An Adapted Savings Algorithm for Planning Heterogeneous Logistics with Uncrewed Aerial Vehicles

Andy Oakey[a], Antonio Martinez-Sykora[b], Tom Cherrett[a]

[a] *Transportation Research Group, University of Southampton, United Kingdom*
[b] *Southampton Business School, University of Southampton, United Kingdom*

Abstract

This paper proposes a new extension to the Sustainable Specimen Collection Problem (SSCP), where medical specimens are transported by vans, bikes, and uncrewed aerial vehicles (UAVs, or drones) from local medical practices/offices to a central hospital laboratory for analysis, employing a two-echelon collection approach. Time restrictions from existing operations and literature are also introduced, with the study being formulated as a weighted multi-objective problem seeking to minimise (i) operating costs; (ii) transit times; and (iii) energy/environmental impacts. A new adaptation of the Clarke and Wright Savings Algorithm is subsequently presented to create collection rounds that leverage each mode's strengths. Subsequently, routes are compiled into workable fixed shifts using a modified bin-packing algorithm in each iteration.

The approach of this study is based on a case study of the UK's National Health Service (NHS), involving the collection of pathology samples using traditional vans operating within fixed time slots. Using case study data from the Solent region (England), a novel test instance generation methodology was also developed, whereby realistic site positioning and origin-destination travel data are captured to enable effective algorithm experimentation. The findings from applying the proposed algorithm to a set of test instances based on this methodology are subsequently discussed, where it was found that the adapted savings and bin-packing approach produced effective solutions quickly, with 90% of all large instances (200 sites) being solved within 15 minutes. Further algorithm developments and the application of the devised problem/methodologies are also discussed.

Keywords: Clarke and Wright, savings algorithm, instance generation, bin-packing, heterogeneous, multi-mode, mixed-mode, uav, drone, pathology

1. Introduction

Approximately 95% of clinical diagnoses within the UK's National Health Service (NHS) depend on timely access to pathology and diagnostics services (NHS, 2014), with a significant proportion of these requiring efficient logistics and routing. Meanwhile, the NHS is looking to become the world's first net-zero emission health provider (NHS, 2020a), requiring substantial changes to the existing logistics practices that are often fossil-fuel-based and less responsive to demand than desired (Oakey et al., 2023). To this end, they are exploring novel logistics modes, such as bicycles (pedal-assisted or otherwise) (NHS, 2020b) and uncrewed aerial vehicles (often referred to as unmanned aerial vehicles, UAVs, or drones), which have seen growing interest in their use for collecting/delivering diagnostic specimens in the medical sector (Otto et al., 2018).

This paper builds on previous work by Oakey et al. (2023), investigating the potential for multiple vehicle modes to work in a two-echelon arrangement to collect patient diagnostic specimens (often referred to as 'pathology samples') from clinics in a community setting and deliver them to a central laboratory at a hospital for analysis.

As outlined in the initial investigations (Oakey et al., 2023), the rationale for investigating a move away from the incumbent logistics modes (diesel-fuelled vans/trucks) was to encourage faster transit times and lower emission deliveries. This study develops the Sustainable Specimen Collection Problem (SSCP) to include UAVs and work shifts (in addition to the two-echelon bike and van arrangement addressed in the original SSCP), where vehicle assets can complete multiple routes to/from the hospital in a given work period (e.g. a morning shift of a few hours).

To the authors' knowledge, presenting the collection problem with this combination of vehicles, objectives, and constraints is a novel contribution, and a robust solution approach was required to obtain good solutions. The column generation method presented by (Oakey et al., 2023) was deemed unsuitable due to the large number of columns that would be generated when vehicle departure times could be varied, resulting in a significant computational demand. To this end, a new adaptation of the Clarke and Wright Savings Algorithm (CWSA) is presented in this study, whereby multiple savings options using different modes are considered, whilst the scheduling aspects of the problem are formulated as bin-packing algorithms, using new adaptations of the first-fit decreasing and best-fit algorithms to efficiently compile routes into work shifts.

The complexity of the proposed problem also introduces challenges in establishing the performance limits under realistic input conditions. Given the importance of representing realistic case study characteristics, a new approach to generate test instances is proposed and implemented, allowing the algorithm to be configured and its performance assessed.

The following section outlines the relevant literature in the context of this multi-faceted problem, relating to heterogeneous problems with UAVs, healthcare routing problems, and the use of savings and packing algorithms in vehicle routing and scheduling.

2. Literature Review

Studies are slowly emerging that demonstrate UAVs are not beneficial in every environment (Goodchild and Toy, 2018) and (Anonymous, 2022), and including them as part of a mixed-mode (heterogeneous) fleet with other vehicles (e.g. bikes and vans) enables their benefits to be leveraged more effectively.

Many studies have investigated operations with a van and a UAV in a multi-echelon arrangement (Murray and Chu, 2015; Benarbia and Kyamakya, 2022), often referred to as the "flying sidekick". Under such an arrangement, the UAV travels with the van and completes sub-tours from that van (the "mothership"), before returning at a later point in the round. Several extensions of these problems have been proposed (Wang and Sheu, 2019), addressing factors such as energy considerations (Kyriakakis et al., 2022), or time dependencies (Wang et al., 2022); however, none of these arrangements account for the likely practicalities imposed by regulator constraints on flight paths or take-off sites (Anonymous, 2022), nor do they investigate additional vehicle modes such as cycles, which have frequently been cited as effective in urban areas (Conway et al., 2017; Gruber and Narayanan, 2019). Özoğlu et al. (2019) proposed an adaptation of the Clarke and Wright Savings Algorithm that was demonstrated to work effectively.

Due to landing and flight restrictions, large UAVs are generally used for return trips, which is the case in this work; however, in Garcia and Santoso (2019) and do C. Martins et al. (2021) the authors explored the more flexible setting where each UAV was able to visit multiple sites, helping to considerably reduce the driven vehicle time. More elaborate settings based on a two-echelon approach are explored in (Hu et al., 2023). More modes of transport are considered in Kovač et al. (2021), for which rail and river freight were used in combination with vans and UAVs in urban logistics, with a "Measurement of Alternatives and Ranking according to Compromise Solution" (MARCOS) Multi-Criteria Decision-Making (MCDM) method being used to identify the most beneficial strategies. A more general problem applied in the context of humanitarian aid is addressed in Scott and Scott (2017) and Romero-Mancilla et al. (2023), where they investigated UAV routing with heterogeneous multi-echelon arrangements using vans and UAVs to deliver humanitarian aid. Both problems addressed a combined routing and facility location problem, using vans before onwards movement by UAV to the customers, and used objectives to minimise either the costs and delivery times, subject to capacity and budget constraints.

The weather resilience of UAVs is a notable challenge highlighted in several studies (Gao et al., 2021; Ranquist, 2017) and (Anonymous, 2022), and in the event of poor weather, a heterogeneous system should be able to adapt plans accordingly by switching to use alternative modes for sites that are not otherwise accessible by air, or scheduling flights for times when conditions are more favourable. To this end, Thibbotuwawa et al. (2020) considered the effects of weather in their formulation, investigating different planning strategies and their effects on factors such as energy consumption and customer satisfaction (a function of delivery priority, demand, and weight). The findings suggested that changing UAV fleet sizes could make services more flexible when wind conditions change; however, it was assumed that flights were always possible in the tested scenarios, and the UAVs could carry out multiple deliveries per route. From the authors' experience implementing flight trials, this is not likely due to the limitations of the technology, where practical payloads require a UAV of significant size; and current regulatory constraints, which restrict take-off/landing locations. With this in mind, the present study seeks to address such limitations and allow for short-notice planning of routes and vehicle allocation in light of traffic and weather forecasts (i.e. in the period directly before implementation on the same day). It should be noted that other modes of transport, such as vans, also face changeable conditions due to traffic or even weather. In this context, Wang et al. (2022) identified a heuristic approach to minimise the impact of traffic on routing vans and UAVs.

In addition to the challenges posed by the heterogeneous VRP when introducing UAVs, there are several factors to account for in healthcare logistics, including time constraints and the multiple

objectives required to satisfy the targets of the stakeholders involved. As was highlighted in (Oakey et al., 2023), there are several similar healthcare logistics VRP studies addressing the movement of diagnostic specimens between medical facilities, subject to constraints around time, cost, and capacity, although heterogeneity was not commonplace. These included Cherrett and Moore (2020); McDonald (1972); Grasas et al. (2014); Smith et al. (2015); Elalouf et al. (2018); Yücel et al. (2013a,b); Zabinsky et al. (2020); Anaya-Arenas et al. (2016); and Naji-Azimi et al. (2016). The full details of each study have been omitted for brevity, though a summary of the most relevant examples is given below, and a comparison of these studies in relation to the present study is shown in Table 1.

The first study investigating the movement of diagnostic specimens was seen in McDonald (1972), who sought to use heuristic methods, including a Gaskell modified saving algorithm (Gaskell, 1967), which builds on the classical VRP approach, the Clarke and Wright Saving's algorithm (Clarke and Wright, 1964). Anaya-Arenas et al. (2016) built on this work to propose the Biomedical Sample Transportation Problem (BSTP), minimising the total distance travelled whilst subject to time windows for collection/delivery, highlighting the importance of reducing the transit time (keeping it lower to 180 minutes). This fact was previously introduced in Wilson (1996), who suggested transport should seek to limit sample damage with punctual delivery and controlled intermediate storage. Zabinsky et al. (2020) proposed an exact method to solve a vehicle routing and scheduling problem for specimen deliveries for specimen deliveries, with an aim to reduce the delivery time across all sites after an initial time point that samples are ready for collection. However, the branch and bound exact method is able to optimally solve small problems (e.g. 9 sites).

The initial problem and investigations proposed in (Oakey et al., 2023) used some of the concepts discussed by McDonald (1972) and Anaya-Arenas et al. (2016), noting that existing operations have limited control over the movement of samples before they are collected by the courier and delivered at the analysis laboratory. This study focused on the UK's National Health Service (NHS), which has ambitions to become net-zero by 2040 (NHS, 2020a); thus, the objectives and arrangement sought to reduce vehicle driving in place of cycle courier use.

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

Author	Problem Title	Key Dynamics	Objective(s)	Formulation, Approach	Est. Runtime*
McDonald (1972)	VRP Case Study - Specimen Collection	Vans Only, Time Constrained	Total Travel Time	Single Obj., Heuristic	Not Stated
Anaya-Arenas et al. (2016)	Biomedical Sample Transportation Problem (BSTP)	Vans Only, Several Time Constraints, Uncapacitated	Travelled Distance	Single Obj., Heuristic	1 hr (limit)
Zabinsky et al. (2020)	Vehicle Routing and Scheduling Gen. Alg. (VeRSA)	Vans Only, Time Constrained, Vehicles Re-Used	Total Duration from Goods Ready to Delivery	Single Obj., Branch and Bound	2 hr (limit)
Oakey et al. (2023)	*Sustainable Specimen Collection Problem (SSCP)*	*Vans and gig-economy cycles, Time Constrained, Multi-Echelon*	*Longest Collection Round Duration, Number of Vehs., Total Travel Time*	*Weighted MOO, Column Gen. with Imp. Heuristics*	*c. 30 secs*
Present Study	*SSCP with work shifts and UAVs*	*Vans, UAVs, and gig-economy cycles, Time-constrained, Multi-echelon*	*Maximum transit time, operating cost, emissions*	*Weighted MOO., Adapted CWSA heuristic with bin-packing*	*Typical <2 mins, Large <~10-20 mins*

Other notable examples of healthcare logistics problems have been discussed by Kergosien et al. (2014); Lodree et al. (2016); Doerner and Hartl (2008); Liu et al. (2013); Osaba et al. (2019); Ouiss et al. (2022), addressing related problems such as the rapid movement of blood stocks in humanitarian disasters, and green vehicle routing for at-home patient care. The objectives in these problems typically related to combinations of cost, time, and emissions. With respect to solution approaches, heuristics were generally most prevalent due to the significant solution spaces and constraints applied to the problems, and the need to solve within short periods prior to implementation.

The scheduling aspect considered in this work could be deemed closely related to a crew scheduling problem, due to drivers and operators being allocated vehicle(s) for an entire shift period. It is worth noting that the number of operators required in any given shift is a key contributor to the total cost. The key variation to this concept is seen in UAV operators managing multiple platforms at once, however, they remain on duty for the entire period. In the literature, it has been noted that such crew scheduling problems can be effectively solved using bin-packing algorithms (Qiao et al., 2010).

The SSCP with work shifts and UAVs is further complicated with variable costs for each route, depending on the assigned position in a shift (due to traffic, etc.). Such behaviour could be likened to a manufacturing makespan problem, which features varying task durations, depending on their position in the manufacturing process (Yu et al., 2013). To this end, bin-packing has been shown to be an effective tool in solving makespan problems (Van De Vel and Shijie, 1991; Coffman et al., 1987); thus,is explored in the present study.

The remainder of this paper outlines the new extension of the SSCP: the "SSCP with work shifts and UAVs", building on Oakey et. al's initial study (Oakey et al., 2023), presenting a heterogeneous problem with three different modes in a two-echelon arrangement. Subsequently, a novel heuristic approach is proposed to solve the problem, using a novel adaptation of the CWSA for route planning with bin-packing to compile effective work shifts. A new test instance generator is also introduced to construct realistic dummy datasets, before test results are presented, using the instances to parameterise and benchmark the algorithm.

3. A Two-Echelon Heterogeneous Vehicle Routing Problem

3.1. Problem Description

This problem further develops the SSCP, where a set of known node locations (clinics) produce diagnostic specimens that require collection and delivery to a single central hospital laboratory node without incurring excessive costs, environmental impacts, or time.

Specimens are routinely used in health services throughout the world to support in the diagnosis of patient medical conditions, with 29% of general practitioner appointments requiring some form of diagnostic test (Ngo et al., 2017). The current method of transport used for these deliveries varies by country. For example, in developed countries such as the UK, fleets of diesel vans are currently used and managed by local healthcare organisational units (known as National Health Service (NHS) trusts). Meanwhile, in some less developed countries such as Rwanda, UAV delivery services are being adopted in select locations in addition to traditional ground transport modes (Sigari and Biberthaler, 2021), and there is interest in adopting similar technologies in more developed nations as well (Cawthorne and Wynsberghe, 2019). This study has been developed to reflect the typical operations in the UK, though the principle arrangement could be adapted for other locations.

As in the initial investigations (Oakey et al., 2023), vans are based at the hospital laboratory and serve multiple clinics in a route before returning to the hospital within a maximum time constraint. Additionally, vans may also be based at the community clinics if facilities and contractual arrangements permit. Bikes operate on a similar basis but can start at any clinic or the hospital and serve only a limited number of sites due to capacity constraints, before delivering their payload back to the same site. Meanwhile, UAVs can only serve one site before returning to the hospital due to payload capacity constraints.

It should be noted that some studies in the literature assume that UAVs can serve multiple sites in one route, e.g. (Wang and Sheu, 2019; Benarbia and Kyamakya, 2022), though this is unlikely given the very limited payloads that UAVs can carry, particularly if industry standard packaging units are used, as highlighted in (Anonymous, 2022). Similarly, other studies have explored the use of UAVs flying from dynamically located positions along van routes (Murray and Chu, 2015), though this is also unlikely due to the regulatory and safety restrictions for such flights to take place.

As in the original problem (Oakey et al., 2023), this study continues to assume that cycling can be undertaken as discrete tasks by gig-economy cyclists outside of mealtime peaks (Lord et al., 2020). Conversely, progressing from the original SSCP study (Oakey et al., 2023), the re-use of vehicles is permitted through introducing work shifts. In business-as-usual NHS operations in the UK, vehicles and their drivers typically operate a morning shift and an afternoon shift, collecting from a given set of sites in each period, with periods being independent of each other. To this end, this study defines shift periods, in which multiple routes (e.g. Hospital→Site A→Site B→Site C→Hospital) can be served by a single vehicle and driver (vans)/operator (UAVs), such that multiple deliveries are made to the hospital in a given period. Due to cycle routes being handled by gig-economy cyclists on an ad-hoc basis, these remain as standalone routes.

Cycling and van routes can be used to complete consolidation deliveries on a local basis, whereby sites are served, and their loads are delivered to another clinic for onward transportation by a (trunk) UAV or van route. Within a shift period, traffic and weather conditions may vary, meaning that travel times can be departure time dependent for vans and UAVs. Travelled trajectories may also vary temporally due to changes in traffic and the associated third-party ground risk and weather conditions (Pilko et al., 2021); thus, distances and associated travel energies are also departure time dependent.

Samples often have a short life-span and should be analysed in a promptly (as stipulated by the regulator) to guarantee quality diagnostic results are obtained (NHS and Sedman, 2020); however, there are sometimes difficulties experienced in realising this (Oakey et al., 2023) and (Anonymous, 2022). As has been noted by Anaya-Arenas et al. (2016) and McDonald (1972), there are many stages in the sample supply chain, though the period in which planners can have greatest influence is in the transportation between clinics and the laboratory. To this end, the proposed problem focuses on the time samples spend in transit (i.e. out of controlled conditions), accounting for transit times as part of the objective and as a maximum time constraint that planners may wish to use.

Furthermore, it is envisaged that logistics planners may want to plan the routes for a given shift with minimal notice to ensure that the strategies match the expected demand for the period in question and maximise the efficiency of the routing. To this end, any approach used for solving the VRP must be capable of identifying effective solutions quickly (i.e. within a few minutes or less).

When considering the combination of consolidation and trunk routes (hereafter referred to as a collection round, Figure 1), the maximum time constraint may be imposed so that there is a limit on the duration between the first collection in a given route combination and the delivery at the hospital. This is similar to the collection round limit in the original investigations, although with an update to consider the transit time only, as opposed to the time to final delivery at the hospital across all rounds, due to the introduction of shift periods and vehicle reuse, which make the original metric less relevant. Similarly, the first objective, relating the problem to patient care levels, seeks to minimise the maximum in transit time. This aligns with other healthcare logistics problems, such as ambulances response times, where metrics seek to ensure operations provide a consistent minimum level of service (NHS, 2018).

