g). Meanwhile, the best objective function value was achieved when the maximum stop parameter was set to 3, though a general trend of a lower objective function result as the maximum stop parameter increased was observed. These results were likely due to shorter routes favouring quicker delivery, meanwhile longer routes favour the use of fewer vehicles and a reduced stem driving total (to and from the hospital site).


In the case of varying the next stop parameters, the best value for the first objective term, u , was achieved when the next stop parameters were set to 1 (75.24 mins) (Figure 12d). The best values for the second and third objective terms, $\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} x_{\bar{r}}$ and $\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} t_{\bar{r}} x_{\bar{r}}$, were seen when next stop parameters were set to 2 (6 rounds, 437.6 minutes) (Figures 12f, 12h). Meanwhile, the best objective function value was achieved when the next stop parameters were set to 4. These findings

were to be expected, as testing a greater number of stop combinations will likely produce more favourable solutions.


The findings presented in Figures 12c-12h suggest that the heuristic parameters can be adjusted to improve computational time and results, though beyond a certain value there is little benefit seen. In the maximum stop parameter, values up to 5 stops generally clustered around with runtimes of approximately 30 seconds, and the difference seen in the objectives (relative to the best solutions) was less than 6%. At values higher than this, runtimes more than doubled on average. A similar trend was seen in the next stop parameter tests, though with a marked improvement in the second/third objective results and a decrement in results in the first objective results between 1 and 2, and a significant increase in runtime above 2 with little benefit in solution quality. Based on these results, it could be suggested that the best outcomes are likely to be achieved with a maximum stop parameter of 5, and next stop parameters of 1 or 2. This is reflected in the subsequent algorithm performance tests (Section 5.5).

Table 8: Mean results of the heuristic performance tests. Each run was completed 10 times. The nearest stop shortlist length was a maximum of 10 surgeries across all tests.


Initial Max Stops	Shortlist Stops Tested	Other Stops in Range Tested	Mean Run Time (mm:ss)	Mean Initial No. Of Veh Routes	Mean u (mins)	Mean $\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} x_{\bar{r}}$ (qty)	Mean $\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} t_{\bar{r}} x_{\bar{r}}$ (mins)	Mean Original Generation	Mean Redist. Longer Heur.	Mean Redist. Shorter Heur.
2	2	1	00:36	87	75.24	8	542.37	12.5%	0.0%	87.5%
3	2	1	00:24	348	77.75	7.1	501.86	18.5%	10.2%	71.3%
4	2	1	00:29	1123	80.28	6.9	491.53	30.4%	5.7%	63.9%
5	2	1	00:36	3252	78.41	6.9	491.09	61.0%	1.4%	37.6%
6	2	1	01:20	8896	81.21	6.6	480.28	95.7%	0.0%	4.3%
7	2	1	03:30	23611	78.99	6.5	478.49	83.3%	0.0%	16.7%
5	1	1	00:24	1221	77.33	7.5	530.97	41.3%	5.0%	53.8%
5	2	2	00:57	6692	79.71	6.7	479.61	64.8%	0.0%	35.2%
5	3	3	02:26	19923	79.39	6.7	477.22	76.2%	0.0%	23.8%
5	4	4	06:30	44334	78.39	7	483.62	97.1%	0.0%	2.9%




(a) VRP Run-time vs. Initial Number of Routes




(b) Route Origins vs. Initial Number of Routes




(c) Max. Stop VRP Run-time vs. u




(d) Next Stop VRP Run-time vs. u




(e) Max. Stop VRP Run-time vs. $\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} x_{\bar{r}}$



(f) Next Stop VRP Run-time vs. $\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} x_{\bar{r}}$



(g) Max. Stop VRP Run-time vs. $\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} t_{\bar{r}} x_{\bar{r}}$



(h) Next Stop VRP Run-time vs. $\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} t_{\bar{r}} x_{\bar{r}}$

Figure 12: Heuristic Performance test results. Symbols indicate the parameter that was varied, e.g., maximum number of stops or next stop options tested in the initial route generation.

5.5. Algorithm Performance

To validate the performance of the algorithm, a series of small test cases were used, allowing full enumeration of all van and cycle routes to ensure the optimal solution was found when using the model presented in Section 4.3. The tests selected 10, 15, or 20 surgeries at random from the Southampton dataset, generating routes for those surgeries only. The 10 surgery and 15 surgery tests were completed 10 times each, and the 20 surgery test was completed 5 times. The algorithm was set to have an initial maximum stop parameter of 4 nodes, nearby next stop shortlist of 5, nearby next stop selection parameter of 2, and other next stop selection parameter of 1.

In all tests the algorithm performed well, reducing the computational time by up to 99% whilst achieving an average 5% difference from the optimal solution (Table 9). As would be expected, the algorithm generally achieved closer to optimal solutions on smaller test cases, though made larger time savings on larger test cases. The reason for this behaviour is that in small cases the original set of routes generated would have been a greater proportion of the fully enumerated set of routes, resulting in a better initial objective function from which to converge and a lower dependency on random selections.

Some tests saw larger deviations from optimality, primarily due to differences in the components of the objective function results (Figure 13). The presented algorithm resulted in a slight bias towards the maximum collection round duration when compared to the optimal solution. Whilst this resulted in improvements in this component of up to 8%, the trade-off in terms of the number of vehicles and driving time resulted in an overall deviation from optimality. To this end, it is likely that the length of the routes initially generated by the algorithm favour this objective and create this skew.

The algorithm was seen to produce consistently good results, with some solutions deviating by less than 1% from optimal (Figure 14). The distribution of results relative to the objective function did vary, suggesting there is scope to improve the consistency of the algorithm in terms of the objective function in the chosen solutions, possibly through further exploration of the weightings and parameterisation. Meanwhile, runtimes became more consistent as the test sizes increased, suggesting that larger use cases will offer reliable savings to planners who may wish to create plans at short notice, similar to how do C. Martins et al. (2021) highlighted the importance of speed in agile vehicle route planning in a related problem.

With a short-notice planning approach, a rapid response time is important. Hence, on investigation of the absolute values produced in the tests (Figure 15), it was seen that enumeration followed an exponential distribution as the case size increases, whilst the algorithm performed in a comparatively more linear relationship. This relationship is key for solving the SSCP, with larger applications (e.g. Southampton case study area) of this problem potentially requiring impractical lengths of time to solve to optimality. It should also be noted that the algorithm reduces the memory requirements for solving when compared to enumeration, meaning that larger case sizes can be processed more universally.

