Land-Air Logistics Optimisation

Andy Oakey

February 27, 2023

This chapter describes the problem and methodology for the carriage of patient diagnostic samples from community clinics to a laboratory for analysis, accounting for realistic time constraints and using practical assumptions. Our definition of land logistics encompasses diesel vans, electric vans and bikes, and air logistics encompasses aerial logistics drones.

1 Background

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

Samples typically have a fairly short time frame in which they must be analysed, generally within the day they are taken, and have specific requirements for storage and transportation (NHS and Sedman 2020). As a result, samples should be delivered to the hospital promptly to enable swift diagnosis and maximise the effective use of laboratory staff undertaking the analyses. The Royal College of Nursing in Wales (2020) identified that community COVID-19 testing was significantly slower than in-hospital testing, with less than a third of community test results being turned around the same day, as opposed to 80% in hospital. Previous studies and anecdotal discussions with hospital staff (McDonald 1972; Allan 2019) have suggested that there may be delaying factors within the surgeries and hospital (e.g. administrative, staff scheduling, etc.), but these are beyond the control of the logistics carrier and are ignored in this research. The period of greatest importance is the time spent travelling in the vehicle as controlled conditions cannot be guaranteed, unlike at the origin surgeries and destination laboratory (Anaya-Arenas et al. 2016); thus, strict constraints on in vehicle times across all surgeries served should be applied. Samples are currently collected as part of scheduled collection periods; a morning shift, and an afternoon shift. During these periods, scheduled sample collections calls are made to surgeries, regardless of whether a collection is required. One area being explored as part of this research is the effect of 'dynamic' collections, where only those sites that have samples to collect are served.

In addition to collection timing, there is an increasing need to reduce congestion and emissions in urban areas which contribute to poor air quality, slow transit times, and anthropogenic climate change (European Commission 2011), with policy makers often stating an aim to move to alternative, more sustainable transport modes (European Commission 2013). Health care providers are responsible for around 5% of the total national carbon dioxide footprint in developed nations (Pichler et al. 2019). Of this, 62% of this contribution can be attributed to medicines, medical equipment, and other supply chain sources (NHS 2020).
The National Health Service (NHS) in the UK has set a goal to be net-zero by 2040, and improving the efficiency of logistics operations will be key to achieving this (ibid.). To support this target,

changes to logistics systems are being explored such as mode-shift and adopting different supply-chain management strategies (NHS 2020). To this end, this research considers how best to integrate multiple modes, including vans, cargo bikes, and drones, into existing operations. Costs are an important factor to consider when implementing such changes, hence the main objective explored in this research is to minmise the cost of operations whilst considering the introduction of these alternative modes.

Other objectives are also explored in this research, including (i) minimising the maximum in-vehicle time; (ii) minimising the number of vans; (iii) minimising the total driving time; (iv) minimising the total fatality risk of the system (for all modes); (v) minimising the total energy requirement of the system; (vi) minimising the total emissions/pollutants resulting from the system. The emissions can also be priced, based on Department for Transport Transport Analysis Guidance (TAG) costings and included in the core cost analyses.

With the use of drones becoming a topic of significant interest in operations research literature, however there are significant practicalities that are neglected by the assumptions in many studies. These include the suitability of sites for drones to safely serve them, the effects of wind on drone trajectories, the effect of third party risk avoidance on drone trajectories, and accurate energy and emission modelling of all modes. Furthermore, cost assumptions are often overly optimistic and assume fully autonomous operations, or ignore key expenses such as the purchase and maintenance of the drone. This study seeks to use more realistic assumptions and apply them to real-world case studies.

2 Problem Description

The complete problem and solution space is described in this section.