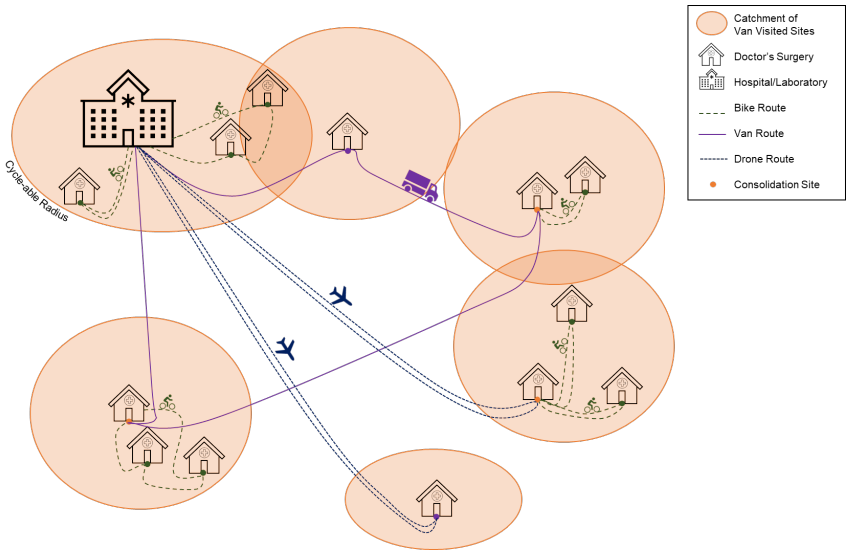

Figure 1: Routing options envisaged in the proposed multi-modal healthcare logistics system.

Where cycling continues to be handled by a third-party gig-economy logistics provider and remains as discrete tasks, a significant level of uncertainty around these routes remained. To this end, collection rounds featuring cyclists account for the full duration from the cycle courier's departure (as opposed to the first collection). Conversely, where route combinations that incurred waiting time were rarely selected in the initial study and introduces delivery/collection time synchronisation risks, consolidation waiting time is not permitted in the developed problem.

It should be noted that using historic sample 'bleed' and 'receipt' time data (i.e. when they were extracted from patients, and when they were checked in to the laboratory, as presented in (Oakey et al., 2023)) is not significantly beneficial, as NHS staff have suggested that appointment times are scheduled around the logistics timetable in the existing system, as opposed to the other way around. Hence, it could be said that there is some flexibility to move appointments to suit vehicle scheduling. Additionally, under the proposed arrangement, shorter work shifts could be adopted to enable a more on-demand style arrangement, as was suggested by (Oakey et al., 2023), whereby only those sites that require a collection are served in a given period.

With respect to the problem's environmental emissions objective, the proposed system should reduce emissions where possible, and the strategic arrangement using consolidation to reduce mileage and the number of vehicles required intrinsically supports this goal. Additionally, any

emissions produced generating the power for electrically-propelled vehicles also need to be accounted for. To this end, a move away from fossil-fuels towards electric-powered vans has been noted to be a committed change in last-mile logistics systems (European Environment Agency, 2020), meaning that further constraints may need to be imposed with respect to vehicle range in some routing problems, as in Green Vehicle Routing Problems (VRP) (Demir et al., 2014). In this particular application, it is anticipated that a vehicle will have sufficient range to complete all of its allocated routes for a given shift before recharging; hence, no overall shift period range limit has been imposed.

Typically, the energy requirements of a given route correlate with the subsequent emissions produced (Krol et al., 2023); hence, the environmental objective of the problem seeks to minimise a function of the energy of the system. This also enables a potential adaptation of the function to normalise the environmental objective using pricing for the impact of emissions, as is often done in transport appraisals (Department for Transport, 2021).

Cost is the final objective of the system, which includes the fixed costs of operations for a given shift period, such as a driver wage for a fixed number of hours, and any standing costs associated with the vehicles, such as insurance. The cost objective also includes the items which vary relative to the length and duration of routes. Van costs typically vary per mile driven (FTA, 2022), whilst the variable costs of UAVs relate to their scheduled maintenance programme based on flight hours (UK CAA, 2020) and (Anonymous, 2022).

UAV staffing is also subject to an operator ratio as a result of multiple UAVs being overseen by a single member of staff (e.g. 1 operator may oversee 20 UAVs (Crosby, 2023)), meaning economies of scale can be possible. Nonetheless, to guarantee the pay of an operator for a shift period, operation costs follow step changes to meet staffing requirements; i.e. 1 UAV = 1 operator, 20 UAVs = 1 operator, 21 UAVs = 2 operators (1:20 ratio).

With the continued assumption of ad-hoc cycling, cycle costs are based on a function of distance, number of stops, and a fixed-fee per route (Stuart, 2023). It should be noted that this is a development on the previous investigations (which were based on a flat rate per route), due to updates to gig-economy company pay structures and surrounding regulations (Lomas, 2023; The Fairwork Project, 2020).

3.2. Integer Linear Programming (ILP) model

As outlined in the problem description, this CVRP defines the selection of a set of vehicle work shifts, each comprising of a set of routes, to collect samples from a known set of nodes before delivering them to a single node for analysis. The following definition can be described as the SSCP with work shifts and UAVs. The formulation is as follows:

Table 2: ILP Notation

Notation	Description	
Timing		
K	Set of discrete time points, $K = \{0, 1, 2, \ldots, \overline{k}_n\}$	(1)
$\overline{k}_n$	Number of time points in K	(2)
a	Number of time intervals in a shift period	(3)
k	Given time point in K	(4)
I_u	Time interval bound by time point k_{u-1} and k_u, containing time point k	(6)
Locations and Modes		
s	Individual surgery (indexed using i and j)	(7)
S	Set of surgeries that require a collection	(8)
H	Target hospital, where vans/UAVs are based. Destination for all samples.	(9)
S'	Set of all nodes, including all surgeries and the hospital	(10)
S^E	Subset of surgeries where vans can also be based	(11)
S^D	Subset of surgeries which can be served by UAVs, $S^D \in S'$	(12)
S^C	Subset of surgeries which can be served by cycles, $S^D \in S'$	(13)
V	Van mode	(14)
D	UAV (drone) mode	(15)
C	Cycle mode	(16)
Costs Parameters		
p^V	Commercial running costs of a van per unit distance (e.g. per kilometre)	(17)
W^V	Standing cost of a van and driver per shift period	(18)
p^D	Commercial running costs of a UAV per unit time (e.g. per flight-hour)	(19)
W^D	Standing cost of a UAV platform per shift period	(20)
W^O	Standing cost of a UAV operator per shift period	(21)
ϕ	Number of UAVs each operator can simultaneously monitor	(22)
p^J	Cost per cycle job	(23)
p^S	Cost per cycle stop after the first pickup	(24)
C^l	Included distance per cycle job	(25)
p^C	Cost per unit distance beyond C^l	(26)
Van Travel Parameters		
$T^V_{k_{u-1}}(i,j)$	Fixed travel time for a van journey between sites i, j, departing at time k_{u-1}	(27)
$T^V_{k_u}(i,j)$	Fixed travel time for a van journey between sites i, j, departing at time k_u	(28)
$t^V_k(i,j)$	Travel time for a van journey between sites i, j, departing at time k, scaled using the fixed travel times for the associated intervals	(29)
$L^V_{k_{u-1}}(i,j)$	Fixed travel distance for a van journey between sites i, j, departing at time k_{u-1}	(30)
$L^V_{k_u}(i,j)$	Fixed travel time for a van journey between sites i, j, departing at time k_u	(31)
$l^V_k(i,j)$	Travel distance for a van journey between sites i, j, departing at time k, scaled using the fixed travel distances for the associated intervals	(32)
$\mathcal{E}^V_{k_{u-1}}(i,j)$	Fixed travel energy for a van journey between sites i, j, departing at time k_{u-1}	(33)
$\mathcal{E}^V_{k_u}(i,j)$	Fixed travel energy for a van journey between sites i, j, departing at time k_u	(34)
$\epsilon^V_k(i,j)$	Travel distance for a van energy between sites i, j, departing at time k, scaled using the fixed travel distances for the associated intervals	(35)

UAV Travel Parameters		
$t^D_k(i,j)$	Travel time for a UAV journey between sites i, j, departing at time k, scaled using the fixed travel times for the associated intervals ($T^D_{k_{u-1}}(i,j)$ and $T^D_{k_u}(i,j)$, analogous to vans, above)	(36)
$l^D_k(i,j)$	Travel distance for a UAV journey between sites i, j, departing at time k, scaled using the fixed travel distances for the associated intervals ($L^D_{k_{u-1}}(i,j)$ and $L^D_{k_u}(i,j)$, analogous to vans, above)	(37)
$\epsilon^D_k(i,j)$	Travel distance for a UAV energy between sites i, j, departing at time k, scaled using the fixed travel distances for the associated intervals ($\mathcal{E}^D_{k_{u-1}}(i,j)$ and $\mathcal{E}^D_{k_u}(i,j)$, analogous to vans, above)	(38)
Cycle Travel Parameters		
$t^C(i,j)$	Travel time for a cycle journey between sites i, j (constant across time periods)	(39)
$l^C(i,j)$	Travel distance for a cycle journey between sites i, j (constant across time periods)	(40)
Routes		
R^V	Set of van routes based at the hospital, H	(41)
R^D	Set of UAV routes based at the hospital, H	(42)
R^C	Set of cycle routes based at any site	(43)
R^E	Set of consolidation van routes based at sites in S^E	(44)
n_r	Number of surgeries visited in a route	(45)
T	Service/dwell time per stop in a route	(46)
$r_{v,k}$	A trunk van route departing at time k	(47)
$t_{r_{v,k}}$	Total duration of route $r_{v,k}$	(48)
$l_{r_{v,k}}$	Total distance travelled by route $r_{v,k}$	(49)
$\epsilon_{r_{v,k}}$	Total energy requirement of route $r_{v,k}$	(50)
$p_{r_{v,k}}$	Total running cost of route $r_{v,k}$, excluding fixed costs	(51)
$r^s_{e,k}$	A consolidation van route based at surgery s, departing at time k, with parameters $t_{r^s_{e,k}}$, $l_{r^s_{e,k}}$, $\epsilon_{r^s_{e,k}}$, and $p_{r^s_{e,k}}$ (analogous to trunk vans)	(52)
$r_{d,k}$	A UAV route, departing at time k, with parameters $t_{r_{d,k}}$, $l_{r_{d,k}}$, $\epsilon_{r_{d,k}}$, and $p_{r_{d,k}}$ (analogous to vans)	(53)
T^D	Downtime duration per UAV route (for battery changes, etc.)	(54)
$r^s_{c,k}$	A cycle route based at surgery s, departing at time k, with parameters $t_{r^s_{c,k}}$, $l_{r^s_{c,k}}$, $\epsilon_{r^s_{c,k}}$, and $p_{r^s_{c,k}}$ (analogous to trunk vans)	(55)
$t_{c_{max}}$	Maximum cycle route duration	(56)
$l_{c_{max}}$	Maximum cycle route distance	(57)
Collection Rounds		
R'	Complete set of all trunk routes	(58)
R''	Complete set of all consolidation routes	(59)
$r_{\alpha,k}$	Individual trunk route, departing at time k	(60)
$r^s_{\beta,\delta}$	Individual consolidation route, based at surgery s, departing at time δ ($\delta \in K$)	(61)
$\overline{r}_k$	Individual collection round, with departure time k (corresponds to collection round's trunk route departure)	(62)
$R''_{r_{\alpha,k}}$	Subset of consolidation routes associated with trunk route $r_{\alpha,k}$	(63)
$\overline{R}$	Set of all collection rounds	(64)
$S_{\overline{r}_k}$	Set of surgeries served by all of the constituent routes of $\overline{r}_k$	(65)

$r^0_{v,k}$	Dummy van route serving only H to enable collection rounds containing only consolidation routes	(66)
$\overline{r}^0_k$	Collection round containing $r^0_{v,k}$ and consolidation cycle routes based at H	
$\epsilon_{\overline{r}_k}$	Total energy requirement of consolidation round $\overline{r}_k$	(67)
$t_{\overline{r}_k}$	Transit time in collection round $\overline{r_k}$ (duration between the first collection and the delivery)	(68)
Variables, Constraints and Objective Parameters		
$x_{\overline{r}_k}$	Binary decision variable to select collection rounds	(69)
$V_{\overline{r}_k}$	Number of van routes (trunk and consolidation) being used in collection round $\overline{r_k}$	(70)
$D_{\overline{r}_k}$	Number of UAV routes being used in collection round $\overline{r_k}$	(71)
A^V_{max}	Maximum number of vans (trunk and consolidation) being used across all collection rounds and time points	(72)
A^D_{max}	Maximum number of UAVs being used across all collection rounds and time points	(73)
$t^{max}_{\overline{r}_k}$	Maximum transit time constraint value	(74)
u	Maximum transit time across all selected collection rounds	(75)
γ	Constant multiplier for converting energy to emissions or priced emissions	(76)
θ_1	Van cost objective function weight	(77)
θ_2	UAV cost objective function weight	(78)
θ_3	Cycle cost objective function weight	(79)
θ_4	Emissions objective function weight	(80)
θ_5	Maximum transit time objective function weight	(81)

Similar to the business-as-usual state, a day is split into multiple shift periods which do not overlap (e.g. morning shift, afternoon shift, etc). To this end, the time range of a given shift period is discretised and partitioned into intervals with approximately equal traffic and weather conditions (Table 2: 1-6). It is worth noting that $\overline{k}_n$ can be large, and for any given time point $k \in K$ there is a unique $u \in \{1, \ldots, a\}$ in such a way that set I_u contains k, i,e., $k_{u-1} \leq k \leq k_u$.

Figure 2 visualises this concept, demonstrating the principles using an example, the time point $k = 3$ falls between k_1 and k_2, hence, $k_{u-1} = k_1$ and $k_u = k_2$.

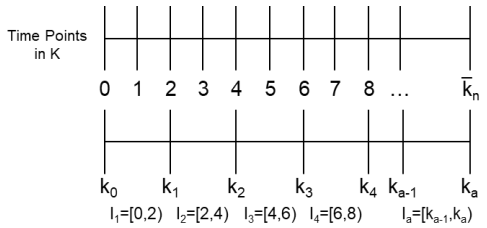

Figure 2: Time point notation visualised

Table 2: 7-13 define the surgeries that require service and the various sets of surgeries relating to the suitabilities of the three modes defined in Table 2: 14-16. It should be noted that sites in $S^E \subseteq S$ must only be used for purpose of consolidation back to the surgery they are based at for contractual and space reasons and all vans must finish at the site from which they are based. For practical reasons (e.g. landing space, staff resource, contracts), UAVs and cycles may be restricted

such that they can only serve a select subset of surgeries. For example, S^C may contain only those sites within large urban areas due to the service areas of gig-economy courier companies

The respective operating characteristics of each mode are defined by Table 2: 17-40, with the van travel times/distances/energies being approximated as a function of the travel times/distances/energies within the travel condition periods that the journey spans. It is assumed that no journey between an OD pair spans more than two travel condition periods, hence, $T^V_{k_{u-1}}(i,j)$ is constant and no period beyond k_u need be considered for each pair departing at time k.

$$t^V_k(i,j) = \min\left(1, \frac{k_u - k}{T^V_{k_{u-1}}(i,j)}\right) \times T^V_{k_{u-1}}(i,j) + \max\left(0, 1 - \frac{k_u - k}{T^V_{k_{u-1}}(i,j)}\right) \times T^V_{k_u}(i,j)$$

$$l^V_k(i,j) = \min\left(1, \frac{k_u - k}{T^V_{k_{u-1}}(i,j)}\right) \times L^V_{k_{u-1}}(i,j) + \max\left(0, 1 - \frac{k_u - k}{T^V_{k_{u-1}}(i,j)}\right) \times L^V_{k_u}(i,j)$$

$$\epsilon^V_k(i,j) = \min\left(1, \frac{k_u - k}{T^V_{k_{u-1}}(i,j)}\right) \times \mathcal{E}^V_{k_{u-1}}(i,j) + \max\left(0, 1 - \frac{k_u - k}{T^V_{k_{u-1}}(i,j)}\right) \times \mathcal{E}^V_{k_u}(i,j)$$

This follows analogously for UAVs, accounting for potential variations in weather and ground risk conditions (as assessed by external pathfinding problems, e.g. (Pilko et al., 2021)). Additionally, it should be noted that the First-In-First-Out (FIFO) constraint of vehicle routing is satisfied by using proportional shares of each time period's travel durations in conjunction with the min-max terms. Meanwhile, cycle travel durations are deemed to be independent of departure time (i.e. constant and unaffected by traffic); thus, the travel duration between a pair of surgeries (i,j) is given as $t^C(i,j) \ \forall k \in K$ and travel distance is given as $l^C(i,j) \ \forall k \in K$. The energy requirement for cyclists is assumed to be zero/negligible for all OD pairs.

It should be highlighted that a per unit time cost is used for UAVs based on the expected part lifespans of the platform being the main driver of the variable costs, as outlined by Anonymous (2022). Additionally, any fixed vehicle fixed costs such as insurance, maintenance, etc. are included in W^V and W^D, whilst the labour cost for the shift period is included in W^V and W^O.

The characteristics of individual routes are listed in Table 2: 41-57, where a trunk van route is defined as $r_{v,k} = (H, s_1, \ldots, s_{n_r}, H) \in R^V$. UAV routes are defined analogously as $r_{d,k} \in R^D$, with the additional constraint that they can only serve one surgery per collection; thus, $n_r = 2 \ \forall r_{d,k} \in R^D$ due to regulation and payload capacity. Similarly, consolidation van and cycle routes are defined by $r^s_{e,k} \in R^E$ and $r^s_{c,k} \in R^C$, where s indicates the surgery where the route is based. For cycling, $n_r \leq 4$ due to capacity constraints (maximum load equal to three surgeries' worth of samples). Additionally, cycles can be based at the hospital if $H \in S^C$, and all routes can only serve those sites where the respective mode is permitted, and vehicles must return to their start location.

