

The Aerial Firefighting Vehicle Routing Problem

Wissam Chalabi
Aviation Operations
Royal Netherlands Aerospace Center
Amsterdam, Netherlands
wissam.chalabi@nlr.nl

Alessandro Bombelli
Faculty of Aerospace Engineering
Delft University of Technology
Delft, Netherlands
a.bombelli@tudelft.nl

Glen Purdam
Fixed wing aerial firefighting pilot
Self-employed contractor
Vestby, Norway
glen.purdam@gmail.com

Abstract—Optimizing the routes of firefighting aircraft can lead to better containment of wildfires, hence yielding great environmental and societal value. In this paper, a novel formulation of the Vehicle Routing Problem (VRP) is customized to address the needs of Aerial Firefighting (AFF). The resulting formulation, named Aerial Firefighting Vehicle Routing Problem (AFFVRP), is a capacitated multi-trip VRP with time-windows and hierarchical objectives. The primary objective is to minimize the time of carrying out all requested water drops (to extinguish wildfires quicker), and the secondary objective is to minimize the total flight time. The multi-trip nuance is adopted to be able to model different aircraft types that might require to revisit the depot for after refueling. Because the model is intended to operate as a decision-making tool to support firefighters, users can input the number and types of aircraft available, the location of the airfield, fires, nearest water body, intensity of each fire, etc. Several random cases and case studied based on real wildfires were solved within the expert-recommended time limit of 5 minutes, yielding good-quality solutions in terms of gap optimality. The problem is scalable and sizes ranging from one to 80 water drops were tested and solved within 22 minutes. Strategic fleet planning is also demonstrated in a case study with the use of Monte Carlo simulation, in order to compare the performance of different fleet options for a given setting. Therefore, the model is not only applicable in live situations, but can also be used as a supportive tool in planning for upcoming fire seasons, or reviewing and learning from past fires.

Keywords—Aerial Firefighting; Vehicle Routing Problem; Min-max problem; Time-windows; Wildfires

I. INTRODUCTION

With the escalating impact of climate change, the frequency and intensity of forest fires are on the rise, influenced by shifts in temperature, precipitation patterns, and vegetation [1]. The global surge in wildfires is evident, leading to extensive destruction of forested regions. Figure 1 illustrates the increased burnt area in European Union countries in 2023 compared to the average in the 2006-2022 period, highlighting the increasing urgency of the issue. This study aims to aid firefighting authorities in mitigating wildfires by optimizing the routing of aerial firefighting aircraft and hence acting as a decision-making tool supporting human operators. This research paper presents a novel adaptation of a Vehicle Routing Problem (VRP) to Aerial Firefighting (AFF), namely the Aerial Firefighting Vehicle Routing Problem (AFFVRP). VRPs adaptations to address disaster

relief are not novel in the literature. We refer readers to [2] for an example focused on crowdsourcing or [3] for an example focused on post-disaster relief distribution. Notwithstanding, to the best of the authors' knowledge, no existing work addresses the unique needs and features of AFF wildfire relief missions. The model considers the wildfires to be known in advance and, given their intensity, a fire-specific amount of water needed to extinguish them is computed. Hence, the primary objective of the model is to minimize the time of carrying out all requested water drops. In AFF operations, high-capacity aircraft (tankers) can be used that are filled with water at the airbase and must discharge all of it at every drop. Conversely, lower-capacity aircraft (scoopers) can also be adopted that can collect water from water bodies and hence do not need to return to the airbase after every drop. To allow tankers to perform several drops, a multi-trip variant of the VRP was adopted. The model has been tested both with random scenarios and with scenarios based on real wildfires. Because of the operational nature of the problem at hand, a maximum computational time of 5 minutes was adopted. Routing solutions across all instances were of good algorithmic quality and were also validated via expert judgment.




Figure 1. Burnt area (ha) in EU countries in 2023 compared to the 2006-2022 average [4].

We believe the AFFVRP contributes to the existing body of literature on optimization-based models addressing disaster relief. Furthermore, it serves as a first step toward the development of decision-making tools to assist AFF human operators and

reduce the environmental and societal impact of wildfires. This paper is structured as follows. The approach and underlying assumptions are presented in section II. section III presents the mathematical formulation of the AFFVRP. The results are discussed in section IV, while conclusions and recommendations are discussed in section VI.

II. MODELING APPROACH

In this section, we define and discuss the main modeling assumptions and simplifications of the AFFVRP.

A. Network representation

The AFFVRP is defined on a directed graph $G = (\mathcal{N}, \mathcal{A})$, where \mathcal{N} is the set of nodes and \mathcal{A} the set of arcs. Three node types are defined: the airfield, the water body, and fires. Note that, with fire nodes, water drop locations are meant. While most water drops are carried out to extinguish fires, sometimes water drops might be performed to lay lines ahead of the fire front. Notwithstanding, we define and model both occurrences as *fire* nodes. The airfield is assumed to be the location of take-off and landing of all available aircraft, and the water body is the nearest water body suitable for water refilling operations. These assumptions represent well real firefighting operations according to expert feedback. Exceptions could be that multiple water bodies are used, if there are several suitable ones near the fire. Multiple airfields can be employed for two key reasons. First, to mitigate congestion by leveraging available alternatives, and second, in scenarios where extensive wildfires demand a substantial fleet, necessitating operations from more than one airfield. This last variation, which is not modeled in the work presented here, bears a similarity with the multi-depot variants of the VRP.

