Effect Oriented Planning of Joint Attacks

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

Abstract We consider tactical planning of a military operation on a large target scene where a number of specific targets of interest are positioned, using a given number of resources which can be for example fighter aircraft, unmanned aerial vehicles, or missiles. The targets could be radar stations or other surveillance equipment, with or without defensive capabilities, which the attacker wishes to destroy. Further, some of the targets are defended, by for example Surface-to-Air Missile units, and this defense capability can be used to protect also other targets. The attacker has knowledge about the positions of all the targets and also a reward associated with each target. We consider the problem of the attacker, who has the objective to maximize the expected outcome of a joint attack against the enemy.

The decisions that can be taken by the attacker concern the allocation of the resources to the targets and what tactics to use against each target. We present a mathematical model for the attacker's problem. The model is similar to a generalized assignment problem, but with a complex objective function that makes it intractable for large problem instances. We present approximate models that can be used to provide upper and lower bounds on the optimal value, and also provide heuristic solution approaches that are able to successfully provide near-optimal solutions to a number of scenarios.

Nils-Hassan Quttineh
Department of Mathematics, Linköping University, Sweden, e-mail: nils-hassan.quttineh@liu.se

Torbjörn Larsson
Department of Mathematics, Linköping University, Sweden, e-mail: torbjorn.larsson@liu.se

Kristian Lundberg
Department of Mathematics, Linköping University, Sweden, e-mail: kristian.lundberg@liu.se

Kaj Holmberg
Department of Mathematics, Linköping University, Sweden, e-mail: kaj.holmberg@liu.se

1 Introduction

Effect Based Operations (EBO) is a military concept which emerged during the 1991 Gulf war for the planning and conduct of operations combining military and non-military methods to achieve a particular effect. The doctrine was developed to take advantage of advancements in weaponry and tactics, from an emerging understanding that attacking a second-order target may have first order consequences for a variety of objectives. The Commander's intent can be satisfied with a minimum of collateral damage or risk to own forces, but EBO planning is complex and hard since it embraces political factors as well as economic.

Despite its complexity, this is not an impossible task. We have been dealing with these challenges on an ad hoc basis throughout history, but we can now use modern technologies and process thinking to provide all ingredients of successful effect based operations.

A network-centric system is a system-of-systems concept where a number of actors are attached to each other in a network sharing information in an adaptable and interoperable manner. Obviously networking enables an enormous rise in accessible information and the intrinsic challenge is the development of systems and functions to shape this information into guidance and control of a variety of operations with multiple objectives. For example, [6] presents an optimization methodology for finding a correct balance between weapons and attack damage assessment sensors.

The above mentioned pinpoints the trend in military operational planning, also at the Swedish military arena. In our case we can use this paradigm shift to put functional and algorithmic requirements on planning of air to ground missions. This leads to adaptation to new doctrines of command and control and to a tool that contains the most of planning experience implemented by planning specialist personnel in cooperation with algorithm experts. Mission performance can be driven to its limits with a model based planning, which simultaneously keeps control of both objective and system performance, which is probably the most cost effective way to gain performance.

1.1 Network Centric Framework

In a network centric framework, a resource is not an entity tightly coupled to a sluggish hierarchical organization but a resource with own intelligence to offer specific effects to a variety of effect customers. Our work does not embrace the full meaning of EBO but is guided by quantifying and responding to effect requests and hence becoming a true entity of a network centric system. In order to understand the paradigm shift in EBO planning or network centric planning, Figure 1 shows the principles of future effect based operations.

Initially an effect must be achieved in order to answer what to do. Thereafter possible systems are considered and how theses systems could manage to do it. The last issue of the effect chain is to decide the resource allocation. As can be noticed,

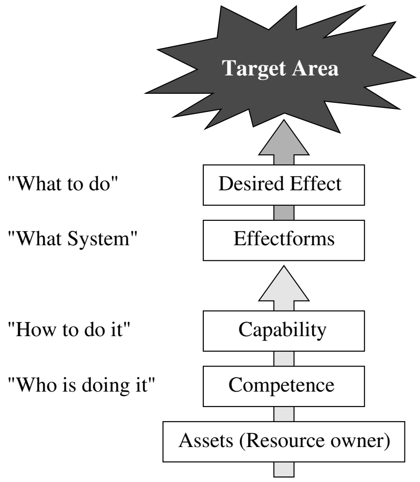

Fig. 1 The effect chain including an EBO principle of a split up of the planning process into stages from the target to allocation of individual platforms.

resource owners are considered in the later planning stages, which is quite a change from traditional planning.

Obviously there are two dimensions in the effect chain, the mission-conduction and the resource owner dimensions. The resource owner dimension keeps and conducts resource supply chains as well as allocation schemes and schedules. The mission-conduction states individual missions and how they shall be implemented.

In order to fulfil requirements on future EBO planning systems, effort must be put on scalable model-based algorithms which promote an easy workflow and a high speed planning performance. Each scenario shall be individually stated by the set of input data, but planning shall always be performed via implemented tactics and knowledge of actual resource performance and mission pattern.

1.2 Mission Planning

An air to ground mission planning system is modular and contains a planning system and weapon systems, hosted by a variety of carriers such as unmanned aerial vehicles or fighter aircraft. In order to perform effect oriented planning in line with Figure 1 we transform the planning process according to Figure 2, where each platform is separated into carrier and weapon performance and tactics producing a certain effect which can be matched with the effect customers needs.

Initially we maximize system effect in the target area by optimally allocating the number of weapons to suppress enemy defense and destroy vital targets. A target area can consist of different ground based targets and sheltering air defense units. Each target has a specific value which indicates its importance. The effect

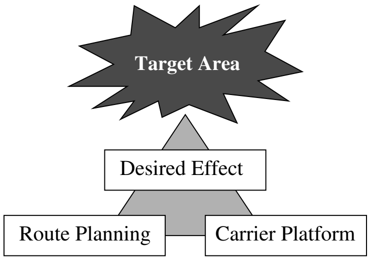

Fig. 2 A resource, a fighter with weapon system, has a relationship between route planning, type of weapon and a set up of tactics which forms the final effect.

oriented weapon allocation of the target area is followed by a search for appropriate platforms, where platform location and scheduling parameters are considered. Each platform must further have a route to the firing position, including tactical features such as hiding and a limited exposure of radar cross section during the flight phase.