Table 9: Mean results of the algorithm performance tests. 10 surgery and 15 surgery tests were completed 10 times, 20 surgeries was completed 5 times. Absolute values show Algorithm / Optimal results, respectively. Percentages show the difference from the optimal solution.

Test	Max. Coll. Round	Qty. Vans	Van Duration	Obj. Fn.	VRP Runtime
10 Surgeries (Avg)	62.95 / 64.27	3 / 2.5	159.63 / 147.06	3.46 / 3.36	0.29 / 5.71
15 Surgeries (Avg.)	70.1 / 76.23	4 / 2.7	242.36 / 196.69	4.24 / 4.03	0.35 / 26.98
20 Surgeries (Avg.)	74.43 / 73.24	4.4 / 3.4	270.04 / 239.28	4.57 / 4.24	0.54 / 301.53
All Tests (Avg.)	68.35 / 70.94	3.71 / 2.78	216.93 / 187.43	4.02 / 3.82	0.37 / 82.16
10 Surgeries (Avg.)	-2%	20%	8%	3%	-69%
15 Surgeries (Avg.)	-8%	50%	24%	5%	-95%
20 Surgeries (Avg.)	2%	30%	13%	8%	-99%
All Tests (Avg.)	-3%	34%	15%	5%	-86%

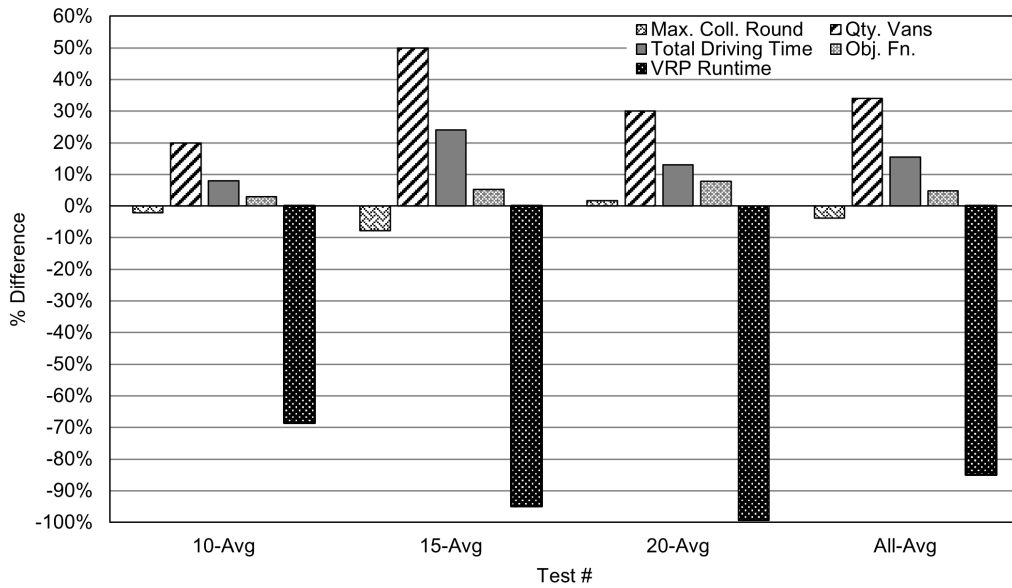


Figure 13: Mean results of each set of algorithm performance test relative to optimal results, comparing the three objectives, overall objective function value, and runtime.

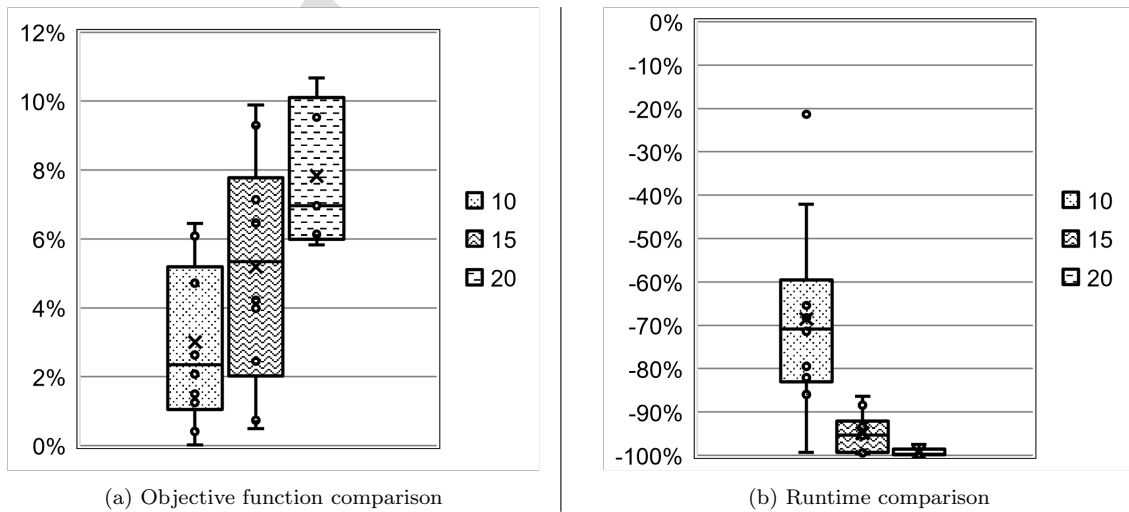


Figure 14: Box plot of algorithm performance test results distribution, relative to full cycle and van enumeration results. Pattern indicates test case size.

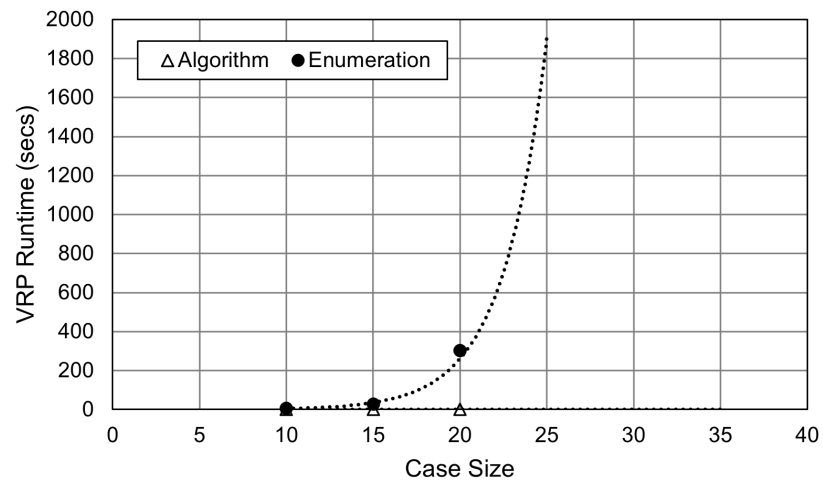


Figure 15: Comparison of absolute computational times with forecast exponential trend lines.

