Proceedings of the 2023 Winter Simulation Conference
C. G. Corlu, S. R. Hunter, H. Lam, B. S. Onggo, J. Shortle, and B. Biller, eds.

SIMULATION-BASED OPTIMIZATION OF AIR FORCE MISSION PLANNING

Mihaela Lechner
Alexander Roman
Tobias Uhlig
Oliver Rose

Chair of Modeling and Simulation
University of the Bundeswehr Munich
Werner-Heisenberg-Weg 39
85577 Neubiberg, GERMANY

Thomas Mayer

ESG Elektroniksystem- und Logistik-GmbH
Livry-Gargan-Straße 6
82256 Fürstenfeldbruck, GERMANY

ABSTRACT

Military planning operations deal with highly dynamic environments and a variety of complex optimization challenges. In order to support decision-makers in this process, innovative concepts are required that can automatically generate applicable solutions for certain aspects of mission planning. Such instruments can simplify the planning process, reduce risks, and lower operating costs. This paper presents a simulation-based optimization framework that addresses three problems in the context of aerial warfare planning: task assignment, scheduling, and route planning. These problems are tackled with interconnected heuristics based on either greedy approaches or genetic algorithms. Additionally, hierarchical task networks are employed to incorporate domain knowledge in form of tactical doctrines into the solution. Our simulation results confirm the viability of the proposed approach for small to medium-sized scenarios. However, further investigation with regard to the evaluation function and the simulation environment is required.

1 INTRODUCTION

A well-planned mission is key to successfully conducting a military operation. However, the planning process of such missions is very complex, since large fleets of different vehicle types with diverse capabilities must be carefully coordinated in time and space. Attack aircraft, for example, need to synchronize with defense suppression fighters when approaching a target to minimize the risk of collateral damage. Considering that these different aircraft types are likely to operate from different locations, fly at different speeds and altitudes, or have different detection ranges, a simple strike mission becomes a highly dynamic and complex optimization problem (Zhang et al. 2020). Moreover, modern technologies such as unmanned aerial vehicles (UAVs) increase tactical diversity by performing operations in areas where human intervention is dangerous.

Although it is a time-consuming and complex process with a large decision space, mission planning is currently a decision-making task usually reserved for human operators (Quttineh and Larsson 2015). However, assisting decision-makers with automated solutions can simplify and speed up the planning process, potentially even computing more effective, less risky, and less expensive mission plans.

An automated planning system would need to support decisions such as the number and type of assets to be used for the attack, the type of payload carried by each vehicle, the point of breach through enemy positions, the attack strategies executed by each team, the sequence of attacks, or the flight routes of each aircraft. These components are highly interconnected, making the entire coordination plan one large coupled optimization problem. Despite the widespread interest in technologies that solve such complex mission planning sub-problems, only few studies create combinatorial systems that provide solutions to multiple aspects of mission planning simultaneously. (For more details on related work, the reader should refer

979-8-3503-6966-3/23/$31.00

to Section 2.) This coupled approach might find ambiguous cause-effect relationships between different components and thus globally-optimal solutions instead of focusing on locally optimal answers. Making decisions in isolation is already a difficult task, but it becomes even more complex when attempting to find an optimal solution for interrelated problems, as noted by Bellingham et al. (2003).

The purpose of this paper is to present a framework that facilitates the creation of military mission plans. The approach creates a joint solution concerning what attack strategy to perform against each enemy position (task assignment), what assets to group in executing each attack strategy and the attack order (scheduling), and what trajectories are best to fly in a given environment (route planning). Interconnected heuristics based on either greedy approaches or genetic algorithms are used to find optimal answers to these questions in terms of damage risk and operational cost reduction. In addition, domain knowledge in the form of hierarchical task networks (HTNs) is included in the decision-making process.

The remainder of the paper is organized as follows. Related work is presented in Section 2. A formal problem description is given in Section 3. Section 4 introduces the proposed method and its components. The experimental setting and the simulation results are discussed in Section 5. Section 6 summarizes our conclusions and future work directions.

2 RELATED WORK

Multiple studies conducted in recent years break down the optimization problem into sub-problems and offer solutions to certain aspects of mission planning. Approaches that solve the task assignment problem in a central manner include linear multi-agent network optimization (Schumacher et al. 2002) or genetic algorithm with a simplified matrix representation of chromosomes (Shima et al. 2006). On the other hand, decentralized approaches can be found in the literature. Jin et al. (2003) propose an approach where aircraft autonomously select tasks in real-time using information about the current situation provided by a central instance. Venugopalan et al. (2015) take inspiration from the team theory and implement an algorithm that uses a search certainty map.

Several traditional and intelligent approaches have been used to solve the scheduling problem with a focus on mission planning. Traditional methods include integer linear programming (Griggs et al. 1997), Lagrange relaxation (Ni et al. 2011), or goal programming (Hocaoğlu 2019). However, due to their good performance, intelligent optimization algorithms are currently widely used in the literature. Sonuc et al. (2017) use parallel simulated annealing to accelerate the computation of pairing a large number of weapons to targets. Kong et al. (2021) apply an improved multi-objective swarm optimization algorithm with high convergence speed. Wang et al. (2021) put forward a genetic algorithm-based multi-objective optimization method to generate the number and the type of weapons engaging with different targets.

Several references use traditional means to solve the route planning problem such as determining polygonal paths through a set of threats using the concept of the Voroni-diagram in combination with algorithms such as Dijkstra (Chandler et al. 2000), Eppstein's k-shortest paths (Beard et al. 2002), or artificial potential fields (Eun and Bang 2004). On the other hand, intelligent heuristics including genetic computation (Gao et al. 2005), simulated annealing (Turker et al. 2016), ant colony optimization (Xin et al. 2021), or deep reinforcement learning (Hu et al. 2021) are used for vehicle routing in mission planning.