These planning aspects are coupled, but with an acceptable loss of generality the effect planning task can be separated from the platform in order to start an overall planning process. Our work addresses a model based approach to rapidly calculate weapon allocation to optimize system effect in an hostile ground based target area. Early work on a similar problem was done by Miercort and Soland in [4], but they consider a less complicated model without intricate dependencies. In a recent paper by Kwon et al. [3], a new weapon-target allocation problem is presented together with a branch-and-price algorithm for solving it. In contrast, Kaminer and Ben-Asher present a model in [2] for maximizing the effectiveness of a defense.

1.3 Paper Overview

In Section 2 we describe the problem at hand, which is basically a weapon-targeting problem, together with some basic concepts that will be used throughout the paper. Section 3 gives a generic mathematical model for the problem. It is straightforward with only simple linear constraints, but comes with a difficult objective function. This section also gives optimistic and pessimistic models that can be used to find upper and lower bounds on the optimal objective value.

In order to use the generic model and solve realistic scenarios, it is necessary to specify how to evaluate a given situation, and especially how the defenders act in different situations. One possible way to do this is presented in Section 4,

Section 5 looks into different heuristic approaches, who cannot guarantee optimality but find high quality solutions for larger scenarios within reasonable time frames. Section 6 contains results for these heuristics. Finally, in Section 7, we present some remarks and conclusions together with suggestions on future work. This paper is based on material that can be found in [5].

2 The Joint Attack Problem

Imagine a large open area, like a desert, where a number of enemy targets are positioned. These can be radar stations or other surveillance equipment, which the attacker wishes to destroy. The targets are however guarded by defenders, like Surface-to-Air Missile (SAM) units. The defenders are also considered to be potential targets for the attack, since the destruction of defenders can improve upon the overall outcome of the attack.

The positions of all targets, both those with and without defense, are known. The set of targets is denoted S, and the subset $\bar{S}$ denotes the targets with defensive capabilities, which are defined by radii of defense and armament. Each target $s \in S$ is given a specified reward r_s, where important targets have higher values.

The attacker's problem is to maximize the expected outcome of a simultaneous attack against the enemy, using at most R identical resources, like aircraft or unmanned aerial vehicles. Each target should be assigned an attack plan which specifies the number of resources to be used against it, and also from which directions.

As illustrated in Figure 3, some targets do not have a defensive system of their own, but depends on the defense of others. Also, the radius of defense for different

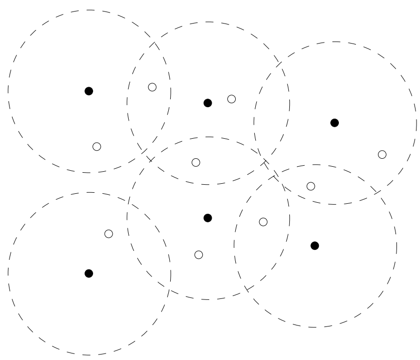

Fig. 3 A possible attack scenario. Some targets, here shown in black, are air defense units. The other targets are radar stations or similar surveillance units who are valuable to destroy.

defenders might overlap. A defender will always protect itself primarily, and then engage resources passing by inside its radius of defense towards other targets.

2.1 Tactics and Angles of Attack

If a target s is attacked, it is done so by a tactic t chosen from a set of tactics, T. In real life there are numerous possible tactics for an attack, but we limit ourselves to tactics using at most 3 resources, as described graphically in Figure 4. The idea behind these tactics is to overload the defensive system of a single defender. This is done either by sending multiple resources from one direction (see tactics 1–3), or by attacking from multiple and evenly spread directions (see tactics 4–5).

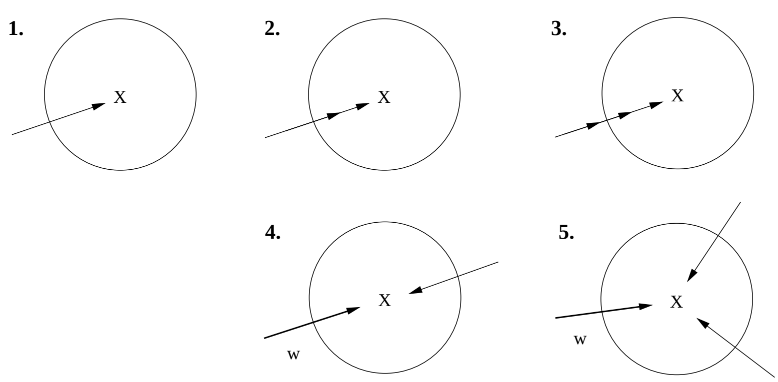

Fig. 4 A graphical description of the 5 tactics considered.

Each tactic t has its own features, such as the number of resources needed, n_t, and the number of attacking directions involved, denoted V_t. The number of resources that is launched from each of the angles V_t is denoted by m_t. Each tactic gives rise to a probability of success, for each of the n_t resources, against a single target s. This probability is denoted p_{st} and might vary between the targets, depending on their respective defensive capabilities.

We consider a coarse angle discretization (every 30 degree), defining a set V of angles. Each tactic $t \in T$ is associated with a reference angle of attack, w, which defines from which direction the attack is launched. For tactics which involve more than one angle of attack (i.e. $V_t > 1$), multiple angles w might give rise to exactly the same attack, since we consider evenly spread angles. To avoid such symmetries, we introduce a subset W_{st} which contains all reference angles w to be used together with tactic t against target s.

For tactics involving multiple angles, we define

$$w_j = w + (j-1) \cdot \frac{2\pi}{V_t}\ , \quad j = 1, \ldots, V_t.$$

We also introduce the concept of an engagement path (s, v), which is the line emanating from target s at angle v. In total, there are $|S| \cdot |V|$ different engagement paths. For a certain tactic and angle, though, only a few of these paths will be used. If there is at least one resource on the path, we call it an active path.

In the following, a reference angle of attack is always denoted w and defined by the set W_{st}, whereas an angle v refers to an individual angle in V used for general discussions involving engagement paths (s, v).

2.2 The Objective

The essence of the attacker's problem is to decide for each target s which tactic t that shall be used (if any) and specify a reference angle of attack w. We therefore introduce the binary variable

$$z_{stw} = \begin{cases} 1 \text{ if target } s \text{ is attacked using tactic } t \text{ from angle } w, \\ 0 \text{ otherwise.} \end{cases}$$

These decisions, at most one for each target s, is defined as an attack plan $\mathbf{z}$. Let $p^{\text{kill}}_{stw}(\mathbf{z})$ be the probability of successfully incapacitating target s when attacked by tactic t from reference angle w. As will be clear from the upcoming analysis, this probability depends on the overall attack plan $\mathbf{z}$, which is a complicating fact.