6. Comparison to Business-as-usual

The results of the computational experiments were analysed through a comparison with the BAU case study to identify the effectiveness of the algorithm. In addition to the three terms of the objective function (u , $\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} x_{\bar{r}}$, and $\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} t_{\bar{r}} x_{\bar{r}}$), eight further Key Performance Indicators (KPIs) have been used to quantify the success of the algorithm. These KPIs include the average driving time of the selected van routes ($\frac{\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} t_{\bar{r}} x_{\bar{r}}}{\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} x_{\bar{r}}}$), the average time to delivery per surgery (i.e., Mean $t_{\bar{r}} x_{\bar{r}} \forall i \in S$, where $S \in S_{\bar{r}}$), the average number of stops per van route, the costs associated with the vehicle (running costs + driver costs), and total operating costs. Where the BAU rounds are more widely spread with intermediate deliveries, the duration from the previous stop at the hospital to the next is used for the equivalent value for $t_{\bar{r}}$.

Whilst the solutions generally present improvements over the BAU in terms of the core objectives, it should be highlighted that this is somewhat expected due to the newly imposed limits on round duration and the introduction of cargo cycles which are absent in the BAU cases. Despite not comparing exactly the same conditions, the results typically show a potential for improvement in the core objectives, though are reflected by significant changes in costs, and make it harder to compare directly. This highlights the possible trade-offs that could occur and the sensitivity towards different BAU inputs. Nevertheless, the comparisons should be interpreted with the imposed changes in mind.

6.1. Southampton

In the Southampton case study (Figure 16a), the algorithm solutions out-performed the BAU case in all areas except total cost, where no cycle costs were incurred due to the operation being van-based only. The average stops per van were considerably reduced; where each vehicle was seen to visit 2x the number of surgeries compared to the algorithm's solutions (5 stops average vs. 10).

With respect to the objective function terms, all of the dominant solutions from the tests out-performed the BAU case. The maximum time to delivery (u) was reduced by 70% in the worst case, and 74% in the best case (285 mins vs. 86 mins vs. 75 mins). The number of rounds (vehicles) used ($\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} x_{\bar{r}}$) was reduced by 30% in the worst case, and 40% in the best case (10 vehs. vs. 7 vehs. vs. 6 vehs.). The total driving time ($\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} t_{\bar{r}} x_{\bar{r}}$) was reduced by 58% in the worst case, and 65% in the best case (1205 mins vs. 509 mins vs. 420 mins).

An average of 49 cycle courier tasks were introduced in each of the dominating solutions, contributing £331 to the cost of the proposed solutions. It should be noted that whilst 49 courier tasks were generated, this does not necessarily correspond to 49 couriers, as this depends on the allocation by the 3PL provider.

The overall cost was increased by between 22% (best case) and 38% (worst case) when compared to the BAU case (£399 vs. £487 vs. £551), with the additional cost of the cycle tasks being partially offset by a 45% (worst case) to 58% (best case) reduction in driving and vehicle costs (£399 vs. £220 vs. £169). The shift towards cleaner transportation was also seen to reduce CO₂ output by up to 43% (162kg vs. 92kg).

It should be remembered that the BAU rounds complete other services, such as internal mail deliveries, and the comparisons in this study were made with respect to a modified version of these rounds in which only specimen collection stops were considered; thus, the comparison is not exact. Nevertheless, this study does highlight the inefficiency of the current route planning with respect to specimens, especially when many of the ancillary services are not required on a daily basis, or are being phased out. Thus, the algorithm outputs still represent the potential savings offered to future rounds.

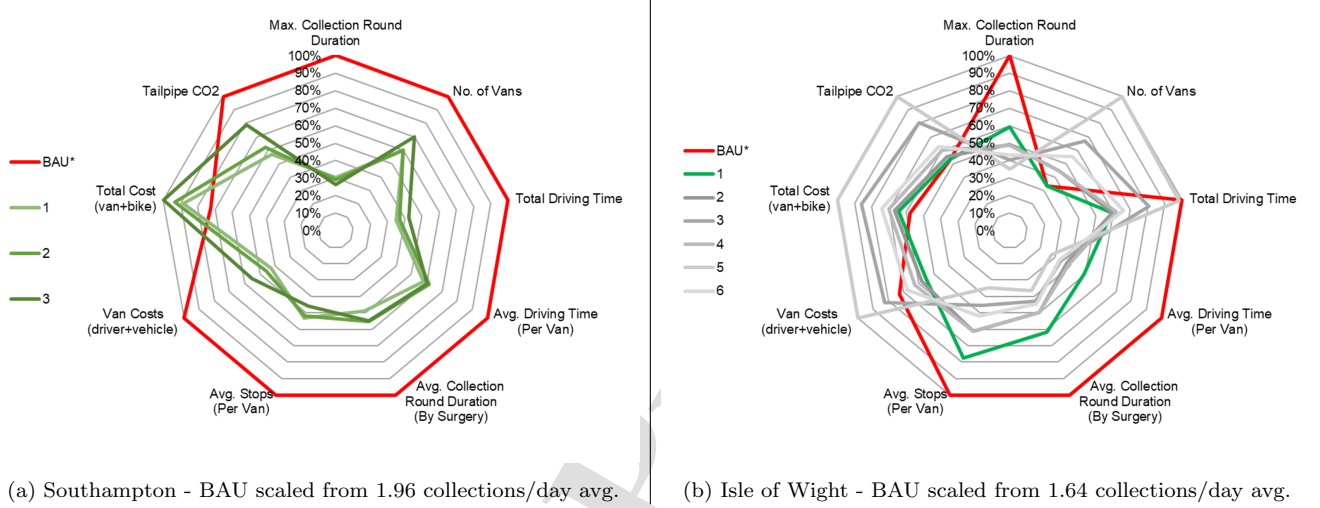


Figure 16: Radar plot KPI Comparison of key solutions from weight sensitivity testing (numbered plots) with BAU activity (red plots). BAU is shown in red. Green plots dominate the BAU in the objective function terms, grey do not. Values scaled with maximums shown at 100%. Smaller is ‘better’ with the exception of ‘Avg Stops Per Van’. BAU activity scaled from multi-collection rounds to compare like-for-like.