Each trunk van route has an associated time, denoted by $t_{r_{v,k}}$ and is calculated by summing the durations of the constituent legs of the route with their individual departure times, where s_i denotes the departure surgery and t^i denotes the arrival time of the i-th leg, with T being embedded

in $t_{r_{v,k}}$:

$$
\begin{aligned}
t^1 &= t^V_k(H, s_1) + k \\
t^i &= t^{i-1} + t^V_{t^{i-1}}(s_{i-1}, s_i) \quad i = 2, \ldots, n_r \\
t^{n_r+1} &= t^{n_r} + t^V_{t^{n_r}}(s_{n_r}, H) \\
t_{r_{v,k}} &= t^{n_r+1} - k + n_r T
\end{aligned}
$$

Each trunk van route also has an associated distance and energy requirement, denoted by l_{r_v} and $\epsilon_{r_{v,k}}$. Where the departure times depend on the previous legs, the totals are also calculated by summing the distances/energies of the constituent legs using the timings calculated above:

$$
l_{r_{v,k}} = l^V_k(H, s_1) + \sum_{i=1}^{n_r-1} l^V_{t^i}(s_i, s_{i+1}) + l^V_{t^{n_r}}(s_{n_r}, H)
$$

$$
\epsilon_{r_{v,k}} = \epsilon^V_k(H, s_1) + \sum_{i=1}^{n_r-1} \epsilon^V_{t^i}(s_i, s_{i+1}) + \epsilon^V_{t^{n_r}}(s_{n_r}, H)
$$

Furthermore, the non-constant (i.e. varying with the length of the route) running cost of a given van route is calculated as $p_{r_{v,k}} = p^V l_{r_{v,k}}$.

Route times, distances, energies (except cycling, where energy is assumed to be zero), and costs (vans only) for the other modes are calculated analogously to trunk vans. For UAVs, the non-constant (i.e. varying by the duration of the route) running cost of a given UAV route is calculated as $p_{r_{d,k}} = p^D t_{r_{d,k}}$, whilst cycling route costs are calculated as a function of the route's stops and distance, $p_{r_{c,k}} = p^J + p^S(n_r - 1) + p^C(l_{r_{c,k}} - C^l)$, where each route's costs increases after one pickup (hence, $n_r - 1$) and for distances beyond C^l.

Building on the original SSCP, the departure times and durations of the route's constituent legs are derived from the departure time of the route such that no waiting time is incurred. Waiting time for any reason (e.g. arriving before consolidation rounds have finished, or waiting for traffic changes) is not permitted. This was changed from the original formulation to provide a greater guarantee with respect to time synchronisation for transshipment.

On arrival back at the hospital, UAVs are also subject to a downtime duration, T^D, in addition to the embedded service time, to allow for battery changes/airworthiness checks to take place prior to their next departure. This means that a UAV cannot be used for T^D minutes on return to the hospital.

Additionally, cycle routes may be subject to a maximum time constraint of $t_{c_{max}}$ and maximum distance constraint $l_{c_{max}}$ to ensure the cycle routes can be managed as discrete gig-economy tasks; $t_{r^s_{c,k}} \leq t_{c_{max}}$ and $l_{r^s_{c,k}} \leq l_{c_{max}}$. The discrete task arrangement of these routes ensures that typical gig-economy cycle operator costs can be accounted for, and cyclist availability is more likely.

Surgeries can be served by:

- Trunk van route or UAV route directly to the hospital; or
- Cycle route directly to the hospital (if the route is based at the hospital); or
- A consolidation van route to a surgery, and then onward by trunk van or UAV; or
- A cycle route to a surgery, and then onward by trunk van or UAV.

Hence, collection rounds are outlined in Table 2: 58-68, where two sets of routes, $R' = \{R^V, R^D\}$ which contains all of the trunk routes (UAVs and vans originating at H); and $R'' = \{R^C, R^E\}$, which contains all of the consolidation routes (bikes and vans). It should be noted that the formulations associated with $r_{\alpha,k}$ and $r^s_{\beta,\delta}$ are analogous to the original route definitions, with k indicating the departure time of the trunk route, δ denoting the departure time of the consolidation route, and s indicating the base of the consolidation route. α and β are indexes of each route type, to differentiate between individual trunk/consolidation routes, respectively.

A **collection round**, $\overline{r}_k = \{r_{\alpha,k}\} \cup R''_{r_{\alpha,k}}$, is defined as the combination of a single trunk route $r_{\alpha,k} \in R'$ with a subset of consolidation routes $R''_{r_{\alpha,k}} \subseteq R''$ such that for any given consolidation route $r_{\beta,\delta} \in R''_{r_{\alpha,k}}$ based at surgery s, it is satisfied that $s \in r_{\alpha,k}$, i.e. any consolidation route in $R''_{r_{\alpha,k}}$ is based at a surgery that is being visited by the trunk route $r_{\alpha,k}$.

Furthermore, other than surgery s where each $r^s_{\beta,\delta}$ begins/ends, the consolidation routes in $R''_{r_{\alpha,k}}$, $\beta \in \{1, \ldots, |R''_{r_{\alpha,k}}|\}$, do not share any other surgeries; i.e. no surgery is served by multiple consolidation routes within one shift period. It should be highlighted that the described problem relates to only one shift period (e.g. a morning), and surgeries may be served by different consolidation routes in separate work shifts.

Note that, since $R''_{r_{\alpha,k}}$ could be empty, then it can satisfied that $R' \subseteq \overline{R}$ and $|\overline{r}_k| = 1$. The set of surgeries served by all of the constituent routes of $\overline{r}_k$ is denoted by $S_{\overline{r}_k}$. Additionally, it should be noted that a maximum of 2 stages (echelons) can occur in a collection round; consolidation to a surgery, followed by trunking to the hospital. Chained consolidation (e.g. bike-bike-trunk van) is not permitted to prevent delays and limit the risks associated with chain of custody over potentially sensitive goods.

Bike routes are also permitted to start and end at the hospital, in order to serve the immediate catchment area of the hospital directly. To account for this in the problem formulation, a dummy van route, $v^0_k \in R^V$, is created. Starting and ending at the hospital, with no intermediate stops ($r^0_{v,k} = (H, H)$) and a travel time of zero ($t_{r^0_{v,k}} = 0$), $r^0_{v,k}$ enables a collection round where surgeries which are cycle served only, $\overline{r}^0_k$. These routes can depart at any time within the shift period so that they are completed by the end time.

The collection round's departure time, k, is given by the trunk route's departure time. Meanwhile, the consolidation routes have their own departure times, and for a consolidation round to be feasible, their departure must result in the consolidation routes in $R''_{r_{\alpha,k}}$ being complete before the trunk route, $r_{\alpha,k}$, serves the consolidation site, i.e. they do not incur waiting time for the trunk route. As previously noted, this builds on the initial modelling to ensure timely transshipment.

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

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

The integer parameter $V_{\overline{r}_k}$ is also defined for each collection round, denoting the number of van routes (trunk and consolidation) being used in a collection round. If there is no active route at time point k, $V_{\overline{r}_k} = 0$. Similarly, $D_{\overline{r}_k}$ defines the number of UAVs (trunk) being used in a collection round.

Vans and UAVs can be reused in the shift period, and in any solution, the number of vans required throughout the entire shift period will never exceed the maximum number in use across all time points in K (e.g. Table 3). This is particularly important when calculating the number of vehicles and drivers/operators that contribute to the cost, with each van being operated by a

single driver who is on-duty and paid for the entire shift period (regardless of how long they are driving routes). Similarly, each UAV is monitored by an operator who is also paid for the entire shift period, but can monitor up to ϕ UAVs simultaneously.

Table 3: Example of the number of vans used in a given solution. Shading indicates van in used. The maximum across all time points in 3.

Time Point	1	2	3	4	5	6	7	8	...	k
Van 1										
Van 2										
Van 3										
# In Use	1	2	**3**	**3**	2	1	1	2	2	2

Hence, the maximum number of vans and UAVs used throughout the entire shift period are given by the sum of the vans/UAVs used across all used collection rounds and time points:

$$A^V_{max} \geq \sum_{\overline{r}_k \in \overline{R}} V_{\overline{r}_{k'}} x_{\overline{r}_k} \;\; \forall \;\; k' \in K \qquad\qquad A^D_{max} \geq \sum_{\overline{r}_k \in \overline{R}} D_{\overline{r}_{k'}} x_{\overline{r}_k} \;\; \forall \;\; k' \in K$$

The number of UAV operators required for a shift period is defined by A^O_{max}; an integer variable constrained such that:

$$A^O_{max} \geq \frac{A^D_{max}}{\phi} \qquad\qquad A^O_{max} < \frac{A^D_{max}}{\phi} + 1$$

Due to the likely contractual arrangement of the drivers/operators, it is assumed that pay is guaranteed for the shift period, regardless of the workload each driver/operator has. Hence, the fixed costs of a solution are given by $W^V A^V_{max}$ and $W^D A^D_{max} + W^O A^O_{max}$ for vans and UAVs, respectively. Bikes are considered to be more ad-hoc in arrangement and can be carried out as standalone discrete tasks (i.e. are not compiled into work shifts) within a given time window.

The priced emissions from the collection system (e.g. greenhouse gas emissions, pollutants) are calculated by the product of the energies of the selected routes and a constant parameter, γ, relating the energies to emissions and their societal impact. In the UK, this can be captured by the energy and emissions pricing values provided by the government for transport appraisals (UK Government, 2022; Department for Transport, 2021).

As in (Oakey et al., 2023), this problem also takes a pseudo-objective of patient care/sample quality into account. In this case, the duration $t_{\overline{r}}$, relates to the sample transit times, i.e. the period from the first collection until delivery, as opposed to in the original formulation, which explored the collection round duration/time after first departure for deliveries to be complete. Cycle routes are an exception to this update, due to routes being served on a more ad-hoc basis with less certainty on timing; hence, the time from the cycle route departure until delivery at the hospital is used in this case. This change accounts for the varying departure times of different routes, meaning that the time to the final delivery from the original formulation was less relevant in the shift context. Subsequently, u is also updated to build on the original formulation, making it the maximum of the transit times across all routes (Constraint 2).

The objective of this problem is to minimise the weighted sum of the operating costs, emissions, and maximum in-transit time for a shift period. To this end, a generalised multi-term objective

function is used to sum the costs of the selected routes and standing costs for each mode, the emissions costs, and the maximum transit time (see Constraint 1). Constants, θ_1, θ_2, θ_3, θ_4, and θ_5 are introduced to allow a Pareto front of solutions to be found by varying the coefficients of each term (Oakey et al., 2023). Similarly, in the real-world application of the problem, the weights could also allow decision makers to easily shift the objective function towards a given aim. It should be noted that in the event that the emissions parameters are not known or transit times are less certain due to more flexible contractual arrangements, the weighting can be set to zero. Furthermore, a weighted objective enables simpler modelling that is more easily transferred to other case study areas.

The problem is constrained such that transit time durations are subject to time constraint of $t_{\overline{r}}^{max}$ minutes or less (e.g. 90 minutes maximum (McDonald, 1972)) to guarantee the timely delivery of samples (Constraint 3). A given collection round duration is calculated as the maximum time between the first collection in any constituent route in $\overline{r}_k$ and delivery to the hospital. If cycle consolidation routes are used in $\overline{r}_k$, the time is measured from the start of the cycle route to account for any uncertainty in the performance of 3PL logistics carriers. The problem assumes an unlimited number of vehicles and operators/drivers are available.

To ensure all sites are served at least once, a further constraint is also added (Constraint 4), whilst $x_{\overline{r}}$ must be binary (Constraint 5).

$$\begin{aligned}\min : \sum_{\overline{r}_k \in \overline{R}} \Bigg(x_{\overline{r}_k} \Bigg(& \sum_{r_{v,k} \in \overline{r}_k \cap R^V} \theta_1 p_{r_{v,k}} + \sum_{r_{e,k} \in \overline{r}_k \cap R^E} \theta_1 p_{r_{e,k}} \\ & + \sum_{r_{d,k} \in \overline{r}_k \cap R^D} \theta_2 p_{r_{d,k}} + \sum_{r_{c,k} \in \overline{r}_k \cap R^C} \theta_3 p_{r_{c,k}} + \theta_4 \epsilon_{\overline{r}_k} \gamma \Bigg) \Bigg) \\ & + \theta_1 W^V A_{max}^V \\ & + \theta_2 (W^D A_{max}^D + W^O A_{max}^O) \\ & + \theta_5 u\end{aligned} \tag{1}$$

$$u \geq t_{\overline{r}} x_{\overline{r}} \quad \forall\ \overline{r} \in \overline{R} \tag{2}$$

$$u \leq t_{\overline{r}}^{max} \quad \forall\ \overline{r} \in \overline{R} \tag{3}$$

$$\sum_{\overline{r}_k ; i \in S_{\overline{r}_k}} x_{\overline{r}_k} \geq 1 \quad \forall\ i \in S \tag{4}$$

$$x_{\overline{r}_k} \in \{0, 1\} \quad \forall\ \overline{r} \in \overline{R} \tag{5}$$

4. An Adapted Savings Algorithm with Bin-Packing Heuristics

With the introduction of temporal variations due to using shift periods, significantly more route options are introduced compared to standard vehicle routing problems. This presents difficulties when using column generation based algorithms, such as in (Oakey et al., 2023), due to the substantially increased memory and processing requirements for computationally solving the derived master problems. Thus, an alternative approach is required to ensure quality solutions are found.

To this end, this paper proposes a novel algorithm to solve the SSCP with work shifts and UAVs; an adaptation of the classical CWSA combined with an adapted first-fit decreasing (AFFD) or an adapted best-fit (ABF) bin-packing algorithm. The adapted CWSA was used in the route generation process to quickly construct effective collection solutions. To enable the fast compilation of multiple routes into reasonable vehicle/work shifts, accounting for vehicle re-use, a one-dimensional bin-packing algorithm was used.

To further enhance processing speeds and keep the computational memory requirements manageable, the upper bounds of the route durations were used to account for the worst-case of traffic conditions when creating the original routes in the CWSA. The bin-packing algorithms subsequently updated these durations to reflect the assigned position in the shift period. Hence, the packing algorithms were variations on classical approaches to account for the variable size of the packed items (Coffman et al., 1984).

After using the adapted CWSA and packing algorithms to construct an initial solution, local search heuristics are used to converge to a local optimum. Meta-heuristics are subsequently applied to allow the solution to escape local optima and explore more of the solution space, before the process is repeated until a final solution is reached. Details of the individual aspects of the algorithm and associated testing are discussed in the remainder of this section, whilst an overview of the full algorithm and its configurable variable parameters are detailed in Section 4.5. Full details of the computational experiments and their results are given in Section 5.

4.1. Route Generation Through Savings Algorithm

As previously outlined, the routes produced in the initial solution are generated using an adaptation of the well-established CWSA. The classical CWSA solution approach starts from a set of routes containing two arcs, such that one route serves each site. The CWSA then progresses and evaluates the objective function savings (typically time) through combining routes, accepting the greatest saving (greedy algorithm) until no further savings can be made within the constraints of the problem.

In the adaptation of the CWSA presented in this study (Algorithm 1), several saving options are tested at each iteration (Table 4), including: (i) the classical option of combining routes, provided they are the same mode and based at the same origin (Figure 3c, Algorithm 1: Line 4); (ii) substituting in a UAV route for a given surgery, removing it from its current route and replacing it with an out-and-back UAV route (Figure 3d, Algorithm 1: Line 6); and (iii) transferring a van/UAV served (trunk route) site to a new or existing van or bike consolidation route, depending on a given catchment area of the consolidation modes (Figure 3e, Algorithm 1: Line 8).

The CWSA adaptation presented by Özoğlu et al. (2019) has some similarities with this approach, with UAVs arcs substituting van arcs, however, their approach is only presented for single van routes, and it is not clear how their adaptation would handle larger, multi-vehicle delivery networks. Furthermore, the sub-tours created by the side-kick UAVs could be likened to the cycle consolidation routes in the present study, suggesting that the CWSA concept provides a good platform for sub-tour related problems.

A further saving option could also be considered, whereby individual sites are tested in all positions within other routes. The savings options could be considered individually (i.e. for a single surgery move), or compounded over multiple iterations over all routes (i.e. multiple moves in a single saving option). This functionality could also be used as a local search heuristic after the solution has converged (see Figure 6c and Algorithm 1: Line 11, detailed further in Section 4.3).

Table 4: Saving options tested in each iteration of the algorithm.

Action	Applies to	Explanation
Combine van routes	All van served nodes	Joins two routes by removing 2 van arcs and replacing with 1 van arc
Introduce bike	All cycleable nodes	Removes 2 trunk arcs and adds 2 bike arcs
Introduce consol. van	All van served nodes up to trunk distance to hospital	Removes 2 trunk arcs and adds 2 consolidation van arcs
Change to UAV	All UAV permitted nodes	Replace 2 trunk van arcs with 2 trunk UAV arcs

Combination options were considered for all routes of the same mode and start location, whilst UAV route substitution and inter-route moves were considered for all sites, and trunk to consolidator options addressed only trunk route (van/UAV) served sites. Furthermore, combination and transfer options were explored using a best insertion approach when adding sites to routes, with all routes being considered for insertion into each other once. e.g. combining route A and route B would be tested by inserting the sites from route B into route A only, not inserting route A into route B. The best insertion was calculated from a time saving perspective (i.e. the greatest time reduction saving option was selected).

Additionally, when considering the trunk to consolidation or inter-route options, if the site being investigated already had consolidation sites associated with it, the associated sites could either be reallocated to other routes (Figure 4), or the investigated site could be excluded from the potential options.

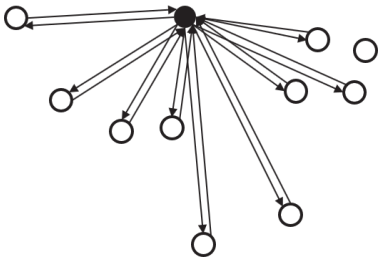

(a) Initial generation of van routes to all sites within range.