2.1 Key Terminology

Surgery = A site producing samples that need to be collected.
Consolidation Surgery/Site = A surgery at which samples are consolidated to and transshipment occurs.
Hospital = The end destination where the analysis laboratory resides and to which all samples must be delivered.
Van = A Light Goods Vehicle (LGV) $<$3.5T Gross Weight (4.25T if electrically powered).
Drone = Uncrewed Aerial Vehicle.
Bike/Cycle = Pedal powered vehicle with limited range and capacity but potential speed advantages in high traffic areas.
Vehicle = One of: (i) a van; (ii) a drone; (iii) a bike.
Driver/Operator = The staff member operating a vehicle (vans), or set of vehicles (drones).
Shift period = The time window in which samples from a set of surgeries must be be collected and delivered to the hospital. No vehicle activity can happen outside of this period.
Constant travel condition period = A time interval within a shift period with constant traffic, weather, and ground risk such that travel times and paths/trajectories are uniform throughout the period.
Path/Trajectory = The series of streets/points in 3D space taken by a vehicle to get from one site to another.
Leg = An individual journey between two sites in a route.
Route = Continuous (i.e., no waiting) closed loop of multiple surgery stops, starting and finishing at the same location, departing at a specific time.
Trunk Route = A route used to service sites directly or to collect from consolidation sites, delivering to the hospital. **Consolidation Route** = A route used to service sites and deliver to a consolidation site for onward transfer by a trunk route. **Collection Round** = The combination of a trunk route and a set of consolidation routes that feed the trunk route.

2.2 Master Problem

The BAU activity presents a problem in which a set of known locations/nodes produce samples that need transporting to a single location/node in reasonable time without incurring excessive cost, congestion, or environmental impact. In this problem, we aim to plan a set of vehicle shifts to collect samples from a known set of surgeries and deliver them to a single hospital laboratory for analysis. The exact numbers of samples produced is not known prior to the collection. **Whilst historic data can suggest sample production rates/trends, these are heavily weighted around the existing van collection schedules, and any new system is likely to cause a shift in these timings regardless.**

A day is split into multiple shift periods which do not overlap (e.g. morning shift, afternoon shift, etc). This problem addresses the planning of an individual shift period in isolation. Surgeries produce samples that require collection and subsequent delivery within a single shift period. Time points in a shift period are discretised as a set of k discrete time points, $K = \{1, 2, 3, \ldots, k_n\}$. Traffic and weather conditions will vary throughout the time period, likely affecting travel times and the trajectories/paths taken; however, many points in K will have approximately equal traffic and weather conditions. Intervals with pseudo constant traffic and weather conditions are identified within K, with the start time points of the intervals being defined as $\{1, k_1, k_2, k_3, \ldots, k_{n-1}\}$. The start of the associated constant travel condition period for a given point, $k \in K$, is defined as k^*, where $k \geq k^*$ and $k < k^* + 1$. For example, the travel durations for a given O-D pair and mode will be uniform for all time points in k^* to $k^* + 1$, where $k^* + 1$ denotes the start of the proceeding time period. In the special case where $k^* = k_{n-1}$, then $k^* + 1 = k_n$, for the purposes of calculating travel times/durations. These features are visualised in Figure 1 and examples are given in Table 1.

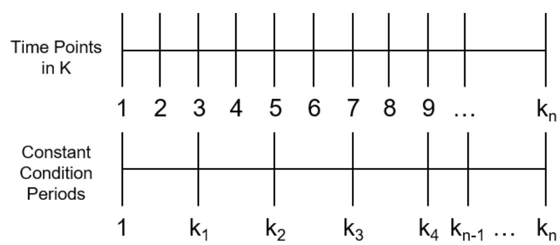

Figure 1: Time point notation visualised

Table 1: Time point and constant travel condition period examples.

Time Point	k^*	$k^* + 1$
1	1	k1
2	1	k1
3	k1	k2
4	k1	k2
k	k^*	$k^* + 1$

We define S as the set of surgeries that require a collection within the given shift period, and H as the Target Hospital to which samples are delivered and vans and drones are based. The service time required at each surgery visit is defined as σ, and captures the time spent completing the collection at each surgery, or delivering at the hospital. The set of all nodes, including the target hospital is defined as $S' = S \cup \{H\}$. Vans can also be based at a further subset of surgeries, $S^E \subseteq S$, though must only be used for purpose of consolidation back to the surgery they are based at for contractual and space reasons. All vans must finish at the site from which they are based.