Several publications solve mission planning with combinatorial optimization methods. Approaches for scheduling and path planning include the conjunction between mixed-integer linear programming and straight-line path approximations (Bellingham et al. 2003), greedy algorithm and Voroni tessellation (Maddula et al. 2004), or Dijkstra algorithm applied in discrete networks and linear bottleneck assignment (Babel 2019). Qie et al. (2019) solve the scheduling and the path planning problem simultaneously by applying a multi-agent deep deterministic policy gradient algorithm.

The literature also provides solutions that attempt to schedule multiple assets to perform tasks on multiple targets while optimizing trajectories to their destinations. Shima et al. (2005) propose a genetic algorithm to collectively execute the task assignment and the scheduling problem. The path planning is built into the fitness function and solved with Dubin's algorithm. Soleyman and Khosla (2020) use hierarchical reinforcement learning to train agents to perform strategic military actions and move around the environment. The assignment of each agent to its nearest target is performed using the Hungarian algorithm.

3 THE MISSION PLANNING PROBLEM

3.1 Problem Description

We consider a suppression or destruction of enemy air defense (SEAD or DEAD) mission in which a group of aerial vehicles aims to cooperatively breach enemy air defense systems, referred to as threats, and engage with predefined targets. While SEAD's goal is to disrupt air defense systems using a combination of electronic warfare, anti-radiation missiles and mere presence, DEAD operations force the physical destruction of threats.

The battlefield is represented as a two-dimensional environment, where assets of both the offensive and defensive teams are distributed in feasible tactical formations. Each side disposes of resources with different abilities. On the one hand, the striking team consists of a variety of aircraft types, ranging from manned to unmanned vehicles, armed with percussion ammunition such as signal-seeking missiles or electric warfare for jamming. The aerial vehicles can change their position in the environment by flying at a constant speed and group together based on predefined attack strategies to engage with targets. When striking, the members of a group must synchronize with each other to minimize the risk of damage. Factors such as the firing range of each asset's weapon or the radar cross-section relative to the target need to be considered. On the defense side, stationary anti-aircraft systems are used to protect high-value unarmed assets, referred to as targets. Each air defense unit is equipped with radars that can detect the position of an attacking aircraft within a unique threat range and strike within a specific lethal range.

The goal of the mission planning system is to automatically generate optimal solutions to the following problems: 1. The task assignment problem allocates each member of the defense team an attack strategy to be carried out against them, taking into account constraints in form of tactical formations, attack angles, timing restrictions, or task precedence. 2. The scheduling problem allocates a group of aircraft to perform each task based on the shortest distance factor. 3. The route planning problem estimates mission plan costs and risks by calculating possible flight trajectories between assets and hostile threats or targets, while complying with spacial, velocity, and safety constraints.

3.2 Notation

We denote $\mathscr{U} = \{U_i | i, N_U \in \mathbb{N} \wedge i \leq N_U\}$ the set of aerial vehicles, $\mathscr{V} = \{V_i | i, N_V \in \mathbb{N} \wedge i \leq N_V\}$ the set of hostile threats, and $\mathscr{W} = \{W_i | i, N_W \in \mathbb{N} \wedge i \leq N_W\}$ the set of targets. $\mathscr{Z} = \mathscr{V} \cup \mathscr{W}$ is used to define the set of length $N_Z \in \mathbb{N}$ that contains all members of the defense side. All resources on the battlefield are modeled as point objects which occupy the location $L_{P_i} = (x_{P_i}, y_{P_i}), \forall P_i \in \mathscr{U} \cup \mathscr{Z}$ in the environment. While threats and targets are stationary, aircraft can move with constant speed σ_{U_i}.

Both the threats $\mathscr{V}$ and the aerial vehicles $\mathscr{U}$ have the ability to use ammunition against their opponents. Air defense systems can detect attacking aircraft at a threat range of up to $\phi^h_{V_i}$ km and strike at a lethal range of up to $\phi^l_{V_i}$ km. Since the positions of the members of the opposing side are known in advance, aircraft within a specified weapon release range ω_{U_i} can attack them. Depending on their capability $\kappa_{U_i} \in \mathscr{K}$, different aircraft types can be used for the attack. Note that for each capability κ only a limited number of resources $N^\kappa_U \in \mathbb{N}$ are available.

3.3 The Mathematical Model

The model of the mission planning problem is inspired by the conceptual model of artificial intelligence (AI) planning as described in Russell and Norvig (2020) and is defined as the 4-tuple $\{\mathscr{S}, \mathscr{A}, \gamma, \mathscr{R}\}$. The following subsections describe each component in detail.

3.3.1 State Space

The state space of length $N_U + N_Z$ can be formulated as $\mathscr{S} = \{L_{U_i} | U_i \in \mathscr{U}\} \cup \{\mathbb{I}_Z(Z_i) | Z_i \in \mathscr{Z}\}$ with

$$\mathbb{I}_Z(Z_i) = \begin{cases} 0, & \text{if enemy asset } i \text{ is destroyed} \\ L_{Z_i}, & \text{otherwise} \end{cases}$$

It contains the coordinates in two-dimensional space of all air vehicles, threats, and targets available in the environment. If a target or threat is destroyed, the coordinate tuple is replaced with the value 0.

3.3.2 Action Space

Each aircraft can either change its position or alter the environment by performing an action. Possible attack strategies on threats and targets are represented as hierarchical task networks (HTNs), an approach inspired by Kiam et al. (2019).

An HTN begins with a goal task (GT) which is progressively broken down into more specific tasks. Depending on the level of abstraction, a distinction is made between compound tasks (CTs), compound task methods (CTMs), and primitive tasks (PTs). CTs are high-level tasks that need to be refined into CTMs. In our case, a CTM represents a specific attack strategy, such as engaging a target with a team consisting of one fighter and one jammer. We denote T_m, $T_m(Z_i)$, $T_m(U')$, or $T_m(U', Z_i)$ as elements of the CTM-set $\mathscr{T}_m$, depending on whether we are referring to a CTM in general, a CTM used against a specific target $Z_i \in \mathscr{Z}$, a CTM performed by a set of resources $U' \subseteq \mathscr{U}$, or a CTM applied by a set of resources against a target. Each CTM is characterized by a set of constraints $C_{T_m} = \{\alpha\} \cup \{(n,\kappa) | n, N_U^\kappa \in \mathbb{N} \wedge T_m \in \mathscr{T}_m \wedge \kappa \in \mathscr{K} \wedge n \leq N_U^\kappa\}$. The tuple (n,κ) denotes the number of assets of a certain capability κ that are needed to fulfill the CTM. α stands for the flight angle between assets.

