Military Aircraft Mission Planning - A Generalized Vehicle Routing Model with Synchronization and Precedence

Nils-Hassan Quttineh, Torbjörn Larsson,
Kristian Lundberg and Kaj Holmberg

March 12, 2013

Abstract

We introduce a military aircraft mission planning problem where a given fleet of aircraft should attack a number of ground targets. Due to the nature of the attack, two aircraft need to rendez-vous at the target, that is, they need to be synchronized in both space and time. At the attack, one aircraft is launching a guided weapon, while the other is illuminating the target. Each target is associated with multiple attack and illumination options. Further, there may be precedence constraints between targets, limiting the order of the attacks. The objective is to maximize the outcome of the entire attack, while also minimizing the mission timespan. We give a linear mixed integer programming model of the problem, which can be characterized as a generalized vehicle routing problem with synchronization and precedence side constraints. Numerical results are presented for problem instances of realistic size.

Keywords: Military operations research, Generalized vehicle routing, Mixed integer programming, Time dependencies, Precedence constraints.

1 Problem Setting

Military mission planning is a very complex task, and it is still mainly performed manually by experienced personnel. This is both time-consuming and the outcomes are probably not even close to being optimal in any sense. In this paper, we address a military mission planning problem where a homogeneous fleet of aircraft should attack a number of ground targets under certain restrictions on the properties of the plan.

The work presented originates from and has been performed in collaboration with an industrial partner.

In general, a military mission might involve various tasks, such as surveillance, backup support, rescue assistance or an attack, where different agendas might be connected to different locations or targets. We consider the situation where there is only a set of targets that shall be attacked.

Except for the specific targets, there are also enemy defense positions, like surface-to-air missiles (SAMs), and protected objects not to be touched by the attacks, like hospitals and schools. We assume that all these objects are stationary with known positions, and that the goal is to find an optimal attack plan where maximal effect is gained within a short timespan.

We also assume that the risk for the aircraft is minimized by restricting them not to fly through defended airspace (or by allowing them to fly only in airspace within an acceptably low level of risk). Weapons may however pass through defended airspace, but at the risk of being shot down.

1.1 Target scene

The geographic area of interest, where targets, defenders and protected objects are situated, is referred to as *the target scene*, which is also defined by a line of entrance and a line of exit for the aircraft. The aircraft fleet is deployed from a base positioned on ground or from hangar ships, usually situated far away from the target scene. They enter the scene at the entry line, and when the mission has been carried out they leave the scene at the exit line and turn back to the base (or some other base). An example of a small target scene is given in Figure 1.

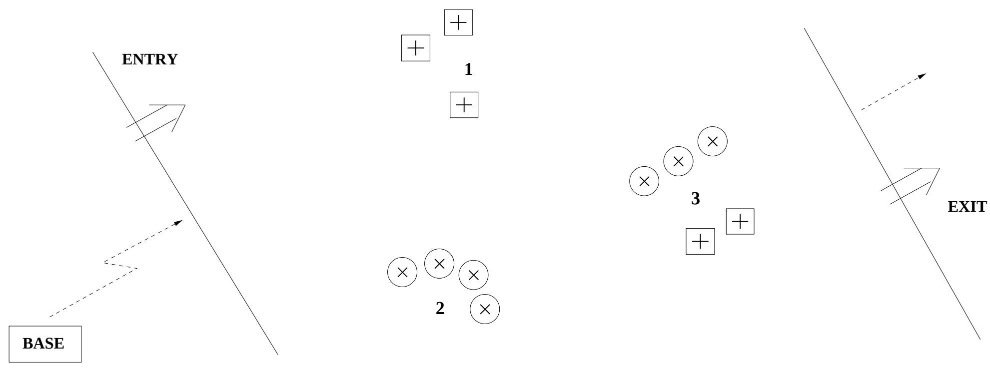

Figure 1: An example of a target scene, including three targets and nearby SAMs ($\times$) and hospitals ($+$) . The entry line is to the *left* and the exit line to the *right*.

The mission time is defined as the time of the first aircraft passing the line of entry until the last aircraft passes the exit line. The diameter of a target scene is usually of the order of 100 km, the distances between targets are of the order of a few kilometers, and the timespan of a mission is of the order of one hour. A large attack would involve 6–8 targets and 4–6 aircraft, and would require several hours, at least, of manual planning just to find a feasible attack plan.

1.2 Attack and illumination

An attack requires two aircraft to team up, where one of them illuminates the target with a laser beam and the other launches the weapon (bomb or missile). We assume that the flight direction of the aircraft is towards the target at the time of the launch. The illumination is required to guide the weapon towards its target, providing high accuracy of the impact, and it needs to be continously visible for the weapon. This means that the aircraft need to rendez-vous not only in time, but also in such a way that the illumination is visible for the attacker at the launch of the weapon. Hence, the illumination begins shortly before the launch of the weapon and has to continue until impact. A typical flight path for an attacker, and a flight path for an aircraft illuminating a target, can be found in Figure 2.

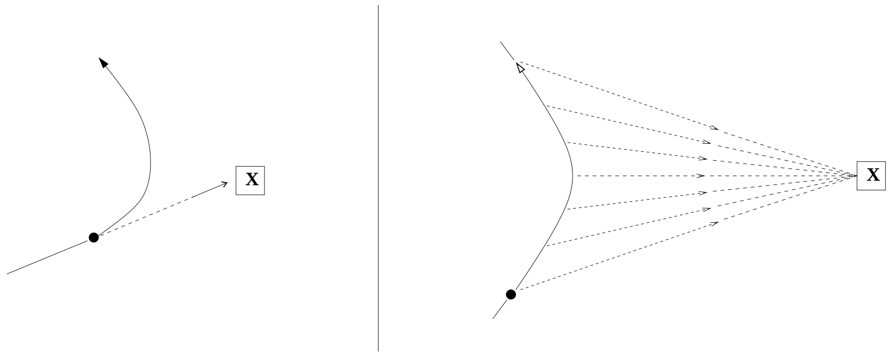

Figure 2: *Left*: Attack path. The aircraft flies towards the target, launches the weapon, and makes an evasive manoeuver. *Right*: Illumination path. From its starting position, the aircraft flies in a parabolic path and illuminates the target continuously.

When a target is attacked, the air around it will be filled with dust and debris, and due to the prevailing wind conditions this might reduce the visibility of nearby targets. Hence it is realistic to assume that some precedence constraints are given, specifying which targets are not allowed to be attacked before other targets.

1.3 Restrictions and limitations

The expected effect of an attack depends, of course, on the kind of weapon being used, which is decided in advance, but also on the direction of the impact and its kinetic energy. The latter factors depend on the velocity and altitude of the aircraft at the time of the launch. Further, if the weapon passes through defended airspace, its expected effect is reduced.

No matter how accurate the attack can be performed, the neighbouring area of the target is always subject to a certain risk of collateral damage, because the weapon can miss its target. This can be due to for example loss of visibility of the illumination, malfunction of the weapon or defense measures.