Three modes are available in this problem, V, which represents a van; D, which represents a drone (UAV); and C, which represents a bike. Subsequently, the set of van routes based at the hospital is defined as R^V, the set of drone routes based at the hospital is defined as R^D, the set of bike routes based at any site is defined as R^C, and the set of van routes based at sites in S^E is defined as R^E.

The commercial running cost of a van per km is denoted by p^V and accounts for the van's variable vehicle costs (e.g. fuel, etc.). The commercial running cost of a drone per hour is denoted by p^D and accounts for the drone's variable vehicle costs (e.g. power, parts/servicing, etc.). A per hour cost is used based on the expected part lifespans of the platform being the main driver of the variable costs. Meanwhile, the commercial running cost of a cycle route is given by a fixed cost per job (route), p^J, and a per km variable cost, p^C, over a minimum distance C^l.

The travel time for a van journey between a pair of surgeries (i, j) departing at time k is defined as $t^V_{i,j,k}$. Due to variations in travel conditions, van travel times are approximated as a function of the travel times within the travel condition periods that the journey spans, where T^V_{i,j,k^*} is the travel time when departing at $k^* \in K$.

$$t^V_{i,j,k} = min(1, \frac{(k^*+1)-k}{T^V_{i,j,k^*}}) \times T^V_{i,j,k^*} + max(0, 1 - \frac{(k^*+1)-k}{T^V_{i,j,k^*}}) \times T^V_{i,j,k^*+1} \tag{1}$$

It is assumed that no journey between an O-D pair spans more than two travel condition periods, hence T^V_{i,j,k^*} is constant and no period beyond k^*+1 need be considered for each pair departing at time k.

At different time points, the chosen path/trajectory of a van may also vary; thus, the distance of the travelled path between a given O-D pair at time point k, $l^V_{i,j,k}$, is approximated in a similar way, where L^V_{i,j,k^*} is the travel distance when departing at $k^* \in K$.

$$l^V_{i,j,k} = min(1, \frac{(k^*+1)-k}{T^V_{i,j,k^*}}) \times L^V_{i,j,k^*} + max(0, 1 - \frac{(k^*+1)-k}{T^V_{i,j,k^*}}) \times L^V_{i,j,k^*+1}$$

Analogously, drone travel times and distances are approximated using the same formulation and are denoted by $t^D_{i,j,k}$ and $l^D_{i,j}$, respectively. Additionally, the energy requirement for a drone to travel between a pair of surgeries (i, j) at departure time k is defined as $\epsilon^D_{i,j,k}$ and is approximated in a similar manner, where ϵ^D_{i,j,k^*} is the travel energy requirement when departing at $k^* \in K$, and a drone's maximum energy capacity is denoted by D^ϵ.

$$\epsilon^D_{i,j,k} = min(1, \frac{(k^*+1)-k}{T^V_{i,j,k^*}}) \times \epsilon^V_{i,j,k^*} + max(0, 1 - \frac{(k^*+1)-k}{T^V_{i,j,k^*}}) \times \epsilon^D_{i,j,k^*+1}$$

Conversely, cycle travel durations are deemed to be independent of departure time, 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 $l^C_{i,j}$ $\forall k \in K$.

For practical reasons (e.g. landing space, staff resource, contracts), drones and cycles may be restricted such that they can only serve a select subset of surgeries. For drones the set of permitted sites is defined as $S^D \subseteq S'$, and for cycles $S^C \subseteq S'$. The set S^C only contains sites within large urban areas due to the service areas of gig-economy courier companies.

Let $v_k = (H, s_1, \ldots, s_{n_v}, H) \in R^V$ be a **trunk** van route departing at time k, where n_v denotes the number of surgeries being visited, $s_i \in S^V$, $\forall i \in \{1, \ldots, n_v\}$. Note that all **trunk** vans are based at the Hospital, H, and must return to H.