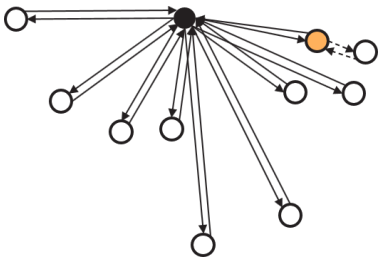

(b) Handling of unallocated sites (consolidate or use UAV).

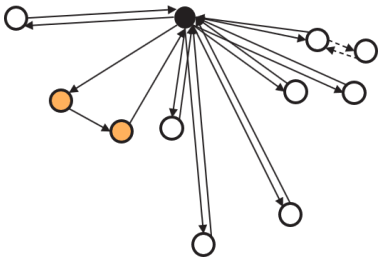

(c) Combine existing routes of the same mode and origin.

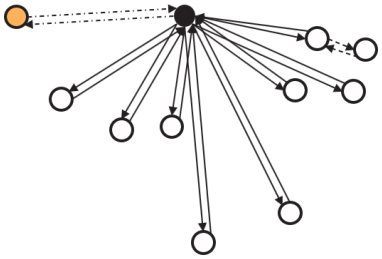

(d) Remove a surgery from a route and replace with a UAV route.

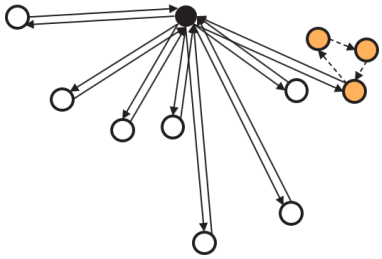

(e) Remove a surgery from a trunk route and add to/create a consolidation route.

Figure 3: Demonstration of the adapted CWSA algorithm. (A)-(B): initial construction, (C)-(E): savings options at each iteration. Solid lines = trunk van, dot-dashed lines = trunk UAV, dashed lines = consolidation.

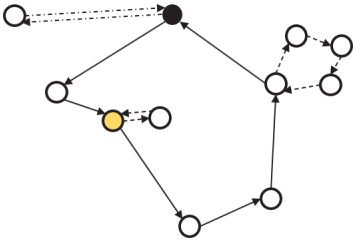

(a) Current solution, with site of interest highlighted.

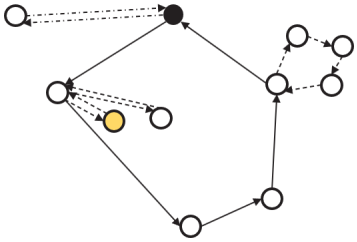

(b) Transfer to consolidation route and reallocate associated consolidator.

Figure 4: Demonstration of reallocation of consolidation routes. Trunk to consolidator saving option shown. Solid lines = trunk van, dot-dashed lines = trunk UAV, dashed lines = consolidation.

Algorithm 1 Adapted Clarke and Wright Savings Algorithm

Input: R: Existing set of routes; R_p: Cost of existing solution; L: Candidate shortlist size; $InterRoute_{CW}$: Inter-route strategy for the savings algorithm; C^{catch}: cycle catchment limit.

1: $improvement = true$
2: **while** $improvement$ **is** $true$ **do**
3: $options = \{\}$ ▷ Create empty list of candidate improvements
4: $options = options \cup combineRoutes(R, R_p, C^{catch})$
5: ▷ Add combination improvements to options list
6: $options = options \cup subInUAV(R, R_p)$
7: ▷ UAV sub-in improvements to options list
8: $options = options \cup trunkToConsol(R, R_p, C^{catch})$
9: ▷ trunk to consol. improvements to options list
10: **if** $InterRoute_{CW}$ **is** $true$ **then**
11: $options = options \cup interRoute(R, R_p, InterRoute_{CW}, C^{catch})$
12: ▷ Add inter-route improvement(s) to options list (multi-move gives one option)
13: **end if**
14: **if** $options \neq \{\}$ **then** ▷ If improvement options found
15: $shortList = \{\}$
16: Sort $options$ by the objective saving (vs. R_p)
17: $shortList \leftarrow options[0 : L]$ ▷ Copying first L elements of $options$ to $shortList$
18: R = Randomly selected option from $shortList$
19: $R_p = R_p - \Delta_p$ ▷ Update current cost
20: **else**
21: $improvement = false$
22: **end if**
23: **end while**
24: **return** R

The classical CWSA is greedy and always selects the saving option of greatest magnitude. To avoid this potentially ineffective behaviour in the adapted algorithm, an element of randomness is used, with options being ranked according to the magnitude of their saving, before an option is selected from a shortlist of the top n options (Algorithm 1, Line 17), in a similar manner to Pichpibul and Kawtummachai (2012). Similar approaches in VRPs have been noted to be highly effective and are well addressed in the literature (Hart and Shogan, 1987). Additionally, if a particular saving option offers significant benefits to the final solution, it will likely continue to feature in the options shortlist and have a greater chance of being selected.

The SSCP also presents additional time constraints that the traditional CWSA does not directly address. With respect to the initial generation, it may be the case that the initially generated van routes cannot serve all of the sites within the given time limits (Figure 3a). In this event, the proposed CWSA adaptation will attempt to reallocate these sites so that they are served by consolidation routes or by UAV routes, keeping the best objective solution after each site is resolved (Figure 3b). In the event a surgery cannot be allocated through these means, the problem is not feasible.

4.2. Shift Compiling Through bin-packing

Subsequent to the creation of effective routes, work shifts must be compiled to account for the number of vehicles and staffing required. In the adapted CWSA, work shifts are compiled each time a saving option is considered, such that the cost associated with each change accounts for any impacts resulting from vehicle assignments.

This study uses bin-packing to complete this task, with routes being the items to be packed, and the work shifts being the bins. The two approaches explored in this study are the adapted first-fit decreasing algorithm and an adapted best-fit algorithm, hereafter referred to as AFFD and ABF, respectively.

For context, the classical version of the First-Fit Decreasing algorithm, has been widely investigated and is generally accepted to be computationally efficient (Johnson, 1973; Dósa, 2007). The method sorts items by decreasing size, before iterating over the existing bins and testing in the next available position, and a similar approach has been adopted in this study. Meanwhile, the ABF presented in this study is understood to be a novel approach, and is not directly based on an existing method.

In the case of the AFFD, routes are sorted and subsequently selected using the upper bound of the route durations (i.e. the sum of the most pessimistic travel times within the shift period), before a revised version of the route with the actual timings at the given time point is assigned to the vehicle shift (bin). Owing to the prominence of this algorithm in literature and for brevity, a pseudocode has not been provided.

Weather and traffic conditions have been noted to vary and be a highly influential factor affecting the use of UAVs (Gao et al., 2021) and (Anonymous, 2022; Anonymous, 2023a). To this end, an alternative packing approach was developed with a view to leverage the periods of shortest travel time. The ABF algorithm opens a new bin (vehicle shift) and iterates over each applicable route at the next available time point (k) in the active bin, calculating the relative (percentage) difference between the upper bound route duration (as previously defined) and the route duration at the time point (k). The route with the greatest variability (i.e. the highest percentage difference) is then assigned; thus, minimising the space a variable item (route) occupies in each bin. If no route fits at the tested time point, a new bin is opened.

The variable scheduling elements have previously been explored with bin-packing algorithms in similar problems relating to makespan and crew scheduling (Qiao et al., 2010; Van De Vel and Shijie, 1991; Coffman et al., 1987), however, the combination of a savings algorithm with bin-packing for vehicle routing with UAVs is a novel approach. Given the adaptable and fast nature of these algorithms, good solutions should be reached within a usable timescale for logistics planners.

With respect to the computational complexity, the algorithm tests a maximum of n items in n positions, meaning it is of quadratic complexity; equal to that of the AFFD algorithm (n^2). The pseudocode of the proposed algorithm is given in Algorithm 2, and a comparison of how the fitting process would work in practice is seen in Figure 5. Furthermore, it should be noted that the classical 'Best-Fit' packing algorithm is not related to the proposed approach.

Algorithm 2 Adapted Best-Fit

Input: R: set of routes to be fitted (built using upper-bounds); M: vehicle mode; k_0: shift start; k_{max}: shift end; T^D: UAV downtime.

$n_r \leftarrow$ number of routes in R
$A \leftarrow (n_r, 1)$ ▷ Array of 1's to track unassigned routes
$b = \{\}$ ▷ Create new empty vehicle
$B = \{\}$ ▷ Create empty set of vehicles
$k = k_0$
while $\sum_{i \in A} i > 0$ **do** ▷ While unallocated routes exist
    $bestRoute = null$
    $bestVariability = 0.00001$ ▷ Set to very small value to ensure any route with 0 variability is still assigned.
    for $a \in n_r$ **do** ▷ Loop over route assign array
        if $A(a) = 0$ **then**
            continue
        end if
        $r = R(a)$ ▷ Isolate a route
        $\hat{r} = reTime(r, k)$ ▷ Create a re-timed variant
        if $k + t_{\hat{r}} > k_{max}$ **then** ▷ If route does not fit vehicle shift, skip
            continue
        end if
        $\Delta = \frac{\hat{r} \text{ duration}}{r \text{ duration}} - 1$ ▷ Calculate variance
        if $\Delta < bestVariability$ **then**
            $bestRoute = \hat{r}$
            $bestVariability = \Delta$
        end if
    end for
    if $bestRoute$ **is** $null$ **then**
        $B = B \cup b$ ▷ Add vehicle to list of vehicles
        $b = \{\}$ ▷ Create a new vehicle
        $k = k_0$
    else
        $b = b \cup \hat{r}$
        $A(a) = 0$
        if $M = D$ **then** ▷ Set next available time, add downtime if UAV
            $k = k + t_{\hat{r}} + T^D$
        else
            $k = k + t_{\hat{r}}$
        end if
    end if
end while
return B

As is discussed further in Section 5, both the ABF and AFFD algorithms were tested for the compilation of trunk van and UAV vehicle shifts (Section 5), and the number of UAV operators was derived from the number of UAVs, using the given operator ratio; e.g. 5:1 ratio: 1 UAV requires 1 operator, whilst 5 UAVs also requires 1 operator.

Meanwhile, the AFFD algorithm was always used for the compilation of consolidation routes (vans and bikes) due to bike travel times remaining constant (unaffected by traffic), and rural areas (where consolidation vans are likely to be based) having low traffic variability. Furthermore, unless they were limited by the start time of the shift, consolidation routes were timed such that they would maintain the maximum permitted in-transit time to reduce the likelihood of transshipment delays, which have been noted to present significant challenges to implementing multi-leg deliveries (Laseter et al., 2018; PharmaAero, 2022). Whilst this may limit the uptake of consolidation routes in some cases where transit times are deemed more important, it is a necessary assumption to ensure service reliability.

An additional variant of the AFFD algorithm (discussed above) was also used to help enhance the computational speed of the algorithm. In this process, previously investigated savings

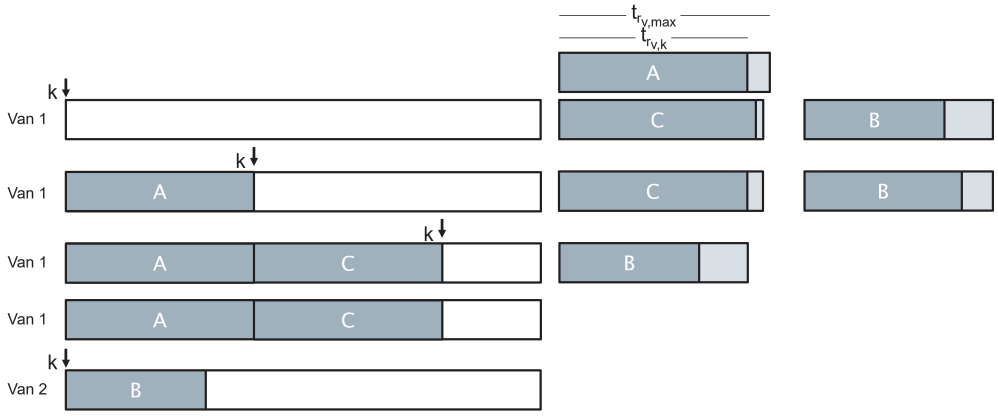

(a) Adapted First-Fit Decreasing Algorithm

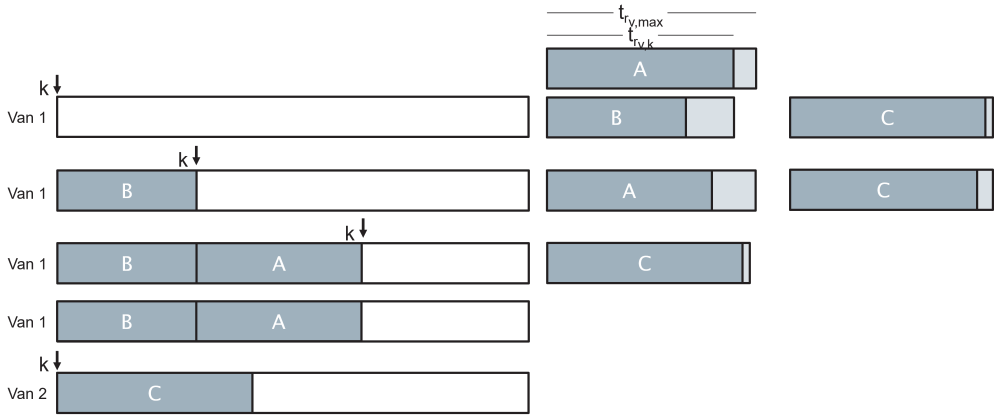

(b) Adapted Best-Fit Algorithm

Figure 5: Comparison of proposed bin-packing algorithms. Dark shading indicates duration of the van route at time point k ($t_{r_{v,k}}$), light shading indicates the upper bound of the van route duration ($t_{r_{v,max}}$) in the shift period.

options were stored for recalling in later iterations, removing the need to recalculate each leg's timings. To check recalled solution savings were still correct after previous iterations potentially re-positioned/removed other routes, the original AFFD approach was used to approximate the solution cost, but without revising and recalculating the allocated route to the assigned time (i.e. retaining the upper bound time). This enabled a rapid check to identify if the previously computed saving was of approximately equal magnitude, and prevented solutions with potentially unrealistic/infeasible costs being carried forward.

4.3. Local Search Heuristics

To resolve any inefficiencies potentially introduced by the savings algorithm, such as routes connecting non-adjacent nodes (i.e. crossing-over paths), local searches are applied to the converged constructed solutions.

A 2-opt local search is applied to all routes longer than 2 stops across all variations of the algorithm. This results in a negligible computational speed decrease, though consistently improves solutions if small inefficiencies arise. Where the 2-opt algorithm evaluates a change to an indi-

vidual route, an approximation of the overall effect on the objective function is required for quick computation of improvements. To this end, a pseudo-objective function is used. For example, in the case of vans:

$$Estimate = \theta_1 p_{r_{v,max}} + \theta_4 \gamma \epsilon_{r_{v,max}}$$

where $r_{v,max}$ indicates the upper bound of the route with respect to durations (i.e. the maximum of duration of all of the included OD pairs in the route in the shift period).

The objective impact estimate is similar for UAV and bike routes, with the formulations referring to the comparable equivalents for those modes. It is assumed that the changes introduced by the 2-opt checks are sufficiently small that required number of vehicles does not change, meaning there is no need for re-compiling work shifts and calculating the complete objective function. Similarly, it is assumed there is a negligible impact on the maximum transit time duration across all routes.

As previously detailed in Section 4.1, the inter-route search algorithm can be applied as either a savings option on a single or multi-move basis, or as an additional local search after the constructed solution has converged, or a combination of both. When applied as a local search tool, the algorithm is applied with multiple moves being compounded in a single change (i.e. looped over all surgeries, cumulative improvement), after to any 2-opt changes have been made. The local searches are visualised in Figure 6, with the improvements being applied to an example solution.

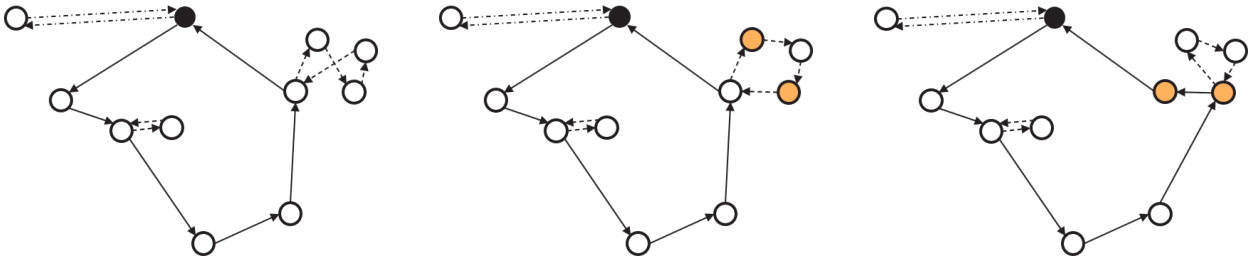

(a) Example converged solution, prior to local search.

(b) 2-opt search, testing within each route.

(c) Inter-route search, testing sites in positions in all routes.

Figure 6: Demonstration of the local search heuristics. Examples shown demonstrate a consecutive improvement from (A) through to (C). Solid lines = trunk van routes, dot-dashed lines = trunk UAV routes, dashed lines = consolidation routes.

4.4. Reset Kick Meta-Heuristic

To avoid cycling around the same locally optimal solution, a meta-heuristic is also introduced after the local search algorithms have been applied, partially destroying the solution in search of a global optima. The proposed meta-heuristic, referred to as the 'reset-kick', takes a given set of surgeries and removes them from their respective routes, before restoring them to new van or UAV out-and-back routes (Figure 7).

In the event that the kick resets routes to UAV service but a site is not permitted to be served by this mode, it is reset to a van route. This approach is similar to the "ruin and recreate" meta-heuristic proposed by Schrimpf et al. (2000), but given the CWSA context and the meta-heuristic's complete restoration to A-B-A routes (as in the initial CWSA solutions), it was defined as a reset kick in this algorithm.