6.2. Isle of Wight

In the Isle of Wight case study, the algorithm’s results also suggested an improvement over the BAU case (Figure 16b), though not to the same extent as the Southampton study. Improvements were consistently observed across all dominant solutions with respect to u (max. time to delivery), $\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} t_{\bar{r}} x_{\bar{r}}$ (total driving time), and average round time. Only one solution was seen to dominate the BAU case in all three objective terms (1 in Figure 16b), whilst the other solutions dominated in terms of u and $\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} t_{\bar{r}} x_{\bar{r}}$, but not $\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} x_{\bar{r}}$ as they exceeded three vehicles in their solutions.

In a dominant solution (1), a 41% improvement in u was observed (125 mins vs. 74 mins), and 41% in $\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} t_{\bar{r}} x_{\bar{r}}$ (360 mins vs. 213 mins), but no change in terms of $\sum_{\bar{r}, \bar{r} \neq \bar{r}_0} x_{\bar{r}}$ was seen.

Six cycle courier tasks were introduced in the dominating solution, contributing £40.50 to the cost. As a result, the overall cost increased by 12% when compared to the BAU case (£112 vs. £125), with the additional cost of the cycle tasks being mostly offset by a 24% reduction in driving and vehicle costs (£112 vs. £85).

The IOW BAU case was not completing any ancillary services and was handling sample collection and delivery only, making their routes more effective for this purpose. This may partly explain the disparity in the algorithm’s effectiveness between the two case studies.

7. Conclusions

This paper presents the Sustainable Specimen Collection Problem (SSCP) for the collection of diagnostic specimens from GP surgeries to a hospital laboratory, and solves it using a weighted objective function and a column generation based heuristic algorithm with some additional improvement heuristics.

The algorithm used was effective in producing solutions that improved on the BAU case with respect to the maximum round duration, the number of rounds (vans), and the total driving time. A weighted multi-objective function with a column generation approach was used to identify the Pareto front, with results suggesting that the approach used was robust for a range of inputs, and could be configured to a decision-maker's needs. A set of solutions was produced, some of which reached the lower bounds of the objective, $\min u$; meanwhile, performance tests indicated that the proposed algorithm was efficient and reduced computational time by up to 99% whilst achieving an average of 5% deviation from optimality.

When using cycle courier consolidation, results suggest that sample delivery times could be improved by between 41% and 74%, based on a comparison with two real-world case studies. Such improvements could be achieved using either the same fleet of vehicles or a reduced fleet. Depending on the locality and selected solution, a reduction in fleet size of up to 40% was possible. A simultaneous reduction in driving time between 41% and 65% was also observed through this model. The proposed system benefits were dependent on the introduction of cycle couriers who offered significant benefits in dense urban areas, with the Southampton case study solutions requiring an average of 49 courier tasks to be completed, whilst IOW solutions required only 6 tasks.

With the requirement to pay for multiple new cycle courier tasks, proposed solutions caused significant price increases of up to 38% in some cases. A proportion of these new rider costs were offset by van and driver reductions, though in reality, these costs are likely to be slightly larger due to the additional requirement to manage the system, as well as fees charged by gig-economy logistics providers. Additional benefits are seen in the reduced CO₂ tailpipe emissions of up to 43%; a key factor in urban areas where there is a need to reduce emissions from logistics activities (European Commission, 2013). Consistent reductions in the number of simultaneously operating vehicles will allow fleet sizes to be reduced, having knock on effects for congestion and operating cost, beyond those explored in the BAU comparisons.

The modelled improvements can offer better utilisation of assets, cleaner transportation, and the potential to improve quality of care in communities through faster and less damaging specimen deliveries, whilst also enabling the possibility of later final collections from surgeries. However, it should be noted that the possible damaging effects (e.g. vibration) of multi-modal transportation on sensitive goods, such as diagnostic specimens, has not been widely explored and should be fully understood before such a heterogeneous model is implemented (Nybo et al., 2019; Oakey et al., 2021). Equally, security of goods and dangerous goods authorisation should be important considerations in the use of multiple modes (Grote et al., 2021; Oakey et al., 2022), as is the case this model.

7.1. Potential for Application

The algorithm described in Section 4 and explored in Section 5 has further application beyond the case studies described in Section 2. The algorithm could be applied to other localities in which

diagnostic specimens need timely delivery. Some areas may have multiple laboratory sites which can be delivered to, though the SSCP model could be adapted to encompass such a variation. Largely the model would remain the same, with the exception of routes being generated to all possible surgeries from each hospital. The constraints in place should subsequently eliminate poor solutions.

The additional costs associated with the introduction of the cyclists may not be the most attractive proposition for decision-makers, though the additional flexibility and speed of delivery offered by this system may be sufficient to warrant the extra cost. Whilst beyond the scope of this paper, future work could account for the cost of introducing the cyclists in the problem to limit the cost increase.

Other situations in which timely delivery is required at a single regional point following collection from a known set of points include delivery of blood donations to a blood manufacturing centre, or ballot boxes during an election or referendum, though security could be a concern in this model due to the additional parties involved.

Should the problem be reversed; i.e. delivering from a single point to multiple sites using a localised consolidation model, it could be seen that the system could be implemented for short term cold-chain logistics where parties are looking to reduce their social and environmental impact. For example, COVID-19 vaccine distribution is carried out using insulated containers with a limited lifespan to local surgeries and vaccination points. However, there is significant pressure to distribute vaccines in a sustainable, affordable and timely manner, particularly in developing countries with reduced access to reliable vehicles (University of Birmingham, 2020). Equally, applications such as hot meal distribution could also follow the reversed SSCP concept.

7.2. Limitations and Further Development

In the problem posed by the SSCP, scheduling was not explicitly considered due to suggestions made by staff (Wessex Academic Health Science Network, 2020); however, should some level of scheduling be addressed, there is scope to reduce the number of rounds required. For example, if collections are between 9 am-12 pm and 3-6 pm, these periods could be divided into 2x 90-minute slots each, meaning that vehicles could work one round and then another on completion, halving the number of required vehicles. Naturally, more complex collection schedules could be produced with further benefits, though there may be larger challenges in practically applying this in terms of management of appointments and opening hours. As mentioned in the problem description, this system could be developed to manage loads day-to-day, though this may be data and decision-maker dependent.

The tests only considered a worst-case scenario in which all surgeries which could potentially require collection were served. Further analysis of a full data set, in which routing and sample data align, could simulate a 'live' scenario in which the routing selections are varied each day depending on the sample production at each surgery. This would identify the efficiency of the proposed system if dynamic (day-by-day) routing was possible. Additionally, in the case studies explored in this paper, surgeries only send samples to one hospital site due to contractual reasons. The proposed problem could be adapted to be capable of multi-hospital delivery, provided routes

from all hospitals are made at the initial route generation phase.