Each **trunk** van route has an associated distance, denoted by l_{v_k} and is calculated by summing the distances of the constituent legs of the route, each departing at time k_s, where s denotes the departure surgery:

$$l_{v_k} = l^V_{H,s_1,k} + \sum_{i=1}^{n_v-1} l^V_{i,i+1,k_{s_i}} + l^V_{s_{n_v},H,k_{s_{n_v}}}$$

Each **trunk** van route also has an associated time, denoted by t_{v_k} and is calculated by summing the durations of the constituent legs of the route, each departing at time k_s, where s denotes the departure surgery:

$$t_{v_k} = t^V_{H,s_1,k} + \sum_{i=1}^{n_v-1} t^V_{i,i+1,k_{s_i}} + t^V_{s_{n_v},H,k_{s_{n_v}}} + n_v\sigma$$

The departure times (k_s) and durations/distance of the route's constituent legs are derived from the departure time of the route such that no waiting time is incurred; e.g. a route departing at time k

will have a first leg departure k, and a second leg departure $k + t^V_{H,s_1,k}$, and so on. Waiting time for any reason (e.g. arriving before consolidation rounds have finished, or waiting for traffic changes) is not permitted. Furthermore, σ is embedded in t_{v_k} and the variable running cost of a given van route is calculated as $p_{v_k} = p^V l_{v_k}$.

Similarly, let $e^s_k \in R^E$ be a **consolidation** van route departing at time k, where each van route starts and finishes at $s \in S^E$. The distances and timings of the consolidation van routes are denoted by $l_{e^s_k}$ and $t_{e^s_k}$, respectively, and are calculated in the same way as trunk van routes, but with the start and finish at surgery s.

Analogously to trunk vans, a **trunk** drone route is defined as $d_k \in R^D$. Further to the definition detailed for van routes, drones are restricted due to regulation and capacity, and can only serve one surgery per collection; thus, $n_d = 2 \ \forall d_k \in R^D$, where the number of surgeries visited is equal to n_d denotes the number of surgeries being visited. Furthermore, the sites visited must be in the list that permit drone service; $s_i \in S^D$, $\forall i \in \{1, \ldots, n_d\}$. Note that all drones are based at the Hospital H and must return to H. On arrival back at the hospital drones and electric vans are also subject to a downtime duration, B, in addition to the embedded service time, to allow for battery changes/charges to take place prior to their next departure. Each **trunk** drone route also has an associated distance and time, denoted by l_{d_k} and t_{d_k}, respectively, and the variable running cost of a given drone route is calculated as $p_{d_k} = p^D t_{d_k}$

Furthermore, the energy requirement to perform drone route is denoted by ϵ_{d_k}. ϵ_{d_k} is calculated as the sum of the energy requirements between the surgeries visited $\epsilon_{d_k} = \epsilon^D_{H,s_1,k} + \epsilon^D_{s_1,H,k_{s_1}}$. To ensure that a drone route can be safely undertaken, an energy limit is applied with a safety factor (FoS); $\epsilon_{d_k} \leq FoS \cdot D^{\epsilon}$.

Likewise, we define a **consolidation** bike route based at surgery s as $c^s_k \in R^C$ with a departure at time point k. The number of surgeries being visited is denoted by n_c, and given capacity constraints for cyclists (maximum load equal to three surgeries' worth of samples), we constrain routes in R^C such that $n_c \leq 4$. It should be highlighted that the bike route should start and end at the same surgery $s \in S^C$.

Each bike route also has an associated distance and time, denoted by $l_{c^s_k}$ and $t_{c^s_k}$, respectively. Additionally, cycle routes are subject to a time constraint of $t_{c_{max}}$ or less to ensure the cycle routes can be managed as discrete gig-economy tasks; $t_{c^s_k} \leq t_{c_{max}}$. The variable running cost of a given cycle route is calculated as $p_{c_k} = p^J + p^C(l_{c_k} - C^l)$

Surgeries can be served by:

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

Hence, we define two supersets of routes: $R' = \{R^V, R^D\}$ which contains all of the trunk routes (drones and vans originating at H); and $R'' = \{R^C, R^E\}$, which contains all of the consolidation routes (bikes and vans). An individual trunk route (drone or van route starting at H) is denoted by $\alpha_k \in R'$, and an individual consolidation route (local van or cycle route) based at s is defined as $\beta^s_\delta \in R''$. The notations associated with α_k and β^s_δ are analogous to the original route definitions, with δ denoting the departure time of the consolidation route.