The probability for a resource to survive the defense of a defender $i \in \bar{S}$ which it passes by on its way towards the target s on path (s,v) is denoted $p_{isv}(\mathbf{z})$, and it depends on what tactics are used against the other targets. Whenever an engagement path (s,v) does not intersect the area of defense for target i, $p_{isv}(\mathbf{z}) = 1$ holds.

The success of an attack against a certain target depends on the following.

1. The number of resources used against the target ($n_t = V_t \cdot m_t$).
2. The target's ability to defend itself against incoming resources (p_{st}).
3. The probability of successfully surviving the defense of every other target which the resource pass by on its way towards the target (p_{isw_j}).

For a given target s, tactic t and angle of attack w, the probability of successfully eliminating target s is

$$p^{\text{kill}}_{stw}(\mathbf{z}) = 1 - \prod_{j=1}^{V_t} \left[1 - p_{st} \prod_{i \in \bar{S} \setminus \{s\}} p_{isw_j}(\mathbf{z}) \right]^{m_t}. \tag{1}$$

The probability of success for a tactic t and angle w against a target s generally depends on which tactics are applied against every other target, that is, the whole attack plan, which means that the probabilities p_{isw_j} are related to each other. This dependence is the core difficulty of the attacker's problem.

The objective is to maximize the expected total reward of the attack, found by multiplying the probability of success of an attack against a target with its reward. Since we want to optimize the total reward of the attack, these expected values should be added. The objective then becomes

$$\max \sum_{s \in S} \left[\sum_{t \in T} \sum_{w \in W_{st}} p^{\text{kill}}_{stw}(\mathbf{z}) \cdot z_{stw} \right] \cdot r_s.$$

For each target $s \in S$, at most one of the decision variables z_{stw}, $t \in T$, $w \in W_{st}$, takes the value one, since it is attacked at most once.

3 Mathematical Models

We here give a generic model for the joint attack problem and two approximate models that can be used to find upper and lower bounds on the optimal value.

3.1 A Generic Model

As stressed above, the probability $p_{isv}(\mathbf{z})$ depends in general on the whole attack plan $\mathbf{z}$, but in the generic model we make no assumptions on the exact nature of this dependence.

$$
\begin{aligned}
\max \sum_{s\in S} & \left[\sum_{t\in T} \sum_{w\in W_{st}} p^{\text{kill}}_{stw}(\mathbf{z}) \cdot z_{stw} \right] \cdot r_s && [GENERIC] \\
s.t. \quad & \sum_{s} \sum_{t} \sum_{w\in W_{st}} n_t \cdot z_{stw} \leq R && (i) \\
& \sum_{t} \sum_{w\in W_{st}} z_{stw} \leq 1, \quad s \in S && (ii) \\
& z_{stw} \in \{0,1\}, \quad s\in S, t\in T, w\in W_{st}
\end{aligned}
$$

It is not necessary to attack all targets. Depending on the rewards specified for the targets, it might be optimal not to do so. Constraint (i) states that we cannot use more resources than we have. Constraint (ii) makes sure that each target is attacked at most once. Both constraints are linear, but the objective is in general non-linear, non-convex and non-separable.

3.2 Optimistic Model

It is possible to construct two auxiliary problems, that provide upper and lower bounds, respectively, on the optimal value of the generic problem. We analyze the expression for $p^{\text{kill}}_{stw}(\mathbf{z})$, under two specific assumptions.

Assume that no target will shoot against resources passing by towards other targets, but just against resources targeting themselves. This means that $p_{isv}(\mathbf{z}) = 1$ would hold for all targets $s \in \underline{S}$, and that $p^{\text{kill}}_{stw}(\mathbf{z})$ would collapse into the quantity

$$
P_{st} = 1 - \prod_{j=1}^{V_t} \left[1 - p_{st} \prod_{i\in \underline{S}\setminus\{s\}} 1 \right]^{m_t} = 1 - (1-p_{st})^{n_t}.
$$

Now the probabilities of success no longer depend on the overall attack plan $\mathbf{z}$. Further, since this expression does not depend on the angle w anymore, we only have to decide which tactic t to use against each target s, if any tactic at all.

We then obtain the optimistic model

$$\begin{aligned}
\max \ & \sum_s \sum_t r_s \cdot P_{st} \cdot z_{st} & & & [OPTIMISTIC] \\
s.t. \ & \sum_s \sum_t n_t \cdot z_{st} \leq R & & & (i) \\
& \sum_t z_{st} \quad \leq \ 1, & & s \in S & (ii) \\
& z_{st} \quad \in \{0,1\}, & & s \in S, t \in T. &
\end{aligned}$$

Solutions to the optimistic model give upper bounds to the original problem, since the values of all coefficients in the objective function are systematically increased. Even more, this is a valid upper bound for all choices of discretization V.

The solution found is also a feasible solution in the original problem, if complemented with an arbitrary reference angle of attack for each tactic used. This means that we can easily calculate a true objective value and also get a lower bound. This bound is only valid for the considered discretization V though.

3.3 Pessimistic Model

In contrast to the assumption made above, we now assume that each target will shoot against all resources passing by, on their paths towards other targets, and with its full defensiv capability. Denote by $\tilde{p}_{isv}$ the resulting probability of surviving the defense from another target. This probability is clearly a pessimistic estimate of the true probability of surviving the defense from this target.

If $p_{isv}(\mathbf{z}) = \tilde{p}_{isv}$ would always hold, then $p^{\text{kill}}_{stw}(\mathbf{z})$ would become the quantity

$$P_{stw} \ = \ 1 - \prod_{j=1}^{V_t} \left[1 - p_{st} \prod_{i \in S \setminus \{s\}} \tilde{p}_{isw_j} \right]^{m_t},$$

and we then obtain the pessimistic model

$$
\begin{array}{llll}
\max & \sum_s \sum_t r_s \cdot P_{stw} \cdot z_{st} & & [PESSIMISTIC] \\
s.t. & \sum_s \sum_t \sum_{w \in W_{st}} n_t \cdot z_{stw} \leq R & & (i) \\
& \sum_t \sum_{w \in W_{st}} z_{stw} \quad \leq \ 1, & s \in S & (ii) \\
& z_{stw} \quad \in \{0,1\}, & s \in S, t \in T, w \in W_{st}. &
\end{array}
$$

The values of $\tilde{p}_{isv}$ might of course be too pessimistic, and hence the solution could provide poor lower bounds on the optimal value of the generic model. Hopefully, though, the structure of the solution (the attack plan $\mathbf{z}$) is close to the optimal one, and by evaluating the true objective one can find a better pessimistic bound.