We refer to the area of unacceptably high risk for collateral damage as the *footprint* of the attack, visualized in Figure 3, and it depends on the altitude and velocity of the attacker. The footprint is in our description of the problem simply given by a straight line from the attack position towards the target, and an angle of maximal deviation from this line. This construction gives an ellipsoid-shaped footprint on the ground.

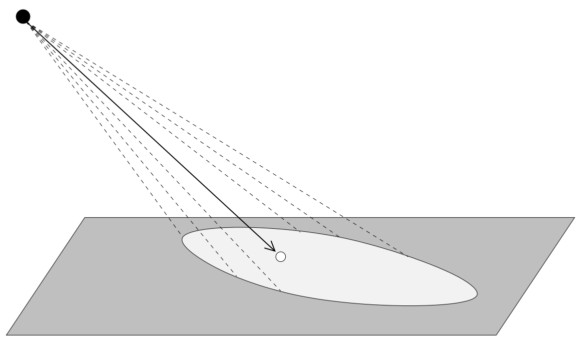

Figure 3: Footprint of an attack position. The black line represents the correct path of the missile, and the dotted lines its maximum deviations.

We define an attack position to be feasible if no protected object is inside its footprint. For a given target, and an aircraft with specified characteristics such as velocity, altitude and armament, one can derive a region of feasible attack positions, referred to as the feasible attack space for that target. Each attack position in this space is associated with a number of feasible illumination positions, where one illumination position can be compatible with multiple attack positions.

Each aircraft has an armament capacity which limits the number of attacks it can perform. In addition to the armament, an aircraft can also carry an illumination laser pod. Without the illumination pod, an aircraft can only perform attacks. An aircraft might also be equipped with the illumination laser pod only, hence only capable of performing illumination.

Note that once the planning of the mission has been made, it is also known how each and one of the aircraft shall be equipped in order to be able to fulfil its tasks during the mission.

1.4 Outline

All in all, the problem is to define an attack sequence for the given fleet of aircraft in such a way that each target is attacked and illuminated exactly once, in a synchronized manner using compatible attack and illumination positions, where the flight path of each aircraft avoids risk from defense systems and the attack sequences must comply with the a priori given precedence constraints. The goal is to maximize the overall expected effect of the attacks on the targets while also minimizing the total timespan of the mission. Due to this bi-objective nature of the problem, we will search for compromise solutions.

The paper is organized in the following way. In Section 2 we discuss modelling issues, and in Section 3 we introduce notations and present the mathematical model. Although our application is very unusual, the model belongs to the familiar class of vehicle routing problems, but enriched with some specific extensions. Each of these extensions appear in other applications, which is also discussed in Section 3. In Section 4 we present some illustrative numerical results, and finally conclusions and future work are discussed in Section 5.

2 Modelling Issues

For a specific type of aircraft and a target requiring a specific type of weapon, one can derive the feasible attack space against the target, here represented by an inner and an outer radius of attack plus an upper and a lower altitude, that is, an attack space that can be visualized as a hollow cylinder. A two-dimensional projection of this cylinder onto the ground is found in Figure 4.

This hollow cylinder is divided into altitude layers and a number of sectors, in which we discretize feasible attack positions. For each and one of these we create compatible illumination positions. Since only feasible attack positions are included, and these depend on the presence of protected objects and air defense, the number of such positions in each sector might vary.

We have chosen to use six sectors together with three discrete possible attack positions in each sector and on each altitude. For the three attack positions we introduce two illumination alternatives. These utilize the same flight

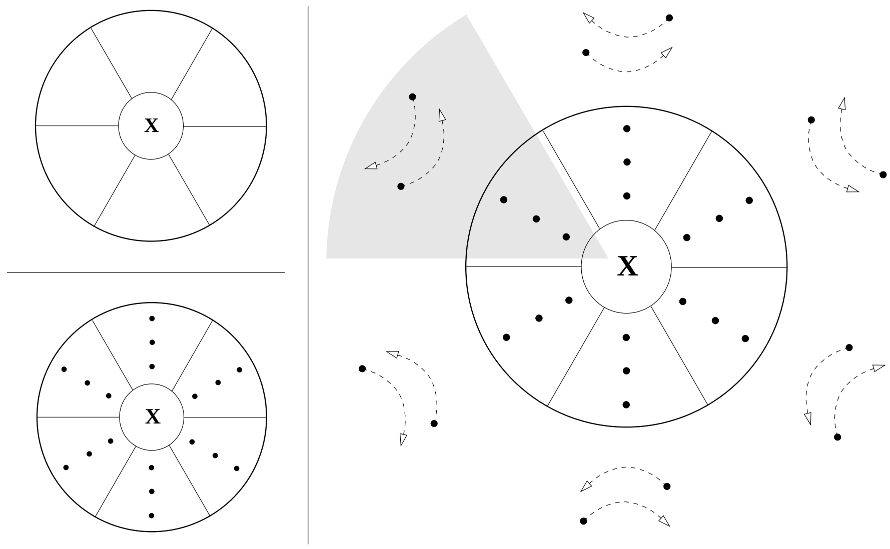

Figure 4: Upper left: The feasible attack space as a hollow cylinder, divided into six sectors. Lower left: A coarse discretization of three attack positions in each sector. Right: Two illumination alternatives for each sector. All such alternatives are compatible with all attack positions in the same sector.

path but differ in flight direction, which is essentially clockwise or counter-clockwise in relation to the target. Each alternative is associated with a position where the illumination begins. See Figure 4 for an illustration of the discretization used.

2.1 Network representation

By performing a discretization of the feasible attack space around each target, representing attack and illumination positions by nodes, and aircraft movements by arcs, the mission planning problem can partly be represented by a network. A dummy origin and a dummy destination are introduced to represent the crossing of the entry and exit lines of the target scene, respectively.

Each target shall be attacked and illuminated exactly once, and an aircraft can not both attack and illuminate a target. Hence, the network only contains arcs between nodes corresponding to different targets, or from the dummy origin or to the dummy destination.

Moreover, no arcs should violate any given precedence constraints between targets. Hence there might for example be arcs from attack nodes for target 1 towards the attack and illumination nodes for target 2, but not the other way around, depending on the given precedence relations.

On a more abstract level, nodes in the network can be clustered and represented as in Figure 5, where each target is associated with two clusters, one containing all attack nodes (A) and the other containing illumination nodes (I). The two illumination nodes in each sector correspond to the two illumination alternatives described above. The exact structure of these clusters is found in Figure 6.