B. Types of aircraft

As anticipated in section I, in this work, we consider two types of aircraft, namely scoopers and tankers. Scoopers can pick up water from a body of water, and continue to perform subsequent requested drops, while tankers must return to the airfield after a drop and refill before they can head out again to another fire (albeit being faster and able to carry more water). Because cruise speed and water capacity are parameters to the model, the AFFVRP can function with a wider variety of vehicles. For example, helicopters can be represented as a scooper variant.

C. Fire intensity

Fire intensity, for the purposes of this research, is defined as the number of scooper drops (because of their lower water capacity) required to extinguish a fire.

D. Subfires

In VRP formulations, a common constraint is that each node is visited only once. An alternative approach, the Split Delivery VRP [5], allows a node to be visited by multiple vehicles so that its demand can be split among them. In the AFFVRP, fires might require more than one water drop to be extinguished. Instead of employing a split delivery approach, we decided to segment each fire into smaller components known as subfires. For instance, a fire with intensity 5 (hence, five scooper water drops are required to extinguish it) can be divided into five subfires, each with a unitary intensity and represented by a distinct node. These nodes, clustered in close proximity, denote the drop locations for that fire, with each having a demand matching a scooper's capacity. A downside of such an approach is that it increases the cardinality $|\mathcal{N}|$ of the set of nodes with respect to a split delivery framework.

E. Tanker multi-trips and subfire drops

Tankers cannot scoop up water from a water body. Once they make a drop, they must return to the airfield to refill the tank with water. The planning horizon for the model assumes a single trip for the scoopers (although with multiple drops), but it does provide the option of multiple trips for tankers. The water capacity of tankers exceeds that of scoopers, meaning a single drop on one subfire might be wasteful. For instance, if the tanker's capacity is 10 units and the demand of each subfire is 5 units, the tanker would waste half its capacity by dropping water on only one subfire. To resolve this, tankers are allowed to make several drops as long as their capacity allows, but only across subfires mapping the same fire.

F. Scooper time matrix and water bodies

Scoopers must visit a water body to refill between consecutive drops. While the sequence of subfires to be visited by each scooper is unknown beforehand, the water body location is a known parameter. Hence, when computing the flight time between subfires for a scooper, the following logic is adopted. For the time matrix of scoopers, the distance between two drop locations is calculated as the distance between the first location and the water body plus the distance from the water body to the second location. Hence, the necessity of visiting the water body is accounted for without unnecessarily complicating routing decisions in the mathematical model.

G. Time windows

Time windows are used to indicate the priority of different fires. A fire does not adhere to a time-window in the traditional sense as in delivery service-relayed VRPs. However, operators might need to prioritize a specific fire, for instance, because it is near a residential area or poses a high risk of fast spreading due to wind patterns. In this case, a time-window can be used to set an earlier upper bound on the maximum time of delivering a

drop, forcing the model to prioritize such a fire over others. The lower bound for all fires is always set to zero, as there is no disincentive to extinguishing fires as early as possible.

H. Hierarchical objectives

Two objectives are identified based on expert advice. Firstly, all water drops must be carried out as soon as possible. Secondly, the total flight time should be minimized. The first objective is more important than the second because it is tied to the success of the mission in terms of safety and disaster relief. The second objective is desirable by operators, to minimize the operation costs. A hierarchical structure is used where the mathematical model is initially solved only accounting for the first objective. Then, the second objective is optimized under the result of the first. In the primary objective, the time of performing the last drop needs to be minimized. However, the time when the drops are carried out is only computed as a result of the model. A min-max approach is then employed via the definition of an auxiliary decision variable that allows to compute, across all sets of feasible routings, the one that minimizes the time when the latest drop is carried out.

III. METHODOLOGY

The AFFVRP is modeled as a Mixed Integer Linear Program (MILP) model. As such, the sets, parameters, and decision variables defining the MILP are defined in Table I., Table II., and Table III. respectively. For the sake of readability, we employ a calligraphic style for sets, upper-case letters for parameters, and lower-case letters for decision variables.

TABLE I. SETS OF THE AFFVRP.