In the experiments, the strategy for selecting kicked nodes and the destination mode were varied. Building on the findings from Oakey et al. (2023), one of the tested strategies kicked sites

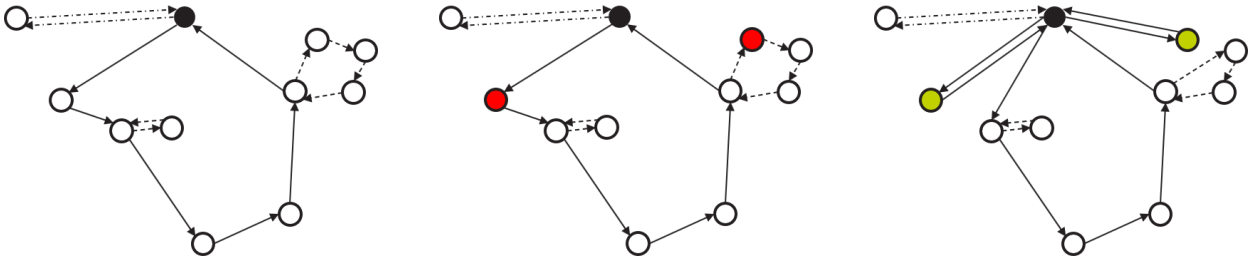

(a) Converged solution, prior to reset kick.

(b) Identifying sites to reset, depending on strategy.

(c) Kicked solution, with sites reset to new routes.

Figure 7: Demonstration of the reset kick meta-heuristic. Example resets sites to van routes, though resetting to other modes was tested. Solid lines = trunk van routes, dot-dashed lines = trunk UAV routes, dashed lines = consolidation routes.

from the longest and shortest routes, which were identified to reduce (a) driver utilisation, and (b) the longest collection round and maximum transit time; all of which are favourable to the objective function. Other tested strategies included randomly selecting a given percentage of sites across all routes.

In the event that a selected surgery was a consolidation site, all of the sites from the associated consolidation routes were also reset, allowing for simpler outcomes that would be unlikely to violate time constraints. The pseudocode for the reset kick is given in Algorithm 3.

Algorithm 3 Reset Kick Meta-Heuristic

Input: R: current set of routes; S^*: set of sites to be reset (built using upper-bounds); S: set of all surgeries, excluding hospital; H: hospital site; M: target mode; M^*: fallback mode if target mode is not permitted.

$R^\mu = \{\}$ ▷ Empty set for kicked solution
$R^* = R$ ▷ Set of existing routes to be updated
for $s \in S^*$ **do** ▷ Surgery loop
  $R^\Delta = \{\}$ ▷ Empty set of routes that are to be added
  $\tilde{R} = \{\}$ ▷ Empty set of routes that have been removed
  for $r \in R^*$ **do** ▷ **Route loop**
    if $s \notin r$ **then**
      continue Route loop
    end if
    if r is a consolidation route **then**
      for $a \in r$ **do** ▷ For each site in r
        if a permits M **then**
          $m_k = \{H, a, H\}$ ▷ Create route using target mode
          $R^\mu = R^\mu \cup m_k$ ▷ Add route to kicked route list
        else
          $m_k^* = \{H, a, H\}$ ▷ Create route using fallback mode
          $R^\mu = R^\mu \cup m_k^*$ ▷ Add route to kicked route list
        end if
      end for
    else
      $r^* = r \setminus s$ ▷ Make updated route without s
      if $r^* \neq \{H, H\}$ **then** ▷ If updated route is not empty
        $R^\Delta = R^\Delta \cup r^*$ ▷ Add to the changed route set
      end if
      if s permits M **then**
        $m_k = \{H, a, H\}$ ▷ Create route using target mode
        $R^\mu = R^\mu \cup m_k$ ▷ Add route to kicked route list
      else
        $m_k^* = \{H, a, H\}$ ▷ Create route using fallback mode
        $R^\mu = R^\mu \cup m_k$ ▷ Add route to kicked route list
      end if
    end if
    $\tilde{R} = \tilde{R} \cup r$ ▷ Add r to removal list
  end for
  $R^* \setminus \tilde{R}$ ▷ Remove old routes featuring surgery s
  $R^* \cup R^\delta$ ▷ Add updated routes, now without s
end for
$R^\mu \cup R^*$ ▷ Add finished list of modified routes
return R^μ

4.5. Algorithm Summary

Where this algorithm uses the well-established principles of the classical CWSA in an adapted form with local search heuristics, effective solutions should be found and local optima should be found quickly. The subsequent reset kick meta-heuristic then facilitates a wider search of the solution space and reduces the likelihood of stagnating at a local optimal solution. The overall algorithm is detailed in Algorithm 4.

As highlighted in the problem description (Section 3), solving speed is an important consideration in this problem, with potential users requiring a solution at short notice to allow for just-in-time deployment of the fleet. To this end, effective calibration of the algorithm is key to ensuring quality solutions without excessive time penalties. The performance of each configurable component of the algorithm is analysed in Section 5.

Algorithm 4 Complete Solution Approach

Input: C^{catch}: cycle catchment size limit; N: reset kick iteration limit; $InterRoute_{local}$: Inter-route strategy for the local search; $InterRoute_{CW}$: Inter-route strategy for savings algorithm; Y: kick strategy; *pack*: packing algorithm; *shortList*: candidate shortlist size; *reallocate*: reallocation strategy

$R = \{\}$ ▷ Empty set of routes
$U = \{\}$ ▷ Empty set of unallocated surgeries
for $s \in S$ **do**
  Define cycle catchments (given by C^{catch})
  if Route to s within constraints **then**
    $R = R \cup \{H, s, H\}$ ▷ Generate routes using van mode (Figure 3a), if within constraints
  else
    $U = U \cup s$ ▷ Add unallocated site to list for resolving
  end if
end for
$R = R \cup resolve(U)$ ▷ Resolve any unallocated sites beyond mode or time constraints using consolidation or UAVs (Figure 3b)
$n = 0$
for $n \in N$ **do**
  if $n \neq 0$ **then**
    $R = resetKick(R, Y)$ ▷ Complete reset kick (on all but first iteration)
  end if
  $improvement = true$
  while *improvement* **is** *true* **do** ▷ Loop until no longer improves and local optima reached
    $R = adaptedCW(R, C^{catch}, InterRoute_{CW}, shortList, reallocate)$ ▷ Run the adapted C&W until no improvement possible (Figure 3, Algorithm 1 and 2)
    $R = 2opt(R)$ ▷ Perform 2-opt local search on all routes (Figure 6b)
    if $Inter - route_{local}$ **is** *true* **then**
      $R = interRoute(R, multi)$ ▷ Perform inter-route (multi-move) local search on all routes (Figure 6c)
    end if
    $improvement = checkImprovement(R)$ ▷ Check for improvement in objective
  end while
  $S^* = kickSurgeries(R, Y)$ ▷ Identify surgeries to be kicked, depending on kick strategy (Figure 7, Algorithm 3)
  $n = n + 1$
end for
return R

5. Computational Experiments and Results

The solution space described by the SSCP with work shifts and UAVs is significant and makes proving optimality of a given solution very difficult. To this end, a set of test instances was generated to help calibrate and test the algorithm's performance, accounting for factors such as traffic variability, UAV landing site suitability, and the spatial distribution of the sites (relative to the road network) needed to be realistic. A new problem generation tool was created and used to complete this task due to the novelty of the problem.

The adapted CWSA/bin-packing algorithm was calibrated using the generated test sets, starting with the most fundamental parameters, such as the shortlist size and packing strategy, exploring the speed and quality of solutions in different configurations. Local search heuristic and kick meta-heuristic experiments followed.

All of the computational experiments in this paper were completed in a Microsoft Visual Studio Code environment running Java (Version 8), with 16 GB of memory and an Intel Core i7-2600 CPU (3.4 GHz).

For reference, it should be noted that Oakey et al. (2023) addressed a different and more simple problem that was solved using a column generation approach that did not reflect the additional modes and associated constraints defined in this work. The original study provided solutions within 5% of optimality with run times up to 99% faster than solving a full enumeration of potential routes

in small cases (<20 surgeries); however, the column generation approach was less beneficial when solving larger scenarios. Nonetheless, the heuristics used in the original approach were notably effective and elements of their operation were explored as part of the metaheuristic element of the present study.

5.1. Test Instance Generator

To allow for a more representative and robust testing of the algorithm, a test instance generation tool[1] was developed in Java, using a locally hosted GraphHopper Routing engine (GraphHopper, 2020) to identify the paths between OD pairs. To support further development of this research, the test instances are available from the GitHub repository detailed above.

The test cases were generated sites using a full list of UK postcodes and their coordinates (Doogal, 2023). GP surgeries were approximately spaced according to population density, which roughly correlated with postcode/street distribution density (Figure 8); hence, sites were chosen at random the list of postcodes to approximately represent a typical spatial distribution of surgeries, using a bounding box to define the limits of the test case. This approach would also work well for other VRPs, where more densely populated areas are typically more likely to see greater demand.

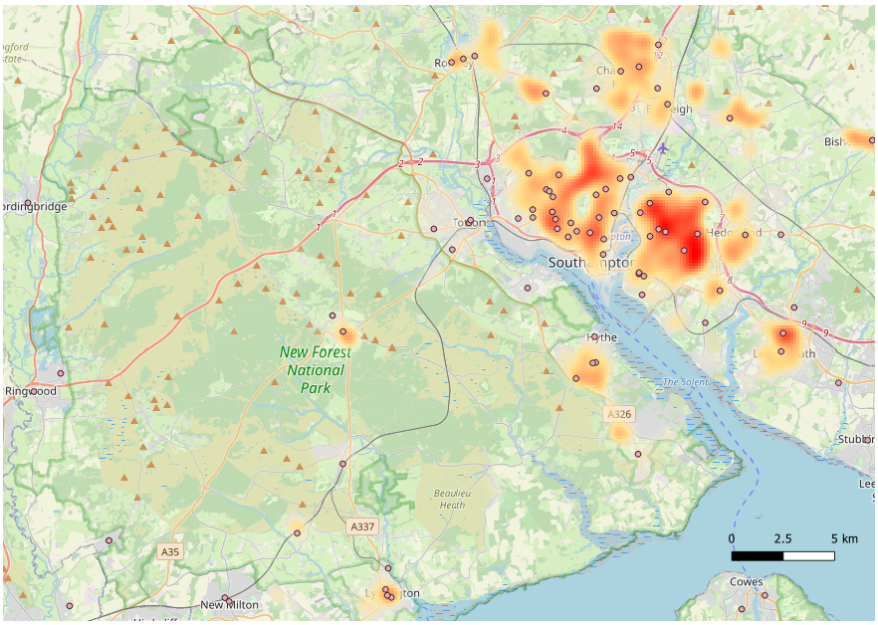

Figure 8: A comparison of GP surgery distribution (points) and "SO" postcode distribution (heat overlay) in the Southampton area. (Base Map © OpenStreetMap contributors).

A set of 40 (4×10) tests were generated, between 50 and 200 sites in size, increasing in 50-site increments between each set of 10 tests. To ensure algorithm performance across locations with different properties, half of the tests were based in the Southampton area (bounding coordinates: (50.782, -1.831), (51.000,-1.172), Figure 9a) with a shift period of 09:00-13:00 (1 hour intervals, 1 minute discretisation), whilst the other half were based in the Birmingham area (bounding coordinates: (52.326, -2.055), (52.558, -2.704), Figure 9b), with a shift period of 13:00-17:00 (1

[1] `URLomittedfordouble-blindreviewing.`

hour intervals, 1 minute discretisation). Delivery locations (i.e. the 'hospitals') were chosen either from a random postcode within the bounding coordinate area, or from the centroid of the selected sites.

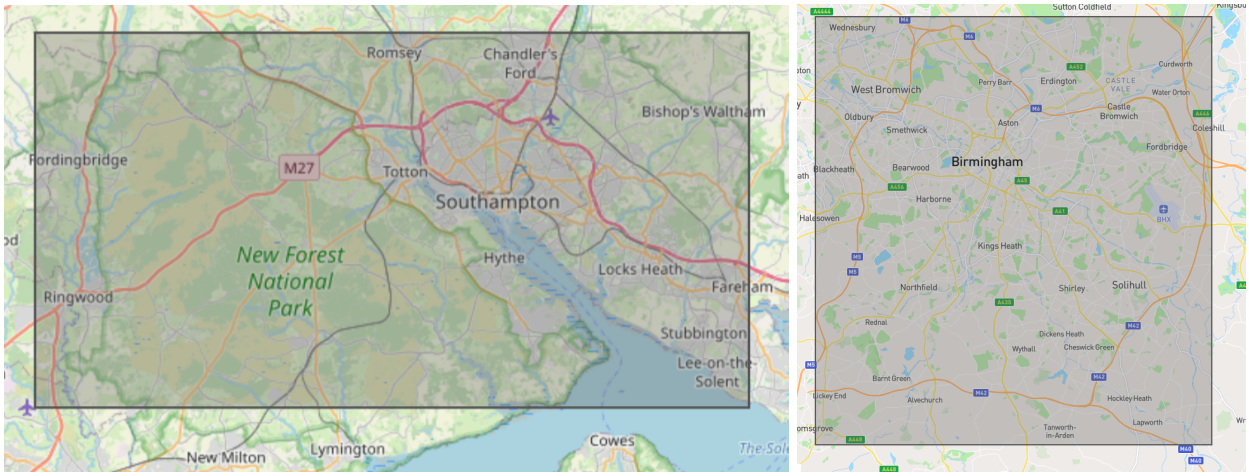

(a) Wider Southampton Area

(b) Birmingham area

Figure 9: Bounding boxes for generating test instances (grey shaded area on map, not to scale). (Base Map © OpenStreetMap contributors).

With respect to site suitability, all sites were assumed to be cycle serviceable ($S^C = S'$). Meanwhile, the UAV suitable sites and trajectory circuity were defined by a normal distribution with a random probability, based on the number of suitable sites in the Southampton area, and a typical trajectory circuity factor from initial analyses of the Southampton area (Anonymous, 2022; Anonymous, 2022a). Consolidation van suitability was assumed to be 5% of all sites. In terms of the test-case objectives, the theta coefficients were set such that the algorithm optimised solely to cost, with equal weighting. This minimised any impacts arising from the OD-time approximations in the test case dataset and avoided the need to accurately calculate emissions/energy.

Traffic variability was captured by applying a normally distributed penalty with random probability to each OD pair, based on UK Department for Transport hourly motor traffic distribution statistics (pre-COVID pandemic) (Department for Transport, 2022), assuming that travel times diminish above 1.25× the average flow. Even though this approach simplifies traffic patterns and may penalise areas that experience lower than average traffic flows, applying a general trend was sufficient to enable robust testing of the algorithm.

5.2. Algorithm Parameter Selection

Several parameters were varied when parameterising and testing the algorithm (Figure 10), including: (i) the packing algorithm used to compile vehicle shifts; (ii) the size of the savings options shortlist; (iii) whether associated consolidation route sites were re-allocated when the base site was moved; (iv) the maximum size of the cycle catchment; (v) how the inter-route neighbourhood search was conducted (within the adapted CWSA, or as a local search); and (vi) the strategy of the reset-kick. Subsequently, combinations of varying parameters were tested to identify the number of iterations required for their application to be most effective.

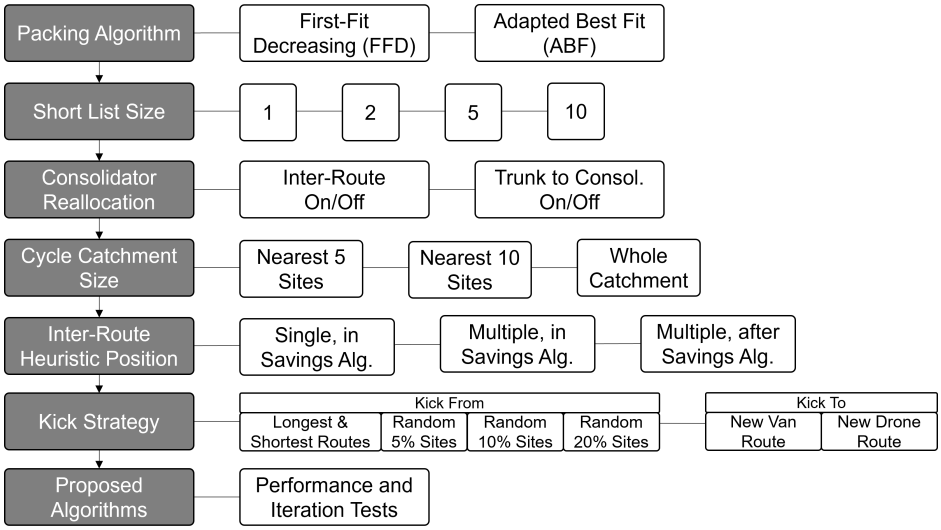

Figure 10: Process/order of algorithm performance testing, with set-up options shown.

The tests for parameters (i)-(v), addressing the initial solution generation, were run for 10 reset kick iterations (i.e. converge and reset-kick 10 times), whilst the kick strategy tests in (vi) were run for 50 iterations to better display the performance of the kick. The final configuration tests were run for 500 iterations to identify the typical run durations to reach the best solutions.

5.2.1. Packing Algorithm for Shift Compilation

To test the performance of the two packing algorithms outlined in Section 4 (AFFD and ABF), the instances were solved under both algorithms, with constant values set for the other parameters.

Both algorithms have a quadratic computational complexity, though the ABF algorithm displayed faster performance (Figure 11a) and the AFFD displayed marginally higher quality results with respect to the objective function. These differences were caused by the packing performance for a given set of routes and the potential savings available at each iteration of the adapted CWSA. In turn, this impacted on the number of iterations within the adapted CWSA loop each packing algorithm typically utilised (i.e. before converging), with the AFFD approach computing an average of 9% more intermediate solutions than the ABF within each CWSA loop before stagnating at a local optima and performing a reset-kick (based on results after 10 reset kick iterations, across all test cases). Thus, the increased iterations reduced the computational speed of the AFFD algorithm relative to the BF algorithm.