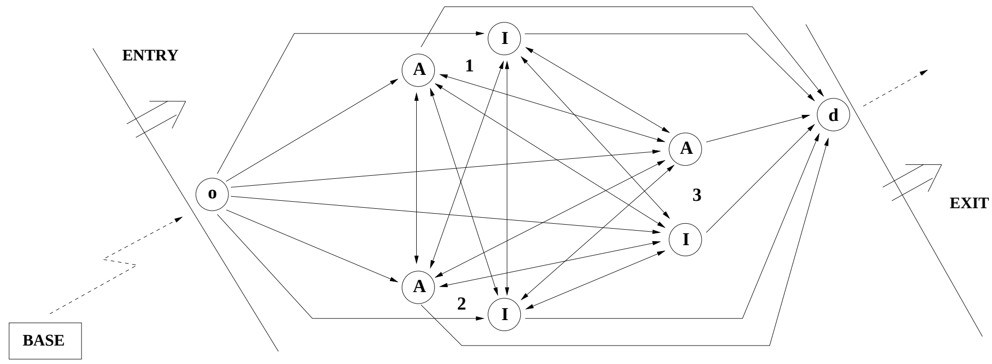

Figure 5: Given an origin (o), a destination (d), and three targets, the aircraft fleet should visit each cluster (A=attack and I=illumination) exactly once. An aircraft is not allowed to both attack and illuminate the same target though.

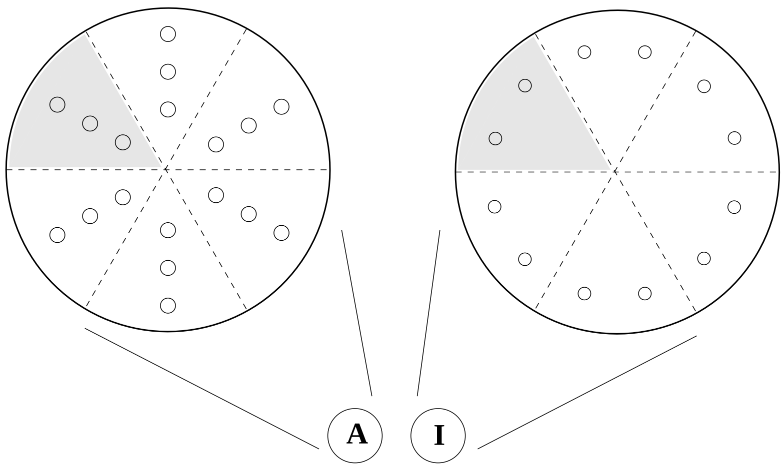

Figure 6: Structure of the clusters. An attack cluster consists of many attack positions, and exactly one should be visited. The same goes for the illumination clusters. Also, the visited attack and illumination positions need to belong to the same sector, here illustrated by the shaded regions.

2.2 Arc costs

For calculating the arc costs in the network representation of the problem, we must find flight distances between all candidate positions. A feasible path between two positions is a path where the restrictions of the aircraft dynamics is taken into account, such as turning radius and other physical limitations. The path also needs to be safe, meaning that the aircraft cannot pass through defended airspace. An illustration of this is found in Figure 7.

In the literature, the problem of finding an optimal flight path from a given starting point to a given destination, while avoiding obstacles, such as defended airspace, is referred to as the Aircraft Routing Problem. This is in itself a difficult optimization problem, but not adressed in this paper, and we refer to [22] and [3] as examples of algorithms that can be used to solve this problem. A closely related routing problem is described in [6] and [19], which gives rise to a shortest-path problem with side constraints.

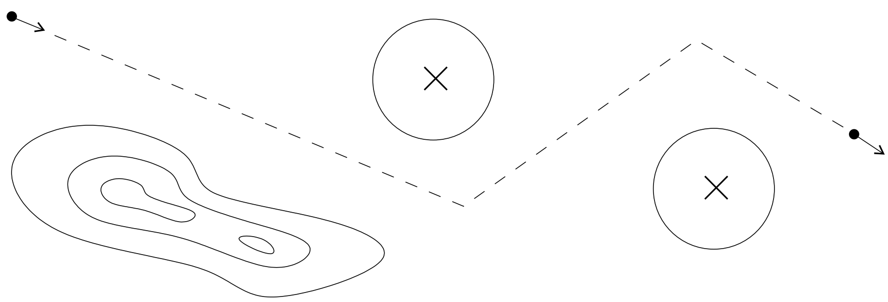

Figure 7: An example of aircraft routing. The problem is to find a flight path from a given starting point to a given destination point, avoiding obstacles and defended airspace.

In our numerical experiments we used a flight path generator provided by our industrial partner. It takes aircraft dynamics into account and is based on a discretization of the airspace and a calculation of a shortest path.

The result of each such routing problem is a feasible path with a length that can be converted into a minimal time required to traverse it. In our network representation, the nodes are associated with *both* a location and a flight direction at the location. The arc lengths and travel times will in general be asymmetric, because of the flight directions and the flight dynamics.

In addition to the time attribute, each arc leaving an attack position also has an attribute that states the expected effect against the target associated with the attack position.

3 Mathematical Description

3.1 Notation

Given is an aircraft fleet $\mathcal{R}$, and a set of targets $\mathcal{M}$ to be attacked. Each target $m \in \mathcal{M}$ is associated with a feasible space of attack positions, discretized into attack positions N_m^A, and their compatible illumination positions, N_m^I. Furthermore, each feasible attack space is divided into sectors, and we let $\mathcal{G}$ denote the set of all sectors for all targets while $\mathcal{G}_m$ is the set of sectors that belong to target $m \in \mathcal{M}$.

Denote by N all positions in the graph, including a dummy origin, o, and a dummy destination, d, and let N^* denote the set of all positions except the origin and the destination. Further, let A denote all arcs in the network and let A_g denote the set of arcs (i, j) such that position j is an attack position in sector g. Similarly, let I_g denote the set of arcs with heads at illumination positions in sector g.

Each aircraft $r \in \mathcal{R}$ is limited to carry at most Γ weapons. Let q_m denote the number of weapons needed towards target $m \in \mathcal{M}$. Let $\mathcal{S}$ denote the set of ordered pairs (m, n) of targets such that target m cannot be attacked before target n. If no precedence relations are given a priori, the set $\mathcal{S}$ is empty.

A parameter c_{ij}^r, $(i, j) \in A$ and $r \in \mathcal{R}$, captures the effect of attacking from a certain position. For arcs (i, j) with $i \in N_m^A,\ m \in \mathcal{M}$, that is, for arcs leaving attack positions, the value of c_{ij}^r is the expected effect of the attack, and otherwise the value is zero.

Further, let T_{ij}^r denote the time needed for aircraft $r \in \mathcal{R}$ to traverse arc $(i, j) \in A$. We also introduce T_{max}, either as a pessimistic estimate of the total mission time or as a given upper time limit for the duration of the mission.

We introduce two types of variables, the binary routing variables x_{ij}^r and the continous time variables t_i, t_m^A, and t_m^I. The routing variable x_{ij}^r equals one if aircraft $r \in \mathcal{R}$ traverses arc $(i, j) \in A$, otherwise it is zero. Variable t_i is the time at which node $i \in N^*$ is visited, by some aircraft, and it equals zero if the node is not visited by any aircraft. The starting time for all aircraft is $t_o = 0$ while the time of the last aircraft to exit the target scene is t_d. Variables t_m^A and t_m^I are the times of the attack and illumination, respectively, of each target $m \in \mathcal{M}$.