An online version of the model could also take account of any routing delays, e.g., cycles delayed to consolidation points, enabling routes to be reconstructed whilst they progress, allowing any waiting vans to continue. Furthermore, whilst more sophisticated optimisation approaches could have been used, the approach tested in this paper appeared to be fairly effective. More thorough exploration of the Pareto front using other methods would be a potential future area of research. The results in this paper present a strong base from which to compare future results with.

The evident improvements offered by the algorithm's outputs would suggest that it is effective in its design, though the sensitivity towards each of the coefficients suggest that some areas will perform differently to others. The greater difference seen in the Southampton case study would suggest that the algorithm performs better in areas that feature greater surgery densities, enabling a greater share of the load to be cyclist consolidated. In even denser areas, such as central London, cycle routes may be capable of completing most or possibly all of the rounds faster than road vehicles. Cycle route heuristics may be needed to account for the larger number of potential cycle routes if enumeration is not feasible.

Whilst tests using the more rural IOW case study showed cycle consolidation can still offer some benefits, should the density of surgeries drop below a point at which cycle consolidation is not possible, it is unlikely this routing algorithm will offer many benefits over the BAU. The model could be further tested and developed using real data from both more dense and more sparse environments to understand the limits of its functionality.

The tests use routing provided by a locally hosted GraphHopper Routing engine which considers traffic as 'free-flowing'. Such a state is evidently not a guaranteed representation of the situations seen in the case studies, particularly in the urban areas which are more prone to congestion (GraphHopper, 2020; TomTom, 2021). Based on the times at which samples are collected, collections are likely to be mid-morning and mid-afternoon; thus, avoiding the peak time congestion. In the Isle of Wight case study and more rural parts of the Southampton study, the free-flow assumption is also less likely to be of issue. Cycling times are also less likely to be affected due to the ability for cyclists to ride through congestion. The modelled results using the GraphHopper routing engine will be slightly biased towards faster travel times, meaning magnitudes of improvements may be slightly inflated; though the general findings will remain largely the same.

In the Southampton case study, ancillary services, such as internal mail deliveries/collections, are disregarded in the design of the routes. Whilst it has been indicated that many of these services are being phased out, there may be need for such services to be considered in the design of routes. Given their lower importance with respect to delivery time, a weighting between load types may be a consideration.

Further to the above limitations and improvements, the introduction of other modes and technologies could be introduced to enhance deliveries further. Drone deliveries of medical goods are becoming increasingly prevalent, particularly in response to the COVID-19 pandemic (Loughran, 2020). They are seen to offer significant benefits in areas where land logistics are difficult or slow, though may present high costs and consume significant quantities of energy if they serve each surgery individually. The consolidation model may open the benefits offered by drones whilst im-

proving practicalities. Range and charging/refuelling, as well as routing and weather would need to be considered. Electric vehicles could also be considered in further developments of the problem. Whilst not a major change, tailpipe emissions will be eliminated, though range limitations would need to be addressed. Electrically assisted bicycles and cargo cycles may also be an area of interest, offering the potential for a wider reach and/or more stops in cycle routes, reducing the need for as many vehicle served sites.

Maximising vehicle asset use may be a further consideration in future model developments. Cargoes of similar sensitive nature heading in similar directions could be handled simultaneously by the vehicle in a variant of the multi-vehicle VRP with Pickup and Delivery (VRPPD) (Desaulniers et al., 2002). Given the many services acting simultaneously but independently in health care, there may be scope to combine logistics movements. Combining non-urgent patient movements with diagnostics specimens and un-processed blood donations could be an example of such a collaboration.

Another area of potential interest may be in the investigation of adjusting opening hours and appointment scheduling to affect the demands produced by surgeries. This may not be practical to fully implement but may have potentially beneficial effects.

Acknowledgements

The authors would like to acknowledge the NHS staff at Southampton General Hospital and St. Mary's Isle of Wight Hospital for their support in this study.

Funding

This research received funding from the EPSRC and was completed as part of the e-Drone project under grant no. EP/V002619/1 and the Department for Transport's Future Transport Zones project overseen by Solent Transport.

Contributions

The authors named in this paper made the following contributions: Conceptualization - T.C., A.O., A.M.S.; Data curation - A.O.; Formal analysis - A.O., A.M.S.; Funding acquisition - T.C.; Investigation - A.O., A.M.S.; Methodology - A.O., A.M.S.; Project administration - A.O.; Resources - A.O., T.C., A.M.S.; Software - A.O.; Supervision - A.O., T.C., A.M.S.; Validation - A.O., A.M.S.; Visualization - A.O.; Roles/Writing - original draft - A.O., A.M.S.; Writing - review & editing - A.O., A.M.S., T.C..

Conflicts of Interest

The authors declare no conflict of interest.

References

- Allan, R., 2019. Southampton Pathology Operations University Project Meeting.
- Allen, J., Bektas, T., Cherrett, T., Bates, O., Friday, A., McLeod, F., Piecyk, M., Piotrowska, M., Nguyen, T., Wise, S., 2018. The Scope for Pavement Porters: Addressing the Challenges of Last-Mile Parcel Delivery in London. *Transportation Research Record* 2672, 184–193. doi:10.1177/0361198118794535. publisher: SAGE Publications Inc.
- Allen, J., Piecyk, M., Cherrett, T., Juhari, M.N., McLeod, F., Piotrowska, M., Bates, O., Bektas, T., Cheliotis, K., Friday, A., Wise, S., 2021. Understanding the transport and CO2 impacts of on-demand meal deliveries: A London case study. *Cities* 108, 102973. doi:10.1016/j.cities.2020.102973.
- Anaya-Arenas, A.M., Chabot, T., Renaud, J., Ruiz, A., 2016. Biomedical sample transportation in the province of Quebec: a case study. *International Journal of Production Research* 54, 602–615. doi:10.1080/00207543.2015.1018455. publisher: Taylor & Francis eprint: <https://doi.org/10.1080/00207543.2015.1018455>.
- Anderluh, A., Hemmelmayr, V.C., Nolz, P.C., 2017. Synchronizing vans and cargo bikes in a city distribution network. *Central European Journal of Operations Research* 25, 345–376. doi:10.1007/s10100-016-0441-z.
- Anderluh, A., Nolz, P.C., Hemmelmayr, V.C., Crainic, T.G., 2021. Multi-objective optimization of a two-echelon vehicle routing problem with vehicle synchronization and 'grey zone' customers arising in urban logistics. *European Journal of Operational Research* 289, 940–958. doi:10.1016/j.ejor.2019.07.049.