Once a CTM is selected for each hostile asset $Z_i \in \mathscr{Z}$, the algorithm can perform a series of low-level PTs specific to each CTM and carry out the attack procedure. Similar to CTMs, we use the notation T_p, $T_p(Z_i)$, $T_p(U_i)$, $T_p(U_i, Z_i) \in \mathscr{T}_p$. The difference is that we are now referring to an individual air vehicle $U_i \in \mathscr{U}$ and not an entire set of resources. Moreover, we denote the PT suggesting the movement of an aircraft to a new location L_{U_i} with $T_p(U_i, L_{U_i})$. In the case of this research, a PT represents a direct action that an air vehicle can perform, such as shooting at a target or moving to a specific location.

Subsequently, the action space is the set containing the PTs performed by the members of the aircraft fleet at a given time and can be formulated as the set $\mathscr{A} = \{T_p | T_p \in \mathscr{T}_p\}$.

The main reason for the integration of HTNs in our solution approach is the structured and formal method to include domain knowledge in the planning process. Furthermore, the search space is drastically reduced and incompatible sequences of actions are avoided. Instead of searching the entire action space, the planner is guided by domain knowledge in solving the planning problem.

3.3.3 Transition Function

The transition between states is captured by the deterministic function $\gamma : \mathscr{S} \times \mathscr{A} \to \mathscr{S}$. To reach the state $S' = \gamma(S, A)$, the set of primitive tasks $A \in \mathscr{A}$ must be applied from state $S \in \mathscr{S}$.

3.3.4 Reward Function

The reward function is used to calculate the costs of a mission plan and is defined as

$$\mathscr{R} = \frac{1}{N_U} \sum_{U_i \in \mathscr{U}} d_{U_i} + \dot{d}_{U_i} + \ddot{d}_{U_i} + \sum_{U_j \in \mathscr{U}} \mathbb{I}_U(U_j) + \sum_{V_k \in \mathscr{V}} \mathbb{I}_{\bar{V}}(V_k), \tag{1}$$

having

$$\mathbb{I}_U(U_j) = \begin{cases} 1, & \text{if asset } U_j \text{ is used in the mission} \\ 0, & \text{otherwise} \end{cases} \quad \text{and } \mathbb{I}_{\bar{V}}(V_k) = \begin{cases} -1, & \text{if threat } V_k \text{ is not attacked} \\ 0, & \text{otherwise} \end{cases}.$$

The first sum of the equation (1) refers to distance and safety costs. The greater the distance flown by the aircraft through safe airspace (d) or through potentially threatening ($\dot{d}$) or lethal ($\ddot{d}$) airspace, the greater the costs associated with the mission plan. The second sum penalizes the planner when using too many resources. Furthermore, the planner is rewarded when engaging with as less threats as possible.

The reward function aims to train the decision-maker to create mission plans that are safe for both the human operators and their equipment while keeping costs down.

3.3.5 Goal

The final goal of the decision-maker is to create a set of mission plans $\mathscr{P} = \{P_i | i \in \mathbb{N}\}$ which are Pareto-optimal with respect to the components of the equation (1). This allows end-users to choose a final solution based on their desirability. One may choose costly mission plans over risky ones, or vice versa.

Each mission plan is an ordered sequence that assigns available resources to execute attacks against targets and threats, plans the strike location, and determines the order in which each enemy asset is attacked: $P_i = \{(T_p(U_i), L_{U_i}) | U_i \in \mathscr{U} \wedge T_p \in \mathscr{T}_p \wedge L_{U_i} = (x_{U_i}, y_{U_i}) \wedge x_{U_i}, y_{U_i} \in \mathbb{R}^+\}$. The routing of the fleet through the environment is not directly part of the mission plan but is necessary for calculating the reward function.

4 SOLUTION APPROACH

A detailed schematic overview of the solution approach is shown in Figure 1. It takes as input the environment presented in Section 3, which contains information about assets, threats, targets, and practicable attack strategies. The task manager, the scheduler, and the path planner then work together to create mission plans. Generated plans are subsequently executed by the simulation manager and evaluated against a set of predefined KPIs. The design of each segment is discussed extensively in Sections 4.1-4.4.

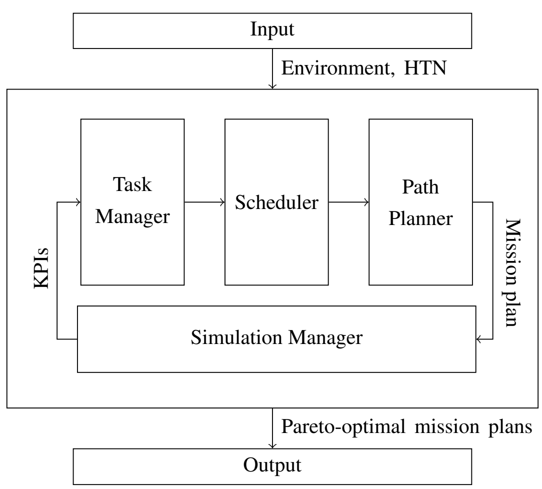

Figure 1: Architecture of the mission planning framework.

4.1 Task Manager

The task manager maps CTMs to targets and threats, defining what attack strategy to use against each member of the defense team. To solve the task allocation problem, we employ a multi-objective evolutionary algorithm (MOEA) provided by the MOEA framework (MOEA 2023).