3.2 The model

The goal is to maximize the expected effect against all targets, while minimizing the time of the last aircraft to pass the exit line. We choose to optimize a weighted combination of these two conflicting objectives, using a parameter $\mu \geq 0$. This yields a solution that is Pareto optimal with respect to the two objectives.

The mathematical model for the Military Aircraft Mission Planning Problem (MAMPP) is given below.

Constraints (1) – (3) ensure that each aircraft enters and leaves the target scene via the dummy nodes, and (4) ensures that each target $m \in \mathcal{M}$ is attacked exactly once, while (5) does the same for the illumination.

$$\max \sum_{r\in\mathcal{R}} \sum_{(i,j)\in A} c_{ij}^r x_{ij}^r - \mu t_d$$

$$s.t. \quad \sum_{(o,j)\in A} x_{oj}^r = 1, \quad r \in \mathcal{R} \tag{1}$$

$$\sum_{(i,d)\in A} x_{id}^r = 1, \quad r \in \mathcal{R} \tag{2}$$

$$\sum_{(i,k)\in A} x_{ik}^r - \sum_{(k,j)\in A} x_{kj}^r = 0, \quad r \in \mathcal{R},\ k \in N^* \tag{3}$$

$$\sum_{r\in\mathcal{R}} \sum_{g\in\mathcal{G}_m} \sum_{(i,j)\in A_g} x_{ij}^r = 1, \quad m \in \mathcal{M} \tag{4}$$

$$\sum_{r\in\mathcal{R}} \sum_{g\in\mathcal{G}_m} \sum_{(i,j)\in I_g} x_{ij}^r = 1, \quad m \in \mathcal{M} \tag{5}$$

$$\sum_{r\in\mathcal{R}} \sum_{(i,j)\in A_g} x_{ij}^r - \sum_{r\in\mathcal{R}} \sum_{(i,j)\in I_g} x_{ij}^r = 0, \quad g \in \mathcal{G} \tag{6}$$

$$\sum_{g\in\mathcal{G}_m} \sum_{(i,j)\in A_g \cup I_g} x_{ij}^r \leq 1, \quad r \in \mathcal{R},\ m \in \mathcal{M} \tag{7}$$

$$\sum_{m\in\mathcal{M}} \sum_{g\in\mathcal{G}_m} \sum_{(i,j)\in A_g} q_m x_{ij}^r \leq \Gamma, \quad r \in \mathcal{R} \tag{8}$$

$$t_i + \sum_{r\in\mathcal{R}} T_{ij}^r x_{ij}^r - T_{max}\left(1 - \sum_{r\in\mathcal{R}} x_{ij}^r\right) \leq t_j, \quad (i,j) \in A \tag{9}$$

$$t_i - T_{max} \sum_{r\in\mathcal{R}} \sum_{(i,j)\in A} x_{ij}^r \leq 0, \quad i \in N \tag{10}$$

$$t_o = 0, \tag{11}$$

$$
\begin{aligned}
\sum_{i \in N_m^A} t_i &= t_m^A, && m \in \mathcal{M} && (12)\\
\sum_{i \in N_m^I} t_i &= t_m^I, && m \in \mathcal{M} && (13)\\
t_m^A &= t_m^I, && m \in \mathcal{M} && (14)\\
t_m^A &\geq t_n^A, && (m,n) \in \mathcal{S} && (15)\\
x_{ij}^r &\in \{0,1\}, && r \in \mathcal{R},\ (i,j) \in A && (16)\\
t_m^A,\ t_m^I &\geq 0, && m \in \mathcal{M} && (17)\\
t_i &\geq 0, && i \in N && (18)
\end{aligned}
$$

Constraint (6) ensures that the attack and the illumination against each target are compatible, that is, that the nodes belong to the same sector. In a sector where no attack is performed, no illumination can be performed either, and vice versa.

Constraint (7) states that each aircraft can visit each target at most once, either for attacking or for illuminating the target. This constraint is actually redundant since the time propagation constraint (9) together with the synchronization constraint (14) make it impossible for an aircraft to both attack and illuminate the same target, but it results in a model with a tighter linear programming relaxation. Constraint (8) is the armament capacity constraint and limits each aircraft to utilize at most Γ weapons.

Constraint (9) propagates time, making sure that if an aircraft r traverses arc (i,j), node j is visited no earlier than the time of the visit to node i plus the time needed to traverse the arc. For $j = d$ the constraint (9) defines the total mission time, t_d, since all aircraft end up at the dummy destination. Note that the construction of this constraint exploits the fact that each arc can be traversed by at most one aircraft, which follows from constraints (3) and (4). Note also that constraint (9) eliminates subtours.

Constraint (10) enforces that $t_i = 0$ holds if node i is not visited by any aircraft, and (11) states that all aircraft start from the origin at time zero.

Constraints (12) and (13) assign the correct times of attack and illumination, respectively, for each target m, and constraint (14) states that these times need to be synchronized. Constraint (15) makes sure that targets are attacked in the prescribed precedence order. The variables t_m^A and t_m^I can be eliminated from the model, but they are used for the sake of readability.

Note that although the network does not contain arcs that violate the precedence relationships, constraint (15) is still needed to fully prevent violation of these relationships. This is so because, without constraint (15), precedence relations between non-consecutive targets can still be violated. Finally, (16) – (18) are definitional constraints.

3.3 Characteristics of the model

The model presented above belongs to the class of Vehicle Routing Problem (VRP). Attack and illumination points are nodes, and paths between such positions are arcs. The aircraft fleet corresponds to resources with capacity constraints on their weapon load, and targets correspond to customers. The model has the following non-standard characteristics.

i) It is generalized in the respect that exactly one node in each cluster shall be visited.

ii) Since the attack and illumination positions for a target need to be compatible, the visits to attack and illumination clusters are coupled by side constraints.

iii) The visits to the compatible attack and illumination nodes for a target are required to be synchronized in time.

iv) The order in time of the visits to the pairs of attack and illumination clusters of all targets are constrained by precedence relations.

In the Generalized Travelling Salesman Problem, the nodes are partitioned into clusters and the salesman shall visit exactly one node in each cluster, at minimum cost. This problem has been studied to some extent, see for example [13] and [16]. The corresponding generalization of the Vehicle Routing Problem (GVRP) has been studied much less. To our knowledge, the first to discuss this problem are Ghiani and Improta [10], who give a transformation to the Capacitated Arc Routing Problem (CARP). Baldacci *et al.* [2] discuss some applications of the GVRP. Formulations and branch-and-cut algorithms for the GVRP are given in the recent paper of Bektaş *et al.* [4]. Decision problems similar to generalized vehicle routing problems arise in school bus routing, which sometimes involves selecting both bus stops and bus routes. A survey on this subject can be found in [17].