Set	Description
\mathcal{H}	Set of original fires, indexed by h
\mathcal{F}	Set of all subfires, indexed by i or j
\mathcal{F}_h	Set of subfires associated with original fire $h \in \mathcal{H}$, indexed by i or j
\mathcal{P}	Set of scoopers, indexed by p
\mathcal{K}	Set of tankers, indexed by k
\mathcal{U}	Set of tanker trips, indexed by u . The first trip is denoted by 0
\mathcal{N}	Set of all nodes, indexed by i or j . The airfield is denoted by 0
\mathcal{A}	Set of all arcs: $\mathcal{A} = \{(i, j) \in \mathcal{N} \times \mathcal{N} : i \neq j\}$
$\mathcal{A}_{\mathcal{K}}$	Set of arcs a tanker can fly
$\mathcal{A}_{\mathcal{P}}$	Set of arcs a scooper can fly

The mathematical formulation of the AFFVRP is:

$$\min Z \quad (1)$$

$$\min \sum_{k \in \mathcal{K}} \sum_{u \in \mathcal{U}} \sum_{(i, j) \in \mathcal{A}_{\mathcal{K}}} T_{ij}^k x_{ij}^{ku} + \sum_{p \in \mathcal{P}} \sum_{(i, j) \in \mathcal{A}_{\mathcal{P}}} T_{ij}^p x_{ij}^p \quad (2)$$

s.t.:

$$\tau_j^{ku} \leq Z \quad \forall k \in \mathcal{K}, u \in \mathcal{U}, j \in \mathcal{F} \quad (3)$$

$$\tau_j^p \leq Z \quad \forall p \in \mathcal{P}, j \in \mathcal{F} \quad (4)$$

$$\sum_{k \in \mathcal{K}} \sum_{u \in \mathcal{U}} \sum_{i \in \mathcal{N} \setminus \{j\}} x_{ij}^{ku} + \sum_{p \in \mathcal{P}} \sum_{i \in \mathcal{N} \setminus \{j\}} x_{ij}^p = 1 \quad \forall j \in \mathcal{F} \quad (5)$$

$$\sum_{j \in \mathcal{F}} x_{0j}^{ku} \leq 1 \quad \forall k \in \mathcal{K}, u \in \mathcal{U} \quad (6)$$

$$\sum_{j \in \mathcal{F}} x_{0j}^p \leq 1 \quad \forall p \in \mathcal{P} \quad (7)$$

$$\sum_{j \in \mathcal{F}} x_{0j}^{k, u+1} \leq \sum_{j \in \mathcal{F}} x_{j0}^{ku} \quad \forall k \in \mathcal{K}, u \in \mathcal{U} \setminus \{0\} \quad (8)$$

$$\sum_{j \in \mathcal{F}} x_{0j}^p - \sum_{j \in \mathcal{F}} x_{j0}^p = 0 \quad \forall p \in \mathcal{P} \quad (9)$$

$$\sum_{i \in \mathcal{N} \setminus \{j\}} x_{ij}^p - \sum_{i \in \mathcal{N} \setminus \{j\}} x_{ji}^p = 0 \quad \forall j \in \mathcal{F}, p \in \mathcal{P} \quad (10)$$

$$\sum_{i \in \mathcal{N} \setminus \{j\}} x_{ij}^{ku} - \sum_{i \in \mathcal{N} \setminus \{j\}} x_{ji}^{ku} = 0 \quad \forall k \in \mathcal{K}, u \in \mathcal{U}, j \in \mathcal{F} \quad (11)$$

$$\sum_{j \in \mathcal{F}} x_{0j}^{ku} - \sum_{j \in \mathcal{F}} x_{j0}^{ku} = 0 \quad \forall k \in \mathcal{K}, u \in \mathcal{U} \quad (12)$$

$$\sum_{k \in \mathcal{K}} \sum_{u \in \mathcal{U}} \sum_{i \in \mathcal{N} \setminus \{j\}} C_k x_{ij}^{ku} + \sum_{p \in \mathcal{P}} \sum_{i \in \mathcal{N} \setminus \{j\}} C_p x_{ij}^p \geq D_j \quad \forall j \in \mathcal{F} \quad (13)$$

$$\sum_{j \in \mathcal{F}_h} D_j x_{0j}^{ku} + \sum_{i \in \mathcal{F}_h \setminus \{j\}} D_j x_{ij}^{ku} \leq C_k \quad \forall k \in \mathcal{K}, u \in \mathcal{U}, h \in \mathcal{H} \quad (14)$$

$$\tau_D^{ku} + R_D + T_{0j}^k \leq (1 - x_{0j}^{ku})M + \tau_j^{ku} \quad \forall k \in \mathcal{K}, u \in \mathcal{U}, j \in \mathcal{F} \quad (15)$$

$$\tau_{D0}^p + R_D + T_{0j}^p \leq (1 - x_{0j}^p)M + \tau_j^p \quad \forall p \in \mathcal{P}, j \in \mathcal{F} \quad (16)$$

$$\tau_i^p + R_i + T_{ij}^p - (1 - x_{ij}^p)M \leq \tau_j^p \quad \forall p \in \mathcal{P}, i \in \mathcal{F}, j \in \mathcal{F} \quad (17)$$

$$\tau_j^p + R_j + T_{j0}^p - (1 - x_{j0}^p)M \leq \tau_{D1}^p \quad \forall p \in \mathcal{P}, j \in \mathcal{F} \quad (18)$$