- Baldacci, R., Mingozzi, A., Roberti, R., Calvo, R.W., 2013. An exact algorithm for the two-echelon capacitated vehicle routing problem. *Operations Research* doi:10.1287/opre.1120.1153.
- Bektaş, T., Erdögan, G., Røpke, S., 2011. Formulations and branch-and-cut algorithms for the generalized vehicle routing problem. *Transportation Science* doi:10.1287/trsc.1100.0352.
- Bernal, N., 2020. Deliveroo says couriers are heroes. In fact, they're fighting to survive. *Wired UK Section: Gig Economy*.
- Biesinger, B., Hu, B., Raidl, G.R., 2018. A genetic algorithm in combination with a solution archive for solving the generalized vehicle routing problem with stochastic demands. *Transportation Science* doi:10.1287/trsc.2017.0778.
- Bredström, D., Rönnqvist, M., 2008. Combined vehicle routing and scheduling with temporal precedence and synchronization constraints. *European Journal of Operational Research* 191, 19–31. doi:10.1016/j.ejor.2007.07.033.
- Buluc, E., Peker, M., Kara, B.Y., Dora, M., 2022. Covering vehicle routing problem: application for mobile child friendly spaces for refugees. *OR Spectrum* 44, 461–484. doi:10.1007/s00291-021-00617-0.
- Buntak, K., Kovačić, M., Martinčević, I., 2019. IMPACT OF MEDICAL LOGISTICS ON THE QUALITY OF LIFE OF HEALTH CARE USERS, in: *Quality Festival 2019*, pp. 1025–1032. doi:https://doi.org/10.24874/PES01.02.109.
- do C. Martins, L., Hirsch, P., Juan, A.A., 2021. Agile optimization of a two-echelon vehicle routing problem with pickup and delivery. *International Transactions in Operational Research* 28, 201–221. doi:10.1111/itor.12796. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/itor.12796.
- Cherrett, T., Moore, A., 2020. Saving the NHS: A case study evaluation of drones and cargo cycles for surgery-to-hospital pathology logistics in Southampton UK, in: *Transportation Research Board (TRB) 99th Annual Meeting*, 12-16 January, Washington D.C.
- Conway, A., Cheng, J., Kamga, C., Wan, D., 2017. Cargo cycles for local delivery in New York City: Performance and impacts. *Research in Transportation Business & Management* 24, 90–100. doi:10.1016/j.rtbm.2017.07.001.
- Cuda, R., Guastaroba, G., Speranza, M.G., 2015. A survey on two-echelon routing problems. *Computers and Operations Research* doi:10.1016/j.cor.2014.06.008.
- Deliveroo, 2020. What kit do you offer? URL: <https://riders.deliveroo.co.uk/en/support/new-riders/what-kit-do-you-offer>.
- Demir, E., Bektaş, T., Laporte, G., 2014. The bi-objective pollution-routing problem. *European Journal of Operational Research* 232, 464–478. doi:https://doi.org/10.1016/j.ejor.2013.08.002.
- Desaulniers, G., Desrosiers, J., Erdmann, A., Solomon, M.M., Soumis, F., 2002. 9. VRP with Pickup and Delivery, in: *The Vehicle Routing Problem*. Society for Industrial and Applied Mathematics. *Discrete Mathematics and Applications*, pp. 225–242. URL: <https://epubs.siam.org/doi/10.1137/1.9780898718515.ch9>, doi:10.1137/1.9780898718515.ch9.

- Doerner, K.F., Hartl, R.F., 2008. Health Care Logistics, Emergency Preparedness, and Disaster Relief: New Challenges for Routing Problems with a Focus on the Austrian Situation, in: Golden, B., Raghavan, S., Wasil, E. (Eds.), *The Vehicle Routing Problem: Latest Advances and New Challenges*. Springer US, Boston, MA. Operations Research/Computer Science Interfaces, pp. 527–550. URL: https://doi.org/10.1007/978-0-387-77778-8_24, doi:10.1007/978-0-387-77778-8_24.
- ECF, 2012. Cargo Bike Crazy: The Potential of Delivering Goods By Bike | ECF. URL: <https://ecf.com/news-and-events/news/cargo-bike-crazy-potential-delivering-goods-bike-0>.
- Eitzen, H., Lopez-Pires, F., Baran, B., Sandoya, F., Chicaiza, J.L., 2017. A multi-objective two-echelon vehicle routing problem. An urban goods movement approach for smart city logistics, in: *2017 XLIII Latin American Computer Conference (CLEI)*, pp. 1–10. doi:10.1109/CLEI.2017.8226454.
- Elalouf, A., Tsadikovich, D., Hovav, S., 2018. Optimization of blood sample collection with timing and quality constraints. *International Transactions in Operational Research* 25, 191–214. doi:10.1111/itor.12354. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/itor.12354>.
- Ettazi, H., Rafalia, N., Abouchabaka, J., 2021. Metaheuristics methods for The VRP in Home Health Care by minimizing fuel consumption for environmental gain. *E3S Web of Conferences* 234, 00094. doi:10.1051/e3sconf/202123400094. publisher: EDP Sciences.
- European Commission, 2011. White Paper: Roadmap to a Single European Transport Area – Towards a competitive and resource efficient transport system. URL: <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:52011DC0144&from=EN>.
- European Commission, 2013. A call to action on urban logistics. URL: <https://ec.europa.eu/transport/sites/transport/files/themes/urban/doc/ump/swd%282013%29524-communication.pdf>.
- Exeter Laboratory, 2017. Turn around times / Cellular Pathology / Exeter Clinical Laboratory International. URL: <https://www.exeterlaboratory.com/cellular-pathology/histopathology-speciality-services/turn-around-times/>.
- Ford, M., 2020. Covid-19 testing 'much slower' in community, says Welsh RCN. URL: <https://www.nursingtimes.net/news/coronavirus/covid-19-testing-much-slower-in-community-says-welsh-rcn-24-07-2020/>.
- FTA, 2020. *Manager's Guide to Distribution Costs 2020*.
- Ghiani, G., Improta, G., 2000. Efficient transformation of the generalized vehicle routing problem. *European Journal of Operational Research* 122, 11–17. doi:10.1016/S0377-2217(99)00073-9.
- Godfrey, K., 2020. Evaluation of the expanded Southampton pilot study (Phase 2) for use of saliva-based lamp testing in asymptomatic populations - Final report. URL: https://eprints.soton.ac.uk/445622/1/Evaluation_of_the_expanded_Southampton_pilot_study_Phase_2_for_use_of_saliva_based_lamp_testing_in_asymptomatic_populations_eprints.pdf.
- GraphHopper, 2020. API - Graph Hopper. URL: <https://docs.graphhopper.com/>.