Various restrictions on times when customers are visited, such as time windows, temporal dependencies, and precedence, are frequently appearing in VRP, with fixed time windows being the most common. Andersson *et al.* [1] study a ship routing and scheduling problem where deliveries of groups of cargoes shall be synchronized within given limits on time differences. Similar synchronization requirements appear in [12], where an aircraft fleet assignment and routing problem is considered; here, certain arrival and departure times may not deviate too much. A recent survey of vehicle routing problems with synchronization contraints is given in [8].

Bredström and Rönnqvist [5] describe a daily homecare planning problem, which is modelled as a vehicle routing and scheduling problem with exact pairwise synchronization of visits (meaning that two staff members are required to visit an elderly person simultaneaously). In addition, this model includes precedence constraints on visits. Redjem *et al.* [18] also consider routing with time windows and synchromized visits for a homecare planning problem.

In [21], vehicle routing with precedence constraints and time windows is considered in order to schedule transportation of live animals to avoid the spread of deseases. A general framework for VRP with time windows and temporal dependencies, including exact synchronization, is given in [7]. In the context of GVRP, the time window extension has been considered by Moccia *et al.* [15], who suggest a meta-heuristic solution method. Their work concerns an application to the design of home-to-work transportation plans.

To the best of our knowledge, the military aircraft mission planning problem presented here has not earlier been approached by operations research techniques, and it has in particular not been modelled as a generalized vehicle routing problem with compatability of visits, synchronization in time, and precedence relations. The work of Schumacher *et al.* [20] is however slightly related to ours. It considers unmanned air vehicle operations, with assignment of multiple tasks against a set of targets, vehicle routing, and precedence between tasks. A MILP model for this problem is presented, together with results for instances with less than four targets, five vehicles, and three tasks per target. These instances are considered to be of realistic size.

3.4 Extending the model

Due to the presence of T_{max} in constraints (8) and (9), the linear programming relaxation is weak. For the specific instances of the problem that we want to solve, it is possible to strengthen the model. We introduce an extra binary variable, u^r_{mn}, that equals one if aircraft r travels directly from target m to target n, and zero otherwise. These variables are defined on the set of ordered pairs $(\mathcal{M} \times \mathcal{M}) \setminus \mathcal{S}$. Defining $N_m = N^A_m \cup N^I_m$, that is, the set of nodes associated with target m, the variables u^r_{mn} and x^r_{ij} are coupled by

$$u^r_{mn} = \sum_{i \in N_m} \sum_{j \in N_n} x^r_{ij}, \qquad r \in \mathcal{R},\ (m,n) \in (\mathcal{M} \times \mathcal{M}) \setminus \mathcal{S}.$$

Let K be the number of nonempty subsets of targets. For each such subset, S_k, we define a subtour eliminating constraint in the variables u^r_{mn}, as given below. By adding all such constraints to the MAMPP model, a stronger model is obtained.

$$\sum_{m \in S_k} \sum_{n \in S_k} u^r_{mn} \;\leq\; |S_k| - 1, \qquad r \in \mathcal{R},\; k = 1, \ldots, K$$

Note that even for a scenario with ten targets, there are only at most $2^{10} - 1 = 1023$ such constraints. (An alternative would be to use constraints similar to the capacity inequalities used for capacitated vehicle routing [14]. These inequalities capture both capacity and connectivity requirements.)

Also, as always in VRP problems where resources are identical, symmetry is an issue. It is possible, with no loss of generality, to add constraints stating that the first target is attacked by a specific aircraft and illuminated by another specific aircraft. For the case of only two aircraft, one can also add constraints forcing them to traverse the targets in the same order, that is, enforcing $u^1_{mn} = u^2_{mn}$.

The aircraft fleet can of course be required to be utilized in different ways. For example, an aircraft r_1 can be dedicated to operate pairwise with another aircraft r_2 throughout the mission, which is now easily modelled as $u^{r_1}_{mn} = u^{r_2}_{mn}$. An aircraft r can also be given a dedicated role throughout the mission, that is, performing attacks *or* illuminations only. This is modelled by $x^r_{ij} = 0$ for all $(i, j) \in I_g,\; g \in \mathcal{G}$, and $x^r_{ij} = 0$ for all $(i, j) \in A_g,\; g \in \mathcal{G}$, respectively. Both assumptions, working pairwise and dedicated roles, can be utilized simultaneously.

In the MAMPP model it is only the total mission time, that is, the maximal mission time among the aircraft, that is minimized. Hence, if an aircraft is not critical with respect to total mission time, the model allows solutions where such an aircraft loiters, instead of heading for the destination node as soon as possible. In order to avoid this, we introduce an individual mission time variable for each aircraft, t^r_d, add the constraints

$$t_i + T^r_{id} x^r_{id} - T_{max}(1 - x^r_{id}) \;\leq\; t^r_d, \qquad r \in \mathcal{R},\; (i, d) \in A,$$

and include the individual mission time variables in the objective function with a small penalty.

4 Empirical Testing

In this section we present numerical results. Throughout all scenarios, we use six sectors with three attack positions and two illumination positions in each, that is, the clusters shown in Figure 6. Only one altitude layer is used. Further, $q_m = 1$ holds for all targets, and we choose the parameter value $\mu = 0.05$ in the objective. We start by analyzing a small-size scenario, in order to verify the mathematical model and illustrate characteristics of solutions. Later we present results for a number of larger scenarios.

The scenario described in Figure 1 is solved for two aircraft, using the network representation shown in Figure 5. This results in a network model including 67 nodes and 2830 arcs. (Some nodes are removed due to the protected objects.) All numerical data used in the scenario were provided by our industrial partner.

The model was implemented using AMPL, see [9], and the tests were performed on a HP DL160 server with two 6-core Intel Xeon CPUs and 72 GB of RAM memory, running Linux, and the MIP solver used was CPLEX/12.3, see [11], for 64 bit environment.

4.1 Default settings

In the baseline case, there are no precedence constraints, and $\Gamma = 3$, allowing one resource to attack all three targets. After AMPL and CPLEX preprocessing, the problem consists of 3249 rows, 4213 columns and 4145 binary variables, and it was solved in 62 seconds. The result is presented in Figure 8.