$$\tau_j^{ku} + R_j + T_{j0}^k \leq (1 - x_{j0}^{ku})M + \tau_D^{k, u+1} \quad \forall k \in \mathcal{K}, u \in \mathcal{U}, j \in \mathcal{F} \quad (19)$$

$$\tau_i^{ku} + T_{ij}^k \leq (1 - x_{ij}^{ku})M + \tau_j^{ku} \quad \forall k \in \mathcal{K}, u \in \mathcal{U}, h \in \mathcal{H} \quad (20)$$

$$x_{ij}^{ku} \in \{0, 1\} \quad \forall k \in \mathcal{K}, u \in \mathcal{U}, (i, j) \in \mathcal{A}_{\mathcal{K}} \quad (21)$$

$$x_{ij}^p \in \{0, 1\} \quad \forall p \in \mathcal{P}, (i, j) \in \mathcal{A}_{\mathcal{P}} \quad (22)$$

$$T_j^0 \leq \tau_j^{ku} \leq T_j^1 \quad \forall k \in \mathcal{K}, u \in \mathcal{U}, j \in \mathcal{F} \quad (23)$$

$$T_j^0 \leq \tau_j^p \leq T_j^1 \quad \forall p \in \mathcal{P}, j \in \mathcal{F} \quad (24)$$

$$\tau_D^{ku} \geq 0 \quad \forall k \in \mathcal{K}, u \in \mathcal{U} \quad (25)$$

$$\tau_{D0}^p \geq 0 \quad \forall p \in \mathcal{P} \quad (26)$$

$$\tau_{D1}^p \geq 0 \quad \forall p \in \mathcal{P} \quad (27)$$

$$Z \geq 0 \quad (28)$$

The objectives are given by equation (1) and equation (2). The former is the min-max objective with higher priority, which minimizes the time of the latest drop. The latter minimizes the total flying time by all active aircraft. With the notation

TABLE II. PARAMETERS OF THE AFFVRP.

Parameter	Description
T_{ij}^k	Time distance between nodes i and j for a tanker
T_{ij}^p	Time distance between nodes i and j for a scooper
D_i	Demand quantity of subfire node i
R_i	Processing time at subfire node i , i.e., the time it takes for the full dropping maneuver
R_D	Processing time at the airfield, assumed to be the same for tankers and scoopers
C_k	Capacity of tanker $k \in \mathcal{K}$
C_p	Capacity of scooper $p \in \mathcal{P}$
$[T_i^0, T_i^1]$	Time-window for subfire i
M	Big arbitrary value

TABLE III. DECISION VARIABLES OF THE AFFVRP.

Variable	Description
x_{ij}^{ku}	Binary. Unitary if tanker k travels from i to j on trip u
x_{ij}^p	Binary. Unitary if scooper p travels from i to j
τ_i^{ku}	Continuous. Time when tanker k starts processing node i on trip u
τ_i^p	Continuous. Time when scooper p starts processing node i
τ_D^{ku}	Continuous. Time when tanker k returns to depot on trip u
τ_D^p	Continuous. Time when scooper p leaves the airfield
τ_{D1}^p	Continuous. Time when scooper p returns to the airfield
Z	Continuous. Time of last drop

$(i, j) \in \mathcal{A}_{\mathcal{K}}$ and $(i, j) \in \mathcal{A}_{\mathcal{P}}$ we imply that tankers and scoopers can use different arc subsets because of the different features of the two aircraft types, as pointed out in section II. Constraints (3) and (4) are the min-max constraints for the tankers and scoopers, respectively, used to determine the time of the latest water drop. Constraint (5) ensures that every subfire is visited exactly once. In such a constraint set and in some others in the formulation, the expression $i \in \mathcal{N} \setminus \{j\}$ is used in the definition of routing variables x_{ij}^{ku} and x_{ij}^p for tankers and scoopers respectively. Although not explicitly stated, it holds that $(i, j) \in \mathcal{A}_{\mathcal{K}}$ in the first case and $(i, j) \in \mathcal{A}_{\mathcal{P}}$ in the second case. The same concept applies to some other constraints in the formulation. Constraint (6) ensures that, on a given trip, tankers leave the airfield at most once. Constraint (7) is the equivalent one for scoopers. Constraint (8) ensures that, for a tanker, trip $u+1$ can only start if trip u has ended. Constraint (9) ensures that scoopers leaving the airfield return to it. Constraint (10) is a classic VRP flow conservation constraint. It ensures that the scooper visiting a subfire node leaves the same node. Constraint (11) is for the tanker conservation flow around subfires: it ensures that tanker entering a subfire on a certain trip, also leaves that subfire during that trip. Constraint (12) does the same but for the airfield. Tankers leaving the airfield return to it for each trip they make. Constraint (13) ensures that the demand of each subfire is satisfied or exceeded, with the latter case being possible if a tanker performs the water drop. Constraint (14) ensures tankers can visit several subfires within the same fire, as long as the total demand of those subfires does not exceed the tanker capacity. Constraints (15)-(20) are time precedence constraints. Constraint (15) ensures that the time of a tanker drop on a certain