A **collection round**, $\bar{r}_k = (\alpha_k, R''_{\alpha_k})$, is defined as the combination of a single trunk route $\alpha_k \in R'$ with a subset of consolidation routes $R''_{\alpha_k} \subseteq R''$ such that for any given consolidation route $\beta^s_\delta \in R''_{\alpha_k}$ based at surgery s, it is satisfied that $s \in \alpha_k$, i.e, any consolidation route in R''_{α_k} is based at a surgery that is being visited by trunk route α_k. Furthermore, other than surgery s where each β^s_δ begins/ends, the consolidation routes in R''_α do not share any other surgeries; i.e. no surgery is served by multiple consolidation routes. We define the set of all collection rounds as $\overline{R}$, and each collection round is denoted by $\overline{r}_k = \{(\alpha_k, R''_{\alpha_k})|\ \alpha_k \in R',\ R''_{\alpha_k} \subseteq R''\}$, $\overline{r}_k \in \overline{R}$.

Note that, since R''_{α_k} could be empty, then it can satisfied that $R' \subseteq \overline{R}$. 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 can occur in a collection round; consolidation to a surgery, followed by trunking to the hospital. Chained consolidation (e.g. bike-local van-bike-trunk van) is not permitted.

Bike routes that start and end at the hospital are permitted in order to serve the catchment area of the hospital directly. To account for this in the model formulation, a dummy van route, $v_k^0 \in R^V$, is created. Starting and ending at the hospital, with no intermediate stops ($v_k^0 = (H, H)$) and a travel time of zero ($t_{v_k^0} = 0$), v_k^0 enables a collection round where surgeries which are cycle served only, $\overline{r}_k^0$. These routes can depart at any time within the collection 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''_{α_k} being complete before the trunk route, α_k, serves the consolidation site, i.e. they do not incur waiting time for the trunk route.

Some example collection round scenarios are given in Table 2 to demonstrate different arrangements of collection rounds.

Table 2: Example collection round scenarios, illustrating potential offsets. Offsets only demonstrated at the first stop, other delays may occur. All consolidation routes deliver to first stop in these examples unless otherwise stated. CR = Collection Round

Scenario	Trunk Route in CR, Dep. Time	Consolidation Routes in CR, $t_{\beta_\delta^s}$	Trunk Time to First Stop	Latest Departure(s) for Compatible Consolidation Routes	Trunk Route Duration	CR Duration
Van + 1 bike, trunk offset	(H,s_1,s_2,s_3,H), 10:00	(s_1,s_4,s_5,s_1), 25 mins	20 mins	09:55	80 mins	85 mins
Drone + 2 bikes, no trunk offset	(H,s_1,s_2,s_3,H), 10:00	(s_1,s_5,s_1), 15 mins; (s_1,s_4,s_1), 10 mins	20 mins	10:05, 10:10	80 mins	80 mins
Drone + 1 local van, trunk offset	(H,s_1,s_2,s_3,s_4,H), 10:00	(s_1,s_5,s_6,s_7,s_1), 30 mins	20 mins	09:50	70 mins	80 mins
Van + no bike, no van delay	(H,s_6,s_7,s_8,H), 10:00	-	30 mins	0 mins	85 mins	85 mins
Bike direct to hospital	(H,H), 10:00	(H,s_9,s_{10},H), 23 mins	0 mins	09:37	0 mins	23 mins

The commercial variable running costs of operating a given collection round is denoted by $p_{\overline{r}_k}$, and is a sum of the running costs ($p_{v_k}, p_{d_k}, p_{c_k}, p_{e_k}$), taken from the respective routes featured in the collection round, $p_{\overline{r}} = p_{\alpha_k} + \sum_{\beta_\delta^s} \in R''_\alpha p_\beta$.