The performance of the packing algorithms themselves depends significantly on the traffic trends for ODs in the problem and the variability from the maximum duration. For example, the van OD travel durations in the Southampton case study (09:00-13:00 shift period) were generally worst in the 09:00 period, and best at 11:00/12:00 (Figure 12a). Such a trend seemed to be typical for most areas (Department for Transport, 2022). The greatest and most consistent travel duration variances were seen on shortest routes, with traffic density generally being more of an issue on shorter OD journeys. To this end, routes containing a majority of short-range OD pairs typically had a greater variance and would be more likely to be fitted first in the ABF algorithm, despite them being largely consistent in their performance throughout the shift period (Figure 12b). Conversely, the AFFD algorithm would fit the longest routes first, which may be more optimal, depending on the typical composition of the routes applied to the route construction.

In conclusion, the different packing behaviours resulted in the AFFD algorithm displaying a better performance (Figure 11b), with increased iterations enabling access to further savings options. If duration variability is more significant and consistent across route lengths, then it may be seen that the ABF performs better, though the difference is likely to remain marginal regardless. The ABF may also be more robust in areas where timings with scheduled connections (e.g. ferries) may affect timings, favouring departures that prevent waiting times and longer durations. For the final algorithms tested in Section 5.3, both options were retained due to the potential trade off between speed and solution quality.

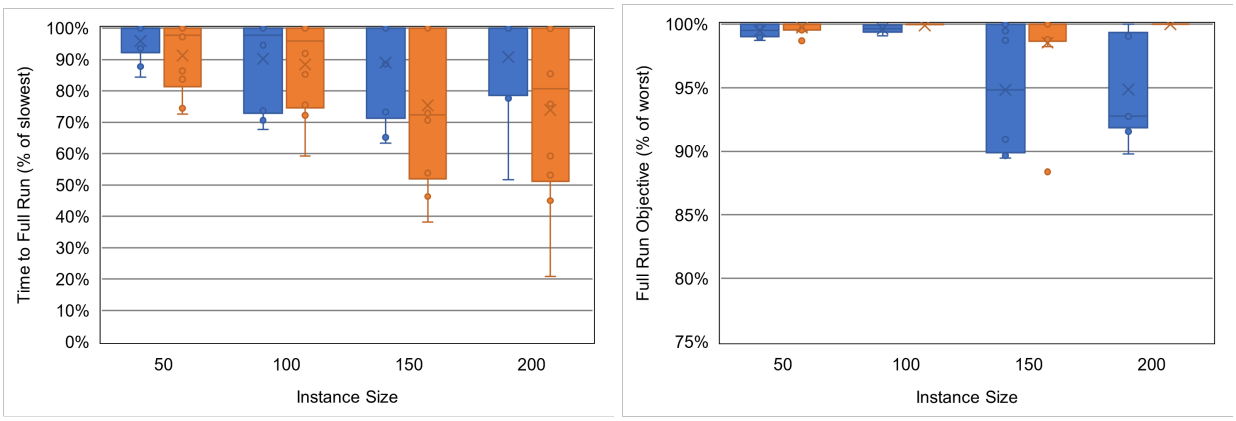

(a) Computational time after 10 reset kick iterations.

(b) Best objective value after 10 reset kick iterations.

Figure 11: Relative performance comparison of the Adapted First-Fit Decreasing (blue) and Adapted Best-Fit (orange) algorithms. Mean of instance results, scaled relative to worst performing in each test (worst = 100%).

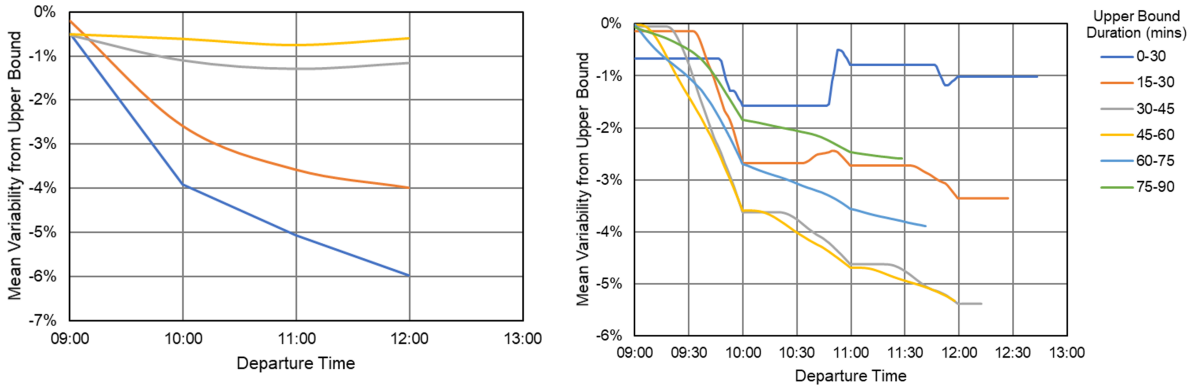

(a) OD-variability, assuming constant travel times for each hour, based on all OD pairs in the Southampton diagnostic specimen case study area.

(b) Route variability, based on a significant set of routes (n=5009) in the Southampton case study area. Discontinuities around each hour result from the steps between constant condition periods.

Figure 12: OD and route variability over a typical morning shift period (09:00-13:00) for sites served in the Southampton diagnostic specimen case study area. Upper bound is the slowest/sum of slowest travel times for each OD-pair.

5.2.2. Introducing Randomness in the Construction Heuristic

To limit any potential impacts of the algorithm being too greedy (i.e. always selecting the best saving within each iteration), the proposed shortlist was designed to introduce an element of randomness into the algorithm. In these performance tests, shortlist sizes of 1 (i.e. fully greedy), 2, 5, and 10 were considered (L in Algorithm 1). It should be highlighted that if a saving option remained feasible and competitive in later iterations of the adapted CWSA, it would recur and still have a small chance of being selected even if the shortlist size limit was set to a larger number.

In the experiments, the greediest approaches, using a shortlist size of 1 or 2 provided the best quality solutions, consistently offering the lowest objective value (Figure A.16b). The likely reason for this is due to the larger shortlists neglecting the best improvements, resulting in an unrecoverable deviation from more optimal solutions. This was particularly noticeable with larger tests cases, where the greater number of required CWSA iterations magnifies any deviations.

In the larger cases, the most greedy shortlist options (1, 2) performed well with respect to computational time (Figure A.16a), despite being less effective in small test cases. A greedy approach enabled effective truncation of searches when many potential savings options were available whilst ensuring saving options continued to occur in smaller tests (i.e. more iterations, slower). In some cases the speed trade-off was significant (up to 50% difference), though the shortlist size of 2 remained the most consistently time efficient option throughout. To this end, the shortlist of 2 and the best quality option of 1 were retained for the final algorithms (Section 5.3).

5.2.3. Reallocation of Consolidation Routes

Reallocation tests explored the effects of permitting reallocation when conducting: (i) trunk to consolidation (TTC) moves, where any consolidation route sites attached to a trunk-served site are also reallocated; (ii) inter-route (IR) moves, also where any consolidation route sites attached to a trunk-served site; or (iii) combinations of (i) and (ii). Results suggested a negative impact in terms of computational time when permitting reallocation (Figure A.17a). This was largely as expected due to reallocation resulting in additional savings options being explored in each iteration.

These tests were run with the IR option set to multi-move (i.e. moves are compounded into a single saving) as a local search, as opposed to being included in the savings options. Hence, the potential time impact arising from introducing reallocation to IR searches was considerably smaller than when introducing it to the TTC search. Not permitting reallocation in both searches was generally more effective in most cases, however, medium sized (100/150) tests were faster when inter-route reallocation was permitted. This was due to the reallocation option encouraging faster convergence without detrimental search time trade-off.

With respect to quality of solutions (Figure A.17b), disabling reallocation was generally more effective, with some select exceptions where other options dominated. The cause is likely due to reallocated options offering benefits on occasion, whilst potentially encouraging more locally optimal moves in most other cases, limiting further progress in the adapted CWSA and a deviation from the optimal solution. Smaller cases did not vary much between settings due to the limited scope for reallocation. Based on these findings, the no reallocation permitted and the inter-route permitted searches were retained for the final algorithms.

5.2.4. Cycling Catchment Size Cap

The maximum cycle catchment size was varied between the nearest 5 sites, nearest 10 sites, and all sites within the cycle-able range (C^{catch} in Algorithm 1). Results generally correlated in terms of computational time and solution quality (Figure A.18), with faster solutions offering

better objective function outcomes. It should be noted that some exceptions to this trend were present, with the correlation being more evident when test cases were larger, due to the density of surgeries being far greater (i.e. catchments were likely to include more sites and reach their catchment limit).

This was largely to be expected, given close-range consolidation offers reasonable speed benefits to the solution, without large cost increases or delays. Hence, longer-distance consolidation was less likely to be selected, whilst also requiring additional time to process. The nearest 5 sites option was retained for all algorithms due to the consistent objective outcomes and computational speed.

Should the algorithm be applied in a real-world context, users with more cycle consolidation opportunities in their area and a stronger environmental objective weighting may benefit from expanding cycle catchments to increase uptake. However, this would likely be detrimental to the computational runtimes.

5.2.5. Use of Inter-Route Search

IR searches were tested in multiple configurations, with combinations of single-move and multi-move (i.e. compounded) savings being considered in the adapted CWSA, whilst multi-move searches were only considered as part of the local search. It should be noted that due to the memory intensity of some of the tests in the 200 site cases, insufficient memory errors often occurred; thus, these cases were omitted from results to maintain consistency.

An anticipated, introducing the single-move IR option to the CWSA and the multi-move IR option to the local search simultaneously resulted in the slowest performance. This was due to the significant number of additional options that required investigation. Similarly, the single-move CWSA option with no local search option was also computationally intensive and recorded slower run times. The remaining options were similar in their computational time, with the compounded CWSA options offering advantages with larger cases.

In small test cases, there was little difference in solution quality (Figure 13b) between all configurations, whilst in larger cases there was a clear benefit in using the IR option, particularly on a single-move basis. Using the multi-move option in the CWSA resulted in very rapid convergence (Figure 13a), with every move replacing a possible saving by another means in the algorithm. This is likely the main cause for the lower solution quality and faster solving speeds when this option is enabled. Single-moves in the savings algorithm took considerably longer to run due to the slower rate of change with each solution progression, though quality is generally better than multi-move. Local searches with the IR option allow minor improvements, with minimal time penalty.

Owing to the significant benefit in solution quality offered by using the local search, the main configurations that were carried forward adopted this set-up; one with a single-move savings algorithm search, and one without. Additionally, the multi-move savings/local search off option was also investigated further for the purpose of fast convergence.

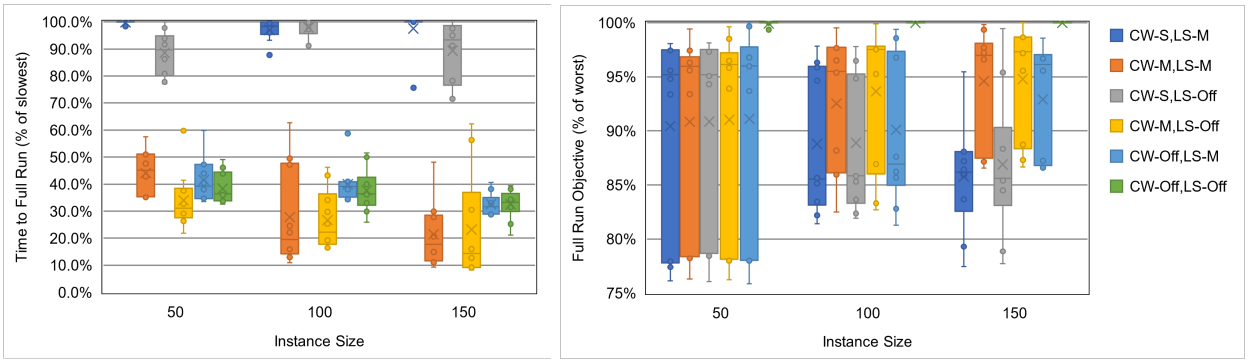

(a) Computational time after 10 reset kick iterations.

(b) Best objective value after 10 reset kick iterations.

Figure 13: Relative performance comparison of the inter-route algorithm usage options. Mean of instance results, scaled relative to worst performing in each instance test (worst = 100%). CW = savings algorithm option, LS = local search option, S = single move, M = multi-move (cumulative), Off = not used.

5.2.6. Meta-Heuristic Kick Strategy

The meta-heuristic experiments were conducted using different kick strategies. They included resetting the sites longest and shortest existing routes to either van-served, or UAV-served routes; or resetting a percentage (5%/10%/20%) of randomly selected sites across all existing routes to either van-served, or UAV-served routes.

Building on the previous tests, these experiments were run for 50 reset kick iterations to better understand the impact of the kick process. The main objective of the reset kicks was to cause sufficient change in the solutions to enable convergence to a different optima without destroying the more optimal parts of solution.

In the strategy tests, there was evident benefit in kicking to van served routes over UAV served routes (Figure 14), whilst the size and origin of the kick was less consistent. In the smaller test cases, the disparity between kicking to vans and kicking to UAVs was less significant due to the number of UAV suitable sites limiting the success of the kicks to UAV-served routes (i.e. if sites weren't UAV suitable, they would reset to a van-served route). Conversely, the van-targeting kicks generally performed better than their respective UAV-targeting equivalent in larger cases. With respect to computational speed, kicking to van-served routes was generally more efficient as well.

With respect to the selection strategy, kicking a random 5% of surgeries typically gave the most consistent outcomes across all test case sizes in terms of both speed and quality. A random 10% kick was successful in some tests, but was not so consistent, meanwhile 20% kicks were evidently destructive to the point where re-convergence was slow and run-times were poor. Kicking from the longest and shortest routes was also fast and gave reasonable solutions, eliminating some of the less effective routes that had a disproportionate impact on the objective function. To this end, the long and short strategy may be effective under a variety of constraints and use cases.

A kick of 5% to van, and long and short to van were both kept for the final algorithms. Additionally, a 10% kick to UAVs was maintained to understand if lower cost UAVs would result in better performance under such a strategy.

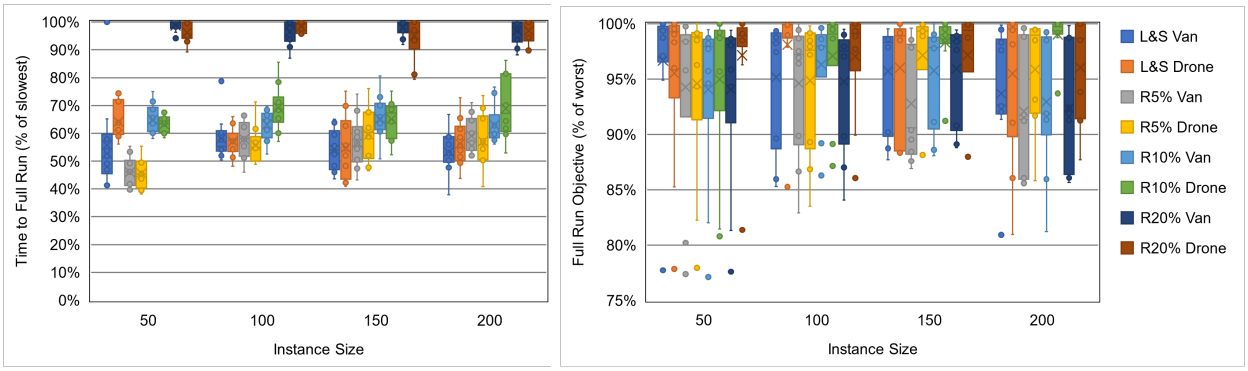

(a) Computational time after 50 reset kick iterations.

(b) Best objective value after 50 reset kick iterations.

Figure 14: Relative performance comparison of the kick-strategy options. Mean of instance results, scaled relative to worst performing in each instance test (worst = 100%). RX% = random X% of sites selected for kick, L&S = sites on longest and shortest routes selected for kick, to van = kicked to van routes (where possible), to UAV = kicked to UAV routes (where possible).

5.3. Defining Final Algorithms and Their Overall Performance

Using the parameter test results, individual algorithms were defined to establish search strategies that performed well under different user circumstances, taking the parameters that offered the best performance with respect to speed and/or quality (green highlights in Table 5). Other parameters that were considered as potentially encouraging greater mode shift were also brought forward for further investigation (yellow highlights in Table 5), with the compounded inter-route search offering the opportunity for several simultaneous moves in a single saving option, and the UAV-targeting meta-heuristic also offering several mode transfers in a single step.

Table 5: Configuration options for all algorithm variables. R/A = Reallocation, IR = Inter-Route, TTC = Trunk to Consolidator, CW = savings algorithm option, LS = local search option, multi = multi-move (compounded), Off = not used. Green highlights options brought forward due to good performance during testing. Yellow highlights indicate other options brought forward to that may enable a more favourable modal shift.

Packing (5.2.1)	Shortlist (5.2.2)	R/A (IR/TTC) (5.2.3)	Catchment (5.2.4)	IR Usage (5.2.5)	Kick Origin (5.2.6)	Kick Destination (5.2.6)
AFFD	1	Off/Off	5 Sites	CW-Off, LS-Off	Long & Short	Van
ABF	2	On/Off	10 Sites	CW-S, LS-Off	Random 5%	UAV
	5	Off/On	All Sites	CW-M, LS-Off	Random 10%	
	10	On/On		CW-Off, LS-M	Random 20%	
				CW-S, LS-M		
				CW-M, LS-M		

Due to the potential time limitations of planning deliveries with minimal foresight of which sites require service, the model would need to be used in a short space of time, e.g. in the morning, prior to commencing any vehicle rounds. Conversely, the model could also be used to complete a longer-term analysis of historic data and identify the typical demands and vehicle requirements for an area, though this will also require reasonable speed due to the potential quantity of data that may be tested.

With this in mind, Algorithm 1 was defined using the parameters that typically offered the best quality outputs, regardless of the computational time; and Algorithm 2 was defined using a slight

compromise in quality with a view to improve the computational speed, using the best time options where quality trade-off was minimal (Table 6). Meanwhile, Algorithms 3 and 4 were constructed to enable a fast comparison of configurations which are likely to encourage greater modal shift.

Table 6: Selected algorithm configurations. R/A = Reallocation, IR = Inter-Route, TTC = Trunk to Consolidator, CW = savings algorithm option, LS = local search option, multi = multi-move (compounded), Off = not used.