- Grasas, A., Ramalhinho, H., Pessoa, L.S., Resende, M.G., Caballé, I., Barba, N., 2014. On the improvement of blood sample collection at clinical laboratories. *BMC Health Services Research* 14, 12. doi:10.1186/1472-6963-14-12.
- Grote, M., Cherrett, T., Oakey, A., Royall, P.G., Whalley, S., Dickinson, J., 2021. How Do Dangerous Goods Regulations Apply to Uncrewed Aerial Vehicles Transporting Medical Cargos? *Drones* 5, 38. doi:10.3390/drones5020038. number: 2 Publisher: Multidisciplinary Digital Publishing Institute.
- Gruber, J., Narayanan, S., 2019. Travel Time Differences between Cargo Cycles and Cars in Commercial Transport Operations. *Transportation Research Record* 2673, 623–637. doi:10.1177/0361198119843088. publisher: SAGE Publications Inc.
- Hà, M.H., Bostel, N., Langevin, A., Rousseau, L.M., 2014. An exact algorithm and a metaheuristic for the generalized vehicle routing problem with flexible fleet size. *Computers and Operations Research* 43, 9–19. doi:10.1016/j.cor.2013.08.017. publisher: Elsevier.
- Kergosien, Y., Ruiz, A., Soriano, P., 2014. A Routing Problem for Medical Test Sample Collection in Home Health Care Services, in: Matta, A., Li, J., Sahin, E., Lanzarone, E., Fowler, J. (Eds.), *Proceedings of the International Conference on Health Care Systems Engineering*, Springer International Publishing, Cham. pp. 29–46. doi:10.1007/978-3-319-01848-5_3.
- Kramer, R., Subramanian, A., Vidal, T., dos Anjos F. Cabral, L., 2015. A matheuristic approach for the pollution-routing problem. *European Journal of Operational Research* 243, 523–539. doi:https://doi.org/10.1016/j.ejor.2014.12.009.
- Landry, S., Philippe, R., 2004. How Logistics Can Service Healthcare. *Supply Chain Forum: An International Journal* 5, 24–30. doi:10.1080/16258312.2004.11517130. publisher: Taylor & Francis eprint: https://doi.org/10.1080/16258312.2004.11517130.
- Liu, R., Xie, X., Augusto, V., Rodriguez, C., 2013. Heuristic algorithms for a vehicle routing problem with simultaneous delivery and pickup and time windows in home health care. *European Journal of Operational Research* 230, 475–486. doi:10.1016/j.ejor.2013.04.044.
- Lodree, E.J., Carter, D., Barbee, E., 2016. The Donation Collections Routing Problem, in: Kotsireas, I.S., Nagurney, A., Pardalos, P.M. (Eds.), *Dynamics of Disasters—Key Concepts, Models, Algorithms, and Insights*, Springer International Publishing, Cham. pp. 159–189. doi:10.1007/978-3-319-43709-5_9.
- Lord, C., Bates, O., Friday, A., McLeod, F., Oakey, A., Cherrett, T., Martinez-Sykora, A., 2020. Written evidence submitted by the FlipGig project (COV0123). Technical Report. URL: https://committees.parliament.uk/writtenevidence/5690/pdf/.
- Loughran, J., 2020. Secure air corridors enable drone deliveries for Covid-19 health-care supplies. URL: https://eandt.theiet.org/content/articles/2020/10/secure-air-corridors-will-allow-drone-deliveries-for-covid-19-healthcare-supplies/.
- Marujo, L.G., Goes, G.V., D’Agosto, M.A., Ferreira, A.F., Winkenbach, M., Bandeira, R.A.M., 2018. Assessing the sustainability of mobile depots: The case of urban freight distribution in Rio de Janeiro. *Transportation Research Part D: Transport and Environment* 62, 256–267. doi:10.1016/j.trd.2018.02.022.

- McDonald, J.J., 1972. Vehicle Scheduling – A Case Study. *Operational Research Quarterly* (1970-1977) 23, 433–444. doi:10.2307/3007958. publisher: Palgrave Macmillan Journals.
- McLeod, F.N., Cherrett, T.J., Bektaş, T., Allen, J., Martinez-Sykora, A., Lamas-Fernandez, C., Bates, O., Cheliotis, K., Friday, A., Piecyk, M., Wise, S., 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. doi:10.1016/j.trd.2020.102311.
- Naji-Azimi, Z., Salari, M., Renaud, J., Ruiz, A., 2016. A practical vehicle routing problem with desynchronized arrivals to depot. *European Journal of Operational Research* 255, 58–67. doi:10.1016/j.ejor.2016.04.007.
- Ngo, A., Gandhi, P., Miller, W.G., 2017. Frequency that Laboratory Tests Influence Medical Decisions. *The Journal of Applied Laboratory Medicine* 1, 410–414. doi:10.1373/jalm.2016.021634.
- NHS, 2017. Pathology networks | NHS Improvement. URL: <https://web.archive.org/web/20200808010235/https://improvement.nhs.uk/resources/pathology-networks/>.
- NHS, 2019. ODS - NHS Trusts and Sites - NHS England Data Catalogue. URL: <https://data.england.nhs.uk/dataset/ods-nhs-trusts-and-sites>.
- NHS, 2020a. Delivering a ‘Net Zero’ National Health Service. URL: <https://www.england.nhs.uk/greenernhs/wp-content/uploads/sites/51/2020/10/delivering-a-net-zero-national-health-service.pdf>.
- NHS, 2020b. Zooming into a greener future: The case for a zero-emissions courier services. URL: <https://www.england.nhs.uk/greenernhs/whats-already-happening/zooming-into-a-greener-future-the-case-for-a-zero-emissions-courier-services/>.
- NHS, 2021. Invitation to Tender to Supply the Pathology Collection Network for The East and South East London NHS Pathology Partnership.
- NHS, Nixon, D., 2015. Transportation of Specimens Procedure. URL: http://www.sthk.nhs.uk/about/Documents/11615_TransportationofSpecimensProcedureTrustDocument.pdf.
- NHS, Sedman, R., 2020. Pathology Specimen Transport. URL: <https://www.cddft.nhs.uk/media/615106/transport%20sop.pdf>.
- Nybo, M., Cadamuro, J., Cornes, M.P., Rioja, R.G., Grankvist, K., 2019. Sample transportation – an overview. *Diagnosis* 6, 39–43. doi:10.1515/dx-2018-0051. publisher: De Gruyter Section: Diagnosis.
- Oakey, A., Grote, M., Royall, P.G., Cherrett, T., 2022. Enabling Safe and Sustainable Medical Deliveries by Connected Autonomous Freight Vehicles Operating within Dangerous Goods Regulations. *Sustainability* 14, 930. doi:10.3390/su14020930. number: 2 Publisher: Multidisciplinary Digital Publishing Institute.
- Oakey, A., Waters, T., Zhu, W., Royall, P.G., Cherrett, T., Courtney, P., Majoe, D., Jelev, N., 2021. Quantifying the Effects of Vibration on Medicines in Transit Caused by Fixed-Wing and Multi-Copter Drones. *Drones* 5. doi:10.3390/drones5010022.