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

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

We also define the integer variable $V_{\overline{r}_k}$ for each collection round, denoting the number of van routes (trunk and consolidation) being used in a collection round. Likewise, $D_{\overline{r}_k}$ defines the number of drones (trunk) being used in a collection round. Vans and drones 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 a the entire shift period (regardless of how long they are driving routes). Similarly, each drone is monitored by an operator who is also paid for the entire shift period, but can monitor up to δ drones 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 drones used throughout the entire shift period are given by:

$$A^V_{max} \geq \sum_{v_k \Rightarrow k, v_k \neq v^0_k} V_{\bar{r}_k} x_{\bar{r}_k} \ \ \forall k \in K$$

$$A^D_{max} \geq \sum_{d_k \Rightarrow k} D_{\bar{r}_k} x_{\bar{r}_k} \ \ \forall k \in K$$

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

$$A^O_{max} \geq \frac{A^D_{max}}{\delta}$$

$$A^O_{max} < \frac{A^D_{max}}{\delta} + 1$$

The standing cost of a van and driver (combined) per shift period is given by W^V, whilst the standing costs of a drone and operator per shift (separate) are given by W^D and W^O, respectively. The costs of vehicle fixed costs such as insurance, maintenance, etc. are included in W^V and W^D, whilst the labour cost for the delivery window is included in W^V and W^O. Due to the contractual arrangement of the drivers/operators, it is assumed that pay is guaranteed for the shift period. Hence, the fixed costs of a solution are denoted by $W^V A^V_{max}$ and $W^D A^D_{max} + W^O A^O_{max}$ for vans and drones, 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 shifts) within a given time window.

The monetised indirect costs from the collection system (e.g. greenhouse gas emissions, pollutants) are calculated as a function of the selected routes and arrival times of samples and are denoted by p^z.

The objective of this problem is to minimise the sum of the operating costs and indirect costs for a shift period. To this end, a multi-term objective function is used to sum the cost of the selected routes, standing costs of each mode, and the indirect costs (Equation 2); meanwhile, the constants θ_1, θ_2, θ_3, and θ_4 are used to balance the different each elements of the cost and allow weighting towards different modes or the indirect costs should a decision maker wish to favour a particular mode. It should be noted that in the event that the indirect costs are not known, the weighting can be set to zero.

The problem is constrained such that collection round durations are subject to time constraint of $t^{max}_{\bar{r}}$ minutes or less (typically set to 90 minutes maximum (McDonald 1972)) to ensure timely delivery of samples (Equation 3). A given collection round duration is calculated as the maximum time between the first collection in any constituent round in $\bar{r}_k$ and delivery to the hospital. If cycle consolidation routes are used in $\bar{r}_k$, 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, meaning the problem is optimisation focused, rather than decision. It should be noted that it is possible for a site to be served more than once by a given mode to ensure time constraints are met without relaxation. The constraints relating to collection round durations must be maintained.

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

$$Min: \theta_1 \sum_{\overline{r}_k} x_{\overline{r}_k} p_{\overline{r}_k} + \theta_2 W^V A^V_{max} + \theta_3 W^D A^D_{max} + W^O A^O_{max} + \theta_4 p^Z \tag{2}$$

$$t_{\overline{r}} \leq t^{max}_{\overline{r}} \tag{3}$$

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

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

3 Solution Approach

The solution algorithm used is an adaptation of the Clarke and Wright Savings Algorithm (Clarke and Wright 1964) (hereafter C&WSA). The main change to the algorithm is that multiple savings options, including the combining of routes, are tested. The initial solution is constructed of multiple routes, with each site within range of the collection round duration being visited by one van route (i.e. one route per site). Those sites beyond the range of the collection round limit are subsequently processed as drone routes or consolidation routes (as a connecting cycle or van), connecting to one of the other van rounds.