A genotype used in the algorithm's genetic code is organized as an acyclic graph and consists of N_Z nodes, each representing a target or threat $Z_i \in \mathscr{Z}$. Every node is characterized by multiple components. The first component of the node constitutes the CTM $T_m(Z_i) \in \mathscr{T}_m$ used for the attack. The second component contains a list of parameters ρ_{T_m} specific to each CTM, such as the bearing angle between an aircraft and a member of the defense side. Finally, a node can be formulated as $(T_m(Z_i), \rho_{T_m})$, $i \in \mathbb{N}$. To allow for genetic diversity in the population, a genotype can undergo a mutation that randomly changes the selected CTM, modifies the bearing parameter, or alters the order of attacks by changing the order of CTMs in the graph.

The generated solution candidates are developed into a complete mission plan by being progressively passed to the scheduler and path planner, which are covered in Sections 4.2 and 4.3. Finally, the simulation manager evaluates the fitness of the candidate based on a set of KPIs, as described in Section 4.4.

4.2 Scheduler

The purpose of the scheduler is to find the best subset of resources to execute the CTMs previously planned by the task manager. The selection is based on capability, availability, and distance to the target or threat. The closest resources of the right type, which are not intended to attack other enemy positions at the same time, are chosen to conduct the CTM. To solve the scheduling problem, the greedy algorithm shown in Figure 2 was developed. The current implementation leads to near-optimal solutions in a reasonable amount of time and with low computational effort.

Get set $M = \{T_m(Z_i) | T_m \in \mathscr{T}_m \wedge Z_i \in \mathscr{Z}\}$ defined by the task manager, which assigns one CTM to every target or threat Z_i.
while there exist unplanned $T_m(Z_i)$ in M **do**
    for all unplanned $T_m(Z_i) \in M$ **do**
        Get the constraint set $C_{T_m} = \{\alpha\} \cup \{(n, \kappa) | n \in \mathbb{N} \wedge \kappa \in \mathscr{K}\}$ of the CTM T_m and the location $L_{Z_i} = (x_{Z_i}, y_{Z_i})$ of target or threat Z_i.
        for all capabilities κ in the constraint set C_{T_m} **do**
            Create set U_κ containing all unplanned assets of type κ.
            if size of $U_\kappa < n$, meaning there are not enough assets available to plan the CTM **then**
                Mark $T_m(Z_i)$ as unplanned.
                continue
            for all $U_i \in U_\kappa$ **do**
                Get location $L_{U_i} = (x_{U_i}, y_{U_i})$ of aircraft U_i.
                Calculate Euclidean distance $dist(L_{U_i}, L_{Z_i})$ between aircraft and enemy asset.
            Plan the n-closest assets of capability κ to perform $T_m(Z_i)$.
        Mark $T_m(Z_i)$ as planned.

Figure 2: Algorithm run by the scheduler to allocate resources to CTMs.

4.3 Path Planner

The path planner determines trajectories between two locations (e.g. the take-off position of an aircraft and the location of a threat). Once more, MOEA framework (MOEA 2023) is used to solve the optimization problem.

A path is modeled as a sequence of waypoints, characterized by coordinate tuples in two-dimensional space. Accordingly, an individual's genotype is formulated as $(L^0_{U_i}, ..., L^f_{U_i})$, where $L^0_{U_i}$ and $L^f_{U_i}$ represent the start and destination locations of a specific aircraft $U_i \in \mathscr{U}$. In the search of optimal trajectories, mutation operators add, remove, or relocate waypoints. Once solution candidates are generated, paths between these nodes are approximated using straight lines. Consequently, KPIs such as path length and trajectory risk are calculated and used by the simulation manager in the evaluation process.

4.4 Simulation Manager

The simulation manager is responsible for executing generated mission plans, observing the behavior of both offense and defense sides, and evaluating each plan against a set of predefined KPIs. These indicators are aimed to minimize costs in terms of the number of resources used and the length of the flight routes. Additionally, they are used to control risk by rewarding decision-makers for engaging with a small number of hostile assets and flying safe routes. The reward function used for performance evaluation is shown in equation (1).

5 SIMULATION RESULTS

5.1 Experimental Setup

The presented solution has been implemented and tested in Java 17.0.2 on a standard PC running Windows 11 with 11th generation Intel Core i7-11370H 3.30 GHz CPU and 16 GB RAM.

We evaluated the feasibility and performance of the solution approach using 350×350 2-dimensional environments. 20 aircraft of three different types take part in the combat mission: 5 $\times$ F-16, 3 $\times$ F-35, and 12 $\times$ remote carriers (RCs). Each vehicle has type-specific features which are described in Table 1.

Table 1: Aircraft characteristics.

Type	Ability	Speed	Weapon release point
F-16	Fight and jam	1100 km/h	35 km
F-35	Fight	1500 km/h	33 km
RC	Fight	277 km/h	15 km

The solution includes domain knowledge in form of an HTN, as presented in Section 3.3.2. All tactical doctrines available for the interaction with hostile assets are illustrated in Figure 3. When it comes to CTMs, three methods are implemented: $\mathscr{T}_m = \{Disengage(V_i),\ AttackWithJamming(Z_i),\ AttackWithRC(Z_i) | V_i \in \mathscr{V} \wedge Z_i \in \mathscr{Z}\}$. The planner can either disengage from enemy installations or attack them. While the striking team can disengage from threats, targets always need to be attacked. The method $AttackWithJamming(Z_i)$ pairs two aircraft, one carrying striking ammunition (fighter) and the other possessing a jamming device (jammer), to fight against enemy $Z_i \in \mathscr{Z}$. On the other hand, manned-unmanned teaming can be performed using the $AttackWithRC(Z_i)$ method, in which four RCs help a fighting aircraft to engage a target or threat. Both methods require an angle of 90° between assets. Each CTM breaks down into a sequence of PTs from the set $\mathscr{T}_p = \{MoveTo(U_i, L_{U_i}),\ Shoot(U_i, Z_i),\ Jam(U_i, Z_i) | U_i \in \mathscr{U} \wedge Z_i \in \mathscr{Z} \wedge L_{U_i} = (x_{U_i}, y_{U_i}) \wedge x_{U_i}, y_{U_i} \in \mathbb{R}^+\}$. Changing the position in the environment, shooting, or jamming are the actions that can be performed by the members of the aircraft fleet.