4 Simulation Details

In order to fully specify the generic model presented in Section 3.1, one needs to describe how the probability $p_{isv}(\mathbf{z})$ depends on the attack plan $\mathbf{z}$. It is obviously a hard task to model a real-life situation. We will here give the assumptions used in our simulation study.

We will analyze the different factors that affect $p_{isv}(\mathbf{z})$, that is, the probability for a resource to survive the defense from another target as it passes by toward its own target, and how it depends on $\mathbf{z}$. To do this, we look into the details of the defensive systems of the targets and define their rules of engagement.

4.1 Specifications of the Defensive System

Since we consider the problem of the attacker, we need to specify a set of deterministic engagement rules for the defenders. Each target with defensive capability is assumed to have a specified number of defensive channels, such as cannons or anti-missile systems. It will primarily defend itself, and any residual defensive channels will be used to defend the other targets, by engaging resources passing by inside its radius of defense. We make the following assumptions for each defender $i \in \bar{S}$.

1. The defender will primarily defend itself.
2. If there are $D_i > 0$ residual defensive channels, then they will be evenly allocated against the active engagement paths that pass by the target.
3. At most F_i channels might be used against a single engagement path.
4. At most G_i different engagement paths might be engaged.
5. All defensive channels should be used if there is something to shoot at.

6. If there are more active paths than defensive channels, one defensive channel is allocated to each path as long as possible with respect to a ranking defined by the distances to the target.

Given an attack plan $\mathbf{z}$, we let auxiliary variables $u_{isv}(\mathbf{z})$ describe how many defensive channels that should be allocated against the resources on each active path (s,v) passing by. The values of these variables will comply with the above rules.

Specifically, the number of resources on each path, denoted N_{sv}, affects the probability of success for each of these resources. We define $K = \max_{t\in T}\{n_t \ : \ V_t = 1\}$ to be the maximum number of resources travelling on a single engagement path. Hence, N_{sv} is in the range $k = 0,\ldots,K$.

We further define the parameter d_{isv} to be the orthogonal distance between a target $i \in \overline{S}$ and the engagement path (s,v). For other targets with positions inside the area of defense of target i, the distance to the mid-point of this path is used. This is illustrated in Figure 5. Each active path is given a rank number, where the path closest to target i gets the highest rank, the second closest path gets the second rank, and so on. Closest path refers to the smallest distance d_{isv} and is thus relative to the target i. This ranking will be used when the defenders cannot engage all paths passing by, but need to prioritize.

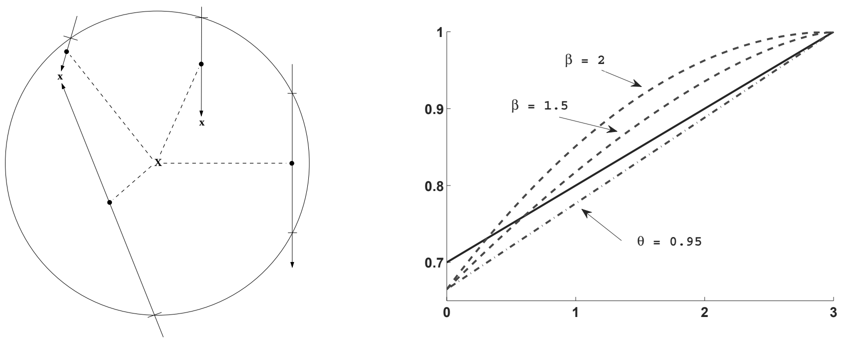

Fig. 5 To the left, an illustration of how the distance between a target and the active engagement path is measured. To the right, an example of how the design parameters β_{ik} and θ_{ik} affect the probability p^k_{isv}.

4.2 Specification of the Objective

The probability for a resource to survive as it passes by target $i \in \overline{S}$ towards target $s \in S$ on path (s,v) is a function of the distance d_{isv} and the number of resources N_{sv} on the path, which are both a direct consequence of the attack plan $\mathbf{z}$. The obvious way to model this dependence would be to demand values for all such combinations as input data, but this is practically impossible. We instead introduce an analytic expression, based on both d_{isv} and N_{sv}.

Let p^k_{isv} be the probability for a resource to successfully pass by one defensive channel of target i. These probabilities are derived from the values of p_{st}, for tactics $t \in T$ where all $k = n_t$ resources are sent from the same angle $(V_t = 1)$. Since this is only relevant for targets in $\overline{S}$, we denote this p_{ik} for all $i \in \overline{S}$ and $k = 1, \ldots, K$.

$$p^k_{isv} \;=\; 1 - \left(1 - \frac{d_{isv}}{\rho_i}\right)^{\beta_{ik}} \cdot (1 - \theta_{ik} \cdot p_{ik})$$

Here, ρ_i is the radius of defense, while β_{ik} and θ_{ik} are design parameters that model the defensive capacities of target i against different numbers of resources k.

The rightmost plot in Figure 5 shows the probability p^k_{isv} on the y-axis as a function of the distance d_{isv} on the x-axis. Here, the probability $p_{ik} = 0.7$ is used, and the solid line corresponds to parameter values $\beta_{ik} = 1$ and $\theta_{ik} = 1$. The dash-dotted line is obtained when the value of θ_{ik} is changed to 0.95. The two dashed curves correspond to the values of 1.5 and 2 respectively for parameter β_{ik}. In all, this expression for p^k_{isv} shows a reasonable behaviour. For $d_{isv} = 0$, its value becomes $\theta_{ik} \cdot p_{ik}$ and for $d_{isv} = \rho_i$ the probability becomes 1. For distances in-between, the parameter β_{ik} is used to model the effectiveness of the defensive system of target i.

Now finally, the probability for a resource to survive as it passes by target $i \in \overline{S}$ towards target $s \in S$ on path (s, v), given the attack plan $\mathbf{z}$, is

$$p_{isv}(\mathbf{z}) \;=\; \prod_{k=1}^{K} \left(p^k_{isv}\right)^{u^k_{isv}}.$$

Here, the auxiliary variable u^k_{isv} equals $u_{isv}(\mathbf{z})$ if $k = N_{sv}$ and zero otherwise. Since $u_{isv}(\mathbf{z})$, and thus also u^k_{isv}, might be greater than one the probability of success decreases with the number of defensive channels assigned to the engagement path. This is realistic as the defensive channels can be seen as independent, and the probability for a resource to survive two channels should be the probability of surviving them both. The general formula (1) now becomes

$$p^{\text{kill}}_{stw}(\mathbf{z}) \;=\; 1 - \prod_{j=1}^{V_t} \left[1 - p_{st} \prod_{i \in \overline{S} \setminus \{s\}} \prod_{k=1}^{K} \left(p^k_{isw_j}\right)^{u^k_{isw_j}}\right]^{m_t}.$$

The values of the variables u^k_{isv} are dependent on the entire attack plan $\mathbf{z}$. Once their values are known, it is however straightforward to evaluate the objective of the generic model.