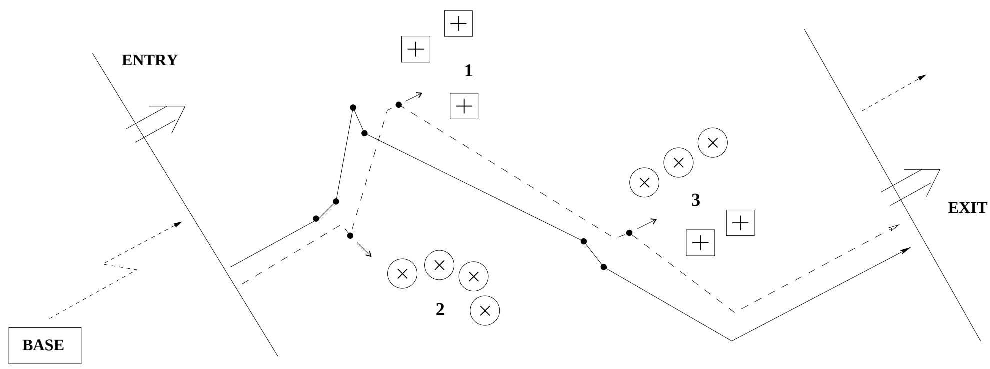

Figure 8: Optimal solution to the example scenario for two aircraft.

The attack sequence becomes 2–1–3, where one aircraft performs all attacks and the other aircraft illuminates the targets, which results in a total mission time of $t_d = 333$ seconds. The expected effect of the attacks against targets 2 and 3 are maximal, among the available attack positions for these targets, while the attack position against target 1 is non-optimal in this respect. The use of an attack position with maximal effect against target 1 would require a longer tour for both aircraft, and therefore this alternative is non-optimal.

4.2 At most two attacks

With the setting $\Gamma = 2$, a single aircraft is not able to attack all targets. As seen in Figure 9, the new solution uses the same attack and illumination positions as before, with the same attack sequence, but the aircraft switch roles against target 2. The total mission time becomes $t_d = 338$ seconds, and the expected effect of the attacks is the same as before.

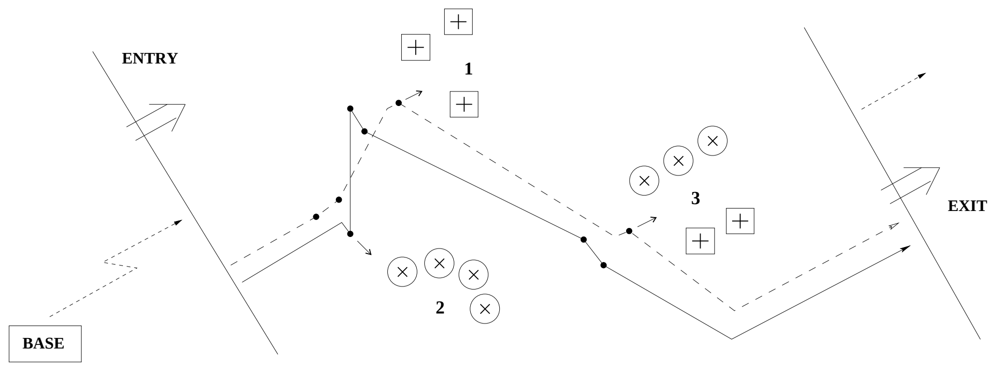

Figure 9: Optimal solution to the example scenario, for two aircraft with at most two attacks each.

After AMPL and CPLEX preprocessing, the problem consists of 3250 rows, 3791 columns and 3723 binary variables, and it was solved in 50 seconds.

4.3 Precedence

For the third case, we set $\Gamma = 3$ again, but impose precedence constraints stating that target 1 must be attacked before both targets 2 and 3. The result is presented in Figure 10. The attack sequence now becomes 1–2–3, which fulfils the precedence relations. Note that the aircraft switch roles even though not forced to do so. The total mission time becomes $t_d = 352$ seconds, and the expected effect of the attacks against targets 1 and 3 are the same as before, while the expected effect against target 2 is lower than before. Comparing the effect of the attack position against target 2 in this solution with the one in previous solutions, one can see that it is lower because of a stronger defense of SAMs from this direction of attack.

After AMPL and CPLEX preprocessing, the problem consists of 2424 rows, 3346 columns and 3278 binary variables, and was solved in 21 seconds.

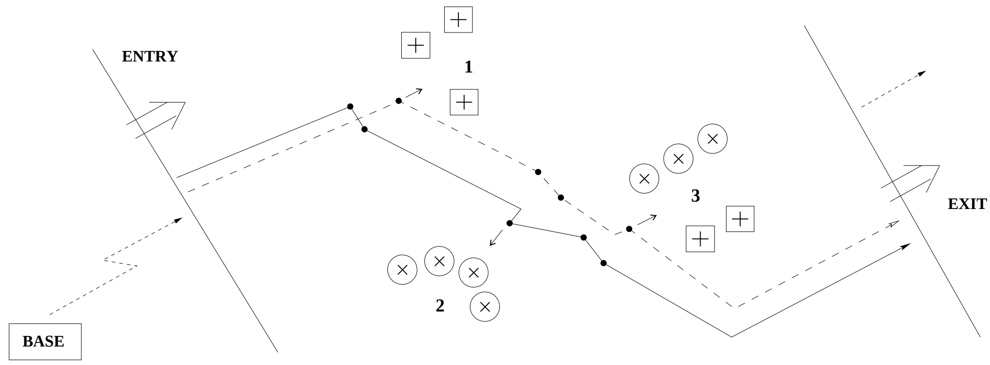

Figure 10: Optimal solution to the example scenario, here with precedence constraints stating that target 1 must be attacked before both targets 2 and 3.

4.4 Dedicated roles

For the fourth case, we specify that one aircraft can only perform attacks and that the second aircraft can only illuminate targets. The precedence constraints are the same as in the previous case, so that target 1 must be attacked before both targets 2 and 3. The result is presented in Figure 11. The attack sequence is still 1–2–3, but the aircraft can not switch roles anymore. The total mission time is again $t_d = 352$ seconds, and the expected effect of the attacks is the same as in the previous case. Targets 1 and 3 are attacked from the same positions while target 2 is attacked from another position.

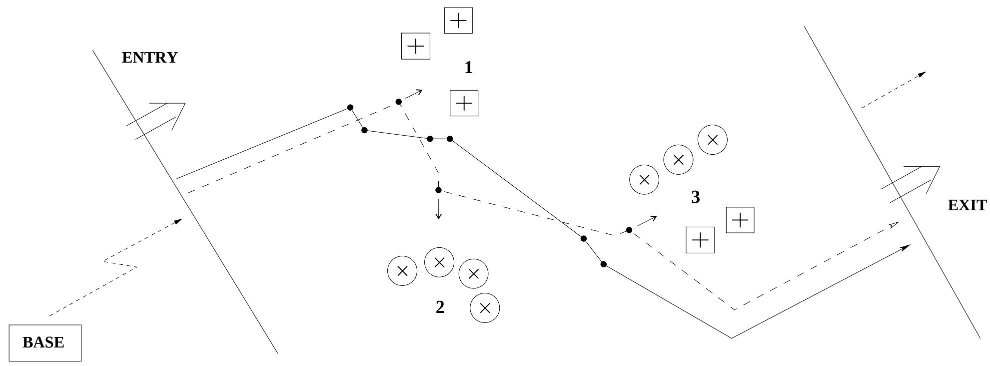

Figure 11: Optimal solution to the example scenario, here with both precedence constraints and dedicated roles.