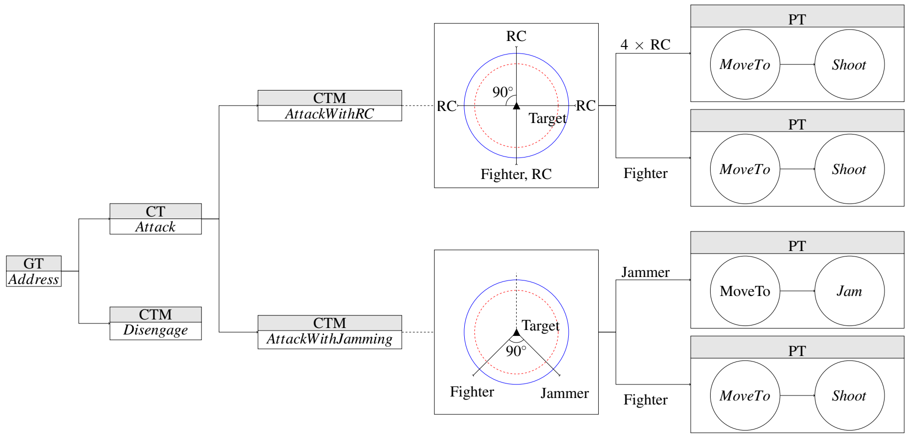

Figure 3: HTN containing tactical doctrines in form of GTs, CTs, and CTMs (levels 1-3). Level 4 constitutes a visual representation of both *AttackWithJamming* and *AttackWithRC* CTMs and their constraints. Level 5 represents the decomposition into PTs.

5.2 Results

Two series of experiments were performed in small and middle-sized environments to demonstrate the feasibility of the solution approach. The first test case consists of a fleet of aircraft trying to break through a small defense team that includes 4 surface-to-air missiles (SAMs) as threats and one target. The threats are organized in a way that protects the target from the attacks. The threat and lethal ranges of the SAMs vary from 16 to 40 km and from 11 to 29 km, respectively. All aircraft are assumed to have the same starting position.

The framework proposes a set of 5 plans as output after creating, simulating, and evaluating more than 500 plans. It takes about 16 minutes for the decision-maker to explore the environment and learn the best solutions for this particular optimization problem. Compared to the beginning of the learning process, the costs of the mission plans calculated according to equation (1) decreased by a factor of 4 on average after the training. This behavior can be seen in Figure 4. It can be noted that the total cost of the plans proposed by the framework varies widely, from about 185 to 1130. This can be argued by the fact that the framework focuses on creating Pareto-efficient plans and not on reducing the total cost function. When closely analyzing the key components of the cost function, namely distance and risk, mission plans with longer flights are preferred to risky trajectories, as presented in Figure 5.

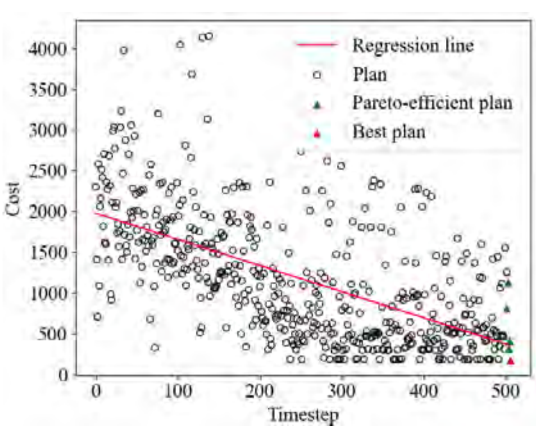

Figure 4: Overall cost per timestep, recorded over 502 timesteps in the basic environment.

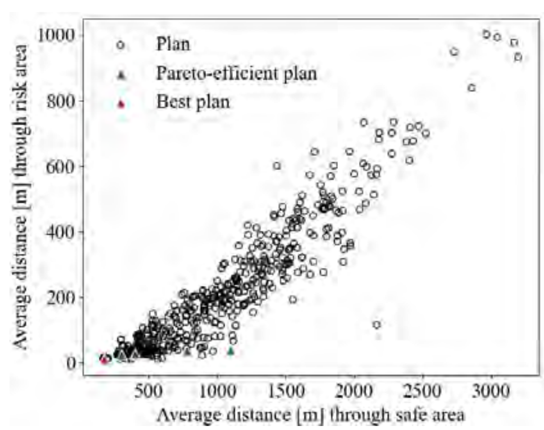

Figure 5: Average distance in safe (d) vs. risky ($\dot{d} + \ddot{d}$) airspace for all 502 plans evaluated.

The overall most cost-effective plan is illustrated in the Figure 6. It brings two air vehicles together to perform the *AttackWithJamming* CTM. An F-35 aircraft takes on the role of jammer, while an F-16 aircraft fires at the target. The planner learns to minimize its costs by disengaging all threats and positioning both attackers as close to the target as possible without exposing the aircraft to great danger. An approximated trajectory connects the starting point and the waypoints. All other plans in the output set are similar variations of this mission plan. Each plan uses the same attack strategy, but places the waypoints in different locations.