Alg. No.	Packing (5.2.1)	Shortlist (5.2.2)	R/A (IR/TTC) (5.2.3)	Catchment (5.2.4)	IR Usage (5.2.5)	Kick Origin (5.2.6)	Kick Destination (5.2.6)
1	AFFD	Shortlist	On/Off	5 Sites	CW-S, LS-M	Random 5%	Van
2	BF	Shortlist	On/Off	5 Sites	CW-Off, LS-M	Long & Short	Van
3	BF	Shortlist	Off/Off	5 Sites	CW-M, LS-Off	Random 10%	UAV
4	BF	Shortlist	Off/Off	5 Sites	CW-M, LS-Off	Random 5%	Van

On testing the two core algorithms (1&2), it was demonstrated that the best solution would be offered by Algorithm 1, as intended (Figure A.19c). Meanwhile, significant speed benefits (at least 50% faster than Algorithm 1 to 500 kick iterations) were made possible using Algorithm 2, demonstrating only a minor loss in quality (<5%) in the majority of cases, particularly with smaller tests which were more typical of real-world case studies (Figure A.19a).

In terms of the required iterations and run-time to find the best solution (within the tested 500 reset kick iterations), test results (Figure A.19b) highlighted that Algorithm 2 converges to its best solution significantly faster than the others. Furthermore, with 90% of solutions (i.e. the 90^{th} percentile) reaching their best solution within 90 (out of 500) reset kick iterations, each iteration was seen to be highly efficient with respect to time and consistency. The equivalent percentile for Algorithms 1, 2, 3 and 4 were 402, 90, 452, and 442 iterations, respectively, suggesting that there may still be scope for improvement if more kicks were permitted, though runtime and memory management may then become a challenge.

The absolute time requirements to solve the SSCP with work shifts and UAVs varied with respect to the size of the test case (Figure 15). The configurations with the shortest computational time to reach their best solution typical exhibited a more linear trend with respect to increasing the test case size, unless a full 500 iteration limit was observed, when an exponential trend was evident (e.g. Algorithms 3, 4). Conversely, the configurations that took longest to reach their best solution were generally more stable per kick iteration but required substantially more iterations to realise their best solution. Nonetheless, 90% of all solutions (i.e. all test cases, regardless of size) were found within 15.5 minutes in all algorithms, and all typical size problems (~50-100 sites, (Oakey et al., 2023)) were solved within 512 seconds, providing the required speed for day-to-day use. In the event that planners are investigating larger cases, it may be advisable to use the faster algorithms (3/4) if planning time is limited, however, this may present a trade-off with respect to the objective function.

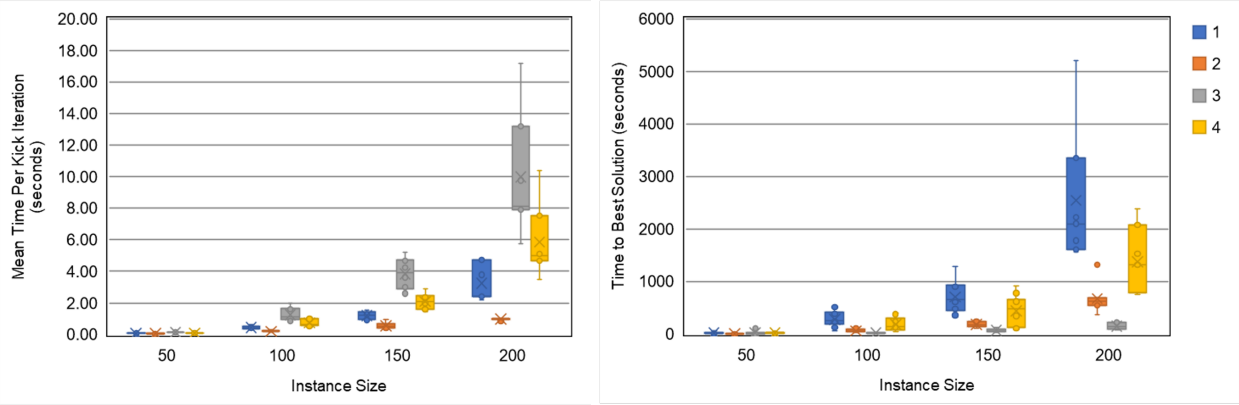

(a) Computational time per reset kick iteration.

(b) Time taken to reach best objective value.

Figure 15: Absolute time performance comparison of the final algorithms. Legend indicates algorithm used, as defined in Table 6.

5.3.1. Meta-Heuristic Sensitivity to UAV Cost Changes

The greatest uncertainty in modelling of real-world scenarios lies with the UAV costs. This is due to the novelty of the technology and limited literature and data covering this area, partially as a result of commercial sensitivities. To this end, a brief sensitivity analysis with UAV costs was undertaken.

To support theoretical scenarios such as cheaper UAVs, where mode-shift may be more likely, Algorithms 3 and 4 were developed and tested. Such an approach was made to ensure that solutions featuring significant number of UAVs were not missed during sensitivity tests, such as in (Anonymous, 2023). To test this concept, the small (50/100) test cases were re-run using these 2 algorithms with UAV costs weight set to 0.1 (other costs set to 1.0), effectively reducing the operating costs of UAV routes to 10% of their value in the other tests.

Findings supported the theory that modal shift was encouraged by the reset kick towards UAV service (Figure A.20), with 20% (4/20) of solutions containing more UAV routes when calculated under Algorithm 3 than Algorithm 4. Whilst this may have benefited the objective function in some solutions, it was also detrimental in others due to the cost challenges associated with operator ratios not being sufficient to outweigh the economy of scale from van-based solutions.

In general, convergence times were reduced under this arrangement (Figure A.20a), likely due to early savings from introducing UAVs limiting objective improvements later on in the solving process and causing stagnation. To this end, the limited site suitability preventing further UAV uptake and operator utilisation will be a key factor in future applications. Moreover, the quality of solutions in small cases was not significantly impacted (<10% in most cases, Figure A.20b).

In summary, Algorithm 3 may be beneficial where cost sensitivities are heightened, such as around the point of cost-parity of UAV routes and van routes, as in (Anonymous, 2023); however, it would be advantageous to explore the kick meta-heuristic further in future studies.

6. Conclusions

Local healthcare logistics typically need to move goods in a timely manner, though environmentally conscious policies are encouraging changes towards green practices. Meanwhile, the use of UAVs is seen as a growing area of interest, particularly in healthcare logistics, though it has

been highlighted that they would only form part of future logistics networks due to operational and regulatory challenges. To this end, this paper has presented an extension of the heterogeneous two-echelon VRP known as the "Sustainable Specimen Collection Problem" (SSCP), introducing UAVs and vehicle reuse as part of scheduled work shifts (SSCP with work shifts and UAVs).

To solve the proposed generalised VRP, a novel heuristic algorithm was introduced, whereby an adaptation of the Clarke and Wright Savings Algorithm is used to create efficient vehicle routes, whilst bin-packing algorithms are subsequently used to compile suitable work shifts from these routes. Additionally, a meta-heuristic based on the ruin and recreate algorithm was proposed to complement the approach.

To enable realistic testing, a new test case generator was also proposed, taking factors such as traffic flow and site distribution densities into account. Using a series of test cases ranging from 50 to 200 sites in size, the component parts of the algorithm were profiled, and effective parameter combinations were defined to ensure computational speed whilst maintaining quality. To this end, the configurations that limited the search space and encouraged smaller incremental improvements were generally the most successful at giving quality results in reasonable timescales. Specifically, it was found that 90% of all solutions (i.e. all test cases, regardless of size: 50/100/200 sites) were found within 15.5 minutes in all of the identified configurations, and all typical size problems (~50-100 sites) were solved within 512 seconds, providing the required speed for day-to-day use in industry.

Whilst the algorithm provides an effective heuristic approach, future work may seek to explore an exact solution and understand the possible optimality gap. Furthermore, the adapted Clarke and Wright Savings Algorithm could be applied to other heterogeneous VRPs, particularly where one mode is likely to serve a substantial proportion of customers.

With respect to the use of UAVs, this paper also identified the challenges that could potentially limit their uptake, relating to cost and operator ratios, when compared to the economies of scale offered by traditional logistics modes. Hence, exploring the sensitivities around these areas and the other objective function components with real case study data could be advantageous. It should be highlighted that such work is already in progress, as published in (Anonymous, 2023), investigating the potential cost requirements for UAV services to become feasible in the region. It should also be noted there are significant contractual and management practicalities to overcome before implementing any solution, many of which are highly specific to the geographies of interest.

In terms of limitations, this study investigates the proposed solution approach using generated datasets based on two UK-based regions that could be deemed representative of many urban and urban-rural areas. This may limit the effectiveness of the algorithm in substantially different areas; thus, understanding how the performance can vary when more generalised datasets based on quantified network characteristics are used would be of significant interest. This would also enable a formulaic rule-of-thumb to be produced for quickly appraising the suitability and expected outcomes under the SSCP with work shifts and UAVs arrangement.

Acknowledgements

This section will be released after peer review.

Funding

This section will be released after peer review.

Contributions

This section will be released after peer review.

Conflicts of Interest

This section will be released after peer review.

References

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. URL: https://doi.org/10.1080/00207543.2015.1018455, doi:10.1080/00207543.2015.1018455. publisher: Taylor & Francis _eprint: https://doi.org/10.1080/00207543.2015.1018455.

Benarbia, T., Kyamakya, K., 2022. A Literature Review of Drone-Based Package Delivery Logistics Systems and Their Implementation Feasibility. Sustainability 14, 360. URL: https://www.mdpi.com/2071-1050/14/1/360, doi:10.3390/su14010360. number: 1 Publisher: Multidisciplinary Digital Publishing Institute.

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. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/itor.12796, doi:10.1111/itor.12796. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/itor.12796.

Cawthorne, D., Wynsberghe, A.R.V., 2019. From HealthDrone to FrugalDrone: Value-Sensitive Design of a Blood Sample Transportation Drone, in: 2019 IEEE International Symposium on Technology and Society (ISTAS), pp. 1–7. doi:10.1109/ISTAS48451.2019.8938019. iSSN: 2158-3412.

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.

Clarke, G., Wright, J.W., 1964. Scheduling of Vehicles from a Central Depot to a Number of Delivery Points. Operations Research 12, 568–581. URL: https://pubsonline.informs.org/doi/abs/10.1287/opre.12.4.568, doi:10.1287/opre.12.4.568. publisher: INFORMS.

Coffman, E.G., Garey, M.R., Johnson, D.S., 1984. Approximation Algorithms for Bin-Packing — An Updated Survey, in: Ausiello, G., Lucertini, M., Serafini, P. (Eds.), Algorithm Design for Computer System Design. Springer, Vienna. International Centre for Mechanical Sciences, pp. 49–106. URL: https://doi.org/10.1007/978-3-7091-4338-4_3, doi:10.1007/978-3-7091-4338-4_3.

Coffman, E.G., Garey, M.R., Johnson, D.S., 1987. Bin packing with divisible item sizes. Journal of Complexity , 406–428URL: https://core.ac.uk/display/82086619?utm_source=pdf&utm_medium=banner&utm_campaign=pdf-decoration-v1, doi:10.1016/0885-064X(87)90009-4.

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. URL: http://www.sciencedirect.com/science/article/pii/S221053951630147X, doi:10.1016/j.rtbm.2017.07.001.

Crosby, I., 2023. 20 Drones, 1 Pilot: Matternet's New Waiver Could Change the Economics of Drone Delivery. URL: https://dronelife.com/2023/02/13/20-drones-1-pilot-matternets-new-waiver-could-change-the-economics-of-drone-delivery/.

Demir, E., Bektaş, T., Laporte, G., 2014. A review of recent research on green road freight transportation. European Journal of Operational Research 237, 775–793. URL: https://www.sciencedirect.com/science/article/pii/S0377221713010175, doi:10.1016/j.ejor.2013.12.033.

Department for Transport, 2021. TAG data book. URL: https://www.gov.uk/government/publications/tag-data-book.

Department for Transport, 2022. Road traffic statistics - About. URL: https://roadtraffic.dft.gov.uk/custom-downloads/road-accidents/reports/53dadb3b-4f3d-40b6-a437-8aee1382d651.

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.

Doogal, 2023. UK Postcodes. URL: https://www.doogal.co.uk/UKPostcodes.

Dósa, G., 2007. The Tight Bound of First Fit Decreasing Bin-Packing Algorithm Is FFD(I) $\leq$ 11/9OPT(I) + 6/9, in: Chen, B., Paterson, M., Zhang, G. (Eds.), Combinatorics, Algorithms, Probabilistic and Experimental Methodologies, Springer, Berlin, Heidelberg. pp. 1–11. doi:10.1007/978-3-540-74450-4_1.

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. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/itor.12354, doi:10.1111/itor.12354. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/itor.12354.

European Environment Agency, 2020. Delivery drones and the environment — European Environment Agency. URL: https://www.eea.europa.eu/publications/delivery-drones-and-the-environment/at_download/file.

FTA, 2022. Manager's Guide to Distribution Costs Oct 2022.

Gao, M., Hugenholtz, C.H., Fox, T.A., Kucharczyk, M., Barchyn, T.E., Nesbit, P.R., 2021. Weather constraints on global drone flyability. Scientific Reports 11, 12092. URL: https://www.nature.com/articles/s41598-021-91325-w, doi:10.1038/s41598-021-91325-w. bandiera_abtest: a Cc_license_type: cc_by Cg_type: Nature Research Journals Number: 1 Primary_atype: Research Publisher: Nature Publishing Group Subject_term: Atmospheric dynamics;Atmospheric science;Climate sciences Subject_term_id: atmospheric-dynamics;atmospheric-science;climate-sciences.

Garcia, O., Santoso, A., 2019. Comparative Evaluation of Drone Delivery Systems in Last-Mile Delivery. Ph.D. thesis. MIT. URL: https://dspace.mit.edu/handle/1721.1/121319. accepted: 2019-06-17T15:22:53Z.

Gaskell, T.J., 1967. Bases for Vehicle Fleet Scheduling. Journal of the Operational Research Society 18, 281–295. URL: https://doi.org/10.1057/jors.1967.44, doi:10.1057/jors.1967.44. publisher: Taylor & Francis _eprint: https://doi.org/10.1057/jors.1967.44.

Goodchild, A., Toy, J., 2018. Delivery by drone: An evaluation of unmanned aerial vehicle technology in reducing CO2 emissions in the delivery service industry. Transportation Research Part D: Transport and Environment 61, 58–67. URL: https://www.sciencedirect.com/science/article/pii/S136192091630133X, doi:10.1016/j.trd.2017.02.017.

GraphHopper, 2020. GraphHopper - RacingBikeFlagEncoder. URL: https://github.com/graphhopper/graphhopper.

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. URL: https://doi.org/10.1186/1472-6963-14-12, doi:10.1186/1472-6963-14-12.

Gruber, J., Narayanan, S., 2019. Travel Time Differences between Cargo Cycles and Cars in Commercial Transport Operations. Transportation Research Record 2673, 623–637. URL: https://doi.org/10.1177/0361198119843088, doi:10.1177/0361198119843088. publisher: SAGE Publications Inc.

Hart, J.P., Shogan, A.W., 1987. Semi-greedy heuristics: An empirical study. Operations Research Letters 6, 107–114. URL: https://www.sciencedirect.com/science/article/pii/0167637787900216, doi:10.1016/0167-6377(87)90021-6.

Hu, Z.H., Li, T., Tian, X.D., Wei, Y.H., 2023. Drone-Based Emergent Distribution of Packages to an Island from a Land Base. Drones 7, 218. URL: https://www.mdpi.com/2504-446X/7/3/218, doi:10.3390/drones7030218. number: 3 Publisher: Multidisciplinary Digital Publishing Institute.

Johnson, D.S., 1973. Near-optimal bin packing algorithms. Thesis. Massachusetts Institute of Technology. URL: https://dspace.mit.edu/handle/1721.1/57819. accepted: 2010-08-31T15:59:47Z.

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.

Kovač, M., Tadić, S., Krstić, M., Bouraima, M.B., 2021. Novel Spherical Fuzzy MARCOS Method for Assessment of Drone-Based City Logistics Concepts. Complexity 2021, e2374955. URL: https://www.hindawi.com/journals/complexity/2021/2374955/, doi:10.1155/2021/2374955. publisher: Hindawi.

Krol, J., Anvari, B., Blakesley, A., Cherrett, T., 2023. Estimation of Energy Usage in Electric and Diesel Vans for Logistics Applications Using Gaussian Processes, in: Proceedings of the 2023 Transportation Research Board (TRB) Annual Meeting, Washington DC.

Kyriakakis, N.A., Stamadianos, T., Marinaki, M., Marinakis, Y., 2022. The Electric Vehicle Routing Problem with Drones: An Energy Minimization Approach for Aerial Deliveries. Cleaner Logistics and Supply Chain , 100041URL: https://www.sciencedirect.com/science/article/pii/S2772390922000142, doi:10.1016/j.clscn.2022.100041.

Laseter, T., Tipping, A., Duiven, F., 2018. The Rise of the Last-Mile Exchange. URL: https://www.strategy-business.com/article/The-Rise-of-the-Last-Mile-Exchange?gko=d0a62. library Catalog: www.strategy-business.com.

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. URL: https://www.sciencedirect.com/science/article/pii/S0377221713003585, 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.

Lomas, N., 2023. European Parliament agrees a way forward on platform workers' rights. URL: https://techcrunch.com/2023/02/02/eu-platform-worker-directive-parliament-mandate/.

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. UK Parliament. URL: https://committees.parliament.uk/writtenevidence/5690/pdf/.

McDonald, J.J., 1972. Vehicle Scheduling – A Case Study. Operational Research Quarterly (1970-1977) 23, 433–444. URL: https://www.jstor.org/stable/3007958, doi:10.2307/3007958. publisher: Palgrave Macmillan Journals.

Murray, C.C., Chu, A.G., 2015. The flying sidekick traveling salesman problem: Optimization of drone-assisted parcel delivery. Transportation Research Part C: Emerging Technologies 54, 86–109. URL: https://www.sciencedirect.com/science/article/pii/S0968090X15000844, doi:10.1016/j.trc.2015.03.005.