The constructed routes use the maximum duration for all arcs in the route based on the time matrices for the shift period, giving an upper bound. With each iteration (including the initial solution), the potential route options are compiled into vehicle shifts based on an adapted best-fit bin packing algorithm. The packing algorithm, hereafter referred to as the adapted best-fit (ABF), selects routes based on the percentage difference between the maximum route duration (i.e. the constructed route), and the duration at the the tested position (the assigned route). Those with a greater difference/improvement over the upper bound are fitted first, encouraging better positioning of more variable routes. Consolidator routes are positioned adjacent to their associated trunk route with no flexibility to change their timing.

To improve computational speed, the proposed savings offered by implementing each change are stored for later recall. Where several routes may have changed in previous iterations, if a stored change is recalled, the realised saving may be less than initially calculated due to time constraints when compiling new shifts. To prevent unexpected deviations from realised savings, a first-fit decreasing packing algorithm (hereafter FFD) is applied to estimate savings, using the constructed route durations only (i.e. based on the upper bound, no recalculation of times). If estimated savings are greater than 5% different from the stored saving, the actual saving is recalculated.

In each iteration of the adapted savings algorithm, three different improvement options are tested: (i) combining routes of the same mode, as in the original C&WSA; (ii) introducing a drone route in place of a van route stop, eliminating the stop from the existing route (or eliminating the route if it is the only stop); and (iii) transferring from a trunk route to a consolidating route (e.g. from a van to a bike that consolidates to a van route). In the trunk to consolidation options, if a surgery is already used as a consolidation site, it cannot be reallocated to a consolidation route.

To reduce the greedy nature of the initial algorithm, a shortlist of the best improvement options was used in each iteration, with a random item being carried forward. Should a change be sufficiently beneficial, it will reoccur for multiple iterations and likely be selected. A local search is also subsequently applied to investigate: (i) a 2-opt algorithm within each route; and (ii) transferring sites between routes.

The use of a kick is also being tested, whereby a given set of sites is removed from the solution and re-introduced as their Hospital-Surgery-Hospital drone routes where possible, or Hospital-Surgery-Hospital van routes if drones are not permitted at that site. The adapted savings algorithm and local searches are then reapplied. Sites are either chosen at random as a given percentage of all sites, or are taken from the longest, or the shortest routes in the best solution. The inter-route local search could also be applied as part of the saving options, with each move being one option, or a full search of all sites being the option considered in the shortlist.

Multiple visits to the same site is not considered by the algorithm. Whilst this does limit the potential ability to combine consolidation van rounds into single shifts, it reduces the solution search space and computational time. This is unlikely to cause significant decrements in solution quality due unless the algorithm is used in a case study area where it is heavily dependant on consolidation (e.g. very long trunk stem mileage). If consolidation van costings are not fixed to the full shift model, their use may also increase, because shift compiling is less of a driving factor in their cost.

4 Integrating With eDrone Modelling Suite

To enable more realistic modelling, the Land-Air Logistics Optimiser integrates with the wider eDrone modelling suite of tools. After a given problem is initially imported, including the sites requiring collection, the destination, available modes, and cost/time constraints/parameters, an enumerated set of O-D pairs for all constant time points is generated for vans and drones. For bikes, only a single time point is required due to cycling times not being affected by traffic.

Further realistic constraints can also be applied at this point, such as gig-economy availability/service areas, captured via the Stuart (gig-economy delivery operator) API (`api-docs.stuart.com`).

The land energy logistics tool (developed by JK) is then queried, using the Google Routing API to capture the routes and travel times/distances/speeds. The emissions and energy analyses are subsequently applied, returning the time, distance, emissions factors, energy requirements, and trajectory for the O-D pair at a given time. For O-D pairs where cycling is permitted, the process is repeated, querying a bike route. Emissions modelling is not required for this mode.