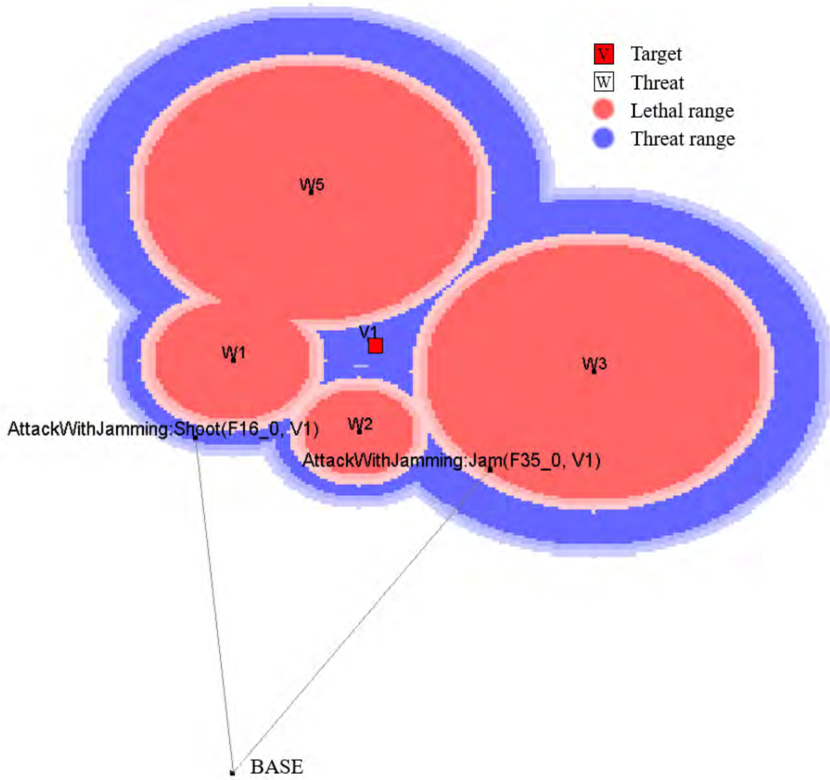

Figure 6: Most cost-effective plan calculated by the framework for the basic scenario. The CTM used is *AttackWithJamming*. The waypoints are positioned in reduced-risk airspace.

The second test case aims to prove the flexibility of the framework to calculate viable mission plans in more complex environments. The fleet of 20 aircraft must hit 7 targets while avoiding 13 threats. The experiment results in a 96-minute training process to output a set of 8 Pareto-efficient plans. The plan with the lowest overall cost is illustrated in Figure 7. The plan's strategy is to breach the lower SAM-belt on the right side by only hitting SAMs with longer lethal and threat ranges (W6, W7, and W8). Short-ranged SAMs, like the W10, W11, and W12, are unable to detect and provide destruction to high-speed air vehicles. This behavior is learned by the decision-maker, who chooses not to shoot at these SAMs even when multiple trajectories pass overhead. The attack sequence is as follows: the first enemy resources attacked are those on the right side of the environment, followed by those in the center, and finally the top left. Hitting the target W6 could have been omitted, as it happens too late in the attack sequence and adds costs without improving the mission plan.

A total of 14 assets are used for the attack in this particular plan. Only 4 assets (1 $\times$ F-16 and 3 $\times$ RC) contribute to one attack only, the rest are utilized in multiple strikes. 7 out of 10 offensive measures are carried out with the *AttackWithJamming* CTM. This behavior is to be expected since the CTM *AttackWithRC* incurs more costs. The total cost caused by this plan is around 1900.

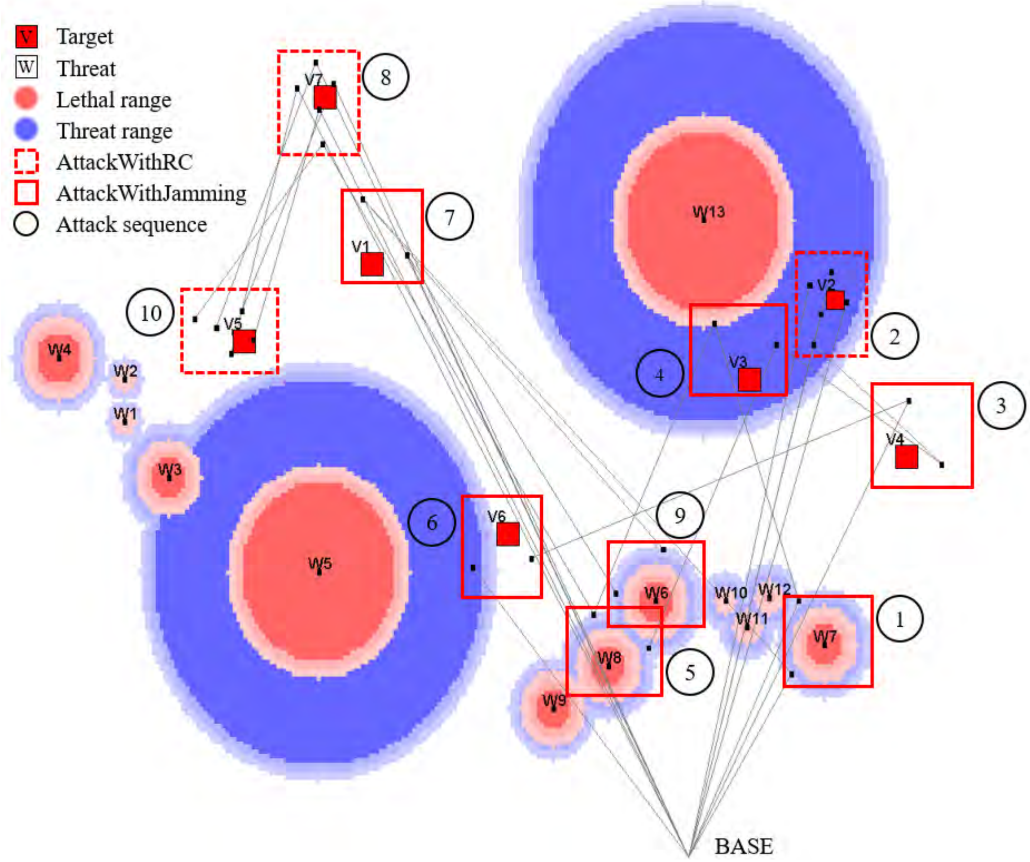

Figure 7: Most cost-effective plan calculated by the framework for the medium scenario. The environment includes 13 threats (W) and 7 targets (V). For the sake of visibility, the labels of the waypoints are not shown. 10 attacks were performed using both *AttackWithJamming* and *AttackWithRC* CTMs. The circled numbers represent the sequence of attacks carried out by the aircraft fleet.