After AMPL and CPLEX preprocessing, the problem consists of 1791 rows, 1090 columns and 1022 binary variables, and was solved in 9 seconds.

4.5 Larger problem instances

We now consider larger problem scenarios, involving up to six targets and eight aircraft, which are solved under various assumptions. For all problem instances, the aircraft fleet is assigned to work in pairs. Problem characteristics and results are presented in Table 1. An upper limit of three hours was set on the solution time, and failure to prove optimality within this time limit is indicated by a dash.

For all problem instances, the ratio between the best and the worst possible expected effect against each target is in the range 2–4. In the table we give the total expected effect that is actually achieved divided by the maximal possible total expected effect. Hence, if maximal effect is achieved for all targets, then this value is 1.000.

To describe given precedence relations between targets, we introduce the notation `{1234}` for no precedences at all among the four targets, while for example `{12|34}` means that targets 1 and 2 must be visited before targets 3 and 4. Further, for example `{1|2|3|4}` means that a totally ordered attack sequence is set in advance.

The solution found for problem instance no. 18 is illustrated in Figure 12. The attackers are shown as dashed lines, while the illuminating aircraft are shown as solid lines. Maximal effect is achieved against all targets except the third one, which receives the effect 0.60 while the maximal possible effect is 0.65. In order to achieve the maximal effect against the third target, an attack from northwest is required, but this would cause a total mission time so long that this alternative becomes non-optimal.

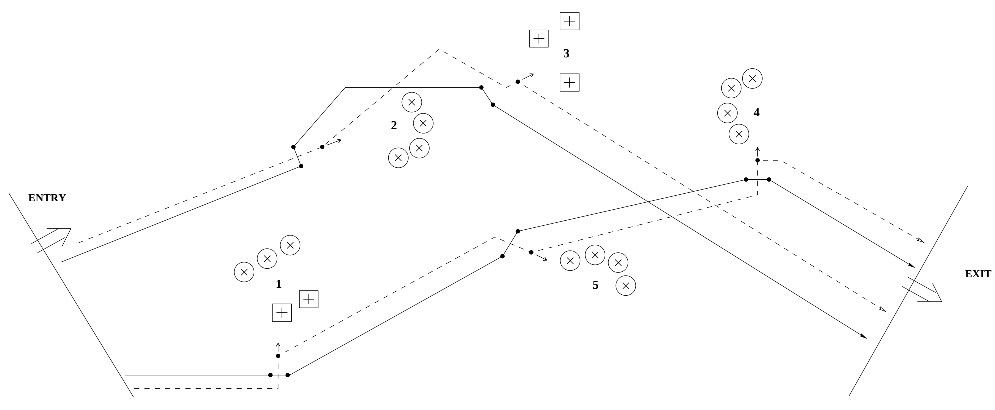

Figure 12: Solution to problem instance no. 18.

As can be seen in Table 1, and as expected, a larger fleet of aircraft enables shorter total mission times and total expected effects that are at least as good.

Table 1: Results for a number of problem instances. Columns $|\mathcal{M}|$ and $|\mathcal{R}|$ state the number of targets and aircraft respectively. *Sequence* defines precedence relations, and Γ states the number of attacks per aircraft. Column *DR* states whether dedicated roles are used (x) or not (-). Column *Bins* gives the number of binary variables, after AMPL and CPLEX preprocessing. Columns *LP* and *IP* state objective values for the root node LP solution and the best optimal integer solution found, while *Time* is the solution time (in seconds). Column *Effect* states total target effect, compared to the maximal possible one, and column t_d states total mission time.

	PROBLEM					SOLUTION					
No.	$\|\mathcal{M}\|$	$\|\mathcal{R}\|$	Sequence	Γ	DR	Bins	LP	IP	Time (s)	Effect	t_d
1	3	2	{123}	3	-	4145	23.173	5.551	62	0.971	333
2				2	-	3723	23.173	5.255	50	0.971	338
3			{1|23}	3	-	3278	22.805	3.014	21	0.857	352
4				3	x	1022	22.745	2.992	9	0.857	352
5	4	2	{1234}	3	-	10799	31.925	1.363	—	1.000	582
6				2	-	9447	31.922	0.915	6169	1.000	591
7			{12|34}	3	-	7094	30.599	1.363	723	1.000	582
8				2	-	6253	30.598	0.915	361	1.000	591
9	4	4	{1234}	3	-	15963	33.277	8.411	4821	1.000	473
10			{12|34}	3	-	11081	32.906	8.411	892	1.000	473
11			{1|2|3|4}	3	-	8128	32.614	7.620	251	1.000	477
12	4	6	{1234}	2	-	7522	34.314	10.919	286	1.000	448
13			{12|34}	2	-	5063	34.246	10.919	115	1.000	448
14			{1|2|3|4}	2	-	3865	33.995	10.489	28	1.000	448
15	5	2	{12345}	3	-	15952	39.876	5.114	—	0.919	625
16			{125|34}	3	-	11259	38.995	1.357	—	0.984	731
17			{1|2|3|4|5}	3	-	3151	38.260	2.254	2163	0.919	678
18	5	4	{12345}	3	x	8537	41.389	12.666	—	0.984	526
19			{125|34}	3	x	6159	41.325	12.666	923	0.984	526
20			{1|2|3|4|5}	2	-	12183	40.443	11.837	4616	0.984	532
21				3	x	4390	40.386	11.838	245	0.984	532
22	5	6	{12345}	2	x	6985	42.680	16.390	2316	0.984	492
23			{125|34}	2	-	18390	42.778	16.390	5687	0.984	492
24				2	x	5063	42.631	16.390	845	0.984	492
25			{1|2|3|4|5}	2	-	14361	42.163	16.150	1704	0.984	492
26	6	4	{2|3|1|4|5|6}	3	x	7249	48.199	9.577	3986	0.987	746
27		6	{2|3|1|4|5|6}	3	x	9030	49.836	13.785	—	0.987	702
28		8	{2|3|1|4|5|6}	3	x	10000	51.230	15.685	6925	0.987	702

Worth mentioning is that solutions with identical effect and total mission time are not always identical, in that the actual routing of the involved aircraft can be very different. This is so because aircraft that do *not* define the total mission time can alter their routes without changing their effects on targets or the total mission time. Also, small differences in objective value for solutions with identical effect and total mission time is due to changes in the individual aircraft mission times.

The LP root node solutions are all found in fractions of a second. In most cases, near-optimal solutions are found rather quickly. However, even for cases that are solved to optimality, the upper bound stays very poor (hundreds of percent above the lower bound) until the very last nodes of the tree search, when it suddenly drops. For the unsolved instances, the final remaining gap is also large, and its value does not really give any indication of the additional time needed to find an optimal solution.