trip is later than the time of leaving the airfield, the processing time at the airfield, and the travel time from the airfield to the drop location. Constraint (16) is the counterpart for scoopers. Constraint (17) ensures that the time of a scooper drop is later than the previous drop operation, in addition to the processing time and the travel time. Constraint (18) enforces the same for the time of a scooper returning to the airfield after the last subfire drop. Constraint (19) ensures time precedence of trips for a given tanker. It ensures that a trip can only start at the airfield after the previous trip has concluded. Constraint (20) is for time precedence of tankers between subfires belonging to the same original fire $h \in \mathcal{H}$. Note that no processing time is included here because when tankers visit multiple subfires, in reality, that represents just one drop that the tanker is performing on the same fire. Finally, constraints (21)-(28) define the nature of the decision variables. In particular, (23)-(24) define the time-windows when the water drop for each subfire can be performed.

IV. RESULTS

This section is structured in three parts. Section IV-A describes the general settings and some computational insights. Section IV-B presents a case study to demonstrate example results and limitations of the model. Finally, Section IV-C shows another use case of the model, namely strategic fleet planning through the use of a Monte Carlo simulation to compare different fleet compositions.

A. General settings and computational insights

According to expert insights, a maximum time of 5 minutes is to be set to run the model and obtain results for a given

wildfire instance. This ensures that operators can access the results before the aircraft take-off, such that the results can be employed without causing any delays to the live urgent response operations. The MILP was programmed in Python and solved with the commercial optimization solver Gurobi. The Gurobi version used is Gurobi Optimizer version 9.5.1 build v9.5.1rc2. All instances were solved using a laptop with an Intel64 Family 6 processor with 6 cores and 12 logical and 16 Gb of RAM. To gain more insight into the performance of the model and the time it takes to obtain useful results, several instances of the simulation with various numbers of fires and subfires were tested, and the performance parameters were recorded and are presented in Table IV.. While this experimental campaign was not extremely exhaustive, and instances of similar size might be characterized by relatively different solution times (to optimality), some conclusions can still be drawn. This investigation used a time limit of 21,600 seconds, or 6 hours, at which point the simulation was stopped and the computational insights recorded. The model limitations are observed here as the optimality gap at the 5-minute time limit is higher than 50% in problems with 15 subfires, which are medium-sized according to expert judgment. For larger problems, the gap increased further, and remains high even after 6 hours.

TABLE IV. COMPUTATIONAL INSIGHTS FOR PROBLEMS OF VARIOUS SIZES.

Num. original fires ($ \mathcal{H} $)	Num. subfires ($ \mathcal{F} $)	Computation time [s]	Optimality Gap [%]	Gap after 5 mins [%]
1	5	1	0	0.0
3	10	69.5	0	0.0
5	15	21,600	39.7	58.1
7	20	21,600	60.7	73.2
8	25	21,600	77.4	81.7
9	30	21,600	85.0	86.9
10	35	21,600	79.2	81.2
10	40	21,600	75.2	89.9
10	45	21,600	78.3	92.1
12	50	21,600	99.9	100.0

B. Case Study

The assumptions used to set up the case study are based on input from an active firefighting pilot who has worked in different regions of the world. A summary of the expert insights is given here before the case study representing the realistic scenario is presented.

Fire size

Most of the fires requiring AFF are still relatively small, the maximum number of fires dealt with by one airbase is around 5. Those fires require a few drops (2 to 4) in order to be kept small (up to about 2 hectares) such that the ground firefighters can control them. However, sometimes fires can get out of control and become medium or large fires that require more AFF involvement. These are classified as medium fires to

keep a simple categorization, although in reality they may be considered large. A medium fire may have up to 20 fire fronts, requiring about 15 to 20 drops. The extremely large fires are referred to as campaign fires. These fires are very large and the main purpose of AFF is to protect property by laying lines ahead of areas with valuable assets, such as homes or residential areas. These fires require hundreds of drops and are outside the current computational capabilities of the presented AFFVRP.

Number of aircraft

Usually, an AFF base has about 2 aircraft on standby, in case a fire starts in the area. If there are already fires in the area, the airfield can have up to 6 aircraft (tankers) on standby. More than 6 is unusual because it can cause traffic issues, and delays when the aircraft need refueling, as well as taking turns to take-off and land. If scoopers are available, up to 6 scoopers can also join the airbase in addition to the maximum of 6 tankers mentioned previously. Scoopers are less of a concern for airfield operations since they do not need to be at the airfield much, as they can scoop up water from water bodies. In addition, they can even refuel at another airfield if necessary. This is because many airfields have the infrastructure for refueling aircraft, but the main airbase is particularly prepared for AFF operations, and tankers need that infrastructure to refill the retardant or water tanks. Furthermore, if helicopters are available, up to about 6 helicopters can also operate out of the same airfield. In summary, the maximum number of aircraft can be up to 6 tankers, 6 scoopers, and 6 helicopters. For the purposes of this paper, that means up to 6 tankers and 12 scoopers (since helicopters are treated as scoopers with different parameters).