- On-Demand Workers Australia, 2018. Snapshot: ON-DEMAND FOOD DELIVERY RIDERS. URL: https://d3n8a8pro7vhmx.cloudfront.net/victorianunions/pages/1411/attachments/original/1517351126/ODFDR_Snapshot.pdf?1517351126.
- Osaba, E., Yang, X.S., Fister, I., Del Ser, J., Lopez-Garcia, P., Vazquez-Pardavila, A.J., 2019. A Discrete and Improved Bat Algorithm for solving a medical goods distribution problem with pharmacological waste collection. *Swarm and Evolutionary Computation* 44, 273–286. doi:10.1016/j.swevo.2018.04.001.
- Perboli, G., Tadei, R., Vigo, D., 2011. The two-echelon capacitated vehicle routing problem: Models and math-based heuristics. *Transportation Science* doi:10.1287/trsc.1110.0368.
- Pichler, P.P., Jaccard, I.S., Weisz, U., Weisz, H., 2019. International comparison of health care carbon footprints. *Environmental Research Letters* 14, 064004. doi:10.1088/1748-9326/ab19e1. publisher: IOP Publishing.
- Pop, P.C., Kara, I., Marc, A.H., 2012. New mathematical models of the generalized vehicle routing problem and extensions. *Applied Mathematical Modelling* doi:10.1016/j.apm.2011.05.037.
- Quadir, A., Underhill, A., Mathur, A., Junaid, M., Wietlisbach, M., Howard, T., 2019. Saving the NHS - Group Design Project. Technical Report. University of Southampton.
- Reid, C., 2018. Confessions of a Deliveroo rider: get fit by delivering fast food. *The Guardian* .
- Semet, F., Taillard, E., 1993. Solving real-life vehicle routing problems efficiently using tabu search. *Annals of Operations Research* 41, 469–488. doi:10.1007/BF02023006.
- Sibilski, L.J., Targa, F., 2020. An old solution to new challenges: The rebirth of the cargo bike. URL: <https://blogs.worldbank.org/transport/old-solution-new-challenges-rebirth-cargo-bike>.
- Smith, A., Colmant, A., Bam, L., H van Vuuren, J., 2015. A new vehicle routing problem with application to pathology laboratory service delivery, in: 2015 Operations Research Society of South Africa Annual Conference. URL: https://www.researchgate.net/publication/328956724_A_new_vehicle_routing_problem_with_application_to_pathology_laboratory_service_delivery.
- Stuart, 2020. Payment structure (rest of UK). URL: <https://web.archive.org/web/20210307002604/https://help-partner.stuart.com/en/articles/1504513-payment-structure-rest-of-uk>.
- The NHS Confederation, Academy of Medical Royal Colleges, National Voices, 2016. Changing care, improving quality. URL: <https://www.nhsconfed.org/-/media/Confederation/Files/Publications/Documents/Changing-care-improving-quality.pdf>.
- TomTom, 2021. Southampton traffic report | TomTom Traffic Index. URL: https://www.tomtom.com/en_gb/traffic-index/southampton-traffic/.
- TravelTime, 2021. Travel And Mapping Software | Travel Time Map | Driving Time Map. URL: <https://traveltime.com/>.

- University of Birmingham, 2020. Vaccine Distribution. URL: <https://www.birmingham.ac.uk/research/quest/sustainable-environments/vaccine-distribution.aspx>.
- Verlinde, S., Macharis, C., Milan, L., Kin, B., 2014. Does a Mobile Depot Make Urban Deliveries Faster, More Sustainable and More Economically Viable: Results of a Pilot Test in Brussels. *Transportation Research Procedia* 4, 361–373. doi:10.1016/j.trpro.2014.11.027.
- Wessex Academic Health Science Network, 2020. GP Focus Group.
- Wilson, M.L., 1996. General Principles of Specimen Collection and Transport. *Clinical Infectious Diseases* 22, 766–777. doi:10.1093/clinids/22.5.766.
- Yücel, E., Salman, F.S., Gel, E.S., Örmeci, E.L., Gel, A., 2013a. Optimizing specimen collection for processing in clinical testing laboratories. *European Journal of Operational Research* 227, 503–514. doi:10.1016/j.ejor.2012.10.044.
- Yücel, E., Salman, F.S., Örmeci, E.L., Gel, E.S., 2013b. A constant-factor approximation algorithm for multi-vehicle collection for processing problem. *Optimization Letters* 7, 1627–1642. doi:10.1007/s11590-012-0578-1.
- Zabinsky, Z.B., Dulyakupt, P., Zangeneh-Khamooshi, S., Xiao, C., Zhang, P., Kiatsupaibul, S., Heim, J.A., 2020. Optimal collection of medical specimens and delivery to central laboratory. *Annals of Operations Research* 287, 537–564. doi:10.1007/s10479-019-03260-9.

Highlights

- New multi-objective problem using cycle consolidation to deliver patient samples
- Potential improvement in maximum delivery time, fleet size and driving time
- Increased operational costs from cycles but potential benefits to healthcare
- Effective algorithm reduced computational time with little change from optimality

Journal Pre-proof

Credit Statement

The authors named in this paper made the following contributions:

Conceptualization - T.C., A.O., A.M.S.;

Data curation - A.O.;

Formal analysis - A.O., A.M.S.;

Funding acquisition - T.C.;

Investigation - A.O., A.M.S.;

Methodology - A.O., A.M.S.;

Project administration - A.O.;

Resources - A.O., T.C., A.M.S.;

Software - A.O.;

Supervision - A.O., T.C., A.M.S.;

Validation - A.O., A.M.S.;

Visualization - A.O.;

Writing - original draft - A.O., A.M.S.;

Writing - review & editing - A.O., A.M.S., T.C..