Precedence between targets has a clear effect on the solution times, which seems reasonable as precedence makes the network sparser and decreases the number of binary routing variables. Similarly, if each aircraft has a dedicated role, many binary variables are removed from the problem, hence decreasing solution times.

We have also made a small investigation of the effect of the subtour eliminating constraints on the solution times. As expected, they are of importance only for instances where the routing aspect is a real issue, that is, when there are few precedence restrictions and each aircraft can visit more than one target during its mission (such as instance no. 7). For such problems the increase in solution time is about 25% on average when the subtour eliminating constraints are not included. On the other hand, for instances with more precedence restrictions and where it is likely that most aircraft visit only one target (such as instance no. 25), the inclusion of the subtour eliminating constraints leaves the solution time virtually unchanged, or even increases it; the reason for the latter is that the model is then cluttered with constraints that are not helpful for the solver.

5 Conclusions and Future Work

We have presented and formulated the military aircraft mission planning problem, involving the routing of an aircraft fleet that shall perform attacks against a given set of ground targets. The airspace is discretized and, as a result, attack positions and illumination positions can be represented by nodes, while a priori generated flight paths between these nodes define arcs.

The network construction yields a mathematical model that is recognized as a generalized vehicle routing problem with several side constraints. Each target needs to be visited in a compatible manner, that is, the attack and illumination positions should match. Further, the attack and illumination actions should be synchronized in time. Finally, precedence constraints restrict the order of the attacks.

Finding an optimal solution through direct application of a general MIP solver to the mathematical model is practical only for scenarios of moderate sizes. Even for problem instances including only five targets, it takes CPLEX several hours to verify optimality, although it is able to find feasible and near-optimal solutions much earlier. Hence, efficient heuristics are needed in order to meet the needs and expectations of this application in a real life setting.

Our model compromises expected target effect and mission time duration through a weighted objective function. Further, it was assumed that risk for the aircraft is avoided by not flying through defended airspace. Risk assessment and avoidance is of course not that simple in real life. A subject for further research is to develop solution methods that are able to provide a decision-maker with multiple solutions that are Pareto optimal with respect to the three objectives of target effect, timespan and risk.

Our ongoing work includes a constructive heuristic where the underlying VRP structure is utilized to provide near-optimal solutions, and a column generation inspired approach which exploits the limited number of targets involved in real world scenarios.

A real life issue when implementing a mission plan is that the generated flight paths must not be in conflict with each other, that is, the aircraft must always be separated by some minimum safety distance. This issue is currently not covered by our model and is therefore also subject to ongoing research.

6 Acknowledgements

The authors thank the editor of the journal and the anonymous reviewers for their helpful comments and suggestions.

The final publication is available at http://www.springerlink.com/openurl.asp?genre=article&id=doi:10.1007/s13676-013-0023-3

References

[1] H. Andersson, J.M. Duesund, and K. Fagerholt: Ship routing and scheduling with cargo coupling and synchronization constraints. *Computers & Industrial Engineering* **61**, 1107–1116 (2011).

[2] R. Baldacci, E. Bartolini, and G. Laporte: Some applications of the generalized vehicle routing problem. *Journal of Operational Research Society* **61**, 1072–1077 (2010).

[3] M.C. Bartholomew-Biggs, S.C. Parkhurst, and S.P. Wilson: Using DIRECT to Solve an Aircraft Routing Problem. *Computational Optimization and Applications* **21**, 311–323 (2002).

[4] T. Bektaş, G. Erdoğan, and S. Røpke: Formulations and Branch-and-Cut Algorithms for the Generalized Vehicle Routing Problem. *Transportation Science* **45**, 299–316 (2011).

[5] D. Bredström, and M. Rönnqvist: Combined vehicle routing and scheduling with temporal precedence and synchronization constraints. *European Journal of Operational Research* **191**, 19–31 (2008).

[6] W.M. Carlyle, J.O. Royset, and R.K. Wood: Lagrangian Relaxation and Enumeration for Solving Constrained Shortest-Path Problems. *Networks* **52**(4), 256–270 (2008).

[7] A. Dohn, M.S. Rasmussen, and J. Larsen: The Vehicle Routing Problem with Time Windows and Temporal Dependencies. *Networks* **58**(4), 273–289 (2011).

[8] M. Drexl: Synchronization in Vehicle Routing – A Survey of VRPs with Multiple Synchronization Constraints. *Transportation Science* **46**(3), 297–316 (2012).

[9] R. Fourer, D.M. Gay, and B.W. Kernighan: AMPL – A modeling language for mathematical programming. *Duxbury Press* (2003).

[10] G. Ghiani, and G. Improta: An efficient transformation of the generalized vehicle routing problem. *European Journal of Operational Research* **122**, 11–17 (2000).

[11] ILOG, Inc. ILOG CPLEX: High-performance software for mathematical programming and optimization, (2012).

[12] I. Ioachim, J. Desrosiers, F. Soumis, and N. Bélanger: Fleet assignment and routing with schedule synchronization constraints. *European Journal of Operational Research* **119**, 75–90 (1999).

[13] G. Laporte, A. Asef-Vazari, and C. Sriskandarajah: Some applications of the generalized traveling salesman problem. *Journal of Operational Research Society* **47**, 1461–1467 (1996).

[14] J. Lysgaard, A.N. Letchford, and R.W. Eglese: A new branch-and-cut algorithm for the capacitated vehicle routing problem. *Mathematical Programming, Ser. A* **100**, 423–445 (2004).

[15] L. Moccia, J-F. Cordeau, and G. Laporte: An incremental tabu search heuristic for the generalized vehicle routing problem with time windows. *Journal of Operational Research Society* **63**, 232–244 (2012).

[16] E. Noon and J.C. Bean: An efficient transformation of the generalized traveling salesman problem. *INFOR* **31**, 39–44 (1993).

[17] J. Park and B-I. Kim: The school bus routing problem: A review. *European Journal of Operational Research* **202**, 311–319 (2010).

[18] R. Redjem, S. Kharraja, X. Xie, and E. Marcon: Routing and scheduling of caregivers in home health care with synchronized visits. *International Conference of Modeling, Optimization and Simulation – MOSIM'12*, (2012).

[19] J.O. Royset, W.M. Carlyle, and R.K. Wood: Routing military aircraft with a constrained shortest-path algorithm. *Military Operations Research* **3**, 31–52 (2009).

[20] C. Schumacher, P.R. Chandler, M. Pachter, and L.S. Pachter: Optimization of air vehicles operations using mixed-integer linear programming. *Journal of Operational Research Society* **58**, 516–527 (2007).

[21] M. Sigurd, D. Pisinger, and M. Sig: Scheduling Transportation of Live Animals to Avoid the Spread of Deseases. *Transportation Science* **38**(2), 197–209 (2004).

[22] M. Zabarankin, S. Uryasev, and R. Murphey: Aircraft Routing under the Risk of Detection. *Naval Research Logistics* **53**, 728–747 (2006).