4.3 An Illustrative Example

Consider a single defender i, as illustrated in Figure 6. We name all paths (s,v) intersecting the area of defense in accordance with their rank, that is, the path with rank 1 is named path 1, and so on. Notice that one of the engagement paths never intersects the area of defense, and it is therefore never considered when the residual defensive channels are assigned. We assume that at most 3 channels might be used against a single engagement path (i.e., $F_i = 3$).

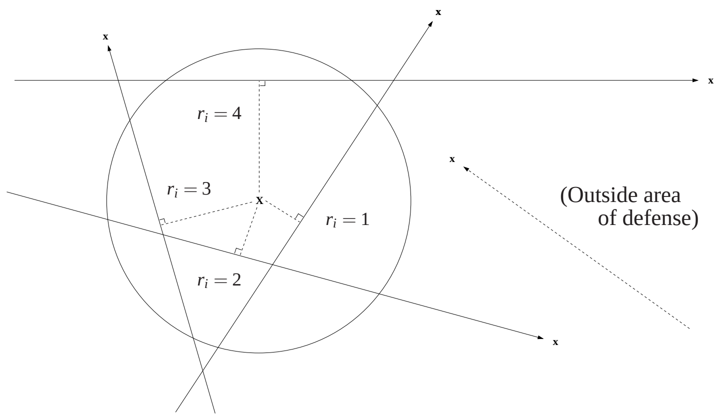

Fig. 6 A situation where multiple engagement paths intersect the area of defense for a target i.

Assume first that at most 4 different engagement paths might be engaged (i.e., $G_i = 4$), and that there are 5 residual defensive channels (i.e., $D_i = 5$). Consider the case where all four paths passing by target i are active (i.e., $B_i = 4$), that is, at least one resource is following each path. Under the given assumptions, all paths should be engaged and first each path gets one defensive channel locked against it. The remaining channel is assigned to the path closest to the target, which is path 1. The variables u_{isv} here take the values $u_{i1} = 2$, $u_{i2} = 1$, $u_{i3} = 1$ and $u_{i4} = 1$.

In the case that B_i or G_i decreases to 3, target i can only engage 3 engagement paths. For $B_i = 4$ and $G_i = 3$, the path most far away will no longer be engaged. The residual defensive channels are then distributed as follows: $u_{i1} = 2$, $u_{i2} = 2$, $u_{i3} = 1$ and $u_{i4} = 0$. If $B_i = 3$ and $G_i = 3$ (or 4), then only three engagement paths are active. Depending on which path that is not active, the other paths are assigned defensive channels like before, with respect to rank. Assume that for example path 2 is not active, in which case we get: $u_{i1} = 2$, $u_{i2} = 0$, $u_{i3} = 2$ and $u_{i4} = 1$.

Finally, if $B_i < 2$, all defensive channels cannot be assigned to an engagement path, since $F_i = 3$. With only one (or none) active path, at most $B_i \cdot F_i \leq 1 \cdot 3 = 3$ channels could be assigned. For example, if only path 3 is active, we obtain: $u_{i1} = 0$, $u_{i2} = 0$, $u_{i3} = 3$ and $u_{i4} = 0$.

5 Heuristic Solution Methods

A problem like this, with only a few constraints (one attack per target and shared resources) and a non-convex objective function, is well suited for meta-heuristics. Throughout this section, we base our work on the following assumptions:

1. The number of available resources is limited, that is, it is not possible to use the maximal number of resources against every target.
2. It is optimal to use all available resources.

The first assumption is reasonable, since otherwise the problem is reduced to choosing between tactics 3 and 5, either assigning all resources on the same path or splitting them on three different paths. (One would however still need to figure out the optimal combination of tactics and angle of attack for each target, and this would be a non-trivial problem.) The second assumption is very reasonable and simplifies the work of defining neighbourhoods and setting up heuristic schemes.

5.1 Local Search

Given a feasible solution to the generic model, $\mathbf{z}$, found by some heuristic scheme, one could try to improve it locally, that is, to perform a local search.

For this problem, where a solution $\mathbf{z}$ states which tactic t and angle w to be used for each target s, it is straightforward to test all feasible angles $w \in W_{st}$ for the assigned tactic t, one target at a time, and save the best improvement (if any). Then, if an improvement is made, one can repeat the same process again (since one target is now attacked from a different angle, and further improvements might be possible) until the process converges.

At the same time as one tests all angles, one can also switch between the tactics that use the same number of resources, hence conserving the overall usage of resources (assumed to be at its upper limit).

A limitation of this local search procedure is that the allocation of resources to targets is never changed. Even so, this procedure has proven to be an effective tool for finding good solutions, for almost any starting solution, as long as the allocation of resources to targets is close to the optimal one.

5.2 A Constructive Heuristic

An intuitive strategy is to iteratively augment a partial solution, adding one extra resource in each iteration. It seems plausible that the optimal solution using, say, 8 resources is close to the optimal solution for 7 resources.

Provided a feasible solution using $k \geq 0$ resources, denoted $\mathbf{z}_k$, we seek a solution $\mathbf{z}_{k+1}$. This is done by considering one target at a time, adding one resource if not $K = 3$ resources are already in use for this target, and then performing a local search. The best such augmentation, over all targets, is saved and returned as the new solution $\mathbf{z}_{k+1}$. The augmentation with one resource at a time is repeated until the available number of resources is reached. The cost of the heuristic will increase with respect to the number of targets, since it performs one local search per target.

Note that this constructive heuristic can be applied to any feasible starting solution. Further, if the initial solution is near-optimal for k resources, then it is likely that the augmented solution is also near-optimal, but now for $k+1$ resources.

As a bonus, this approach will generate Pareto-like solutions, stating the expected outcome of an attack for different numbers of resources, which also yields marginal values for additional resources with respect to the expected outcome. This information is useful when choosing the number of resources to use for an attack. As will be seen in the forthcoming results, the gain in expected outcome of an additional resource decreases as a function of the number of resources already in use.

5.3 Simulated Annealing

A popular meta-heuristic, which is easy to implement, is simulated annealing. The basic idea, which makes it a meta-heuristic and not a local search method, is to accept solutions which are non-improving in order to escape local optima. This is done by chance, and the probability to accept a non-improving value is related to the change in objective value from the current solution to the new one.