Scooper operations

There are two relevant insights concerning the operation of scoopers. Firstly, scoopers tend to work in pairs or groups of 3. The reason is to increase efficiency. In a pair, twice the capacity is achieved for the same operational effort. Secondly, scoopers can deliver 100 or more drops if the water body is nearby. They can manage the scooping operation quickly (under a minute) and return to the fire. Depending on the location of the water body, this could mean up to a drop every 5 minutes. Furthermore, aircraft can operate for up to 12 hours a day, limited by darkness and weather conditions among other factors.

Computation time requirements

The model presented in this paper can have multiple use cases. One of the uses is live during a fire, to determine the routes for the available aircraft. This use case is highly time-sensitive. From the moment that a fire is reported and known to the airbase, it usually takes about 20 minutes before the aircraft are in the air. Hence, the maximum computation time should not exceed this. However, having preliminary results after 5

minutes is preferable, so that this information is available at the right time in the decision-making process for operators.

Insights into AFFVRP modifications to tackle large-scale cases

Fires of this scale require several firefighting aircraft, and hundreds of drops per aircraft per day. Thus, modeling the problem can involve thousands of nodes if approached in this way. However, the fire is dynamic and the drop requests keep changing. The number and location of drops can change on an hourly basis, or even quicker in some cases. Hence, a more realistic approach to represent and apply the problem is to optimize a snapshot of the problem at a certain moment, and then re-run with the updated information at certain time increments in a rolling-horizon approach. As an example, consider a large fire that lasts for three days. At the first hour of the first day, the fire is still small, perhaps only one or two aircraft are available and deployed, and only a few drop requests are identified. By the first hour of the third day, reinforcements would have arrived and the fleet expanded to several aircraft. The wind may have changed direction and the fire may have advanced towards a residential area, so the requested drop locations change accordingly. This means that large fires cannot and should not be modeled at once with a full fleet and all the drop requests at once. An incremental approach is more useful with several runs of the model, where the inputs are continuously updated and the model re-run. The result is updated with optimized routes at every increment. This approach has the benefit of solving the computation time issue. The incremental approach divides this problem into smaller problems, which can each be solved in minutes or seconds. Hence, for large-scale fires, this approach features advantages both computation-wise (smaller problems vs. a large one) and application-wise (updated information is used dynamically to revise aircraft routes according to the changes in the fire fronts).

Case Description

In July 2022, two large fires south of Bordeaux, in the department of Gironde in France, burnt more than 20,800 hectares (ha) of land. One of the fires was in La Teste-de-Buch and the other in Landiras, with the two areas about 50 km apart. The fires started in the afternoon of July 12th 2022, and progressed aggressively for the following days. The aircraft types used to assist in fighting the fires were the Canadair CL-415 and CL-215 as scoopers (Capacity of 6,140 and 4,800 liters, respectively) [6], and the Bombardier Dash 8 400-AT (capacity of 10,000 litres) [7].

Case results

Two instances from the fire operation were selected to demonstrate the use of the model for large fires. The assumed inputs

are summarized in Table V.

TABLE V. INPUTS OF THE BORDEAUX CASE STUDY SIMULATIONS.

Parameter	Value
Scooper capacity	5,000 [litres]
Tanker capacity	10,000 [litres]
Scooper speed	6 [km/min]
Tanker speed	10 [km/min]
Planning horizon	ca 1 [hour]
Required number of scooper drops for simulation #1	[Teste-de-Buch, Landiras] [7, 13]
Required number of scooper drops for simulation #2	[Teste-de-Buch, Landiras] [30, 50]
Preferred Time limit	5 [minutes]

The fire situations on July 16th and July 19th, 2022 are found through the EFFIS system and mapped accordingly. The number and types of mobilized aircraft are obtained from the press releases of the local authority [8]. The subfires, which can be seen as the drop locations, are split among the two large fire areas. In the first fire situation (July 16th), the Teste-de-Buch (west) fire has 7 drop requests and the Landiras (east) has 13. The resulting time carrying out all requested drops is 76.05 minutes. The resulting routes are visualized in Figure 2. Note that the tanker is carrying out the drops at Landiras, which are further from the water body, while the scoopers mainly tend to the drops at Teste-de-Buch, which are near the water body. This verifies the model as it is a logical assignment, as scoopers can make use of the water body, and the tanker has a higher cruise speed. Therefore, this routing leverages the strengths of the different aircraft types.

In the second fire situation (July 19th), the fires are larger and out of control. The wind is moving south and hence the fires are spreading in that direction. This case is the largest instance of all attempted case studies so far with a total of 80 drop requests. The fire at Teste-de-Buch (west) is estimated to have 30 drop requests, while the fire at Landiras (east) is estimated to have 50. According to the local press release [8], the number of mobilized scoopers on this day was 8, and the number of mobilized tankers was 2. For this fleet combination, the resulting time to carry out all requested drops is 98.7 minutes. The routes are shown in Figure 3. At this size of the simulation, involving 80 subfires and 10 aircraft, the model shows some limitations. A feasible solution was not obtainable within the preferable time limit of 5 minutes. This solution was obtained with a time limit of 1,300 seconds, meaning in just under 22 minutes. This may still be acceptable as it takes at least about 20 minutes before the aircraft are in the air. Nonetheless, this demonstrates the size of the problem at which the model surpasses the limit of the acceptable time limit range.