If an O-D pair can be served by a drone, each possible time point is queried using the Air Energy and Risk Route Optimisation (AERRO) tool (developed by AB/AP/JK), identifying a drone trajectory between the points that is optimised in terms of both energy and risk. As with the land energy logistics tool, the AERRO tool returns a time, distance, emissions factors, energy requirements, and trajectory for the O-D pair at a given time. If a drone cannot be used at a given time due to adverse weather, the tool will return a failure.

The data is subsequently fed into the LALO tool, where the Adapted Clarke and Wright Savings Algorithm with the Adapted Best-Fit Packing Algorithm is used to identify a solution.
The objectives of the algorithm can be changed to optimise towards alternative objectives, as detailed in Section 1

References

Allan, Rick (Dec. 2019). *Southampton Pathology Operations University Project Meeting.*

Anaya-Arenas, Ana Maria et al. (Jan. 2016). "Biomedical sample transportation in the province of Quebec: a case study". In: *International Journal of Production Research* 54.2. Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/00207543.2015.1018455, pp. 602–615. ISSN: 0020-7543. DOI: 10.1080/00207543.2015.1018455. URL: https://doi.org/10.1080/00207543.2015.1018455 (visited on 12/22/2021).

Cherrett, Tom and Amy Moore (Jan. 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.* Paper 20-04349. 12-16 January, Washington D.C.

Clarke, G. and J. W. Wright (Aug. 1964). "Scheduling of Vehicles from a Central Depot to a Number of Delivery Points". In: *Operations Research* 12.4. Publisher: INFORMS, pp. 568–581. ISSN: 0030-364X. DOI: 10.1287/opre.12.4.568. URL: https://pubsonline.informs.org/doi/abs/10.1287/opre.12.4.568 (visited on 02/27/2023).

European Commission (Mar. 2011). *White Paper: Roadmap to a Single European Transport Area – Towards a competitive and resource efficient transport system.* URL: https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:52011DC0144&from=EN (visited on 11/05/2019).

— (Dec. 2013). *A call to action on urban logistics.* URL: https://ec.europa.eu/transport/sites/transport/files/themes/urban/doc/ump/swd%282013%29524-communication.pdf (visited on 11/22/2019).

Ford, Megan (July 2020). *Covid-19 testing 'much slower' in community, says Welsh RCN.* en. URL: https://www.nursingtimes.net/news/coronavirus/covid-19-testing-much-slower-in-community-says-welsh-rcn-24-07-2020/ (visited on 01/29/2021).

McDonald, J. J. (1972). "Vehicle Scheduling – A Case Study". In: *Operational Research Quarterly (1970-1977)* 23.4. Publisher: Palgrave Macmillan Journals, pp. 433–444. ISSN: 0030-3623. DOI: 10.2307/3007958. URL: https://www.jstor.org/stable/3007958 (visited on 02/01/2021).

Ngo, Andy, Paras Gandhi, and W Greg Miller (Jan. 2017). "Frequency that Laboratory Tests Influence Medical Decisions". In: *The Journal of Applied Laboratory Medicine* 1.4, pp. 410–414. ISSN: 2576-9456. DOI: 10.1373/jalm.2016.021634. URL: https://doi.org/10.1373/jalm.2016.021634 (visited on 12/23/2021).

NHS (Oct. 2020). *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 (visited on 04/14/2021).

NHS and David Nixon (2015). *Transportation of Specimens Procedure.* URL: http://www.sthk.nhs.uk/about/Documents/11615_TransportationofSpecimensProcedureTrustDocument.pdf (visited on 01/29/2021).

NHS and Rebecca Sedman (Apr. 2020). *Pathology Specimen Transport.* URL: https://www.cddft.nhs.uk/media/615106/transport%20sop.pdf (visited on 01/29/2021).

Pichler, Peter-Paul et al. (May 2019). "International comparison of health care carbon footprints". en. In: *Environmental Research Letters* 14.6. Publisher: IOP Publishing, p. 064004. ISSN: 1748-9326. DOI: 10.1088/1748-9326/ab19e1. URL: https://iopscience.iop.org/article/10.1088/1748-9326/ab19e1/meta (visited on 04/14/2021).