6 CONCLUSION AND OUTLOOK

The importance of a well-planned military mission cannot be overstated as it can significantly impact the success or failure of an operation. However, this involves coordinating a large fleet of diverse vehicles. Often these assets operate from different locations, fly at different speeds or altitudes, and assume various roles, increasing the tactical diversity of a military operation.

This paper proposes an automated mission planning framework that aims to assist the human decision-maker in planning a SEAD or DEAD military operation by providing a set of Pareto-efficient solutions to three coupled problems: task assignment, scheduling, and route planning. These sub-problems are tackled with strongly interrelated heuristics based on either greedy approaches or on genetic algorithms. In addition, domain knowledge organized as an HTN helps accelerate the learning process. The results show that the framework is able to generate feasible solutions to the three optimization problems within a short period of time, especially in small to medium-sized environments.

Nevertheless, further improvements can be added to the solution approach. The subsequent research must refine the evaluation function and the KPIs in cooperation with a subject matter expert (SME). Whether these indicators lead to valid military mission plans has not yet been investigated. It must be analyzed what adjustments need to be made to configuration parameters or performance measures to meet the needs of SMEs. Therefore, extensive experiments have to be performed. Moreover, a more realistic assessment of missions is possible by expanding the simulation environment to a 3-dimensional scenario.

The incorporation of factors such as fuel consumption, flight altitudes, or weather conditions could further increase the realism of the environment. To achieve this, we intend to incorporate external combat simulation software such as CPE (MatrixGames 2023) or FLAMES (Ternion 2023) into the solution.

Although improvements need to be addressed, this study represents an initial step towards developing more efficient and effective mission planning processes.

ACKNOWLEDGMENTS

We thank Airbus Defence and Space, especially Fabian Miehe, for assistance in identifying the most relevant and pressing issues in this research field. We are particularly grateful for their generous support and willingness to share their knowledge and expertise. This study is supported by dtec.bw - Digitalization and Technology Research Center of the Federal Armed Forces via the project MissionLab. dtec.bw is funded by the European Union – NextGenerationEU.

REFERENCES

Babel, L. 2019. "Coordinated Target Assignment and UAV Path Planning with Timing Constraints". *Journal of Intelligent & Robotic Systems* 94(3-4):857–869.

Beard, R. W., T. W. McLain, M. A. Goodrich, and E. P. Anderson. 2002. "Coordinated Target Assignment and Intercept for Unmanned Air Vehicles". *IEEE Transactions on Robotics and Automation* 18(6):911–922.

Bellingham, J., M. Tillerson, A. Richards, and J. P. How. 2003. "Multi-Task Allocation and Path Planning for Cooperating UAVs". In *Cooperative Control: Models, Applications and Algorithms*, edited by S. Butenko, R. Murphey, and P. M. Pardalos, Volume 1, 23–41. Boston, MA, USA: Springer US.

Chandler, P., S. J. Rasmussen, and M. Pachter. 2000. "UAV Cooperative Path Planning". In *AIAA Guidance, Navigation, and Control Conference and Exhibit*, 1255–1265. Denver, CO, USA: American Institute of Aeronautics and Astronautics.

Eun, Y.-J., and H. Bang. 2004. "Cooperative Control of Multiple UCAVs for Suppression of Enemy Air Defense". In *AIAA 3rd "Unmanned Unlimited" Technical Conference, Workshop and Exhibit*, 1–14. Chicago, IL, USA: American Institute of Aeronautics and Astronautics.

Gao, X.-G., X.-W. Fu, and D.-Q. Chen. 2005. "A Genetic-Algorithm-Based Approach to UAV Path Planning Problem". In *Proceedings of the 5th WSEAS International Conference on Simulation, Modelling and Optimization*, 503—507. Corfu, Greece: World Scientific and Engineering Academy and Society (WSEAS).

Griggs, B. J., G. S. Parnell, and L. J. Lehmkuhl. 1997. "An Air Mission Planning Algorithm Using Decision Analysis and Mixed Integer Programming". *Operations Research* 45(5):662–676.

Hocaoğlu, F. M. 2019. "Weapon Target Assignment Optimization for Land Based Multi-Air Defense Systems: A Goal Programming Approach". *Computers & Industrial Engineering* 128:681–689.

Hu, Z., X. Gao, K. Wan, Y. Zhai, and Q. Wang. 2021. "Relevant Experience Learning: A Deep Reinforcement Learning Method for UAV Autonomous Motion Planning in Complex Unknown Environments". *Chinese Journal of Aeronautics* 34(12):187–204.

Jin, Y., A. A. Minai, and M. M. Polycarpou. 2003. "Cooperative Real-Time Search and Task Allocation in UAV Teams". In *42nd IEEE International Conference on Decision and Control*, 7–12. Maui, HI, USA: IEEE.

Kiam, J. J., E. Besada-Portas, V. Hehtke, and A. Schulte. 2019. "GA-Guided Task Planning for Multiple-HAPS in Realistic Time-Varying Operation Environments". In *Proceedings of the Genetic and Evolutionary Computation Conference*, 1232–1240. Prague, Czech Republic: ACM.

Kong, L., J. Wang, and P. Zhao. 2021. "Solving the Dynamic Weapon Target Assignment Problem by an Improved Multiobjective Particle Swarm Optimization Algorithm". *Applied Sciences* 11(19).

Maddula, T., A. A. Minai, and M. M. Polycarpou. 2004. "Multi-Target Assignment and Path Planning for Groups of UAVs". In *Recent Developments in Cooperative Control and Optimization*, edited by R. Murphey, P. M. Pardalos, S. Butenko, R. Murphey, and P. M. Pardalos, Volume 3, 261–272. Boston, MA, USA: Springer US.

MatrixGames 2023. "Command Professional Edition". https://command.matrixgames.com/?page_id=3822. Accessed 29th March 2023.