Also, in order to assure finding a local optimum, the probability of accepting worse solutions decreases over time. This is done by a temperature parameter, which decreases as the iterations goes by. A simulated annealing approach is successfully used for a weapon-target allocation problem in [1].

In order to apply a simulated annealing approach, we need to define a neighbourhood for a solution $\mathbf{z}$. Under the assumptions stated above, all we need is to work with feasible attack plans $\mathbf{z}$ that use all available resources. Hence we define five neighbourhoods of an attack plan $\mathbf{z}$, denoted $N_k(\mathbf{z})$, in the following ways.

1. The angle of attack w is changed for one target s and tactic t in the attack plan, that is, $z_{stw} \rightarrow z_{st\bar{w}}$.
2. The tactic against one target is changed by switching between one angle and multiple angles, that is, $z_{stw} \rightarrow z_{s\bar{t}\tilde{w}}$. If necessary, the reference angle w is adjusted. For example, instead of two resources attacking from the same angle, they attack from different angles. Notice that the number of resources involved in the attack is still the same though.
3. Pick two targets at random and switch their tactics and angle of attack. For example, variables $z_{s_1t_1w_1}$ and $z_{s_2t_2w_2}$ become $z_{s_1t_2w_2}$ and $z_{s_2t_1w_1}$ instead.
4. Pick two targets at random and exchange their angle of attack. For example, variables $z_{s_1t_1w_1}$ and $z_{s_2t_2w_2}$ become $z_{s_1t_1w_2}$ and $z_{s_2t_2w_1}$ instead.
5. Pick two targets at random, which do not use the same number of resources, and change to new tactics which increase/decrease the number of resources used

respectively. For example, one target is changed to be attacked by two resources instead of one, while another target is attacked by two resources instead of three.

The use of multiple neighbourhoods provides diversity to the search, and by repeatedly changing between them all feasible solutions can be reached. Notice that neighbourhood N_5 is crucial, since without it the number of resources allocated against each target would remain fixed to that of the initial solution throughout the search.

The implemented simulated annealing heuristic consists of outer and inner iterations. At the end of each outer iteration the temperature is decreased (from the initial temperature 0.9 and with the cooling factor 0.7). In each outer iteration, we cycle once through the different neighbourhoods, according to the sequence $\{5,2,1,3,4,1,5,2,1\}$. For each of these, we perform 100 evaluations of neighbours. During the search, we keep track of the overall best found solution.

6 Numerical Experiments

The optimistic and pessimistic models presented in Sections 3.2 and 3.3, respectively, are easily solved using a linear integer programming solver, in our case CPLEX. They provide upper and lower bounds on the true optimal value, and these are found in fractions of a second. In order to improve the lower bound, the pessimistic solution provided by the solver is simply evaluated using the true objective function. This step improves the bound significantly and is also instant. Moreover, if a local search, as described above, is performed from the pessimistic solution, an even better solution can be found. This is fairly inexpensive and improves the bound in most cases.

The constructive heuristic is initiated with the locally improved pessimistic solution obtained for $k = 2$ resources. It is then applied to find a solution with the available number of resources. The procedure should generate near-optimal solutions to the cost of at most one application of the local search procedure for each target and each new resource.

The simulated annealing method is applied as described in Section 5.3. This is a fairly time-consuming method, but is likely to produce the solutions of best quality.

6.1 Case 105

The test case, called Case 105, includes 2 targets with defense and 5 other targets, which are positioned as shown in Figure 7. One unit step in the picture corresponds to 1 km. The targets with defense are positioned 10 km apart, and each of them has a defensive radius of 10 km. The distances between the targets are 300–500 meters. When modelling the problem, a coarse angle discretization of 12 angles is used.

We define three different reward settings for the targets. In setting I, $r_s = 0$ for $s \in \overline{S}$ and $r_s = 1$ for $s \in S \setminus \overline{S}$, that is, there is no reward for the defenders and the

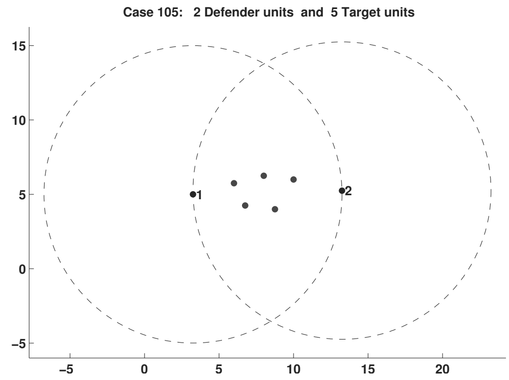

Fig. 7 Test case 105, with 2 defenders and 5 other targets.

same reward for every other target. Although this setting does not reward the targets with defense, it might still be optimal to attack the defenders in order to reduce their defensive capabilities and thus increase the overall reward of the attack. In reward setting II, $r_s = 1$ for $s \in \bar{S}$ and $r_s = 2$ for $s \in S \setminus \bar{S}$, so that the defenders are also considered valuable but only second to the other targets. In setting III, $r_s = 1$ for $s \in \bar{S}$ and $r_s = 5$ for $s \in S \setminus \bar{S}$, which differentiates the two types of targets more. Below, we present and analyze the result for the different reward settings.

6.2 Results for Case 105

In Figure 8 we see a graphical representation of the best found solution for Case 105 with reward setting III and 14 resources available. Both defenders, numbers 1 and 2, are attacked by tactic 5 which means 3 resources from different directions. Targets 5 and 6 are attacked using tactic 4, where 2 resources attack from opposite directions. Target 3 is attacked using tactic 2, that is, 2 resources from the same direction, indicated by the dashed line. Finally, targets 4 and 7 are attacked by single resources.

The solutions are not always intuitive at first glance. For example, one of the attack paths toward target 1 intersects the defensive area of target 2 for a long distance, and vice versa. Is it not better to attack with all 3 resources from the same angle and avoid the defense of the other defenders? The explanation is logical. Consider the resource attacking defender 1. By travelling inside the defensive area of defender 2, some of the defender's defensive capability will be allocated against this resource. As one of three resources taking part of the attack against target 1, the total expected probability of success will be quite high even though this specific resource faces great danger. In this way, the defensive capabilities available for target 1 to use against other resources are reduced, and the overall expected outcome will gain.