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. URL: https://www.sciencedirect.com/science/article/pii/S0377221716302260, 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. URL: https://doi.org/10.1373/jalm.2016.021634, doi:10.1373/jalm.2016.021634.

NHS, 2014. National Pathology Programme. URL: https://www.england.nhs.uk/wp-content/uploads/2014/02/pathol-dig-first.pdf.

NHS, 2018. Ambulance Quality Indicators: Data specification for Systems Indicators. URL: https://www.england.nhs.uk/statistics/wp-content/uploads/sites/2/2018/07/20180525-Ambulance-System-Indicators-specification.pdf.

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, Sedman, R., 2020. Pathology Specimen Transport. URL: https://www.cddft.nhs.uk/media/615106/transport%20sop.pdf.

Oakey, A., Martinez-Sykora, A., Cherrett, T., 2023. Improving the efficiency of patient diagnostic specimen collection with the aid of a multi-modal routing algorithm. Computers & Operations Research , 106265URL: https://www.sciencedirect.com/science/article/pii/S0305054823001296, doi:10.1016/j.cor.2023.106265.

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. URL: https://www.sciencedirect.com/science/article/pii/S2210650217307630, doi:10.1016/j.swevo.2018.04.001.

Otto, A., Agatz, N., Campbell, J., Golden, B., Pesch, E., 2018. Optimization approaches for civil applications of unmanned aerial vehicles (UAVs) or aerial drones: A survey. Networks 72, 411–458. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/net.21818, doi:10.1002/net.21818. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/net.21818.

Ouiss, M., Ettaoufik, A., Marzak, A., Tragha, A., 2022. A Parallel Genetic Algorithm for Solving the Vehicle Routing Problem with Drone Medication Delivery, in: Saeed, F., Al-Hadhrami, T., Mohammed, E., Al-Sarem, M. (Eds.), Advances on Smart and Soft Computing, Springer, Singapore. pp. 225–233. doi:10.1007/978-981-16-5559-3_19.

PharmaAero, 2022. WP3: The Use of Drones in Transporting Vaccines From the Factory to the Patient. Technical Report. URL: https://pharma.aero/wp-content/uploads/2022/11/White-Paper_UAV-Project_WP-3.pdf.

Pichpibul, T., Kawtummachai, R., 2012. An improved Clarke and Wright savings algorithm for the capacitated vehicle routing problem. ScienceAsia 38, 307. doi:10.2306/scienceasia1513-1874.2012.38.307.

Pilko, A., Sóbester, A., Scanlan, J.P., Ferraro, M., 2021. Spatiotemporal Ground Risk Mapping for Uncrewed Aerial Systems operations, in: AIAA SCITECH 2022 Forum. American Institute of Aeronautics and Astronautics. AIAA SciTech Forum. URL: https://arc.aiaa.org/doi/10.2514/6.2022-1915, doi:10.2514/6.2022-1915.

Qiao, W., Hamedi, M., Haghani, A., 2010. Algorithm for Crew-Scheduling Problem with Bin-Packing Features. Transportation Research Record 2197, 80–88. URL: https://doi.org/10.3141/2197-10, doi:10.3141/2197-10. publisher: SAGE Publications Inc.

Ranquist, E., 2017. Exploring the Range of Weather Impacts on UAS Operations, AMS. URL: https://ams.confex.com/ams/97Annual/webprogram/Paper309274.html.

Romero-Mancilla, M.S., Hernandez-Ruiz, K.E., Huerta-Muñoz, D.L., 2023. A multiobjective mathematical model for a humanitarian logistics multimodal transportation problem. Journal of Humanitarian Logistics and Supply Chain Management URL: https://doi.org/10.1108/JHLSCM-01-2023-0004, doi:10.1108/JHLSCM-01-2023-0004.

Schrimpf, G., Schneider, J., Stamm-Wilbrandt, H., Dueck, G., 2000. Record Breaking Optimization Results Using the Ruin and Recreate Principle. Journal of Computational Physics 159, 139–171. URL: https://www.sciencedirect.com/science/article/pii/S0021999199964136, doi:10.1006/jcph.1999.6413.

Scott, J.E., Scott, C.H., 2017. Drone Delivery Models for Healthcare, p. 8.

Sigari, C., Biberthaler, P., 2021. Medical drones: Disruptive technology makes the future happen. Der Unfallchirurg 124, 974–976. URL: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC8554499/, doi:10.1007/s00113-021-01095-3.

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, 2023. Earnings structure (All cities except London) | Stuart Help Center. URL: https://web.archive.org/web/20230615122923/https://help.stuart.com/en/articles/6999668-earnings-structure-all-cities-except-london.

The Fairwork Project, 2020. The Gig Economy and Covid-19:Looking Ahead. Technical Report. URL: https://fair.work/wp-content/uploads/sites/131/2020/11/COVID-19-Report-September-2020.pdf.

Thibbotuwawa, A., Bocewicz, G., Radzki, G., Nielsen, P., Banaszak, Z., 2020. UAV Mission Planning Resistant to Weather Uncertainty. Sensors 20, 515. URL: https://www.mdpi.com/1424-8220/20/2/515, doi:10.3390/s20020515. number: 2 Publisher: Multidisciplinary Digital Publishing Institute.

UK CAA, 2020. CAP722 Unmanned Aircraft System Operations in UK Airspace – Guidance.

UK Government, 2022. Greenhouse gas reporting: conversion factors 2022. URL: https://www.gov.uk/government/publications/greenhouse-gas-reporting-conversion-factors-2022.

Van De Vel, H., Shijie, S., 1991. An Application of the Bin-Packing Technique to Job Scheduling on Uniform Processors. The Journal of the Operational Research Society 42, 169–172. URL: https://www.jstor.org/stable/2583183, doi:10.2307/2583183. publisher: Palgrave Macmillan Journals.

Wang, Y., Wang, Z., Hu, X., Xue, G., Guan, X., 2022. Truck–drone hybrid routing problem with time-dependent road travel time. Transportation Research Part C: Emerging Technologies 144, 103901. URL: https://www.sciencedirect.com/science/article/pii/S0968090X2200314X, doi:10.1016/j.trc.2022.103901.

Wang, Z., Sheu, J.B., 2019. Vehicle routing problem with drones. Transportation Research Part B: Methodological 122, 350–364. URL: http://www.sciencedirect.com/science/article/pii/S0191261518307884, doi:10.1016/j.trb.2019.03.005.

Wilson, M.L., 1996. General Principles of Specimen Collection and Transport. Clinical Infectious Diseases 22, 766–777. URL: https://doi.org/10.1093/clinids/22.5.766, doi:10.1093/clinids/22.5.766.

Yu, X., Zhang, Y., Huang, K., 2013. Minimizing the Makespan for Scheduling Problems with General Deterioration Effects. Mathematical Problems in Engineering 2013, e218981. URL: https://www.hindawi.com/journals/mpe/2013/218981/, doi:10.1155/2013/218981. publisher: Hindawi.

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. URL: http://www.sciencedirect.com/science/article/pii/S0377221712008090, 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. URL: https://doi.org/10.1007/s11590-012-0578-1, 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. URL: https://doi.org/10.1007/s10479-019-03260-9, doi:10.1007/s10479-019-03260-9.

Özoğlu, B., Çakmak, E., Koç, T., 2019. Clarke & Wright's Savings Algorithm and Genetic Algorithms Based Hybrid Approach for Flying Sidekick Traveling Salesman Problem. Avrupa Bilim ve Teknoloji Dergisi , 185–192URL: https://dergipark.org.tr/en/pub/ejosat/issue/49069/637816, doi:10.31590/ejosat.637816. publisher: Osman SAĞDIÇ.

Anonymous, 2022. Details omitted for double-blind reviewing.

Anonymous, 2022a. Details omitted for double-blind reviewing.

Anonymous, 2023. Details omitted for double-blind reviewing.

Anonymous, 2023a. Details omitted for double-blind reviewing.

Appendix A. Additional Parameterisation Results Plots

Mean of instance results, scaled relative to worst performing in each test (worst = 100%).

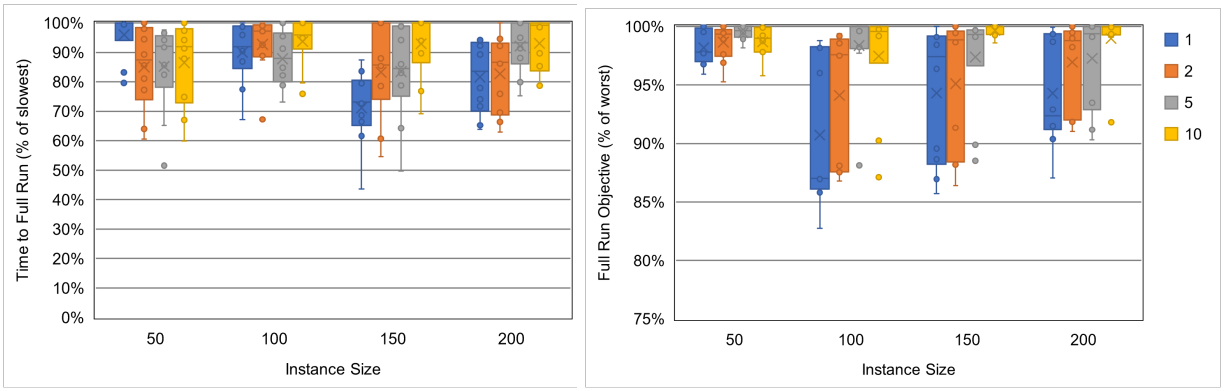

(a) Computational time after 10 reset kick iterations.

(b) Best objective value after 10 reset kick iterations.

Figure A.16: Relative performance comparison of the savings candidate shortlist length options (indicated by legend).

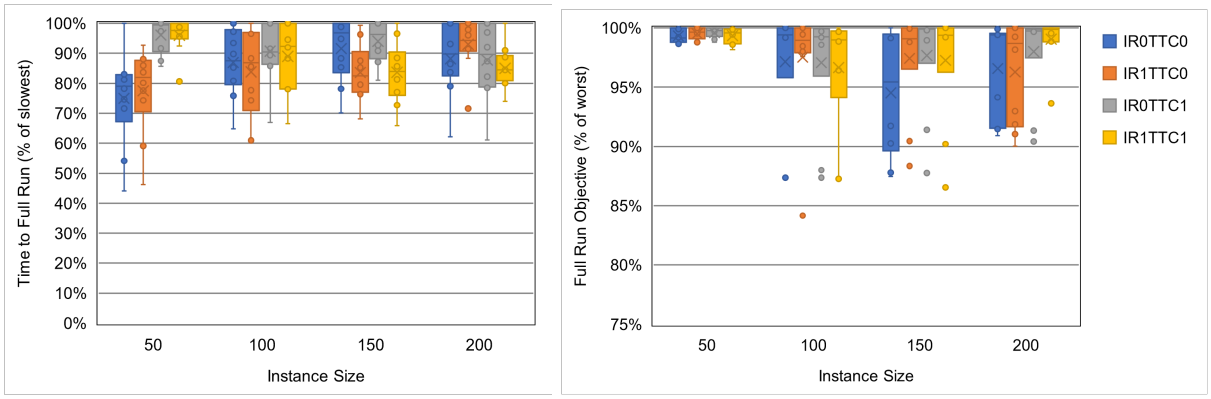

(a) Computational time after 10 reset kick iterations.

(b) Best objective value after 10 reset kick iterations.

Figure A.17: Relative performance comparison of the reallocation options. IR = inter-route, TTC = trunk to consolidator, 0 = reallocation off, 1 = reallocation on.

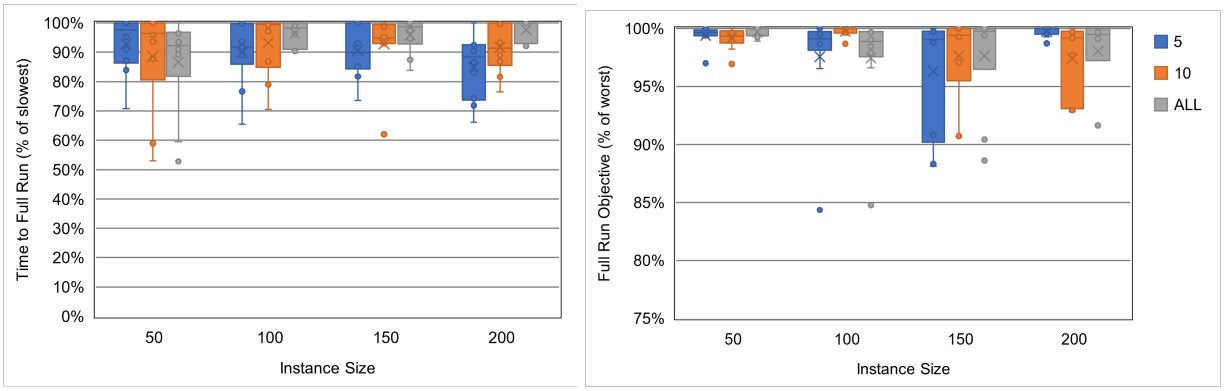

(a) Computational time after 10 reset kick iterations.

(b) Best objective value after 10 reset kick iterations.

Figure A.18: Relative performance comparison of the cycle catchment size options (limit indicated by legend).

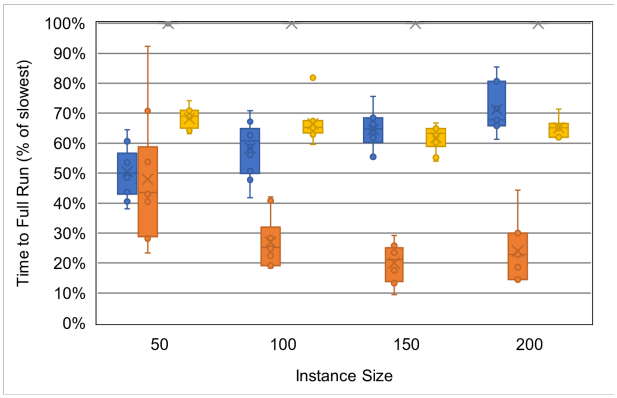

(a) Computational time after 500 reset kick iterations.

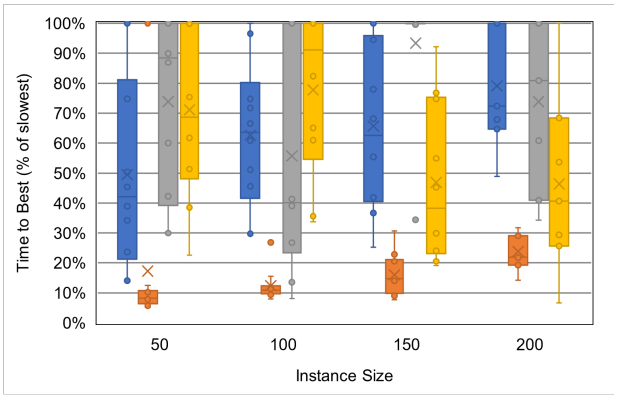

(b) Time taken to reach best objective value.

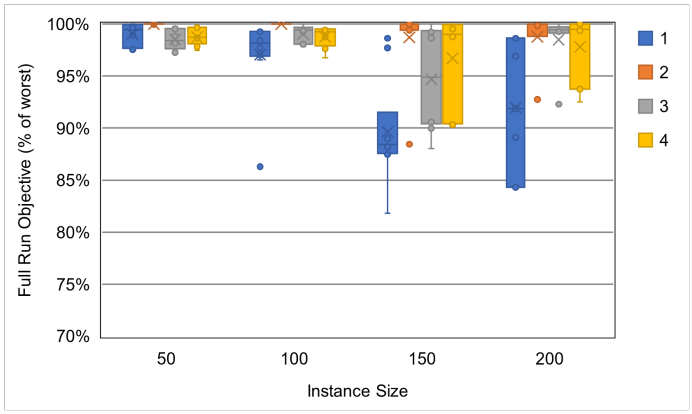

(c) Best objective value after 500 reset kick iterations.

Figure A.19: Relative performance comparison of the final algorithms. Legend indicates algorithm used, as defined in Table 6.

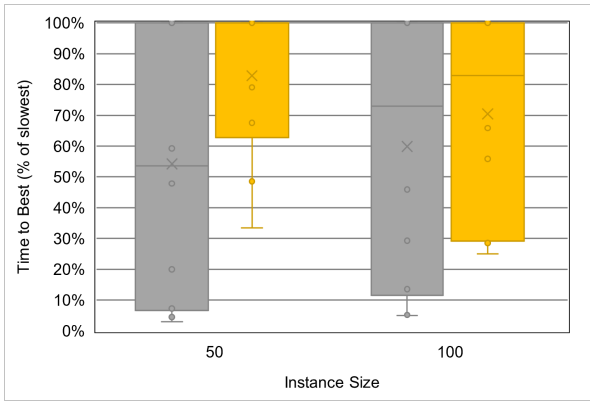

(a) Time taken to reach best objective value.

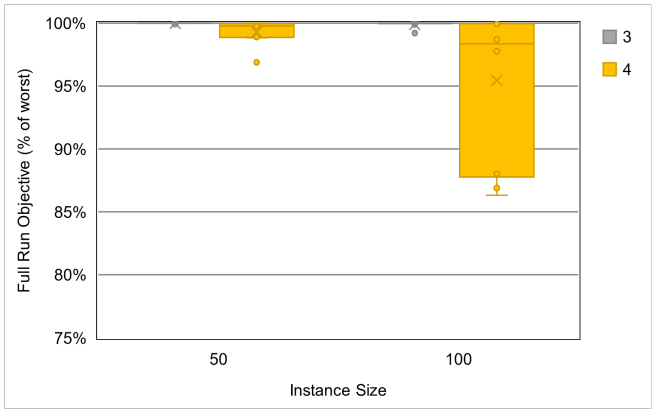

(b) Best objective value after 500 reset kick iterations.

Figure A.20: Relative performance comparison of Algorithms 3 and 4 under reduced UAV costs. Legend indicates algorithm used, as defined in Table 6.