MOEA 2023. "MOEA Framework". http://moeaframework.org. Accessed 29th March 2023.

Ni, M., Z. Yu, F. Ma, and X. Wu. 2011. "A Lagrange Relaxation Method for Solving Weapon-Target Assignment Problem". *Mathematical Problems in Engineering* 2011:1–10.

Qie, H., D. Shi, T. Shen, X. Xu, Y. Li, and L. Wang. 2019. "Joint Optimization of Multi-UAV Target Assignment and Path Planning Based on Multi-Agent Reinforcement Learning". *IEEE Access* 7:146264–146272.

Quttineh, N.-H., and T. Larsson. 2015. "Military Aircraft Mission Planning: Efficient Model-Based Metaheuristic Approaches". *Optimization Letters* 9(8):1625–1639.

Russell, S., and P. Norvig. 2020. *Artificial Intelligence: A Modern Approach*. 4. edition ed. Hoboken: Pearson.

Schumacher, C., P. R. Chandler, and S. J. Rasmussen. 2002. "Task Allocation for Wide Area Search Munitions". In *Proceedings of the 2002 American Control Conference (IEEE Cat. No.CH37301)*, Volume 3, 1917–1922. Anchorage, AK, USA: IEEE.

Shima, T., S. J. Rasmussen, and A. G. Sparks. 2005. "UAV Cooperative Multiple Task Assignments Using Genetic Algorithms". In *Proceedings of the 2005, American Control Conference, 2005.*, 2989–2994. Portland, OR, USA: IEEE.

Shima, T., S. J. Rasmussen, A. G. Sparks, and K. M. Passino. 2006. "Multiple Task Assignments for Cooperating Uninhabited Aerial Vehicles Using Genetic Algorithms". *Computers & Operations Research* 33(11):3252–3269.

Soleyman, S., and D. Khosla. 2020. "Multi-Agent Mission Planning with Reinforcement Learning". In *AAAI Symposium on the 2nd Workshop on Deep Models and Artificial Intelligence for Defense Applications: Potentials, Theories, Practices, Tools, and Risks*, 51–57. virtual: AAAI.

Sonuc, E., B. Sen, and S. Bayir. 2017. "A Parallel Simulated Annealing Algorithm for Weapon-Target Assignment Problem". *International Journal of Advanced Computer Science and Applications* 8(4):87–92.

Ternion 2023. "FLAMES Simulation Framework". https://flamesframework.com. Accessed 29th March 2023.

Turker, T., G. Yilmaz, and O. K. Sahingoz. 2016. "GPU-Accelerated Flight Route Planning for Multi-UAV Systems Using Simulated Annealing". In *Artificial Intelligence: Methodology, Systems, and Applications*, 279–288. Varna, Bulgaria: Springer International Publishing.

Venugopalan, T., K. Subramanian, and S. Sundaram. 2015. "Multi-UAV Task Allocation: A Team-Based Approach". In *2015 IEEE Symposium Series on Computational Intelligence*, 45–50. Cape Town, South Africa: IEEE.

Wang, C., J. Gao, N. Lv, Y. Cao, S. Zhao, and J. Wu. 2021. "Multi-Objective Optimization of Weapon Target Assignment Based on Genetic Algorithm". In *2021 International Conference on Computer, Internet of Things and Control Engineering (CITCE)*, 29–34. Guangzhou, China: IEEE.

Xin, C., Q. Luo, C. Wang, Z. Yan, and H. Wang. 2021. "Research on Route Planning Based on Improved Ant Colony Algorithm". *Journal of Physics: Conference Series* 1820(1):012180.

Zhang, L. A., J. Xu, D. Gold, J. Hagen, A. Kochhar, A. Lohn, and O. Osoba. 2020. *Air Dominance Through Machine Learning: A Preliminary Exploration of Artificial Intelligence–Assisted Mission Planning*. Santa Monica, CA, USA: RAND Corporation.

AUTHOR BIOGRAPHIES

MIHAELA LECHNER is a Ph.D. candidate in the Modeling and Simulation Chair of the University of the Bundeswehr Munich. Her research interest lies in the design and analysis of both statistical and simulation-based models. The main areas of application are logistics and military mission planning. Her e-mail address is mihaela.hanea@unibw.de.

ALEXANDER ROMAN is currently pursuing a Ph.D. in the Modeling and Simulation Chair at the University of the Bundeswehr Munich. He is primarily focused on designing and evaluating artificial intelligence and simulation models, with a specific emphasis on their use in military mission planning. His email address is alexander.roman@unibw.de.

THOMAS MAYER is working as data architect for ESG, a leading German technology and innovation partner for defense and public safety. He holds a Ph.D. degree in computer science from the University of the Bundeswehr Munich. His work focuses on the application of AI in military decision making and planning processes. His email address is thomas3.mayer@esg.de.

TOBIAS UHLIG is a postdoctoral researcher at the University of the Bundeswehr Munich, Germany. He holds a M.Sc. degree in computer science from Dresden University of Technology and a Ph.D. degree in computer science from the University of the Bundeswehr Munich. His research interests include operational modeling, logistics, natural computing and heuristic optimization. He is a member of the ASIM and he is one of the founding members of the ASIM SPL workgroup on the Investigation of Energy-related Influences in SPL. His email address is tobias.uhlig@unibw.de.

OLIVER ROSE holds the Chair for Modeling and Simulation at the Department of Computer Science of the University of the Bundeswehr Munich, Germany. He received a M.S. degree in applied mathematics (1992) and a Ph.D. degree in computer science (1997) from Würzburg University, Germany. His research focuses on the operational modeling, analysis and material flow control of complex manufacturing facilities, in particular, semiconductor factories and assembly systems. He is a member of INFORMS Simulation Society, ASIM (German Simulation Society), and GI (German Computer Science Society). Currently, he is member of the board of the ASIM and the ASIM representative at the Board of Directors of the WSC. His email address is oliver.rose@unibw.de.