Figure 2. Resulting routes for the July 16th fire.

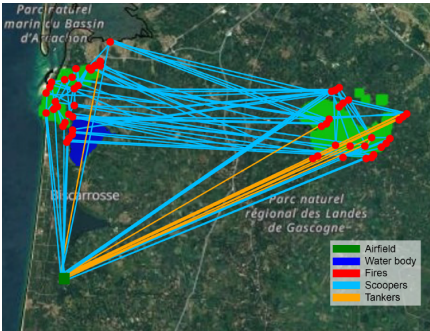


Figure 3. Resulting routes for the July 19th fire.

C. Strategic fleet planning with Monte Carlo simulation

The model can also be used for strategic fleet planning for a given geographical region. This is demonstrated with a Monte Carlo simulation. Random fires are generated within the boundaries of a given area. It is assumed that the existing fleet consists of 1 scooper and 1 tanker, and the study is to compare the performance of adding another scooper vs adding another tanker. Thus, *fleet A* is composed by 2 scoopers and 1 tanker, and *fleet B* by 1 scooper and 2 tankers. The results in Figure 4 show that, on average, fleet B completes the mission sooner than fleet A and hence is to be preferred for the area of interest.

V. VALIDATION TEST

The ultimate aim of the model is to craft a decision-support tool that can help firefighting authorities contain fires more efficiently. The AFFVRP model has been formulated and programmed to produce efficient AFF routes, but the question that needs to be answered is if these routes are more efficient than what human operators can come up with in a real or simulated situation. To investigate this, a validation test is designed and taken by an expert in the field (an active firefighting pilot), and the results are compared with the outcome of the AFFVRP model based on the same input. The hypothesis is that the model will outperform the human, especially as the cases increase in

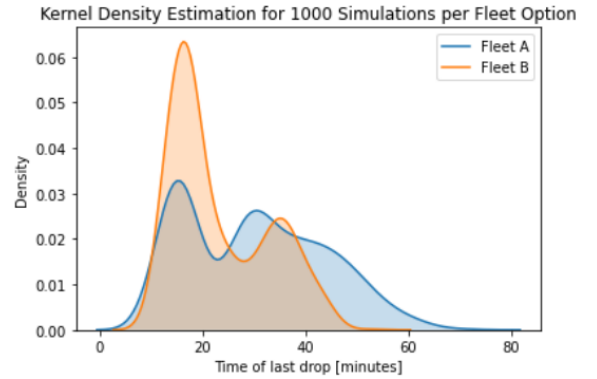


Figure 4. Density estimation comparing the performance of Fleet A vs Fleet B in a Monte Carlo simulation of a total of ,2000 fire situations

complexity. The available fleet consists of 2 scoopers (denoted p_1 and p_2) and 2 tankers (denoted k_1 and k_2). The scoopers have a capacity of 5 [k litres] and a speed of 5 [km/min]. The tankers have a capacity of 10 [k litres] and a speed of 10 [km/min]. All aircraft need 10 minutes at the airfield before take-off, and 1 minute for each drop maneuver. The time for scoopers to collect water from the water body is considered negligible. All subfires are considered to have equal urgency (thus no time windows are used), and the only objective considered is the primary objective, that is to minimize the time of the last drop, or to satisfy all drop requests at the earliest time possible. The graphics used to represent the situation are simple and clear. Furthermore, the test layout and some examples are explained thoroughly to the test taker before the test takes place. The test taker is also asked to use a timer and attempt to solve each case within 5 minutes. The same time limit is used when the model is run to solve the given cases. The test contains 6 cases of varying complexity, the simplest of which contains 2 main fires and 6 subfires, and the most complex contains 7 main fires and a total of 23 subfires. The results are reported in Table VI. for cases 1-3 and Table VII. for cases 4-6. The main hypothesis was that the model will provide more efficient routes, and earlier last drop times, than the human expert. This is shown to be true in all cases except case 1, where they both provided the same solution, as it was a simple case and the optimal route was almost trivial to find. Another hypothesis is that the advantage of the model will grow in correlation with case complexity. As the cases become more complex, it was expected that the human will have a harder time finding good routes, resulting in a larger difference in the time of the last water drop. In this validation test, this hypothesis is not verified. An explanation for this is that with growing complexity, the solution space increases as well, and the model takes longer to find superior solutions. Hence, the hypothesis is still expected to hold if no time limits are imposed (or more relaxed ones are used), but for the used time limit of 5 minutes, the model also

struggles to find good solutions as the optimality gap is still large. This is best seen in case 6, the most complex case of the test, involving 23 subfires and 4 aircraft. The resulting time of the last drop by the expert was 61.3 minutes and by the model 58.4 minutes, as shown in Table VII.. This difference is smaller than the difference observed in the less complex cases. However, if the time limit is extended, the solution does improve further, yielding a last drop time of 54.9 minutes for case 6 when the time limit is extended to 10 minutes.