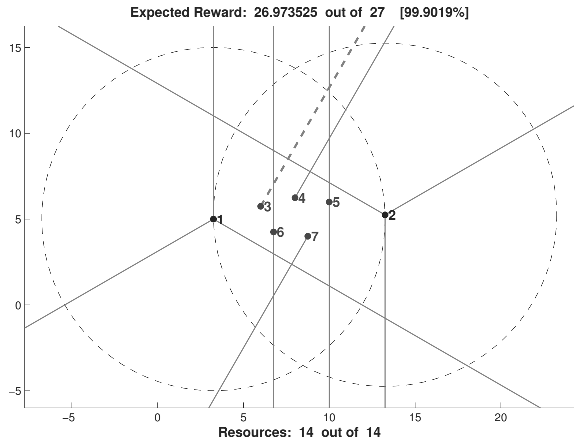

Fig. 8 Test case 105, with 2 defenders and 5 other targets. Best solution for 14 resources.

Figure 9 shows a graphical representation of the best found solution for the same case but with 17 resources available. The objective value is improved somewhat.

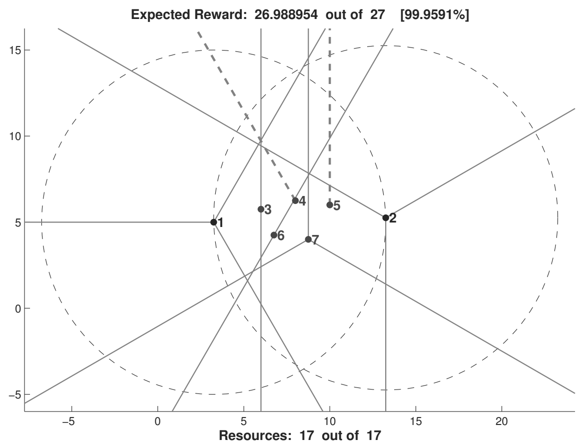

Fig. 9 Test case 105, with 2 defenders and 5 other targets. Best solution for 17 resources.

The use of reward setting I (i.e. reward 0 for defenders and reward 1 for other targets), render the result seen in Figure 10. The x-axis represents the number of resources available and the y-axis the corresponding objective values. The two outer dash-dotted lines represent the upper and lower bounds, respectively, found by

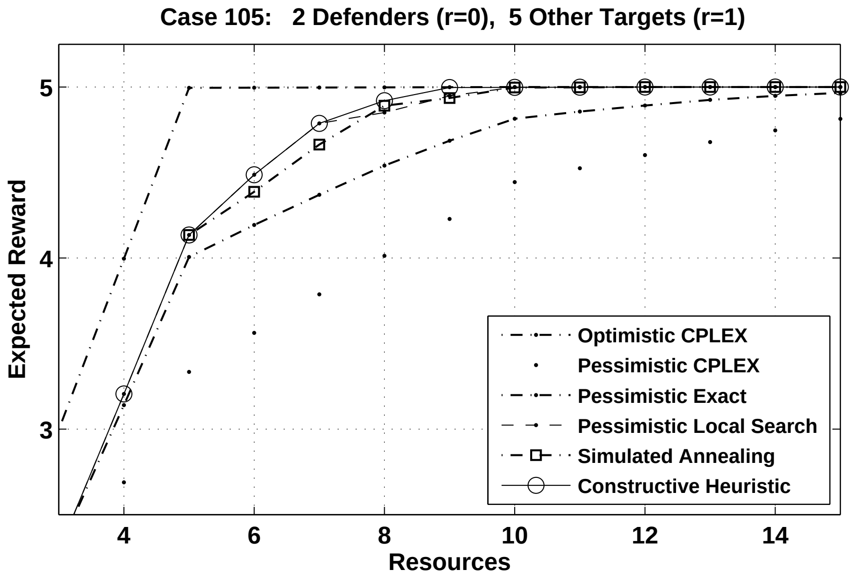

Fig. 10 Test case 105, with 2 defenders and 5 other targets.

CPLEX, where the pessimistic solutions have been evaluated using the true objective function. The single dots represent the pessimistic values given by CPLEX. The dashed line with dots is the locally improved pessimistic solutions. We can see that the improvement is substantial for most numbers of resouces. The dash-dotted line with squares shows the best found solutions from the simulated annealing heuristic. The solid line with circles shows the result of the constructive heuristic. These solutions are in general the best ones found, but sometimes simulated annealing solutions are equally good.

For reward settings II and III, a similar behaviour can be observed in Figure 11 and Figure 12, respectively. Obviously, the objective values differ due to the different reward settings, but the overall trend is the same.

We conclude this section with some remarks. The behaviour is very similar for the different reward settings. The optimistic and pessimistic bounds are not tight for 5–10 resources, but a local search from the pessimistic solution improves the situation. For Case 105, with only 7 targets, using more than around 15 resources is not very interesting, and, as can be seen in the graphs, the optimistic and pessimistic bounds are then tight.

The simulated annealing algorithm performs very well, and provides solutions comparable with the constructive heuristic approach, but it requires comparably long time even for a moderate number of resources. Mostly, the constructive heuristic finds the best found solution, and it is beaten by the simulated annealing method on only single occasions, but it requires even more time than the latter algorithm when considering many resources.

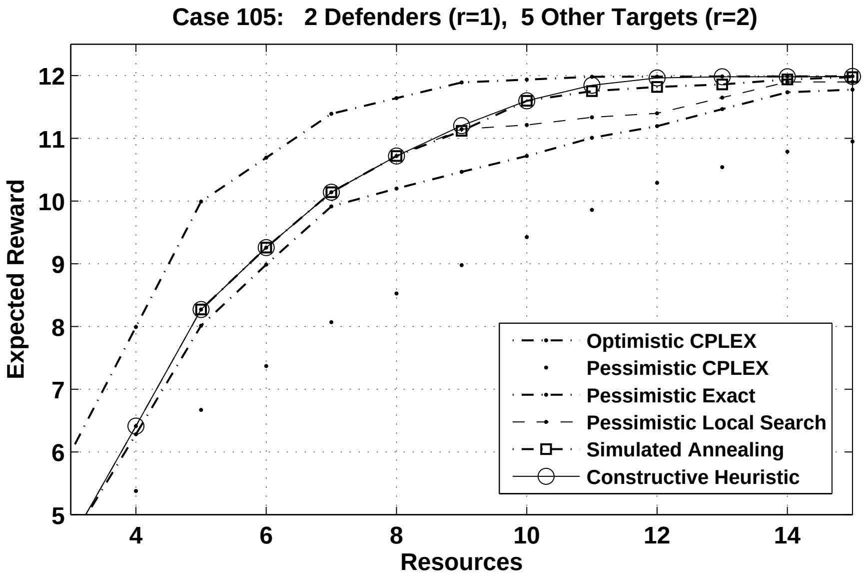

Fig. 11 Test case 105, with 2 defenders and 5 other targets.