TABLE VI. VALIDATION TEST RESULTS COMPARING PERFORMANCE OF EXPERT HUMAN AND AFFVRP MODEL FOR CASES 1-3. TIME OF LAST DROP IN MINUTES.

	Case 1		Case 2		Case 3	
	Expert	Model	Expert	Model	Expert	Model
Overall	14.0	14.0	30.9	26.9	31.1	27.8
p_1	14.0	14.0	30.9	26.9	23.0	23.0
p_2	14.0	14.0	25.9	23.2	23.0	18.9
k_1	12.0	14.0	12.0	27.0	14.5	14.5
k_2	14.0	12.0	14.0	14.0	31.1	27.8

TABLE VII. VALIDATION TEST RESULTS COMPARING PERFORMANCE OF EXPERT HUMAN AND AFFVRP MODEL FOR CASES 4-6. TIME OF LAST DROP IN MINUTES.

	Case 4		Case 5		Case 6	
	Expert	Model	Expert	Model	Expert	Model
Overall	53.9	42.2	48.7	43.6	61.3	58.4
p_1	53.9	37.7	48.7	41.3	61.3	52.5
p_2	40.9	38.3	43.7	39.0	61.0	52.8
k_1	43.7	34.9	40.5	43.6	49.9	58.4
k_2	29.5	42.2	29.5	33.7	49.9	46.8

VI. CONCLUSIONS AND RECOMMENDATIONS

The AFFVRP model makes use of a combination of features from various types of well-studied VRPs. It is a capacitated multi-trip VRP with time windows and hierarchical objectives. Two types of aircraft are used which can represent most fire-fighting aircraft by simply changing the capacity and speed. The two types are tankers and scoopers. The main difference is that scoopers can make use of natural bodies of water to refill their tank. Helicopters can also be represented by this type. Hence, the model is flexible and can capture significantly different real-life situations. While the model itself should undergo algorithmic enhancements to tackle larger instances, results are encouraging as the AFFVRP already defines an effective decision-support tool to assist human decision-makers in planning faster wildfire containment routing strategies. In particular, a validation via expert judgment, albeit limited in size, showed that non-optimal solutions obtained within the recommended 5 minute time limit can still outperform manually-obtained solutions by expert operators. To enhance and extend this work, several recommendations are suggested. Firstly, expanding the problem setting is

paramount. The current AFFVRP utilizes a single airfield and considers only one (nearest) water body. For more comprehensive coverage, especially in addressing large fires, the model can be extended to involve multiple airbases handling the same fires and accommodating situations where multiple accessible water bodies are feasible. This inclusion would enhance the model's realism. Secondly, the aircraft selection should be broadened. The current simulation allows the use of only one type of scooper and one type of tanker. Expanding this capability to model various scoopers, helicopters, and tanker types concurrently for a given fire simulation would provide greater flexibility for incorporating and analysing diverse fleets. Lastly, for creating a split delivery VRP, the current discretization approach splits large fires into subfires, allowing an aircraft to visit each subfire once. However, this approach poses a drawback, as aircraft with a capacity larger than the smallest subfire risk wasting part of their capacity. Besides further developments of the problem definition, improving the solution quality via tailored algorithms is also recommended. The optimality gap with a 5 minute time limit was larger than 80% for problems comprising of 25 subfires and more. For non-urgent use of the model, this is acceptable as better solutions, and eventually an optimal solution, are found when more time is afforded. However, for the live application where a 5 minute time limit is a constraint, a faster convergence towards quasi-optimal solutions is recommended. Using suitable heuristics to find such solutions within this time limit, can make this work more useful in real-life applications.

REFERENCES

- [1] Dennison, P. G. (2009). "Climate change and forest fires." *Global Change Biology*, 15(8):2279– 2288.
- [2] J. Parappathodi and C. Archetti, "Crowdsourced humanitarian relief vehicle routing problem," *Computers & Operations Research*, vol. 148, paper 105963, 2022.
- [3] X. Wei et al., "An integrated location-routing problem with post-disaster relief distribution," *Computers & Industrial Engineering*, vol. 147, paper 106632, 2020.
- [4] European Forest Fire Information System (EFFIS), Current Statistics Portal. Accessed on Jan 17, 2024.
- [5] M. Dror and P. Trudeau. "Savings by split delivery routing." *Transportation Science*, pages 141–145, 1989.
- [6] Viking Air (2023), CL-415 aircraft. Accessed last on May 10, 2024.
- [7] Conair (2021). Dash 8-400at type 2 airtanker. Accessed on: Mar 3, 2023.
- [8] Préfète de la Gironde. Communiqué de presse. Accessed on: Mar 3, 2023.