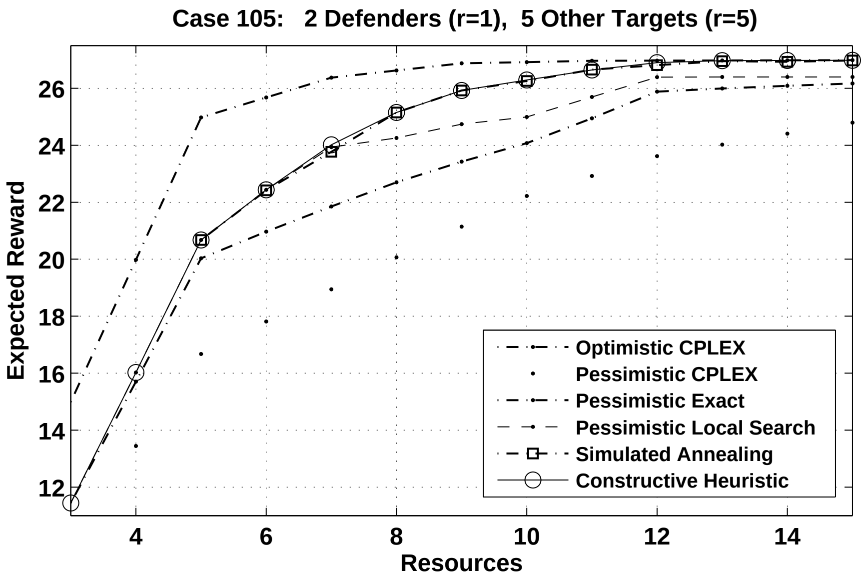

Fig. 12 Test case 105, with 2 defenders and 5 other targets.

6.3 Results for Larger Instances

In addition to Case 105, a number of different cases have been studied. These ranges from 7 to 21 targets. Case 105 is a good representative for all of them, with respect to the behaviour of the heuristic solution approaches. In Table 1, we give mean objective values for 12 different cases, with a varying number of resources. Here, the objective values are normalized with respect to the optimistic value found for each case.

Table 1 Normalized mean objective values for each method. Best values are in boldface and second best values are emphasized.

	Resources					
Method	5	10	15	20	25	30
Opt. CPLEX	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000
Pess. Exact	0.6702	0.7691	0.8399	0.8894	0.9209	0.9463
Pess. Local	*0.6893*	0.8300	0.9023	0.9522	0.9737	0.9834
Constr. Heur.	**0.6999**	**0.8599**	**0.9382**	**0.9852**	**0.9941**	**0.9988**
Sim. Ann.	0.6845	*0.8545*	*0.9369*	*0.9840*	*0.9918*	*0.9940*

For larger instances, with 10–20 targets, the quality of the optimistic and pessimistic bounds are not as good as for smaller instances. We suspect that the pessimistic bound is tight for up to 10 resources, and that the strength of the optimistic bound improves with an increasing number of resources. For instances where 10–20 resources are available, none of the bounds seems to be tight.

The constructive heuristic is the most stable of all solution methods, providing high quality solutions for all different scenarios and reward settings. The simulated annealing method is also very successful. Because of the long calculation times required for a single run of the simulated annealing method, it is only competitive with the constructive heuristic approach when seeking a single solution for a specific and quite large number of resources. Otherwise, the constructive heuristic provides both better calculation times and solution quality, with the important extra feature of providing a range of solutions, one for each number of resources. In all, the constructive heuristic is the clear winner.

7 Conclusions and Future Work

We have introduced and defined a mission planning problem. A generic mathematical model of the problem is presented, and the complex objective function is analyzed in detail. The generic model can be approximated in order to derive optimistic and pessimistic models. Such models are an important tool since they provide upper and lower bounds on the optimal value, hence limiting the uncertainty of the quality of solutions.

However, in order to solve problem instances of realistic sizes, it is necessary to use heuristic methods. We have proposed a constructive heuristic method and a simulated annealing heuristic to solve this difficult problem. The methods were tested on a set of problem instances, and the results are very promising. The constructive heuristic method has good solution times, while solution times are relatively long for the simulated annealing algorithm.

All methods are generic and can handle different scenarios for the defender's strategy. It is sufficient to provide a black-box function to call whenever the objective needs to be evaluated. Hence, if the assumptions in Section 4 are inadequate, or needs to be modified in any way, the given framework will still be applicable.

This paper has focused on the development of a planning system only considering target scene parameters such as target location and defense system description, and how the defense reacts upon attack. Resource performance is certainly included in the analysis but just in the sense of a static set-up of effect-on-target as a function of tactics, and the ability to survive in a surface-to-air defense system environment. This approach complies with future command and control doctrines which promote a separation of effect planning and resource allocation planning.

To extend the mission scope we can include planning aspects of the platform. Route planning can be conducted in a flexible way with its own objectives to conclude the overall mission success. Obvious aspects are minimizing radar cross section exposure during route phase, and minimize time to target, that is, to explore hiding possibilities or by clever surveillance tactics during the cruise phase. An obvious continuation from our work within this paper, is to investigate the coupling between route and effect planning. If this is solved properly, a large step is taken to control and comprise vital aspects of ground attack planning.

Further, firing platforms must not be given in advance, instead maximizing the effect of the target area can be the driver to find the best platforms from a larger set. Based on this fact, future work could address at least two obvious scenarios. The first is when the target scene is known and there is a predefined number of platforms where route planning is included in the overall mission. A second scenario is to consider when several platforms are available. In this case we must allocate good firing units from a set of platforms but also decide firing position and route planning.

References

1. Ciobanu, C., and Marin, G.: On Heuristic Optimization. Analele Stiintifice ale Universitatii Ovidius Constanta **9**, 17–30 (2001)
2. Kaminer, B.I., and Ben-Asher, J.Z.: A Methodology for Estimating and Optimizing Effectiveness of Non-Independent Layered Defense. Systems Engineering **13**, 119–129 (2010)

3. Kwon, O., Lee, K., Kang, D., and Park, S.: A Branch-and-Price Algorithm for a Targeting Problem. Naval Research Logistics **54**, 732–741 (2007)
4. Miercort, F.A., and Soland, R.M.: Optimal Allocation of Missiles against Area and Point Defenses. Operations Research **19**, 605–617 (1971)
5. Quttineh, N.-H.: Models and Methods for Costly Global Optimization and Military Decision Support Systems. Linköping Studies in Science and Technology, Dissertations, No. 1450, Linköping University (2012)
6. Yost, K.A., and Washburn, A.R.: Optimizing Assignment of Air-to-Ground Assets and BDA Sensors. Military Operations Research **5**, 77–91 (2000)